



# Developing the Fringe Routing Protocol

by  
**Deb Shepherd**

a thesis submitted to Victoria University of Wellington  
in fulfilment of the requirements for the degree of  
Masters of Engineering  
in  
Network Engineering

2011

extending the work begun by  
**Don Stokes**  
Knossos Networks Ltd

supervised by  
**Andy Linton and Ian Welch**  
Victoria University of Wellington

---

for ian  
an engineer of the old-fashioned kind  
wish you could read this

with thanks to  
Andy Linton  
*Victoria University of Wellington*  
for your light hand of supervision  
DonStokes  
*Knossos Networks*  
Ian Welch  
*Victoria University of Wellington*  
Peter Komisarczuk  
formerly of *Victoria University of Wellington*  
now at *Thames Valley University*  
John Rumsey  
*Knossos Networks*  
and  
Stephen Marshall  
*husband extraordinaire*  
for much more than can be listed here

---

## Table of Contents

---

Abstract.....	5
<b>Chapter One</b>	
Introduction .....	7
<b>Chapter Two</b>	
Background .....	11
Reference implementation.....	12
The Quagga Software Routing Suite.....	12
Tools .....	14
FRPsniffer.....	15
Resources and documentation .....	16
<b>Chapter Three</b>	
Specifying the Fringe Routing Protocol .....	19
The goals of the Fringe Routing Protocol.....	20
FRP gateways .....	21
FRP route forwarding.....	24
FRP router configuration .....	29
FRP message formats .....	32
FRP packet exchange .....	40
<b>Chapter Four</b>	
Designing the Quagga FRP daemon .....	47
Quagga Architecture .....	48
Initial process .....	51
Code structure.....	53
Program flow .....	57

## Chapter Five

Implementing and testing the Quagga FRP daemon .....	61
Daemon start up .....	61
Peers.....	66
Events .....	69
Polling.....	70
Triggered updates .....	76
Incoming packets.....	77
Outgoing packets .....	94
Testing .....	102

## Chapter Six

Conclusions and future work .....	105
Bibliography .....	107
Appendix A.....	113
Initial Fringe Routing Protocol specification from Don Stokes .....	113
Appendix B.....	119
Slides from the FRP presentation at the NZNOG conference .....	119
Appendix C.....	123
Changes made to the Quagga code to hook in a new daemon.....	123
Appendix D.....	127
The Quagga FRP daemon implementation code.....	127
The FRPsniffer implementation .....	129
The original implementation of FRP by Don Stokes.....	131

---

# Developing the Fringe Routing Protocol

---

## Abstract

---

An ISP style network often has a particular traffic pattern not typically seen in other networks and which is a direct result of the ISP's purpose, to connect internal clients with a high speed external link. Such a network is likely to consist of a backbone with the clients on one 'side' and one or more external links on the other. Most traffic on the network moves between an internal client and the external world via the backbone.

But what about traffic between two clients of the ISP? Typical routing protocols will find the 'best' path between the two gateway routers at the edge of the client stub networks. As these routers connect the stubs to the ISP core, this route should be entirely within the ISP network. Ideally, from the ISP point of view, this traffic will go up to the backbone and down again but it is possible that it may find another route along a redundant backup path.

Don Stokes of Knossos Networks has developed a protocol to sit on the client fringes of this ISP style of network. It is based on the distance vector algorithm and is intended to be subordinate to the existing interior gateway protocol running on the ISPs backbone. It manipulates the route cost calculation so that paths towards the backbone become very cheap and paths away from the backbone become expensive. This forces traffic in the preferred direction unless the backup path 'shortcut' is very attractive or the backbone link has disappeared.

It is the analysis and development of the fringe routing protocol that forms the content of this ME thesis.



---

# 1. Introduction

---

The basic underlying architecture and protocols of the Internet have been developed and successfully used for several decades. As the system evolves, it is inevitable that new ideas and concepts emerge requiring new ways of doing things, and that better ways of handling existing concepts appear. Such improvements are typically created in order to streamline or change the behaviour of one small part of the larger system, often in fairly narrow circumstances.

This project examines a new protocol that does exactly that. The Fringe Routing Protocol (FRP) is designed by Don Stokes of Knossos Networks ([knossos.net.nz](http://knossos.net.nz)), a commercial ISP in downtown Wellington, New Zealand. Don devised FRP to solve a particular issue common to ISP style networks — an ISP tends to want traffic between it's connected client stub networks to move via the core backbone rather than via any alternative back routes added in to the network for reliability and redundancy.

An ISP designs their network primarily to connect their customers to the outside world. The core of the network is a highspeed, high capacity backbone capable of delivering traffic for the ISP itself and for its customers, quickly and efficiently. The outer fringe will have upstream (external) gateway connections to multiple layer-2 providers for redundancy and robustness. The inner fringe holds a multitude of internal client gateways and routers connecting layer-3 client stub networks to the backbone.




Figure 1.1  
 showing a simplified ISP network  
 ISP infrastructure is blue, client stub networks are grey  
 default routes are solid lines, backup routes are dotted  
 external gateway are labeled xG, internal (client) gateways are labeled iG,  
 routers are labeled R, client stub networks are labeled A–F

The nature of an ISP network — to provide connectivity to customers — means that the main flow of customer generated traffic through the network core is either from the outside world to the client's network, or from the client to the outside world.

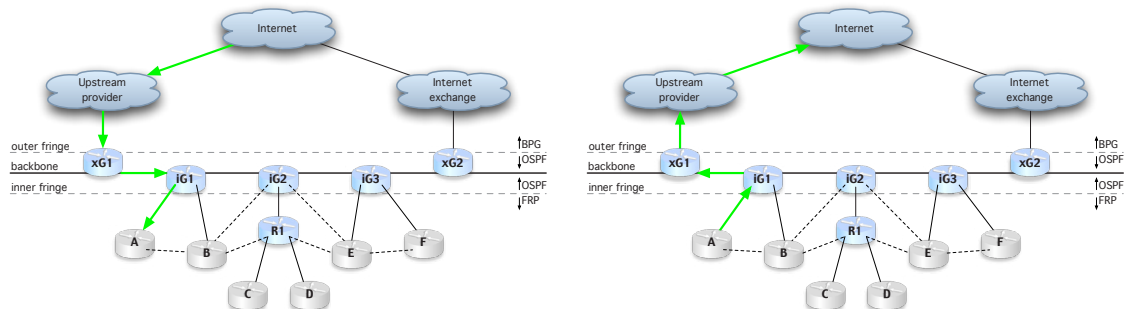


Figure 1.2  
 the main flow of client based network traffic through an ISP network

Considerably less traffic is likely to be generated between individual ISP customers and the traffic that is generated stays entirely within the ISP's network. This means that the ISP is able to impose some level of control over the routing of that traffic. The left hand diagram in figure 1.3 shows the preferred route for routing traffic between customer A and customer B, but A and B are also connected by a backup link, allowing traffic to take that path if it chooses.



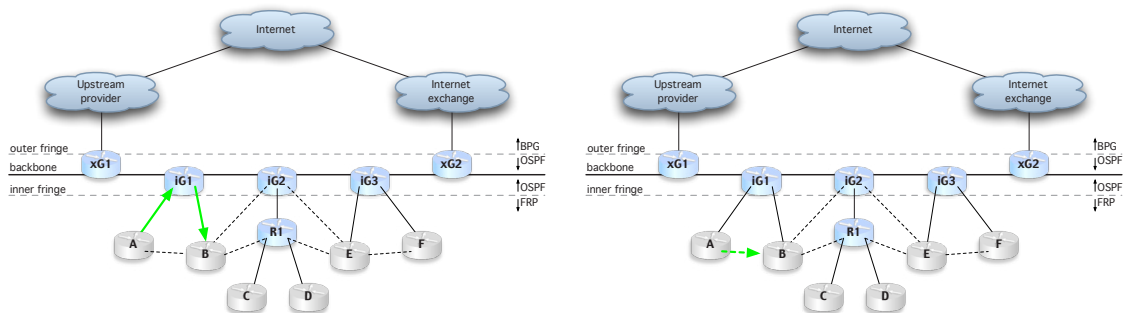


Figure 1.3  
traffic between clients of the same ISP

This is the scenario that FRP is designed to address, a lightweight solution to routing traffic at the inner fringe of the ISP network and an alternative to the commonly used but more heavyweight protocols such as the Border Gateway Protocol (BGP) or Open Shortest Path First (OSPF).

An initial implementation of the fringe routing protocol has been developed at Knossos Networks as a proof of concept. That version was written, from scratch, specifically for the Knossos network and is currently running successfully.

This ME project involved taking FRP and creating a new, independent, implementation to complement the first, this time in a completely different development environment. Of the number of possible choices, the *Quagga software routing suite* was mandated. The process was to take the initial version as a reference and reverse engineer it with the aid of the accompanying notes. A specification of the protocol was then developed, which became the basis for the development of the *Quagga FRP* routing daemon.

In order to create a new *Quagga* protocol daemon, the intricate complexities of the *Quagga* suite need to be understood. It was necessary to also reverse engineer a number of the other *Quagga* daemons while building the new one as this process is not well documented.



---

## 2. Background

---

The project of developing a working version of the fringe routing protocol starts with the provided reference implementation and accompanying notes (see Appendix A). The brief set of notes provided is sparse but concise, the barest beginnings of a specification, and although padded out with a couple of face-to-face sessions, initial lack of knowledge limited the usefulness of these conferences.

The notes are essentially complete in that they address all the significant architectural components of the protocol and contain all the necessary detail. But understanding needs to be teased out, there being no explanation as to why elements are necessary or why certain factors need to be handled in a certain way. It is very much an architectural description of intent rather than an engineering framework for implementation. Included are a brief list of goals, the basics of the route forwarding algorithm, a short explanation on the importance of gateways, the message formats, a concise description of the batch processing, and an example of how the sequence numbering works.

So the initial part of the project became the process of understanding the sources provided, and an attempt at a first draft of a specification was written. The draft specification was created by reverse engineering the prototype to the point where the new version could be created from the specification alone, using it as a buffer between the code of the two separate implementations.




Figure 2.1  
*The ideal FRP development process*

## Reference implementation

---

The Knossos Networks reference implementation (see Appendix B) is a stand-alone system, combining the implementation of a basic, single threaded, Unix based router with the implementation of the fringe routing protocol. The code generally tries to keep these two functions separate but sometimes the two do become combined. At times the notes and the code do not match and the implementation significantly expands on, or even introduces new functionality.

The first step to understanding the workings of the system was to deconstruct it and fit it to the description of the protocol provided. The code is spread across two header and three code files. The main program loop plus supporting functions resides in a file named `main.c`, which is in turn supported by `route.c` and `parse.c`.

One header contains all defines, variables and data structures for the router in general. The other specifies those needed explicitly for the fringe routing protocol making it an invaluable resource. It was used extensively to define and build the required message and packet formats and many of the data structures. The main program loop is not multi-threaded so everything happens in a linear sequence, repeated indefinitely. This is one of the major differences between the implementations as the *Quagga* version is multi-threaded.

The essence of the FRP decision algorithm is in `route.c`. This only contains two functions, one to add an IP route to the routing table and another labelled as the ‘main routing engine’ [58]. `parse.c` handles numerous additional functions including adding peers, handling addresses, handling access control lists and parsing the configuration file.

## The Quagga Software Routing Suite

---

One of the major outcomes of the FRP project is the development of a fringe routing protocol routing daemon using the *Quagga software routing suite*. *Quagga* is a popular open source alternative to *Cisco* or *Juniper* routers. The major difference, apart from price, is that *Cisco* and *Juniper* produce highly optimised dedicated hardware routers whereas *Quagga* is a software suite that allows a standard Unix based desktop machine to operate as a router, using a very similar command structure to *Cisco* routers. As the ability to hook an entirely new routing protocol seamlessly into the inner workings of the router

was required, the open source route was an obvious choice — access to *Cisco* or *Juniper* at that level is virtually impossible.

The selection of *Quagga* from among the available choices was mandated. This springs from the fact that not only do *Knossos Networks* use *Quagga* in house, the Wellington Internet exchange (WIX) and *CityLink*, the company that run it, utilise *Quagga* extensively on small routers at the edge of their network.



*The quagga (Equus quagga quagga) is an extinct subspecies of the plains zebra*  
— Wikipedia (<http://en.wikipedia.org/wiki/Quagga>)  
image is in the public domain

*Quagga* is (a very much alive) fork of the older (and now essentially extinct) GNU Zebra project. Zebra began in 1996 under the auspices of Kunihiro Ishiguro [<http://www.zebra.org/history.html>] with the intention of creating a modular, TCP/IP based software routing engine made freely available under the GNU General Public License [<http://www.zebra.org/what.html>]. Zebra's modularity is achieved by creating a distinct daemon for each of the supported protocols — RIPv1, RIPv2 and OSPFv2 and BGP-4. This allows individual protocols to be modified separately, taking down the affected process while leaving the rest of the system on line. Likewise, the failure of any one daemon will only take out that individual process [<http://www.zebra.org/features.html>].

In 2003, *Zebra* forked into two separate projects, the original *Zebra* and a new project called *Quagga* [<http://www.quagga.net/news2.php?y=2003&m=8>]. The last *Zebra* release was zebra-0.95a in late 2005 [<http://www.zebra.org/index.html>]. As of early 2011, the latest *Quagga* release is quagga-0.99.17 [[www.quagga.net](http://www.quagga.net)].

*Quagga* is written in the C programming language and is firmly based on the *Zebra* model. It currently supports the BGP-4 (rfc1657, 1771, 1965, 1997, 2545, 2796, 2842, 2858), ISIS, OSPFv2 & v3 (rfc1850, 2328, 2370, 3101), OSPF6 (rfc2740), RIPv1 & v2 (rfc1058, 1724, 2082, 2453), and RIPv6 (rfc2080) routing protocols [<http://www.quagga.net/docs/docs-multi/Supported-RFC.html>]. *Quagga* is available for various Unix and Linux operating systems including FreeBSD, NetBSD, Debian, Ubuntu, Gentoo, and MacOSX. [[www.quagga.net/about.php](http://www.quagga.net/about.php)]

Somewhat surprisingly given that it is an open source project, *Quagga* proved to be a stable software platform with a well written development environment. The area in which it failed to deliver was the one which became a defining factor for this project. *Quagga* completely lacks clear, useful documentation. While the code is, on the whole, well written and well commented, it is a large complex system which is extremely difficult to understand sufficiently well to start to develop within it.

### **FRP in Quagga**

Implementation started by experimenting with turning a copy of an existing *Quagga* protocol daemon into something that pretended to be a FRP daemon. Although this was useful, the next step was to start again and build an empty shell of a daemon that worked correctly inside *Quagga* but did nothing beyond that. The FRP implementation was then added to the shell.

As already mentioned, the intention was to develop the daemon entirely from the newly written specification but this proved to be impossible and, eventually, reference to the code of the prototype version became necessary in order to answer the many questions that arose. However, only one small piece of the prototype code was included in the *Quagga* daemon, the rest is written from scratch.

## **Tools**

---

This project was conceived and developed entirely in a Unix operating system environment and ideally FRP should run on any ‘flavour’ of Unix. In the development of the FRP protocol to this point in time, a number of different Unix platforms have been used. The original implementation was developed and currently runs on FreeBSD. In the scope of the current project, this original version has been ported to NetBSD and to MacOSX, simply for convenience.

One of the problems with understanding and developing in a computer network environment is the fact that it is not possible to actually see the data that is travelling through the network. Tools are required to look at different parts of the environment at different points in time to see what is actually happening. Unix supplies a good range of useful command line tools including *netstat*, *route*, *ip*, and *ifconfig*, which are good for looking up information on sockets, interfaces, addresses, protocols and the kernel routing table. *ping* is useful for checking that machines acting as routers are up and visible and that traffic is actually routing through the network. *tcpdump* [<http://www.tcpdump.org>] captures network traffic over a period of time, printing out useful information and providing the ability to trace packets and analyse the data they contain.

*WireShark* [<http://www.wireshark.org>] is similar to *tcpdump* but comes as a standalone package with a graphical user interface and the ability to visually drill down through the contents of a packet. *WireShark* helpfully sections up each packet into its component headers and payload. It also makes traffic filtering easier than in *tcpdump*. *WireShark* is a cross platform, Unix based open source application using *tcpdump*'s *pcap* library to handle packet capture.

---

## FRPsniffer

---

It became clear however, that it would be useful to have a packet sniffer tool that collected only FRP packets and knew how to correctly extract data from them. From this idea, *FRPsniffer* was designed and implemented — which also provided a useful first attempt at understanding and working with the FRP message and packet structure. Although *tcpdump* and especially *WireShark* are useful tools, this dedicated FRP traffic capture tool became more useful still. It was extensively used to watch FRP behaviour and the traces shown later in this document all come from *FRPsniffer*.

*FRPsniffer* is written in C and compiled with the MacOSX Xcode compiler. As this is essentially a Unix GNU Compiler Collection (gcc) compiler, the code is portable to any Unix based environment. The sniffer is written around the same library as *tcpdump* and *WireShark*, *tcpdump*'s *libcap* [<http://www.tcpdump.org>]. This is an open source C/C++ library providing a packet capture application programming interface (API). *FRPsniffer* is hardcoded to promiscuously collect all ethernet based UDP traffic on FRP's port and to display it broken down into FRP's packet and message structure.

## Resources and documentation

---

To continue the theme of paucity in documentation, the standard technique of conducting a literature search for academic, peer reviewed papers proved fruitless. The following is a brief look at the sources that proved to be the most useful and the most regularly consulted during the completion of this project. For a complete list, refer to the bibliography on page 107.

### General background

In order to even start on this project, a good general knowledge of computer networks and their architecture needed to be moved on to the next level and the inevitable gaps needed to be plugged. A couple of general networking texts

Kurose, J. F. and Ross, K. 2002 Computer Networking: a Top-Down Approach  
Featuring the Internet. 2nd. Addison-Wesley Longman Publishing Co., Inc.  
Tanenbaum, A. S. 1985 Computer Networks. Prentice Hall PTR.

provided the starting point, followed by several more detailed works including

Perlman, R. 2000 Interconnections (2nd Ed.): Bridges, Routers, Switches, and  
Internetworking Protocols. Addison-Wesley Longman Publishing Co., Inc.  
Stallings, W. 1993 Data and Computer Communications. 4th. Prentice Hall PTR.

### Routing

The purpose of this project is the development of a routing protocol. A number of books provide good, general background reading on the topic of routing but the most useful was

Parkhurst, W. 2004 Routing First-Step. Cisco Press.

which provided a solid base to build on. The other useful volumes are

Beijnum, P. I. and Beijnum, I. V. 2002 BGP. O'Reilly & Associates, Inc.  
Malhotra, Ravi. 2002. IP routing. O'Reilly Media, Inc.  
Medhi, D. and Ramasamy, K. 2007 Network Routing: Algorithms, Protocols, and  
Architectures. Morgan Kaufmann Publishers Inc  
Moy, J. T. 1998 OSPF: Anatomy of an Internet Routing Protocol. Addison-Wesley  
Longman Publishing Co., Inc.

The definitive source of information on routing protocols and the topics surrounding them is the Internet Engineering Task Force (IETF) Request for Comments (RFC) series.

Carpenter, B., Ed., "Architectural Principles of the Internet", RFC 1958, June 1996.

Bush, R. and D. Meyer, "Some Internet Architectural Guidelines and Philosophy",  
RFC 3439, December 2002.

provide general conceptual discussion on Internet design and Architecture, while

Hedrick, C., "Routing Information Protocol", RFC 1058, June 1988.



gives a very full and useful description of the distance vector algorithm and realities of implementing it, and

Fuller, V. and T. Li, “Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan”, BCP 122, RFC 4632, August 2006.

contains critical information on current IPv4 address space usage.

Bradner, S., “The Internet Standards Process – Revision 3”, BCP 9, RFC 2026, October 1996.

Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels”, BCP 14, RFC 2119, March 1997.

Postel, J. and J. Reynolds, “Instructions to RFC Authors”, RFC 2223, October 1997.

has useful information on how to write and submit a RFC or any other documentation and, although obsolete,

Hinden, R., “Internet Engineering Task Force Internet Routing Protocol Standardization Criteria”, RFC 1264, October 1991.

is interesting in that it still has some useful insight into what should be in a RFC for a routing protocol.

## FRP

Specific information on the FRP protocol was obtained from the supplied specification notes and the code of the original implementation, both written by Don Stokes of Knossos Networks. Additional information came from talking to Don and from listening to him present the FRP protocol at the New Zealand Network Operators Conference in Wellington, February 2011 (see Appendix B).

## Unix networking

The creation of a routing daemon for a Unix environment required a good working knowledge of the Unix socket API and how it works. Two sources proved to be invaluable. The book

W. R. Stevens *et al*, *Unix network programming: the sockets networking API*, 3<sup>rd</sup> ed. Boston, MA: Addison–Wesley, 2004.

provided a good, in depth recap of the topic and filled in many gaps but it was the online resource

B. J. Hall. (2009). *Beej’s guide to network programming: using Internet sockets (version 3.0.14)*. [Online]. Available: <http://beej.us/guide/bgnet/>

that proved the most valuable as it provided succinct descriptions of the different system calls, functions and structures in a logical manner that made it a very easy to use reference manual.

## Quagga

The Quagga software routing suite package is complex and difficult to come to terms with. The Quagga package comes with a manual

K. Ishiguro *et al.* (2006). *Quagga: a routing software package for TCP/IP networks (Quagga 0.99.4)*. [Online]. Available: <http://www.quagga.net/docs.php> which concentrates on installing, building and configuring Quagga plus providing a list of user commands. It also contains useful (if basic) sections on Quagga's architecture, supported RFCs and the Zebra protocol.

When it came to the Quagga daemon itself, the best resource was the actual program code Available: <http://www.quagga.net/download/>

and the comments contain within it. The comments are generally adequate, although often brief and sometimes incomplete Two additional documents, although unfinished and very dated, provided invaluable background material and advice. The first,

Y. Uriarte. (2001). *Zebra for dummies*. [Online]. Available: <http://www.quagga.net/zh.html>

proved particularly useful as it provides a 'how to' guide to creating a new Quagga routing daemon from scratch, including information on mandatory inclusions that is not available from any other source. The second,

Pilot. *Zebra hacking notes: for GNU Zebra 0.93b (rev 1.6)*. [Online]. Available: <http://quagga.net/faq/zebra-hacking-guide.txt>

is more a list of notes, not as helpful as the Dummies guide but useful as an addendum to it.

## Network packet sniffing

When writing the FRPsniffer support tool to collect and display traffic and packet contents in order to watch FRP actually running, the following article was very useful

L. M. Garcia, "Programming with libpcap: sniffing the network from our own application," in *Hackin9*, vol. 3, no. 2, pp. 38-46, 2008.

as was this documentation from the *tcpdump and libpcap* website

T Carstens. *Programming with pcap* [Online]. Available: <http://www.tcpdump.org/pcap.html>

---

### 3. Specifying the Fringe Routing Protocol

---

This chapter defines and specifies the fringe routing protocol. It looks at the goals it is expected to achieve, and the network environment it is designed to work in. The components of the protocol are specified: the importance of gateways; route forwarding; router configuration; message formats; and packet exchange.

The fringe routing protocol aims to send client to client traffic in an ISP style network up to the backbone and down again rather than across secondary backup links — that is from A to iG1 to B, rather than directly from A to B. To understand FRP however, it is necessary to move from looking at traffic delivery in the forwarding plane — how actual traffic is directed around the network — to studying the advertisement of preferred, or ‘best’ routes between routers in the control plane.




Figure 3.1  
showing a simplified ISP network  
ISP infrastructure is blue, client stub networks are grey  
default routes are solid lines, backup routes are dotted  
external gateway are labeled xG, internal (client) gateways are labeled iG,  
routers are labeled R, client stub networks are labeled A-F

A primary concept of FRP is that of gateways. A FRP gateway is one that connects a customer stub network to the ISP backbone, so iG1, iG2 and iG3 are all potential FRP gateways. Each FRP router must have a single designated gateway and prefers to only advertise routes it has discovered if they lead towards its gateway. The only exception is if a link has gone down or if an alternative route is extremely cheap. The result of this carried through to the forwarding plane is that traffic is generally delivered up to the backbone, via the FRP gateway, and back down again.

## **The goals of the Fringe Routing Protocol**

---

The primary goal of the fringe routing protocol is to simplify the management of an ISP style public network. The protocol achieves this by having the following characteristics:

1. Traffic is routed, by preference, towards the backbone via a designated gateway.  
Traffic is only routed away from the backbone in exceptional circumstances.
2. FRP is intended to be subordinate to the existing IGP, which is still used at the network core. FRP supplies reduced routing complexity at the network's inner fringe.
3. FRP attains reduced complexity by operating as a lightweight protocol. It achieves this by:
  - a. remaining quiescent until an FRP gateway route to the backbone is established;
  - b. and therefore only updating the kernel routing table or advertising routes if there is a path to a gateway;
  - c. keeping routing table sizes to a minimum;
  - d. keeping traffic down to keepalives if there are no route updates to be exchanged.
4. FRP must work on public access networks and so therefore must be secure. This security is based on a configured peer relationship.

## FRP gateways

A key part of FRP is the use of gateways. FRP is designed to forward traffic up on to the ISP backbone for delivery — to destinations beyond the ISP network, to the ISP itself, or to another client of the ISP. FRP facilitates this delivery by utilising FRP designated gateways sitting on the backbone.

Every router using the fringe routing protocol needs to discover the ‘best’ path to one of the FRP backbone gateways. This becomes the router’s designated gateway. It shares knowledge of this gateway, including the path to it, with its peers using special gateway messages. In the context of FRP, a peer is any directly connected router designated as such in the host’s configuration.

### Establishing gateway routes

FRP specifies that if a router has no gateway, then it is quiescent. However routers are only told if they themselves are a gateway, not what their ‘nearest’ or ‘best’ gateway is. Therefore any router that is designated as a FRP gateway needs to tell its peers that they can use it as such. This allows any connected peer on the FRP network to eventually compile a path to gateway for themselves. The newly discovered gateway and its path are then shared with peers and the propagation continues.




Figure 3.2  
A and B gateway via G1  
C and D gateway via G2

In figure 3.2, two FRP routers, G1 and G2, come on line and discover that they are gateways. They do not acknowledge each other as potential gateways because their own gateway route, being 0, is cheaper.

G1 and G2 duly inform each of their peers, A, B, C, and D, of this status so that each peer now has a path to gateway set. A and B set their gateway route via G1. C and D set theirs via G2.




Figure 3.3  
*B advertises G1 to E and F*  
*C advertises G2 to F*  
*D advertises G2 to H*  
*E gateway via G1, nexthop B*  
*F chooses say G1, nexthop B over G2, nexthop C*  
*H gateway via G2, nexthop D*

In figure 3.3, B now advertises G1 to E and F, while C advertises G2 to D and F, and D advertises G2 to C and H. Consequently, E has a gateway route via G1 with a nexthop of B and H has a gateway route via G2 with a nexthop of D.

Both C and D ignore the advertisements from each other as they already have a cheaper route to G2.

F is given two gateway routes to choose from, one via G1, nexthop B and another via G2, nexthop C. F needs to pick the cheapest of these two — let's say F, B, G1.




Figure 3.4  
*F advertises G1 to I*  
*H advertises G2 to J*  
*I gateway via G1, nexthop F*  
*J gateway via G2, nexthop H*

Next F advertises G1 to I, because it is cheaper than G2, and H advertises G2 to J.

I now has a gateway route via G1, nexthop F, and J now has a gateway route via G2, nexthop H.

Convergence is reached, leaving the setting of gateways complete, until a change in network topology forces a node to re-advertise and the process begins again.




Figure 3.5  
F to D via B, G1, G2, D rather than via C, D

FRP traffic is weighted to prefer gateway routes over other routes, so F will prefer to send traffic for D via F, B, G1, G2, D rather than F, C, D. In this example, route F, C, D will be seen as a redundant or back route because FRP prefers to pass traffic towards the gateway and F's gateway nexthop is B.




Figure 3.6  
F to D via B, G1, G2, D rather than via C, D unless C, D is very cheap

There are only two reasons that F will route traffic to D via C. One is if the route from F to C to D is very cheap, an unlikely occurrence because then the cheapest gateway is more likely to be G2. The second is if the gateway weighting changes so that G2, nexthop C, becomes F's gateway route. The route from F to D then becomes F, C, G2, D.

### Exchanging gateway routes

An FRP router is told in its configuration if it is a gateway or not. On start up, it works through the list of peers it is provided with and communicates with each one. If the router is a gateway, it uses this initial exchange to inform its peers of its gateway status and they adjust their path to gateway appropriately, sharing any new gateway routes with all their peers.

This interaction results in all 'alive' peers exchanging their current path to gateway. Routers that are not a gateway use these to reassess the cheapest gateway path for themselves,

then sharing this with all their peers. Gateway route exchanges are triggered whenever a router changes their gateway status or recomputes a new path to gateway.

### **Avoiding loops in gateway paths**

A common problem experienced by routing protocols is the creation of loops in the route paths discovered. This same issue can potentially occur in FRPs path to gateway discovery. It is to avoid this problem that a FRP host sends its entire gateway path to its peers. On receiving a path to gateway message, each router checks to see if its own IP address is in the path list. If it is, then the gateway path is excluded as a “*router will not learn a path that contains itself*” [58]. This avoids counting to infinity by ensuring that loops are not formed.

## **FRP route forwarding**

---

The primary function of any router, or host, is to develop a forwarding table to facilitate the correct and efficient delivery of network traffic at the forwarding plane level. The forwarding plane encompasses the ‘front end’ of the router, the mechanism by which it accepts a packet for delivery and decides, by consulting its forwarding table, which of its peers is the ‘nexthop’ in the chain of routers the packet will pass through to reach its destination.

The primary function of any routing protocol is to provide the host with a routing table containing the ‘best’ routes according to that routing protocol. A host may employ more than one routing protocol, so the forwarding table is built by combining all the routing tables submitted by all the different protocols, and choosing the ‘best’ of ‘the best’ routes. This ‘back end’ of the router is the control plane.

FRP is designed to be subordinate to the existing interior gateway protocol (IGP) (goal 2, see page 20), typically OSPF or IS-IS, with the intention of simplifying route advertisement at the inner fringe of the network. The core of the network still requires the heavyweight complexity of the existing well-known, well-understood and well-tested protocols. However, the potentially large number of routes produced by these protocols do not necessarily need to be passed on to the client stub networks at the fringe when a single default route may be all that is necessary. So FRP takes over from OSPF et al, establishing a single gateway and telling the relevant client network to push all traffic towards it (goal 1, see page 20).



An issue that has not yet been addressed is where the FRP should sit in the table of administrative distance that is utilised to allow a router to decide between different routes to the same destination provided by different routing protocols. For the purposes of testing, FRP was given an arbitrarily high rank so that it automatically trumped every other protocol. On reflection, it seems that a high position on the list is perhaps the correct answer, given that FRP is designed to simplify the number of routes at the fringe. This does, however, stand out against the idea of FRP being subordinate to the core network IGP.

Protocol	Administrative distance
Directly connected route	0
Static route	1
External BGP	20
OSPF	110
IS-IS	115
RIP	120
Internal BGP	200

Figure 3.7  
An abbreviated look at the administrative distance table [18]

In common with many other routing protocols, FRP builds a routing table, or routing information base (RIB), by selecting the ‘best’ route to each known destination from among all the routes provided by all its neighbours, or peers. This table of ‘best’ routes is then shared with all peers, excluding in each case any routes that the peer is nexthop for. Meanwhile, the host’s peers are still passing on their ‘best’ routes, triggering changes in the host’s RIB as new and ‘better’ routes are supplied until convergence is reached — that is all hosts have completed their RIBs and no more routing information is passed on. From this point, routes are only exchanged when changes to the network topology are discovered and shared.

Where FRP differs to other protocols is in the mechanisms it employs to build the RIB and in how it distributes routes to peers, in order to achieve the goals specified in its design.

## Quiescence

A FRP router remains quiescent until it is able to determine a path to a FRP gateway (goal 3a, see page 20). A quiescent router exchanges configuration information and receives gateway path and route update information from other peers but does not send out routing updates until a gateway path is established. Should a node lose its gateway path and not be able to establish a new one, it returns to a quiescent state.

## Establishing routes to destinations

FRP employs a distance vector algorithm to decide which route to a destination is the 'best'. Distance vector is iterative, distributed and self terminating. It is also asynchronous in that peers do not have to synchronously exchange data, although FRP overrides this in its message exchange process. A distance vector based protocol receives routing information from directly attached peers and then redistributes new routing information back to its peers.

The distance vector algorithm specifies that a node establish a route for each known peer, setting the cost to the cost of the link between them and setting the nexthop as that peer. This information is shared with all peers. Using the similar information delivered from those peers, identify any new destinations and create a route to them, specifying the peer that supplied the route as the nexthop. The cost of the route is calculated by adding the cost of the link to the nexthop peer to the cost of the route as advertised. If more than one peer offers a route to a destination, select the one with the lowest cost. Share this new information with all peers. Continue to update and share whenever new routing information is received or whenever the network topology changes.

FRP modifies the behavior of the algorithm by adding the criteria that traffic should be preferentially forwarded towards the host's gateway. This creates the need for a more complex check to determine the 'best' route and brings the FRP decision algorithm in to play. FRP peers exchange information on the cost of the link between them, using this to assign a cost to each hop and so building up a metric for each route, as per the normal distance vector method. Additional to this is the metric associated with the host's gateway route and a flag that is added to each route that leads towards that gateway.

For a route to be included in the host's forwarding table, and therefore to be used as a route for delivering traffic, it must be cheaper than the combined gateway costs of the two nodes. FRP prefers to route traffic via the gateways and the backbone rather than over routes between peers. As a consequence, very few routes will be included in the FRP forwarding table as the FRP algorithm ensures that the gateway costs are almost always lower. Therefore, most traffic is directed towards the host's gateway and the backbone.

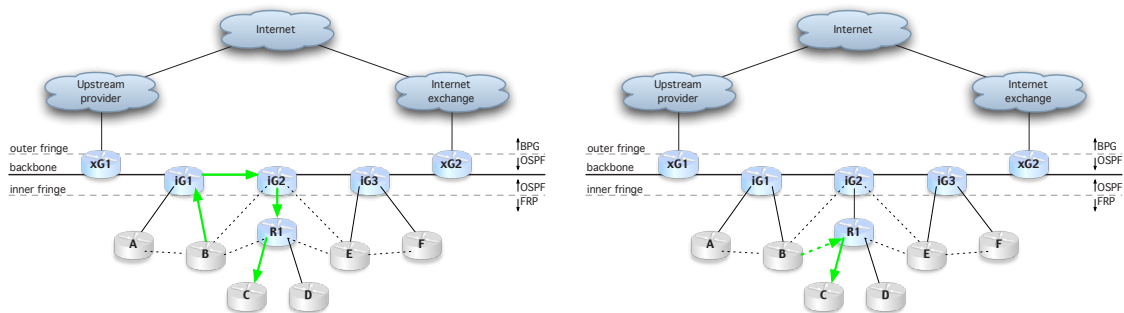


Figure 3.8  
The example below illustrated

### For example:

B's gateway is iG1, C's gateway is iG2 via R1. Should B route traffic for C via its designated gateway, across the backbone, and down to C (B, iG1, iG2, R1, C); or is it cheaper to send it via the backroute (B, R1, C). FRP prefers the first.

### The FRP algorithm

A route leading away from the gateway is included in the RIB if:

$$\text{route-cost} + \text{link-cost} < \text{route-gw-cost} + \text{target-gw-cost} + \text{is-gw-route}$$

where:

- route-cost is the accumulated sum of the costs of intervening links
- link-cost is the cost of the link to remote-node
- route-gw-cost is the cost from the route's originating node to its gateway, as expressed in the routing update
- target-gw-cost is the cost from the remote-node to its gateway, as expressed in its link advertisements
- is-gw-route equals 1 if the route is being passed toward the gateway, 0 otherwise

[57]

Note that the cost of traversing the backbone is 0.

This same algorithm is also used to determine if a route should be advertised.

### Split horizon

Having placed a route in the RIB, when should a FRP router advertise that route to its peers? The first caveat is that a route is never advertised back to the peer that provided it. This is the standard distance vector technique of avoiding counting to infinity — two peers each specifying that the other is the nexthop of a particular route.

## Advertising routes

The second caveat is that the FRP algorithm is applied using the host's knowledge of the peers gateway cost. The route is only forwarded to the peer if the calculation indicates that it is cheaper for the peer to use that route than it is to send traffic via its own designated gateway.




Figure 3.9  
The example below illustrated

### For example:

R1 is a router on the ISP fringe. It is not a gateway — its gateway is iG2 and its peers are B, C, D, E.

B's gateway is iG1, C&D's gateway is iG2 via R1, E's gateway is iG3.

*Does B pass the route B,A to R1?*

$$\text{route-cost} + \text{link-cost} < \text{route-gw-cost} + \text{target-gw-cost} + \text{is-gw-route}$$

$$B,A + B,R1 < A,iG1 + R1,iG2 + \text{not-a-gateway}$$

The hop count only scenario, assume all links = 1.

$$1 + 1 < 1 + 1 + 0$$

The weighted links scenario, assume solid (preferred) routes = 1 and dotted (redundant) routes = 2.

$$2 + 2 < 1 + 1 + 0$$

*Does R1 pass route R1,B,A to C?*

$$R1,B,A + R1,C < A,iG1 + C,R1,iG2 + \text{not-a-gateway}$$

The hop count only scenario, assume all links = 1.

$$2 + 1 < 1 + 2 + 0$$

The weighted links scenario, assume solid (preferred) routes = 1 and dotted (redundant) routes = 2.

$$4 + 1 < 1 + 2 + 0$$

In both cases, the hop count comes out equal so the backbone route is the favoured one, although another hop between R1 and its gateway would mean that the redundant route was preferred. With the link weighting, the backbone route is clearly favoured over the alternative.

## FRP router configuration

A router requires a certain amount of data available to initially get up and running. Typically this might include address, interface and network information; the rules for accessing and communicating with neighbours or peers; output locations for debugging, status and logging data; and in the case of FRP, gateway information.

### Configuration file

Each FRP router needs a configuration file. It stores the information required during the start up process and is where the mandatory and default operating values are set. Although each individual implementation of FRP will need to handle configuration differently in line with the system being developed for, the file that came with the prototype implementation provides an indication of what is important to FRP. It is clear that this example is a work in progress as different methods, some of which are commented out (#), are used to specify the same data. The original prototype of FRP sets all configurable values in this file, these values can not be changed once the router is up and running.

```
gateway always
debug 2
pidfile frpd.pid
#secret fred
statusfile frpd.status
#acl peers permit ip 192.168.151.0/24
#acl accepts permit ip 192.168.0.0/16-24
#listen 192.168.151.110 secret fred poll 0.5 retry 0.1 fail 2
#listen 192.168.151.110 accept-connect peers
#listen 192.168.151.110 accept accepts
listen 192.168.151.110
#peer 192.168.151.91 secret fred
#peer 1.1.1.2
override static [57]
```

In this example, the basic FRP information is

- `gateway always`, indicating that this router is a gateway
- `secret fred`, setting the secret used by this router
- `listen 192.168.151.110 secret fred poll 0.5 retry 0.1 fail 2`, setting the local address, secret, poll time, retry time and fail time
- `listen 192.168.151.110`, the address of this router
- `peer 192.168.151.91 secret fred`, a peer known to this router — its address and secret

The configuration file lines

```
#acl peers permit ip 192.168.151.0/24
#acl accepts permit ip 192.168.0.0/16-24
#listen 192.168.151.110 accept-connect peers
#listen 192.168.151.110 accept accepts
```

indicate that the original implementation allowed the use of access lists to control which peers a router can communicate with. Given that the use of ‘secrets’ mandates the explicit setting up of peers by an administrator (see *Peers and Secrets* below), this may not actually be necessary.

Some output information is also specified to support multiple levels of debugging and the creation of two output files, one is used to store the process id and the other a complete copy of the entire FRP routing table.

## User commands

The original prototype of FRP sets all configurable values in the configuration file. The *Quagga* implementation requires the mandatory basics to be set there but other mandatory and optional values are pre-set as defaults in the daemon code. These values may be reset either within the configuration file or via a user command line tool while the router is up and running.

Consequently, the *Quagga* implementation of FRP needs to identify the list of configuration settings required, the default value for each setting, and whether the setting is to be turned in to a command to allow an administrator to live change the setting.

## Peers and secrets

FRP peers need to be known and set up in advance. An access control list can be set up to specify which addresses are acceptable as peers but the real restraint is the need to know each individual peer's specific security information before any messages can be exchanged. Therefore each peer requires a secret to be stored in the configuration file and loaded on start up, or provided via the command line with all the other information if a peer is added while the router is up and running.

The only way to achieve this is for the secret to be set manually at each end. Consequently the secret is never sent via the network except as part of the security hash used to encode and decode packets, a good security feature if the secrets are exchanged via some trusted and secure means. This is helped by the fact that FRP is designed to be used in isolated pockets within confines of an ISP network so it may often be the case that one administrator is setting all the secrets anyway, and even if there are more than one, then they probably still work for the same organisation.

Other information also needs to be stored for each peer — the peer's address, the cost of the link between router and peer, and the poll and retry times to be set for that peer.

## Sequence numbers

Each FRP router contains a random number generator that provides a sequence number each time a new conversation is started. A router keeps track of the sequence currently in use with each peer (*lseq*), incrementing the previously used number by one before embedding it in a new packet.

Each host also maintains a copy of the peer's current sequence number (*rseq*), and the two are used in tandem to catch gaps where packets have gone missing and to ensure that packets are processed in the correct order. When a packet arrives from a peer, the sequence numbers are checked to ensure that the expected number is in fact the one used. If this is not the case, then a problem has obviously occurred and the packet is not processed.


A packet arriving from a peer stating that the host's latest sequence number is 0 indicates the start of a new conversation with that peer. Consequently, if a sequence number ever naturally wraps to 0, it is automatically reset to 1.

## FRP message formats

---

Often the concept of a ‘packet’ and a ‘message’ is essentially the same but FRP clearly differentiates between the two. This is a direct result of the fact that FRP allows multiple messages to be sent in a single packet, potentially reducing the amount of traffic flowing between FRP routers but increasing the complexity of building and processing packets.

A single FRP message begins with a message header providing the length of the entire message in bytes and a code indicating what type of message this is. The header is followed by the fields specific to that message. A FRP packet consists of a packet header containing security and sequencing information followed by zero or more messages in the above format. Packets containing just the packet header and no messages are only used in two special circumstances; generally at least one message is attached. Some message types may occur more than once in a packet.



*Figure 3.10*  
*A representation of a packet containing multiple messages*

Although there is no requirement for most messages to be sent in any particular order, the reality is that the linear execution of code tends to ensure that they are. Each message is however a complete and self-sufficient bundle and once extracted, can be processed independently.



## The FRP packet header

The packet header sits at the beginning of each packet sent by an FRP router to one of its peers. There is exactly one packet header per packet sent, regardless of how many messages the packet contains.

Size: 128 bits / 16 bytes

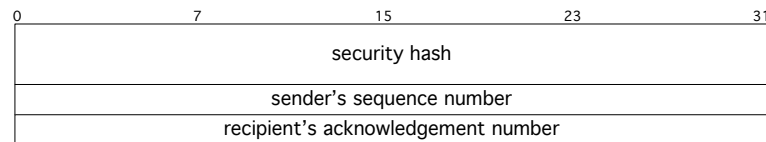


Figure 3.11  
Specification of a packet header

security hash	64 bits / 8 bytes	a hash of the local router's secret plus the contents entire packet
sender's sequence number	32 bits / 4 bytes	the current sequence number of the local (sending) router
recipient's acknowledgement number	32 bits / 4 bytes	the latest received sequence number from the remote router that is at the other end of the current conversation

## FRP message headers

A FRP packet can contain multiple messages. Each message must be prefixed with exactly one message header specifying the type and length of the message.

Size: 16 bits / 2 bytes





Figure 3.12  
Specification of a message header

message length	8 bits / 1 byte	the actual length of this message, including the message header, in bytes
message type	8 bits / 1 byte	the type of the message that follows specified in one of the following hexadecimal codes: 0x01 control message 0x41 IPv4 configuration message 0x42 IPv4 path to gateway message 0x43 IPv4 route update message 0x61 IPv6 configuration message 0x62 IPv6 path to gateway message 0x63 IPv6 route update message

## Session control messages

Control messages inform the receiving peer about the type of conversation in progress.

Size: 32 bits / 4 bytes



*Figure 3.13*  
*Specification of a control message*

message header: message length	4						
message header: message type	0x01						
control type	the type of this control message is specified using one of the following numerical codes: <table><tr><td>1</td><td>POLL</td></tr><tr><td>2</td><td>ACK</td></tr><tr><td>3</td><td>NAK</td></tr></table> <p>POLL is used to send a message checking that a peer that has not been heard from for a period of time is still alive</p> <p>ACK is used to indicate that the local router received the last message of a conversation from the remote peer and is now closing the conversation down</p> <p>NAK is used to indicate receipt of a malformed packet</p>	1	POLL	2	ACK	3	NAK
1	POLL						
2	ACK						
3	NAK						
control parameters	not currently used						

## Static configuration messages

Configuration messages tell other peers about the state of the sending router. One is always sent at start up and then again at any other time that the router state changes.

Size: 96 bits / 12 bytes




Figure 3.14  
Specification of an IPv4 configuration message

message header: message length	12
message header: message type	0x41
cost of link	16 bits / 2 bytes the cost of the link between the local router and the remote peer
poll time	16 bits / 2 bytes the amount of time to wait without hearing from the remote peer before sending a POLL control message to see if the peer is still alive
retry time	16 bits / 2 bytes the amount of time to wait for an acknowledgement from the remote peer before resending the last packet
router-ID of peer	32 bits / 4 bytes the unique id of the remote peer, typically it's IPv4 address

Size: 192 bits / 24 bytes




Figure 3.15  
Specification of an IPv6 configuration message

message header: message length	24
message header: message type	0x61
cost of link	16 bits / 2 bytes the cost of the link between the local router and the remote peer
poll time	16 bits / 2 bytes the amount of time to wait without hearing from the remote peer before sending a POLL control message to see if the peer is still alive
retry time	16 bits / 2 bytes the amount of time to wait for an acknowledgement from the remote peer before resending the last packet
gateway address to use	128 bits / 16 bytes the unique id of the remote peer, typically it's IPv6 address

## Path to gateway messages

If a FRP router is not told that it is a designated gateway router, it needs to find out the path to its nearest gateway. The peers of the router pass it ‘path to gateway’ messages telling it of the best path currently known by that peer. The local router can use this to choose the best gateway path for itself, which is then passed on to all its own peers (except the one who originally told it about the gateway path).

Size: from 64 to 2016 bits / from 8 to 252 bytes




Figure 3.16

Specification of an IPv4 path to gateway message

message header: message length	1+1+2+4 + 0 to 244 in 4 byte increments (ie: between 0 and 61 additional IPv4 addresses)
message header: message type	0x42
cost from peer to gateway	16 bits / 2 bytes cost of the link between the advertising router and it's gateway if the advertising router is a gateway, the cost is 0 if the advertising router does not know of a gateway, the cost is infinity (which is 0xffff in FRP)
path from peer to gateway	32 bits / 4 bytes x n (1 ≤ n ≤ 62) a list of 1 to 62 IPv4 addresses specifying the path taken by this router to reach it's designated gateway the first address in the list is always this router's address

Size: from 160 to 7968 bits / from 20 to 996 bytes

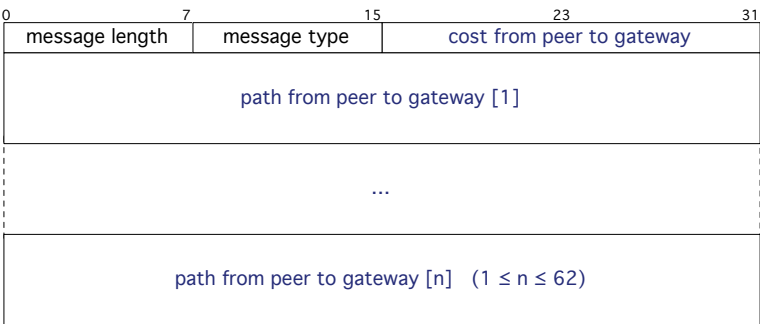


Figure 3.17  
Specification of an IPv6 path to gateway message

- message header: message length    1+1+2+16 + 0 to 980 in 16 byte increments  
(ie: between 0 and 61 additional IPv6 addresses)
- message header: message type    0x62
- cost from peer to gateway    16 bits / 2 bytes  
cost of the link between the advertising router and it's gateway  
if the advertising router is a gateway, the cost is 0  
if the advertising router does not know of a gateway, the cost is infinity  
(which is 0xffff in FRP)
- path from peer to gateway    128 bits / 16 bytes x n (1 ≤ n ≤ 62)  
a list of 1 to 62 IPv6 addresses specifying the path taken by this router to reach it's designated gateway  
the first address in the list is always this router's address

## Route update messages

Once the local router establishes a path to a gateway, it starts to exchange routing information with other FRP peers using route update messages. It utilises the information provided by peers via these messages to build and rebuild its own routing table, selecting the best route to each known destination to store and to share with its peers.

Size: 96 bits / 12 bytes

0	7	15	23	31
message length		message type	update type flag	prefix length
route cost		cost from originator to gateway		
IP prefix				

*Figure 3.18*  
*Specification of an IPv4 route update message*

message header: message length	12
message header: message type	0x43
update type flags	8 bits / 1 byte
	the update flags are specified by setting one or more of the following bits:
	0x01 begin
	0x02 commit
	0x04 null
	0x08 update
	0x10 delete
	0x80 gateway
prefix length	8 bits / 1 byte
	length of the prefix mask (a number between 0 and 32)
route cost	16 bits / 2 bytes
	cost of the new route being advertised
cost from originator to gateway	16 bits / 2 bytes
	cost of the link between the advertising router and it's gateway
IP prefix	32 bits / 4 bytes
	destination address or prefix

Size: 96 bits / 12 bytes

0	7	15	23	31
message length	message type	update type flag	prefix length	
route cost		cost from originator to gateway		
IPv6 prefix (truncated)				

Figure 3.19  
Specification of an IPv6 route update message

message header: message length	12
message header: message type	0x63
update type flags	8 bits / 1 byte
	the update flags are specified by setting one or more of the following hexadecimal codes:
	0x01 begin
	0x02 commit
	0x04 null
	0x08 update
	0x10 delete
	0x80 gateway
prefix length	8 bits / 1 byte
	length of the prefix mask (a number between 0 and 128)
route cost	16 bits / 2 bytes
	cost of the new route being advertised
cost from originator to gateway	16 bits / 2 bytes
	cost of the link between the advertising router and it's gateway
prefix	32 bits / 4 bytes
	destination prefix

### Batch processing of update messages

Updates typically happen in batches as a router sends its current routing table to all peers, one update message at a time. The first message of an update batch carries a BEGIN flag and the final message of the batch carries a COMMIT flag. Messages in between the BEGIN and the COMMIT are assumed to belong to the same batch. Any individual route that is also a gateway route carries a GATEWAY flag as well.

There are instances where update messages do not occur as a batch. A router can tell a peer that it has nothing to share by sending a single update carrying the BEGIN + NULLRT + COMMIT flags. Modification to or deletion of an individual entry in the routing table carry BEGIN + UPDATE + COMMIT or BEGIN + DELETE + COMMIT.

To avoid processing out-of-date routes, each time a new BEGIN is received, any previously unCOMMITTED batches are discarded in favour of the new arrival. If the batch sequencing is interrupted or lost for any reason, the current batch is likewise discarded.

## FRP packet exchange

---

In common with any other routing protocol, FRP has a set of rules for the exchange of packets between FRP routers to ensure the secure delivery of messages in the correct order. The diagrams show only the packet header and the message headers. The colour coding for the different message types is used consistently throughout this document.

### The packet header

The packet header holds the information that is unique to the entire packet — the sender's sequence number (or local sequence number or lseq), the recipient's acknowledgement number (or remote sequence number or rseq), and a security code created by using the SHA library to create a hash of the entire packet plus the sender's secret. The hash is duplicated by the recipient on receipt of the message. If the two do not match, the packet is discarded. This process is followed every time two FRP routers exchange packets.

A new header and the subsequent hash are produced each time a packet is constructed, although the hash is actually the last element to be generated as it requires the packet contents to be in place before it can be created.

### Initial handshake

When a router comes online for the first time, it works through the list of peers provided in the configuration file. By sending a SYN packet to each one, the router determines which peers are 'alive' and lets each one know that it also is now 'alive'.

To start the initial conversation with a new peer, the initiating router creates a packet header setting the sequence number to the randomly generated starting sequence number for this peer and the acknowledgement number to 0, as the peer's current sequence is not yet known. The 0 indicates to the remote peer that this is the start of a new conversation.

This is a SYN packet and is specified as a null packet, so no messages are added to the packet header. Packet creation is now complete so the security hash can be computed and the packet sent to the remote peer.



*Figure 3.20  
The start of a new conversation*



The remote peer takes receipt of the packet, checking both the security hash and that the initiating router is one of its peers. It responds with a null ACK packet, setting the sequence number to the next available sequence number for this router and the acknowledgement number to the sequence number sent in the SYN packet.



Figure 3.21  
Handshake is complete

### Batches of messages

The initiating router receives the ACK packet from the peer and replies, incrementing the sequence number and building a new packet setting the acknowledgement number to the one sent by the peer in the ACK packet. This response packet may contain more than one message, as the host router will scan through the information held on this peer creating messages for each of the ‘send message’ flags currently set.

Following the initial handshake, a configuration message is added to the packet header, sending the peer the initiating router’s poll and retry times, address, and the cost of the link between them. A path to gateway message is also added, sending the peer information about the router’s gateway, which, at this point, will be either “I have no gateway” or “I am a gateway”.




Figure 3.22  
Sending more than one message in a packet

The responding peer receives the configuration and gateway messages and replies in kind by adding a configuration message and then a gateway message to the return packet. Incidentally, it is at this point that a router can begin to build a path to gateway if it needs to, using the path to gateway information it has just received.

## Triggered updates

In addition to the configuration and gateway messages, the remote peer may have route updates to send the initiating router. Initially these updates are a complete exchange of routing tables. Subsequent information exchanges are triggered by one of the hosts receiving an update from elsewhere and propagating the route information on to all its peers.

The remote peer can add a batch of update messages, with the appropriate flags set, to the packet at this point.

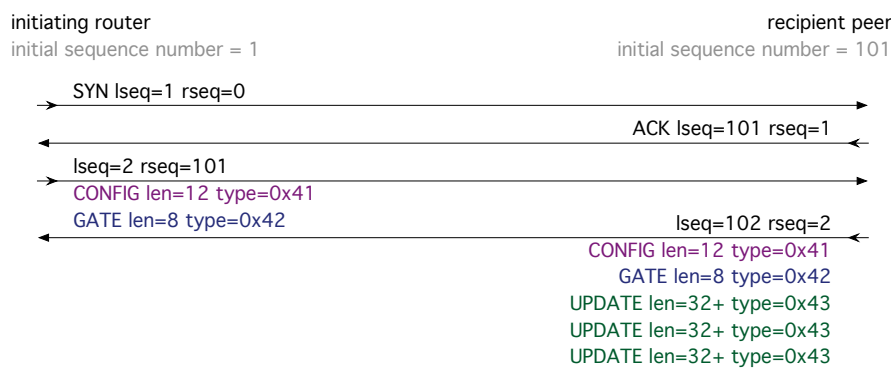


Figure X 3.23  
Sending config, gateway and update messages  
in response to receiving a config and a gateway message

## Closing the conversation

The initiating router receives and processes the messages. If it has nothing to send in return, it closes the conversation by sending a packet header plus a control message with the type field set to ACK, letting the peer know that this exchange is over.

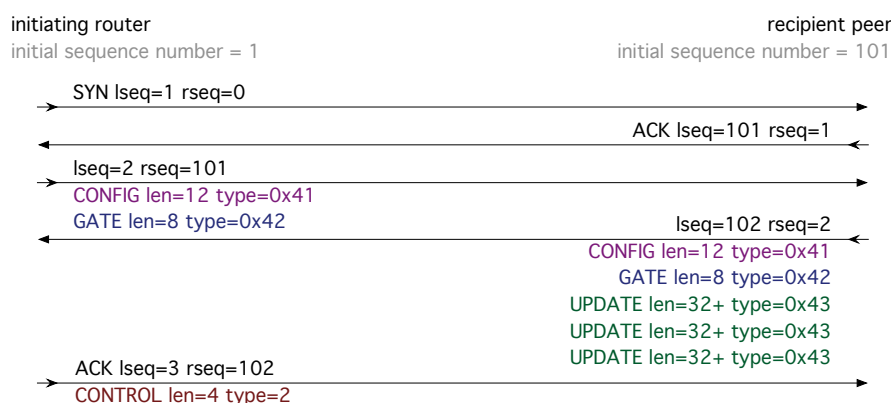


Figure X 3.24  
Closing a conversation

Alternatively, if the initiating router also has route information to share, it may send its own batch of route update messages back to the remote peer. The peer is then the one to conclude the conversation by sending a packet header plus a control message with the type field set to ACK, letting the router know that this exchange is over.

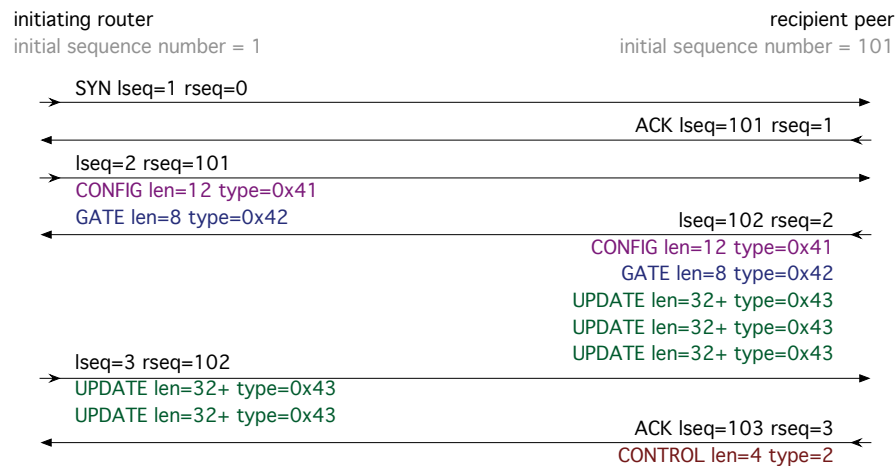


Figure X 3.25  
A complete conversation

### Starting a subsequent conversation

Next time the router needs to communicate with this peer, it starts the conversation process again by sending a SYN packet with the next number in the sequence used for this peer and the acknowledgement number set to 0, indicating the start of a new exchange. The peer ACKs this SYN, providing it's next sequence number. The initiating router then proceeds to send a packet of messages. The peer may respond with messages of it's own but eventually one or the other will send a final ACK to close the conversation again.




Figure X 3.26  
Starting a subsequent conversation

## Polling

If a peer has not been heard from for a pre-specified period of time (set as a configuration variable), then a host sends a poll message, essentially saying “Are you still there?”. The host still needs to conduct the standard SYN/ACK handshake before sending a control message with the POLL control type (1) set. The correct response is a control message with the ACK control type (2) set.





Figure X 3.27  
A poll conversation

## Dealing with failure

A common delivery issue in many protocols is the problem of packets arriving out of order, especially when relying on UDP as the transport protocol as UDP has no delivery guarantees. The synchronous nature of the FRP message exchange provides it with the robustness it needs by insisting on a TCP style requirement that a host pause to receive an acknowledgement from the peer before sending the next packet. The synchronicity rises from FRP sending multiple messages in one packet rather than each individual message on its own.

A potential failure occurs when a router disappears and so stops replying to messages. This particular failure cannot be recovered from unless the peer reappears sometime in the future. All that can be done in the interim is to make sure the peer really has disappeared, rather than just becoming slow to respond, and then note the fact that it has gone. If a host does not receive a response to any sent packet within a pre-specified (configurable) period of time, the packet is resent and then resent again until the retry time runs out. At this point the peer is declared ‘dead’.



*Figure X 3.28  
The peer has 'died'*

When the peer reappears, an initial handshake exchange updates the information held in the host's list of peers and conversations continue as usual.

Other common problems in packet delivery are detected as sequence number mismatches. In a similar situation to the one above for example, if the peer was merely having network issues causing it to slow down, it may get a message from the initiating router more than once. If the message is a SYN packet, each new one will be treated as a new conversation request, so the previous conversation request is dropped. Hopefully an ACK will get through before the host declares the peer 'dead' but if it doesn't an initial handshake will take place next time the peer has a communication for the host.

Alternatively, the host's repeated message could occur in the middle of a conversation. In this case, the first message received by the peer is replied to and any subsequent ones dropped as having the incorrect 'rseq'. This is because the peer is expecting an incremented sequence number but the host is still ACKing the original until the peer's reply arrives — at which point a normal conversation continues.

Another possible cause of failure is a corrupt packet, either by damage to the packet in transit or by more malignant means. In this case the problem is detected when the receiving router attempts to validate the packet. As the initial security hash performed by the sending host utilises both its secret and the contents of the packet, any deliberate attempt to modify the packet will only succeed if the third party knows the secret. Otherwise, the recomputation of the hash by the recipient fails and the packet is ignored and discarded.



---

## 4. Designing the Quagga FRP daemon

---

This chapter covers the architecture of the *Quagga* system, the initial process of designing and building the *Quagga* FRP daemon, the architecture and structure of the daemon and the program flows necessary for the correct application of the FRP algorithm and associated operational requirements.

The design problem is to, within the confines of the *Quagga* software routing suite, create a fringe routing protocol daemon that:

- can bootstrap itself, — ie: get up and running using a configuration file and can inject any routes in that configuration into the kernel;
- can be told a set of peers via the configuration file or via a terminal, can send messages to those peers, and can receive messages from those peers;
- can extract information from the received peer messages, can make the correct decisions about building a routing table using the data in those messages, and can distribute changes in the routing table back to the peers;
- can use the routing table to correctly route traffic.

## Quagga Architecture

---

Quagga is not a router in the physical sense of the word, in that it is not a dedicated piece of hardware, optimised to perform one specific task as quickly and efficiently as possible. It is instead, a modular, open source software suite for Unix and Linux platforms. In place of the specialist hardware, the package allows a standard Unix OS desktop to operate as a router using Quagga to facilitate the routing functions.

The Quagga command set and the syntax it uses is very similar to that used by Cisco for their routers and so is familiar to most users, although Quagga does not have the same extent of functionality provided by Cisco routers [<http://sourceforge.net/apps/mediawiki/quagga/>].

At Quagga's core is the Zebra daemon providing an abstraction layer using a client/server model to interface between the individual protocol daemons and the Unix kernel [[www.quagga.net/about.php](http://www.quagga.net/about.php)]. Each protocol daemon is a standalone process, maintaining its own state and routing table consistent with its protocol's algorithm. Zebra is responsible for handling the injection of routes from each individual daemon into the kernel routing table and for the redistribution of routes between the individual protocol daemons. The Zebra daemon also manages other shared functions that rely on cooperation with the kernel, such as using interfaces and sockets.

Each Quagga daemon is run as a separate and independent process on the host operating system. The Zebra process must be running before any of the protocol processes will run correctly. Individual protocol daemons may be taken up and down without affecting any other Quagga process.

Figure 4.1 shows the way Quagga separates the two major functions of a router, that of routing traffic from that of creating and maintaining a table of destinations and the directly connected nexthops. The top part of the diagram shows Quagga enabling the Unix system at the forwarding plane level where incoming traffic is accepted, the destination address consulted, and the nexthop determined. The underlying Unix kernel's routing table is used as the forwarding information base (FIB) and is consulted to determine which interface is used to send outgoing traffic on its way.



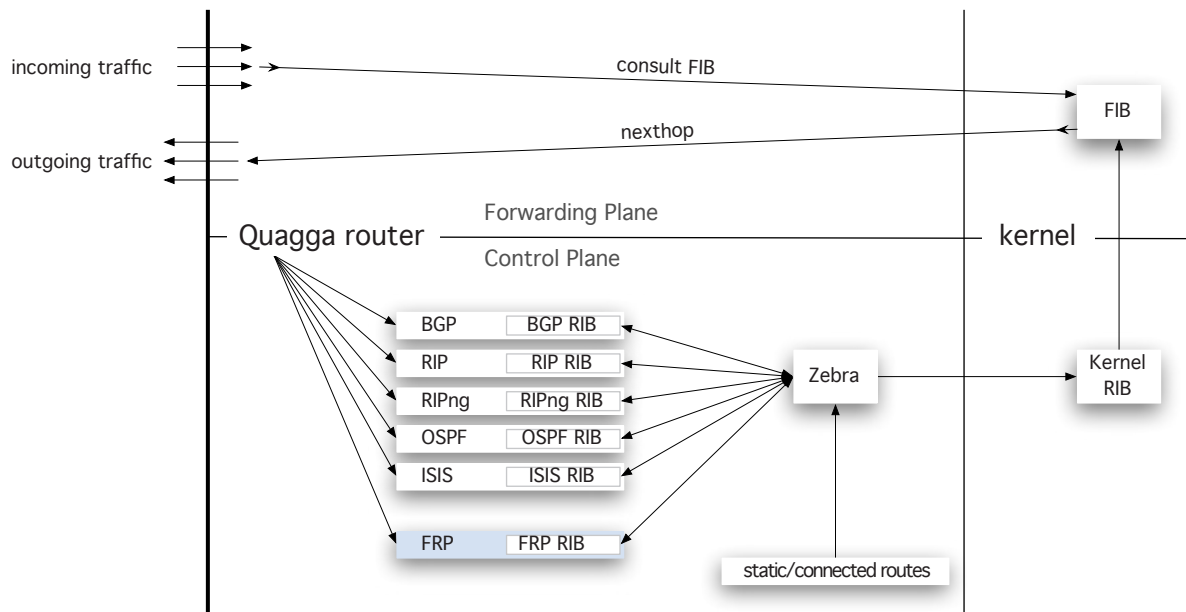


Figure 4.1  
The architecture of a Quagga software router

The control plane part of the diagram shows the modular approach of *Quagga*. The individual protocol daemons, each with their own individual routing table, are clients of the Zebra server daemon that facilitates access to the Unix kernel. Protocol daemons provide their chosen ‘best’ routes to the kernel and to each other via the Zebra daemon. Zebra also allows user defined static and connected routes to be defined, duplicating some of the functions that can also be performed directly via the Unix OS.

## Communicating with Zebra

Communication between Zebra and the protocol daemons is provided by *Quagga*’s Zserv API running over a TCP or Unix stream to the client daemons [www.quagga.net/about.php]. Knowledge of the message protocol used for this communication is not required in order to create a new routing daemon to add to the suite.

Zserv supplies zclient functions to the client daemons, allowing them to make calls to API functions that provide access to standard kernel mechanisms. Each new protocol daemon must also implement the callback functions necessary to allow Zebra to communicate with that daemon, as Zebra makes these calls to the daemon assuming that they have been implemented. These API and callback functions handle updating the router id, detecting, adding and deleting interfaces, taking interfaces up and down, adding and deleting addresses to interfaces, discovering addresses attached to interfaces, and adding and deleting routes in the kernel’s routing table [60].

## Configuration and commands

Like any other router, *Quagga* daemons cater to the custom needs of the network they are running on and consequently need to be configured. This can happen in one of two different ways. A minimal amount of information is required for loading at start up and this is stored in a simple, text based configuration file. Additional data may be added either to the configuration file or entered by an administrator while the daemon is running. The current running configuration may be written out to the configuration file at any point.

Re-configuring the daemon while running is facilitated by the built in virtual teletype terminal (VTY), which provides a command line interface (CLI) for issuing daemon commands. The Zebra command set provides the standard settings for the routing system — setting static and connected routes for example. A basic set is provided for each new daemon and these are fully expandable and customisable to fit the needs of the protocol. The VTY commands are loaded during daemon initialisation.

*Quagga* has an elegant solution for handling these configuration commands. A macro has been defined which makes adding new commands reasonably straightforward by providing a standardised framework. The format in which a command is executed is identical whether it is issued by a user or comes from the configuration file, so both methods use the same piece of code. Commands can be created anywhere in the code base and so are typically defined in the same file as the feature they relate to.

## Threads and events

*Quagga* is a multi-threaded, event driven system. On initialisation, a daemon spawns a master thread within an infinite while loop which responds to events. Each daemon has its own list of events, typically the arrival of a packet, a change to the routing table, and the triggering of a poll timer but there is scope for other events to be included.

Within the main daemon engine, the initialisation sequence triggers the first instance of each event. The event thread is created and the appropriate function attached. The thread then sleeps until that event is triggered at which point the thread wakes and the code is executed. The first action within the function is to create a new event thread, attach the function to it, and set to wait for the next trigger.

## Library support

The Quagga suite supplies a good set of library files to provide support for protocol daemon development [www.quagga.net/about.php]. These include `zclient.h/c`, which provides the Zserv API mentioned above and `zebra.h` with the global defines, macros, external variables and prototypes made available by the Zebra daemon.

Other library files provide support for:

- the management of memory, signal handling, threads, logs, and privileges;
- creating user commands, extracting them from the configuration file, and executing them in the VTY
- working with networks, interfaces, and sockets;
- all manner of routing related functions like route filtering and distribution, and working with prefixes, routemaps and tables;
- the standard program requirements of buffers, hash tables, linked lists, strings, vectors and the network.

## Initial process

---

The first step towards building a FRP daemon was to get to grips with the *Quagga* code and see how a protocol daemon hooked in to the Zebra daemon and the associated libraries. RIPng is the smallest of the existing *Quagga* protocol daemons and the name also supplies a good, unique search string. So RIPng was duplicated and turned in to an instant 'FRP' daemon simply by running a global find/replace of '*ripng*' with '*frp*' within the duplicated directory, then recompiling and working through the errors generated.

Once the new 'FRP' daemon compiled in isolation, the next step was to get the Zebra daemon to recognise it. This became an extensive hunt through the Zebra and library code looking for every instance of RIPng, duplicating the necessary code and updating it for the new daemon — which at this point mainly meant changing RIPng to FRP. Later on, it required analysing what was happening at that point and deciding what modifications were required for the real FRP daemon. It also became apparent that looking at more than one protocol was necessary, so RIP and BGP searches were conducted as well. Once again, this method eventually came down to constantly recompiling and working through the errors generated. The overall technique yielded a good indication of where the Zebra daemon, the library code, and the `make` and `configure` files need to be modified, adjusted and added to in order to support a new protocol daemon.

The other major unknown was how a *Quagga* daemon behaved at the basic level before any routing protocol specific code was added. The initial idea was to take the RIPng version of FRP and remove all the RIPng specific code, leaving a basic daemon shell for the FRP protocol to be built on top of. However, the fact that RIPng is an IPv6 only protocol started to cause problems. This left a choice between RIP, which as it is IPv4 only, was also not ideal but easier to work with than RIPng, and BGP, which was considerably more complex. RIP was the obvious choice and a new ‘FRP/RIP’ daemon was created and compiled. The technique ultimately failed however, as it became too difficult to separate the RIP protocol specific code from the generic daemon code.

Going through the above process was necessary but had left the code base in an irreparable mess. A clean start was required. This time the modifications to the Zebra daemon and the libraries were systematically worked through (see Appendix C) before a new daemon was built from the ground up using the *Zebra for dummies* [60] document as a guide and referring to the RIP, RIPng and BGP code as necessary. The daemon was created in two stages, the first being the creation of a shell and the second being the addition of the FRP protocol.

In the creation of the shell, the RIP, RIPng and BGP code base became an extremely useful resource and in many cases, code for basic functionality was lifted directly from one or another of them. The RIP/RIPng combination was invaluable as a guide to the differences between IPv4 and IPv6 whereas BGP showed how to combine IPv4 and 6 in one daemon.

When reading the FRP daemon code included in this document, note that as *Quagga* only uses the `/*comment*/` style of comments, `/*comment*/` denotes code taken directly from an existing daemon and `//comment` indicates new code.

## Code structure

The Quagga protocol daemons mostly adhere to a common code file structure, although there is some diversity between them. Inside the main Quagga directory, a `lib` directory holds shared code and a `zebra` directory stores the implementation of the Zebra daemon. Each individual Quagga protocol daemon has its own directory called `protocol`. The initial code file is called `protocol_main.c` that calls `protocold.c`, and it is this file that contains the core daemon code. Both files use the same header, `protocold.h`, which holds all the core datastructures and declarations. Other files are created as necessary using a `protocol_` prefix. Interaction with Zebra happens inside `protocol_zebra.c` and other common suffixes appear in multiple daemons — `_debug`, `_interface`, `_peer`, `_route` and `_routemap` for example. The FRP daemon was designed to fit in to this structure.

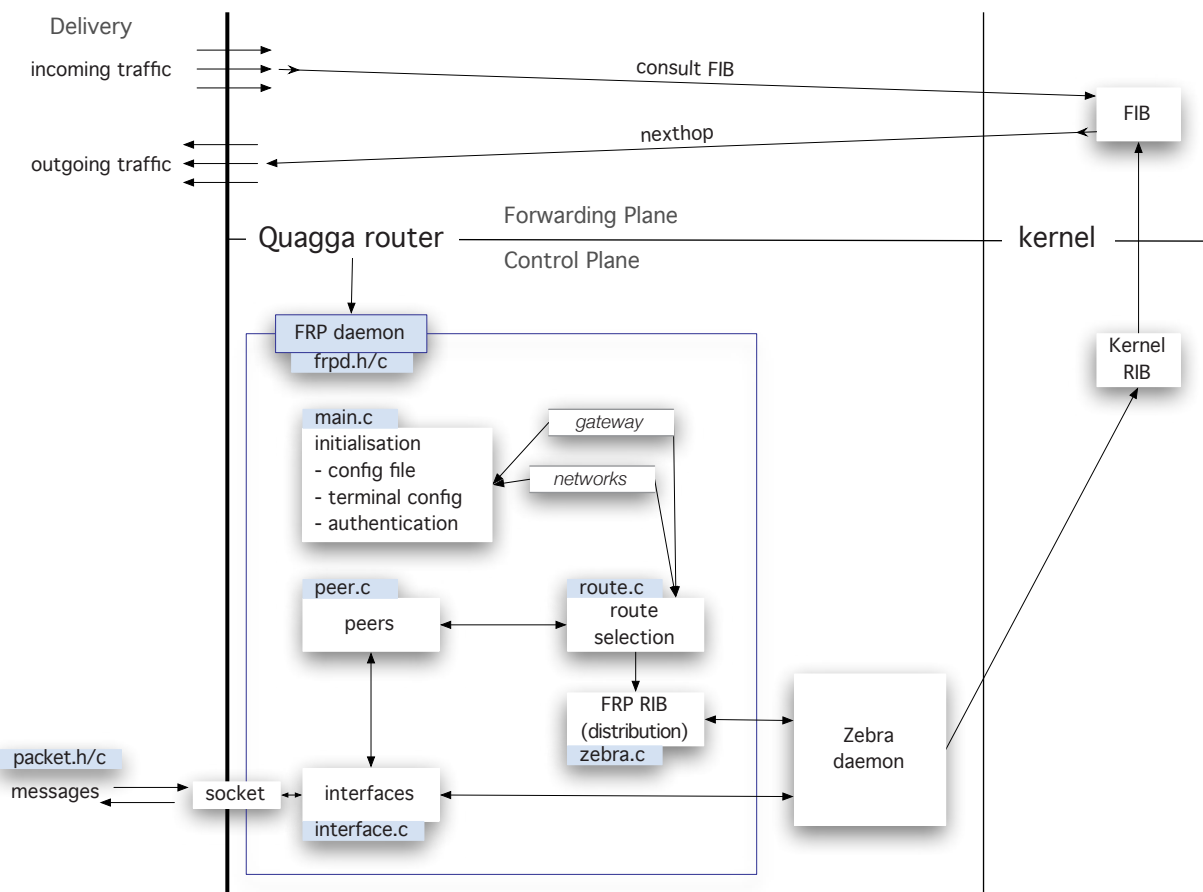


Figure 4.2  
The FRP daemon architecture  
(note that the file name `frp_` prefixes have been dropped)

**frp\_main.c**

The *protocol\_main.c* files are not only common to all Quagga daemons but the code is very similar in each case. It is here that the set up and initialisation of the system occurs. The main program:

- ‘includes’ the *frpd.h* file that sets up many of the data structures and global variables
- sets the daemon’s privileges
- initialises debug logging
- extracts and executes the command line arguments used when the daemon was booted
- creates and sets up the master thread
- calls the system level initialisation functions — zebra privileges, signal handlers, commands, VTY, memory
- calls the FRP and Zebra (*zclient*) daemon initialisation functions
- installs and sorts the VTY commands
- reads the config file and executes the commands
- changes to running in daemon mode
- creates VTY socket and starts the socket listeners
- creates the pid file
- starts up the main program loop, an infinite `while` loop in which the master thread continuously fetches and executes the next thread

**frp\_zebra.c**

The primary function of the `frp_zebra.c` file is to initialise `zclient` and to set up the callback functions for the Zserv API. This file is present in all the Quagga protocol daemons. The callback functions implemented in FRP are the following:

```
int (*interface_add) (int, struct zclient *, uint16_t);
int (*interface_delete) (int, struct zclient *, uint16_t);
int (*interface_up) (int, struct zclient *, uint16_t);
int (*interface_down) (int, struct zclient *, uint16_t);
int (*interface_address_add) (int, struct zclient *, uint16_t);
int (*interface_address_delete) (int, struct zclient *, uint16_t);
int (*ipv4_route_add) (int, struct zclient *, uint16_t);
int (*ipv4_route_delete) (int, struct zclient *, uint16_t);
```

which need to be mapped to the functions that actually implement them in the appropriate files. Two are implemented here in the function `frp_zebra_read_ipv4`.

```
zclient->ipv4_route_add = frp_zebra_read_ipv4;
zclient->ipv4_route_delete = frp_zebra_read_ipv4;
```

Two more callbacks need to be added to complete the daemon:

```
int (*ipv6_route_add) (int, struct zclient *, uint16_t);
int (*ipv6_route_delete) (int, struct zclient *, uint16_t);
```

**frp\_interface.c**

Much of the interface code is common to all Quagga daemons and handles anything to do with network interfaces. The interface initialisation function is defined here, as is the VTY `enable FRP network` command. The rest of the functions define the many requirements of using interfaces including bringing them up and down, creating, enabling and checking them, attaching addresses, networks and peers to them, and looking up information about them.

The interface callback functions are implemented and linked to the Zebra daemon here:

```
zclient->interface_add = frp_interface_add;
zclient->interface_delete = frp_interface_delete;
zclient->interface_up = frp_interface_up;
zclient->interface_down = frp_interface_down;
zclient->interface_address_add = frp_interface_address_add;
zclient->interface_address_delete = frp_interface_address_delete;
```

**frp\_debug.h and frp\_debug.c**

The Zebra daemon and the Quagga libraries handle a wide range of debugging options that can be hooked into as necessary. Output to screen, terminal or log file is handled by embedding `zlog_debug` statements into the daemon code. Individual protocol daemons can also specify their own custom debug statements to be called and displayed via the `zlog_debug` system. FRPs small amount of debugging code resides here and consists predominantly of VTY commands used to turn various debug levels on and off.

**frpd.h**

In common with other Quagga prototype daemons, the main repository of FRPs external and global includes, `#defines`, variables, data structures, and function prototypes. One of the key structures is `struct frp` which holds information about the FRP router including its routing table and its gateway status. Another is `struct frp_peer` which stores data on a single peer, each one being held in a linked list called `frp_peers`. There is only one `struct frp` and only one list `frp_peers`. Both are specified as prototypes here.

Also defined here is `struct frp_route` which stores a single route given by a peer, a collection of which makes up the stored copy of that peer's forwarding table, and `struct frp_info` which holds meta data about the router's routing table and is used to handle access to it via the systems set up by Quagga.

**frpd.c**

This is the main daemon file containing much of the code specific to running the daemon. `frpd.c` handles the daemon's internal setup and configuration including creating the `struct frp` instance, creating and setting up the socket for FRP to run on, and adding the FRP specific commands to the VTY interface. This file also manages re-writing the configuration file on demand, once gain via the VTY interface.

The bulk of the code stored here however, controls the mechanics of running the fringe routing protocol. FRP's implementation of Quagga's triggered event driven system is defined here, as are the functions that handle each individual event.

**frp\_packet.h**

The repository of all information relating to FRP messages and packets. This file defines all the different sizes, types and flags required, and creates all the data structures needed to build a FRP packet. Data structures are defined for a packet header, a message header, and for each of the different message types.



**frp\_packet.c**

Drawing on the data defined in `frp_packet.h`, `frp_packet.c` defines the specific message and packet handling functions. These are called on to make up individual messages dependent on the message type provided and to build FRP packets out of one or more messages and a packet header. The `frp_send_packet` function, which is mostly made up of code taken from other protocol daemons, also resides here.

**frp\_peer.c**

Functions for the creation, initialisation, handling and support of peers, including the associated VTY commands. Where overlaps occur, code has been lifted from other protocol daemons. Of particular interest are the two security functions, `dohash` and `checksecure`, which are taken directly from the original Knossos implementation of FRP.

**frp\_route.c**

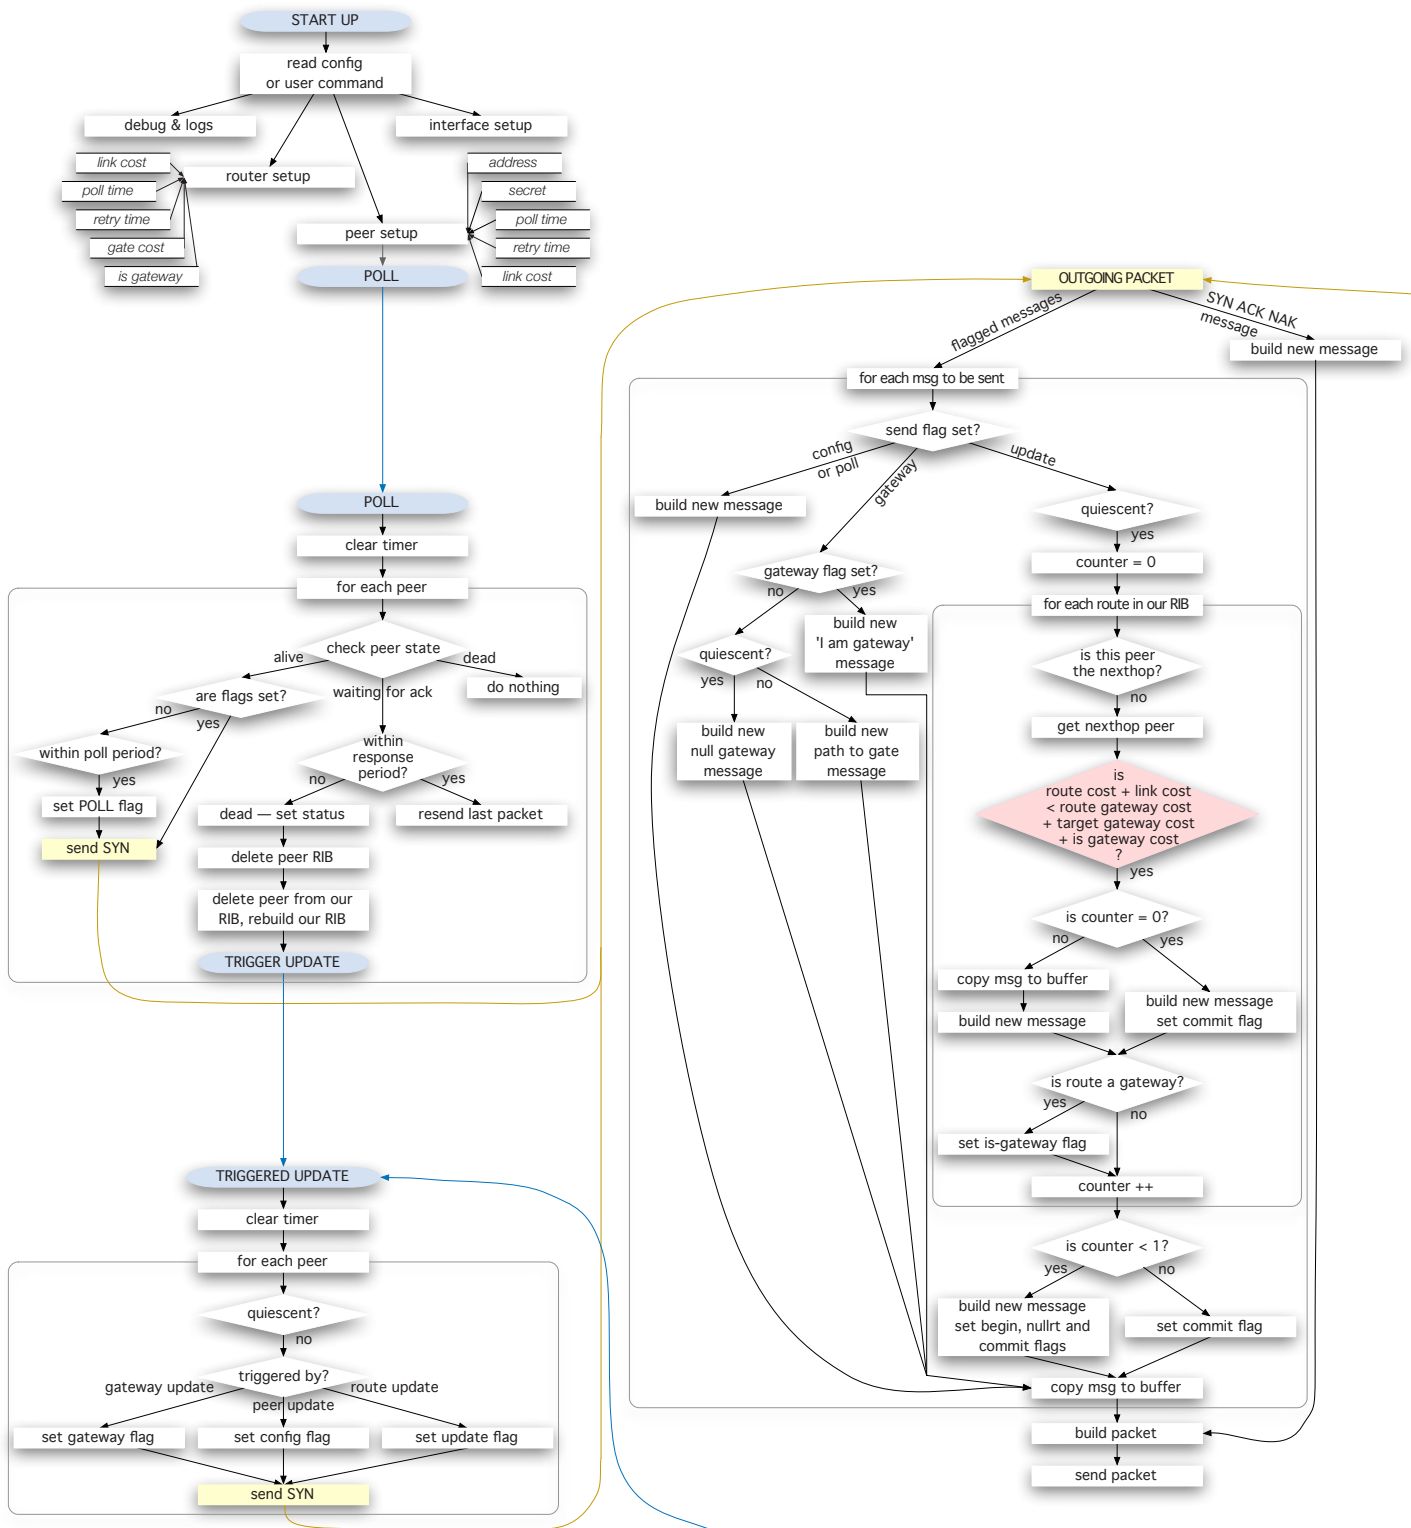
Quagga supplies many of the mechanisms required to handle the FRP routing table. The FRP daemon only needs to set up the necessary data structures and make calls to the relevant library functions. `frp_route.c` holds the functions that address the specific FRP routing mechanisms and apply the FRP decision algorithm.

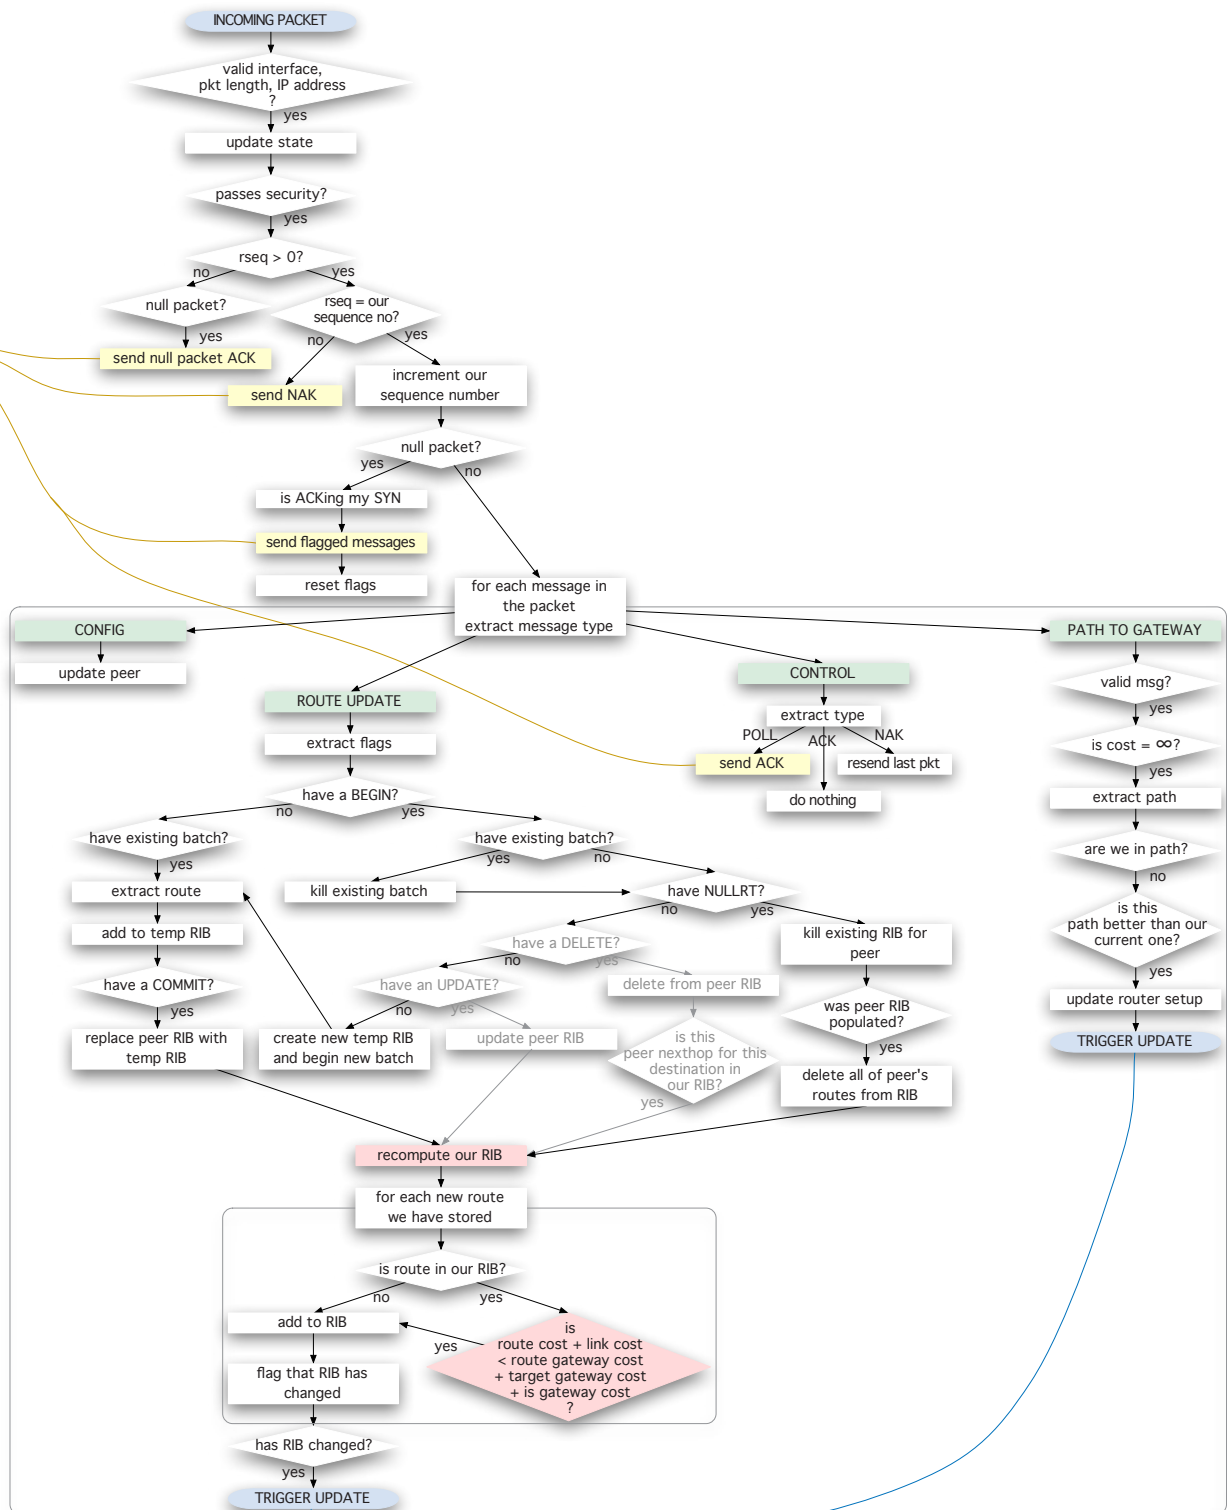
---

## Program flow

Once the empty daemon shell was working correctly, the fringe routing protocol itself could be built into it, necessitating a complete change of direction in the development phase. The Quagga system is a framework providing an API for basic or common router functionality — the mechanics of building the routing table for example. Any FRP specific functionality — deciding what actually goes into the routing table for instance — needed to be written from scratch. The design for this functionality is shown over the page. It shows the program flows necessary for the correct application of the FRP algorithm and associated operational requirements. The implementation of this FRP daemon design is described in the following section.

Figure 4.3  
FRP daemon program flows







---

## 5. Implementing and testing the Quagga FRP daemon

---

This chapter describes in detail the implementation and testing of the fringe routing protocol as a *Quagga* routing daemon. It covers the daemon start up, peers, events, polling and triggered updates. It describes how an incoming packet can trigger the sending of the different types of FRP messages — control, configuration, gateway, and route updates — and describes how route updates force a recomputation of the routing table. Finally, the process of building outgoing messages and packets is explained. The IPv4 version of the daemon is used as the example throughout the section.

### Daemon start up

---

On start up, the FRP daemon, like all the other *Quagga* protocol daemons, is either launched from a startup script or from the Unix command line using the standard Unix launch command plus any of the additional arguments built in to the FRP implementation. For example:

```
sudo ./frpd -d -f /opt/local/etc/quagga/frpd.conf
```

Here, `-d` is specifying that the daemon should be launched in daemon mode and `-f` is providing the path to the configuration file. The start up arguments are implemented in `frp_main.c` and allow customisation of various settings including the VTY port and address, and the user and group names.

```
struct option longopts[] =
{
    { "daemon",      no_argument,      NULL, 'd'},
    { "config_file", required_argument, NULL, 'f'},
    { "pid_file",    required_argument, NULL, 'i'},
    { "dryrun",      no_argument,      NULL, 'C'},
    { "help",        no_argument,      NULL, 'h'},
    { "vty_addr",    required_argument, NULL, 'A'},
    { "vty_port",    required_argument, NULL, 'P'},
    { "retain",      no_argument,      NULL, 'r'},
    { "user",        required_argument, NULL, 'u'},
    { "group",       required_argument, NULL, 'g'},
    { "version",     no_argument,      NULL, 'v'},
    { 0 }
};
```

The header file `frpd.h` contains most of the daemon's external function prototypes and global variables. It also specifies a number of defines to hold various initialisation values required to set up the daemon as a FRP router.

```
/* FRP version number. */
#define FRP_VERSION 1
/* Default config file name */
#define FRP_DEFAULT_CONFIG "frpd.conf"
/* FRP ports */
#define FRP_PORT_DEFAULT 343
#define FRP_VTY_PORT 2609
// frp router defaults
#define FRP_DEFAULT_COST 1
#define FRP_DEFAULT_POLL 5
#define FRP_DEFAULT_RETRY 1
#define FRP_INFINITY 0xffff
// set to determine how many polls to wait before declaring a peer
'dead'
#define FRP_PEER_DEAD 5
```

The version number is for compatibility with other *Quagga* daemons but is not really used in FRP at present. The name of the configuration file follows existing *Quagga* convention and is stored in the default directory set by *Quagga*. This is different on each platform and is configurable in *Quagga* itself.

The FRP port and the VTY port are necessary settings. 343 is the port number chosen by Don and used in his implementation of FRP. *Quagga* is currently using ports 2600 to 2608 for Zebra and other protocol demons so port 2609 was chosen for FRP. It needs to be noted that neither of these port choices are official IANA assigned numbers.

There are a number of parameters required to set up a FRP router and these are stored in the data structure `frp` defined in `frpd.h`.

```
struct frp
{ int version;
  /* frp output buffer */
  struct stream* obuf;
  /* frp socket */
  int sock;
  // secret
  const char* secret;
  // cost of the link (16 bits)
  u_short cost;
  // poll time (16 bits)
  u_short poll;
  // retry time (16 bits)
  u_short retry;
  /* frp routing information base (linked list) */
  struct route_table* rib;
  /* frp only static routing information (linked list) */
  struct route_table* routes;
```

```

/* frp neighbors (linked list) */
struct route_table*      neighbors;
/* gateway types
#define FRP_GATEWAY_ALWAYS      0
#define FRP_GATEWAY_YES        1
#define FRP_GATEWAY_NO         2
// is this peer a gateway?
enum flag                is_gateway_flag;
/* frp gateway path (linked list) */
struct list*             gateway_path;
// which peer is the current next hop to the current gateway?
// - use this to trigger update if this peer changes its path to
gateway
struct frp_peer*         gateway_nexthop;
// number of hops to gateway
int                      gateway_cost;
/* frp threads */
struct thread*           t_read;
struct thread*           t_poll;
struct thread*           t_update;
struct thread*           t_update_interval;
int                      update_trigger;
/* timer values. */
unsigned long            update_time;
unsigned long            timeout_time;
unsigned long            garbage_time;
};

```

This data structure stores the basic information the router needs to interact with peers; the secret to be exchanged; the poll interval for contacting peers to check they are still there and the retry interval for contacting peers that have not responded for a period of time. It stores gateway information; if this router is a gateway, the current path to the currently designated gateway (if there is one), the nexthop peer in that path and the cost to reach it. The threads and the timers are used to handle events. Note that `struct route_table* routes` is required so that Zebra can handle static routes and that `struct route_table* neighbors` has been left in at present because it is not clear whether Zebra also requires it to be here, despite the fact that the FRP daemon is handling peers in a different manner.

The main function in `frp_main.c` initialises the daemon and sets many of its defaults, setting up the basic router configuration by reading and executing the commands written in the configuration files. The basic, non-FRP router configuration requirements are set via the Zebra daemon.

For example:

```

! Zebra configuration
hostname sapphire
password zebra
enable password zebra
log file zebra.log
!
interface en0
 ip address 10.0.1.20/24

```

At the barest minimum, the FRP configuration file must contain a host name for the router and a password for accesses that host. Zebra typically encrypts passwords but the ones shown here are in clear text. A destination for logging output (typically stdout) is likewise usually specified. To set up the FRP router, specify the network range(s) it can route traffic to, state whether it is a gateway (default is no) and set the gateway cost to 1 if it is (default is infinity, or 0xffff if it isn't). State what secret this router will share with its peers. If the linkcost (1), poll (60sec) and retry (60sec) times are different to the defaults, set them here as well. Finally, set up any known peers in the format:

```
neighbor address secret theirsecret
```

For example:

```
! FRPd configuration
hostname sapphire
password zebra
log file frp.log
log stdout
!
router frp
  network 10.0.1.0/24
  network en0
  gateway yes
  gatecost 0
  secret sapphire
  cost 2
  poll 30
  retry 30
  neighbor 10.0.1.50 secret artemis
  neighbor 10.0.1.60 secret emerald
```

Note that the American spelling of 'neighbor' has been used to stay consistent with Quagga usage.

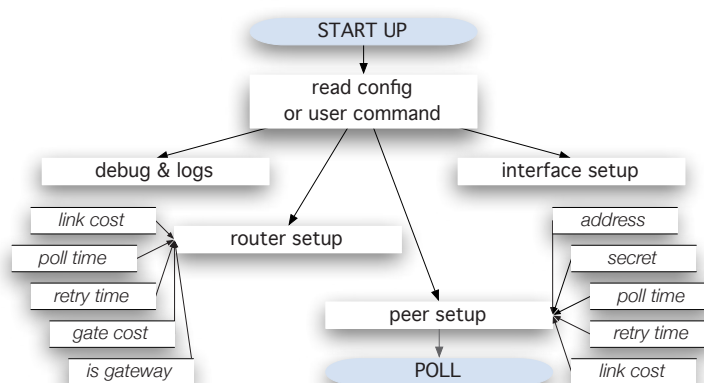


Figure 5.1  
Start up flow



The main function also calls the functions that initialise the FRP `zclient` and set up the Zebra callback functions, as well as installing all the FRP specific commands in the VTY. Debugging is initialised via a call to a function in the `frp_debug.c` file, as is interface initialisation via `frp_interface.c`. In both these cases, the code is essentially the same as all the other *Quagga* protocol daemons, suitably modified for FRP. Peer initialisation, beyond that handled by the configuration file, is simply setting up data structures, including the struct `frp_peer` that holds the list of peers, and VTY commands at this point in the process. Additional information about peers is gathered immediately after start up is complete by triggering a POLL event that works through the peer list contacting each one.

The following commands were modified or added to the daemon menu system (highlighted in red):

```
frpd(config)#
  banner      Set banner string
  debug       Debugging functions (see also 'undebug')
  enable      Modify enable password parameters
  end         End current mode and change to enable mode.
  exit        Exit current mode and down to previous mode
  help        Description of the interactive help system
  hostname    Set system's network name
  line        Configure a terminal line
  list        Print command list
  log         Logging control
  no          Negate a command or set its defaults
  password    Assign the terminal connection password
  quit        Exit current mode and down to previous mode
  router      Enable a routing process
  service     Set up miscellaneous service
  show        Show running system information
  write       Write running configuration to memory, network, or terminal
frpd(config)# router frp
frpd(config-router)#
  cost        Cost of link
  end         End current mode and change to enable mode.
  exit        Exit current mode and down to previous mode
  gateway     Is this router a FRP gateway? (yes/no)
  help        Description of the interactive help system
  list        Print command list
  neighbor    Specify a neighbor router
  network     Enable routing on an IP network
  no          Negate a command or set its defaults
  poll        Poll / keepalive frequency
  quit        Exit current mode and down to previous mode
  retry       Timeout to failure after acked packet
  secret      Specify secret
  show        Show running system information
  write       Write running configuration to memory, network, or terminal
frpd(config-router)#
```

## Peers

---

Peers (or neighbours) are other FRP routers directly connected to the host router. Therefore peers are known not only by their network address but also by the interface they are connected on. Due to the security of the shared secrets, peers are only ever set up by manual means — there is no automated polling required.

Two pieces of information about a peer are required: the address (which must be matched to an interface) and the secret, which is to be shared for the security hashes during packet exchange. A peer must be specified in either the start up configuration or via the command line, as the secret cannot be entered in any other way. Any additional information is sent by the peer during the first communication exchange and then stored.

A list of all peers known to a host is stored in the linked list

```
struct list* frp_peers = NULL;
```

which is created in `frp_peer.c` and initialised

```
frp_peers = list_new ();
```

in the function `frp_peer_init` in the same file. A prototype

```
extern struct list* frp_peers;
```

is listed in `frpd.h` to ensure there is only one list in existence.

`frp_peers` holds a list of peer records, one for each peer known. These records are data structures containing all the information held on the peer.

```
struct frp_peer
{ /* peer address */
    struct in_addr    address;
    const char*      secret;
    u_short          cost;
    u_short          poll;
    u_short          retry;
    // our current sequence number with this peer
    u_int32_t        lseq;
    // this peer's current sequence number
    u_int32_t        rseq;
    // current path to gateway for this peer
    struct list*      gateway_path;
    // cost of gateway, ie: number of hops in path to gateway
    int              gateway_cost;
    // temporary storage for new incoming routes
    struct list*      rib;
    struct list*      temp_rib;
    // latest packet sent to this peer
    u_char*          packet_latest;
    u_int8_t         packet_latest_length;
    u_int32_t         packet_latest_lseq;
```

```

// latest timer timestamps
time_t      time_latest_packet;
time_t      time_last_heard;
time_t      time_sent_config;
/* timeout thread */
struct thread* t_timeout;
// flags
enum flag    flag_alive;
enum flag    flag_send_syn;
enum flag    flag_send_poll;
enum flag    flag_send_config;
enum flag    flag_send_gateway;
enum flag    flag_send_update;
enum flag    flag_awaiting_ack;
};

```

Many of the variable names make it obvious what information is being stored. Of particular interest is the use of `lseq` and `rseq`. These are the terms used in the initial FRP specification and stand for ‘local sequence’ and ‘remote sequence’. Local always equates to the router whose perspective is current, the host or ‘me’ or ‘us’; and remote equates to the peers of this router, or ‘them’. Thus, `lseq` is ‘our’ sequence and `rseq` is ‘their’ sequence. This perspective shifts as the focus shifts to another router or host. Each host maintains a separate local sequence, or `lseq`, with each individual peer.

The linked list `rib` is the routing table as supplied by the peer. `temp_rib` is the temporary storage used when changes to that routing table are received from the peer. The new routing table is built piece by piece in `temp_rib` and then copied into `rib` when confirmed as complete and correct.

The latest packet sent to a peer is stored so that it can be easily resent if a ‘no acknowledgement’ (NAK) is received indicating that the packet has not arrived at the peer, or if the peer does not acknowledge (ACK) the packet within a specified period of time.

The timers are required to implement polling and keepalives and the timeout thread is part of *Quagga*’s thread management system.

The flags are an enumerated boolean

```

enum flag
{ OFF,
  ON,
};

```

which is turned ON to indicate which message types are ready to be sent the next time a packet is put together, except for `flag_alive` which is turned OFF to indicate that a peer has not been heard from for a certain period of time and subsequently pronounced ‘dead’.

## The FRP daemon in action

On start up, a FRP host works through its list of peers passing the necessary configuration information. The following *FRPsniffer* trace shows a standard configuration exchange between two *Quagga* FRP daemons, 10.0.1.20 and 10.0.1.50. Both peers are quiescent so no route information is exchanged.

20 starts the sequence with the initial handshake by sending a SYN packet header with no message (a null packet), telling 50 what 20's current sequence number is and indicating that is a new conversation by setting the recipient's sequence number to 0. Note the security hash which 20 has created using its secret.

50 continues the initial handshake by sending a null ACK packet header saying "acking your 1, my next sequence number is 2". Note that for 50 to reach the point of sending this response, the initial packet security hash must have been successfully reversed using 50's stored copy of 20's secret. This new packet is encrypted with 50's secret.

The initial handshake is now complete so 20 can send the configuration message that triggered this conversation. First comes the packet header — note the incremented sequence number; "acking your 2, my next sequence number is 2". Then comes the config message (type 0x41) providing the cost, poll and retry values, plus 20's id which in this case is its address.

50 receives the message and discovers that their 'send config' flag is set for 20. The packet is constructed — header plus config message containing data — and sent.

```
FRP PACKET HEADER
  source address: 10.0.1.20 (a.0.1.14)
  destination address: 10.0.1.50 (a.0.1.32)
  security hash: ??;??
  sender's seq no: 1 (0x1)
  recipient's ack no: 0 (0x0)
--SYN Packet
```

Waiting for incoming FRP packet

```
FRP PACKET HEADER
  source address: 10.0.1.50 (a.0.1.32)
  destination address: 10.0.1.20 (a.0.1.14)
  security hash: ??'L?q
  sender's seq no: 2 (0x2)
  recipient's ack no: 1 (0x1)
  Null Message: 0x0
--ACK Packet
```

Waiting for incoming FRP packet

```
FRP PACKET HEADER
  source address: 10.0.1.20 (a.0.1.14)
  destination address: 10.0.1.50 (a.0.1.32)
  security hash: d?????N
  sender's seq no: 2 (0x2)
  recipient's ack no: 2 (0x2)
```

```
IPV4 Config message: 0x41
  cost: 3
  poll: 60
  retry: 60
  router-id: 10.0.1.20
```

Waiting for incoming FRP packet

```
FRP PACKET HEADER
  source address: 10.0.1.50 (a.0.1.32)
  destination address: 10.0.1.20 (a.0.1.14)
  security hash: q>??r???
  sender's seq no: 3 (0x3)
  recipient's ack no: 2 (0x2)
```

```
IPV4 Config message: 0x41
  cost: 10
  poll: 60
  retry: 60
  router-id: 10.0.1.50
```

Waiting for incoming FRP packet

20's response contains a packet header and a control message of type 'ack' indicating that, as far as 20 is concerned, this conversation is over.

```
FRP PACKET HEADER
  source address: 10.0.1.20 (a.0.1.14)
 destination address: 10.0.1.50 (a.0.1.32)
  security hash: ???H?
 sender's seq no: 3 (0x3)
 recipient's ack no: 3 (0x3)

Control message: 0x1
--Control ACK, param = 0
```

## Events

FRP's event handler, function `frp_event` in `frpd.c`, recognises three different events: a poll, an update triggered by a change to the routing table and an incoming packet. These are defined in `frpd.h` as an enumerated list.

```
enum frp_event
{ FRP_EVENT_INCOMING,
  FRP_EVENT_UPDATE,
  FRP_EVENT_POLL,
};
```

Outgoing packets are not handled as an `frp_event` because they are created as required in response to one of the three listed events.

As in the other *Quagga* protocol daemons, a call to each event is triggered during initialisation and an event thread is created and packaged up with the appropriate function code and data. Subsequent event calls create a new thread for the next event of that type to use. This process starts in `frp_event` where, to use `FRP_EVENT_INCOMING` as an example, a new read thread is created and given all the necessary data required to run independently — the thread that spawned it, the code it will run and the current socket.

```
void frp_event (enum frp_event event, int sock)
{ switch (event)
  { case FRP_EVENT_INCOMING:
    // create a new read thread and tell it to run frp_incoming_
    // packet()
    frp->t_read = thread_add_read (master, frp_incoming_packet,
    NULL, sock);

    break;
  case FRP_EVENT_UPDATE:
    ...
  case FRP_EVENT_POLL:
    ...
  default:
    ...
  }
}
```

Consequently, `frp_incoming_packet` is required to be an independent piece of code, the only parameter taken being the thread. The first thing these thread called functions must do is create a new thread which immediately blocks and sits waiting for the next event of the correct type to trigger, at which point the process starts again. The original thread meanwhile executes the code in the function passed to it, thus completing the current event.

```
int frp_incoming_packet (struct thread* t)
{  variables etc ...

    /* fetch socket then register myself */
    sock = THREAD_FD (t);
    frp->t_read = NULL;
    /* add myself to the next event */
    frp_event (FRP_EVENT_INCOMING, sock);

    // read the packet from the socket
    ...
}
```

## Polling

---

The FRP daemon keeps a single poll timer that activates on a regular basis triggering a POLL event. Each time this happens, the list of peers is traversed and each one individually checked to see if a POLL message needs to be sent. The criteria to be met are that the peer is alive, it is not waiting for other messages to be received or sent, and it has been silent for more than the individual poll period for that peer.

The original FRP specification states

*Keep traffic down to keepalives (5 sec?) if no updates.*

and

*Poll / keepalive frequency in tenths of seconds. Minimise received value with local configuration.*

The current is 60 seconds as this is the most convenient frequency for debugging purposes. Quagga uses seconds so this FRP daemon also uses seconds. At present, the implementation only uses the host router's poll period, this needs to be changed to compare the two potentially different values and choose the lowest one as per the specification. The modification requires the router to store a little more state on each individual peer than happens at the moment. Both the router and the peer poll settings are customisable.

Polls are triggered within the FRP event handler. If there is already a POLL event active, it is cancelled and a new POLL thread is created.

```
case FRP_EVENT_POLL:
    if (frp->t_poll)
    { thread_cancel (frp->t_poll);
      frp->t_poll = NULL;
    }
    frp->t_poll = thread_add_timer (master, frp_poll_peers, NULL,
                                   (unsigned long)frp->poll);
    break;
```

When a POLL event is issued, the function `frp_poll_peers` in `frpd.c` is executed. First a new POLL thread is created and set to wait for the next POLL to occur. Then the poll timer is cleared and reset. The list of peers in `frp_peers` is iterated through and the state of each peer is checked. This information is held in two flags, `flag_alive` and `flag_awaiting_ack`. Flags are boolean types that can be either ON or OFF.

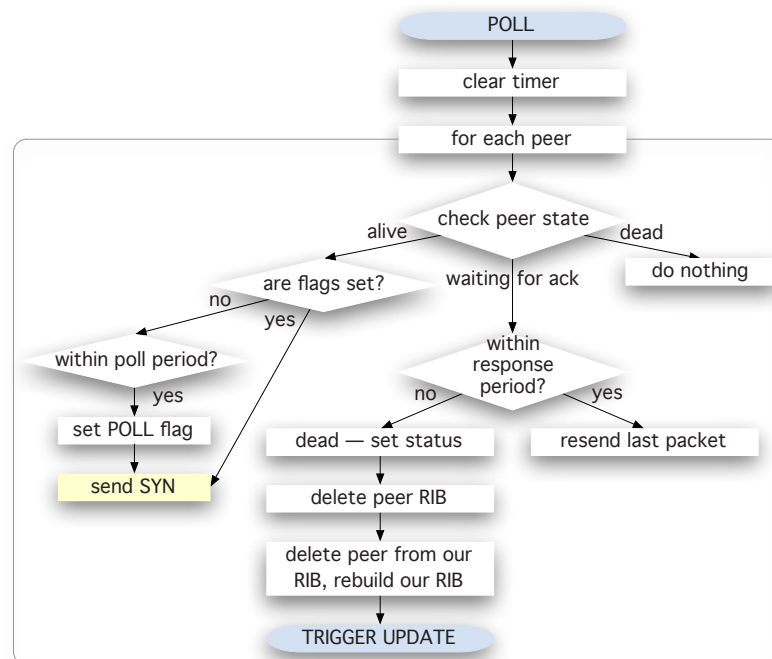


Figure 5.2  
Polling flow

If `flag_alive` is OFF, then the peer has not responded for the specified number of retries and has been declared dead. So do nothing.

If `flag_alive` is ON and `flag_awaiting_ack` is OFF, then the peer has simply not had any information to share within the last poll period and this router's poll timer has triggered before the peer's has. The first thing to do is determine if any outgoing messages are

waiting to be sent by checking the `flag_send_xxx` flags. If none of these are currently set, check that the poll period for this peer has expired and, if it has, turn on `flag_send_poll`. Finally, send a SYN packet to start the outgoing packet sequence. The outgoing packet will be made up of messages corresponding to any of the flags that were set.

If `flag_alive` is ON and `flag_awaiting_ack` is ON, check to see if the retry timer is still within the retry time. If it is, resend the last packet — which will be the one that the outstanding ACK is for. To facilitate this, the last sent packet is stored in `packet_latest`. Send it exactly as is with the same sequence number as previously.

If the retry time has over-run the specified retry period, then declare the peer dead and clean up. Leave the peer record in the peers list (it may return) but set `flag_alive` to OFF. Delete the peer's routing table — if the peer does return, it will send a new set of routes at that point. Work through the router's RIB and delete every route that uses the dead peer as a nexthop, then re-compute the table. Finally, trigger an UPDATE event to indicate that the routing table has changed.

### The FRP daemon in action

The following *FRPsniffer* trace shows a standard poll exchange between two Quagga FRP daemons, 10.0.1.20 and 10.0.1.50.

50 starts the sequence with the initial handshake by sending a SYN packet.

```
FRP PACKET HEADER
  source address: 10.0.1.50 (a.0.1.32)
  destination address: 10.0.1.20 (a.0.1.14)
  security hash: I??? {yk
  sender's seq no: 4 (0x4)
  recipient's ack no: 0 (0x0)
--SYN Packet
```

Waiting for incoming FRP packet

20 responds with a null packet ACK.

```
FRP PACKET HEADER
  source address: 10.0.1.20 (a.0.1.14)
  destination address: 10.0.1.50 (a.0.1.32)
  security hash: \???J??x
  sender's seq no: 4 (0x4)
  recipient's ack no: 4 (0x4)
  Null Message: 0x0
--ACK Packet
```

Waiting for incoming FRP packet

The initial handshake is now complete so 50 can send the poll message that triggered this conversation. First comes the packet header, followed by the control message of type 'poll'.

```
FRP PACKET HEADER
  source address: 10.0.1.50 (a.0.1.32)
  destination address: 10.0.1.20 (a.0.1.14)
  security hash: q?G??N??
  sender's seq no: 5 (0x5)
  recipient's ack no: 4 (0x4)
```

```
Control message: 0x1
--Poll, param = 0
```



The response from 20 contains a packet header and a control message of type 'ack' indicating that, as far as 20 is concerned, this conversation is over.

Waiting for incoming FRP packet

#### FRP PACKET HEADER

source address: 10.0.1.20 (a.0.1.14)  
 destination address: 10.0.1.50 (a.0.1.32)  
 security hash: ?0?)??  
 sender's seq no: 5 (0x5)  
 recipient's ack no: 5 (0x5)

Control message: 0x1

--Control ACK, param = 0

Some times when a host polls a peer, that peer has disappeared for some reason. The configuration information provided to each host at start up specifies how long that host will wait before re-polling and how many polls will be sent before declaring the peer 'dead'. The following trace shows 10.0.1.50 polling 10.0.1.20 after 20 has been shut down.

50's poll timer triggers and so sends a SYN to 20. A copy of the packet is placed in the latest packet storage for peer 20. 50 waits.

#### FRP PACKET HEADER

source address: 10.0.1.50 (a.0.1.32)  
 destination address: 10.0.1.20 (a.0.1.14)  
 security hash: &4[1??  
 sender's seq no: 24 (0x18)  
 recipient's ack no: 0 (0x0)

--SYN Packet

Waiting for incoming FRP packet

Waiting for incoming FRP packet

Waiting for incoming FRP packet

Waiting for incoming FRP packet

Waiting for incoming FRP packet

Waiting for incoming FRP packet

Waiting for incoming FRP packet

No reply is received from 20 within the retry time so 50 retrieves and sends the same packet again. 50 waits

#### FRP PACKET HEADER

source address: 10.0.1.50 (a.0.1.32)  
 destination address: 10.0.1.20 (a.0.1.14)  
 security hash: &4[1??  
 sender's seq no: 24 (0x18)  
 recipient's ack no: 0 (0x0)

--SYN Packet

Waiting for incoming FRP packet

Waiting for incoming FRP packet

Waiting for incoming FRP packet

Waiting for incoming FRP packet

Waiting for incoming FRP packet

Waiting for incoming FRP packet

Waiting for incoming FRP packet

No reply is received from 20 within the retry time so 50 retrieves and sends the same packet a third time. 50 waits.

In this example, the number of retries was set to 3, so when no reply is received from 20 within the retry time, 50 declares 20 to be dead.

Time passes and 20 is brought back up again.

20 starts the usual configuration handshake with a SYN packet.

50 ACKs.

```
FRP PACKET HEADER
    source address: 10.0.1.50 (a.0.1.32)
destination address: 10.0.1.20 (a.0.1.14)
    security hash: 84[1??]
    sender's seq no: 24 (0x18)
    recipient's ack no: 0 (0x0)
--SYN Packet
```

Waiting for incoming FRP packet
Waiting for incoming FRP packet
Waiting for incoming FRP packet
Waiting for incoming FRP packet
Waiting for incoming FRP packet
Waiting for incoming FRP packet
Waiting for incoming FRP packet
Waiting for incoming FRP packet
Waiting for incoming FRP packet
Waiting for incoming FRP packet
Waiting for incoming FRP packet
Waiting for incoming FRP packet
Waiting for incoming FRP packet
Waiting for incoming FRP packet
Waiting for incoming FRP packet
Waiting for incoming FRP packet
Waiting for incoming FRP packet
Waiting for incoming FRP packet
Waiting for incoming FRP packet
Waiting for incoming FRP packet
Waiting for incoming FRP packet
20 Restarted
Waiting for incoming FRP packet
Waiting for incoming FRP packet

```
FRP PACKET HEADER
    source address: 10.0.1.20 (a.0.1.14)
destination address: 10.0.1.50 (a.0.1.32)
    security hash: ??;??
    sender's seq no: 1 (0x1)
    recipient's ack no: 0 (0x0)
--SYN Packet
```

Waiting for incoming FRP packet

```
FRP PACKET HEADER
  source address: 10.0.1.50 (a.0.1.32)
destination address: 10.0.1.20 (a.0.1.14)
  security hash: ?NT:Q??
  sender's seq no: 25 (0x19)
  recipient's ack no: 1 (0x1)
    Null Message: 0x0
--ACK Packet
```

Waiting for incoming FRP packet

20 sends a configuration message.

```
FRP PACKET HEADER
  source address: 10.0.1.20 (a.0.1.14)
destination address: 10.0.1.50 (a.0.1.32)
  security hash: ?J??c?
  sender's seq no: 2 (0x2)
recipient's ack no: 25 (0x19)
```

```
IPV4 Config message: 0x41
  cost: 3
  poll: 60
  retry: 60
router-id: 10.0.1.20
```

Waiting for incoming FRP packet

50 responds in kind with its own configuration message

```
FRP MESSAGE
  source address: 10.0.1.50 (a.0.1.32)
destination address: 10.0.1.20 (a.0.1.14)
  security hash: S80?f
  sender's seq no: 26 (0x1a)
recipient's ack no: 2 (0x2)
```

```
IPV4 Config message: 0x41
  cost: 10
  poll: 60
  retry: 60
router-id: 10.0.1.50
```

Waiting for incoming FRP packet

and 20 closes the conversation with a control ACK.

```
FRP PACKET HEADER
  source address: 10.0.1.20 (a.0.1.14)
destination address: 10.0.1.50 (a.0.1.32)
  security hash: $???/?
  sender's seq no: 3 (0x3)
recipient's ack no: 26 (0x1a)
```

```
Control message: 0x1
--Control ACK, param = 0
```

## Triggered updates

An UPDATE event is triggered within the FRP event handler to indicate that a change has been made to the hosts routing table, path to gateway, or configuration and this new information needs to be propagated out to Zebra and to the host's peers. Quagga uses an update timer to try to batch updates to a certain extent — “after a triggered update is sent, a timer should be set for a random interval between 1 and 5 seconds. If other changes that would trigger updates occur before the timer expires, a single update is triggered when the timer expires” [35] and this has been retained in the FRP daemon.

```
case FRP_EVENT_UPDATE:
    if (frp->t_update_interval)
        frp->update_trigger = 1;
    else if (! frp->t_update)
        frp->t_update = thread_add_event (master, frp_update_peers,
                                          NULL, 0);
    break;
```




Figure 5.3  
Triggered updates flow

When a UPDATE event is issued, the function `frp_update_peers` in `frpd.c` is executed. First a new UPDATE thread is created and set to wait for the next UPDATE to occur. Then the update timer is cleared and reset. The list of peers in `frp_peers` is iterated through, checking to see if the peer is quiescent. If it is, then no path to gateway currently exists for that peer and consequently, no update is needed. Otherwise, the correct flag is set for the type of update triggered, Quagga waits it's random period, then a SYN packet is created and sent.

## Incoming packets

Dealing with an incoming packet is the most complex part of the FRP demon, mainly because of the many decisions to be made at each step. It is within this sequence that the different types of message are dealt with and the correct responding message sent in reply.

### FRP packets

Note that the terms ‘packet’ and ‘message’ have different meanings in relation to FRP communications. A message is a single announcement of a set of data in a rigid, pre-specified format. A packet is the unit of transfer of data across the network and consists of 0 or more messages plus a packet header.

Quagga uses a union to store the contents of an incoming packet in a buffer. When combined with pointers to mark how much of the data has been extracted, this provides the flexibility to use the one buffer type for each incoming packet and to decide what underlying structure to give those contents based on what is extracted.

```
/* buffer to store frp data */
union frp_pkt_buf
{ char buf[FRP_PKT_MAXSIZE];
  struct frp_pkt_hdr          frp_pkt_hdr;
  struct frp_msg_hdr          frp_msg_hdr;
  struct frp_msg_control      frp_msg_control;
  struct frp_msg_ipv4config   frp_msg_ipv4config;
  struct frp_msg_ipv4gateway  frp_msg_ipv4gateway;
  struct frp_msg_ipv4update   frp_msg_ipv4update;
};
```

Initially, the contents are treated as a `frp_pkt_hdr`. If dealing with a null packet, the process ends here. Otherwise the first message header is extracted with provides the length type of that message, allowing it to be extracted using the correct underlying structure once again. This process is repeated until the buffer is empty.

### Incoming packet events

The `frp_incoming_packet` function in `frpd.c` is triggered by a packet arriving via the socket set up by the daemon during initialisation. The function code is divided in to three sequential parts, the mechanics of receiving a packet, extracting data and validity and security checking, and determining what type of messages have arrived and where they fit in the current conversation. `frp_incoming_packet` calls on an appropriate separate function (`frp_incoming_xxx_msg`) to handle each of the different types of incoming message. The diagram below concentrates on the second and third of these three parts.

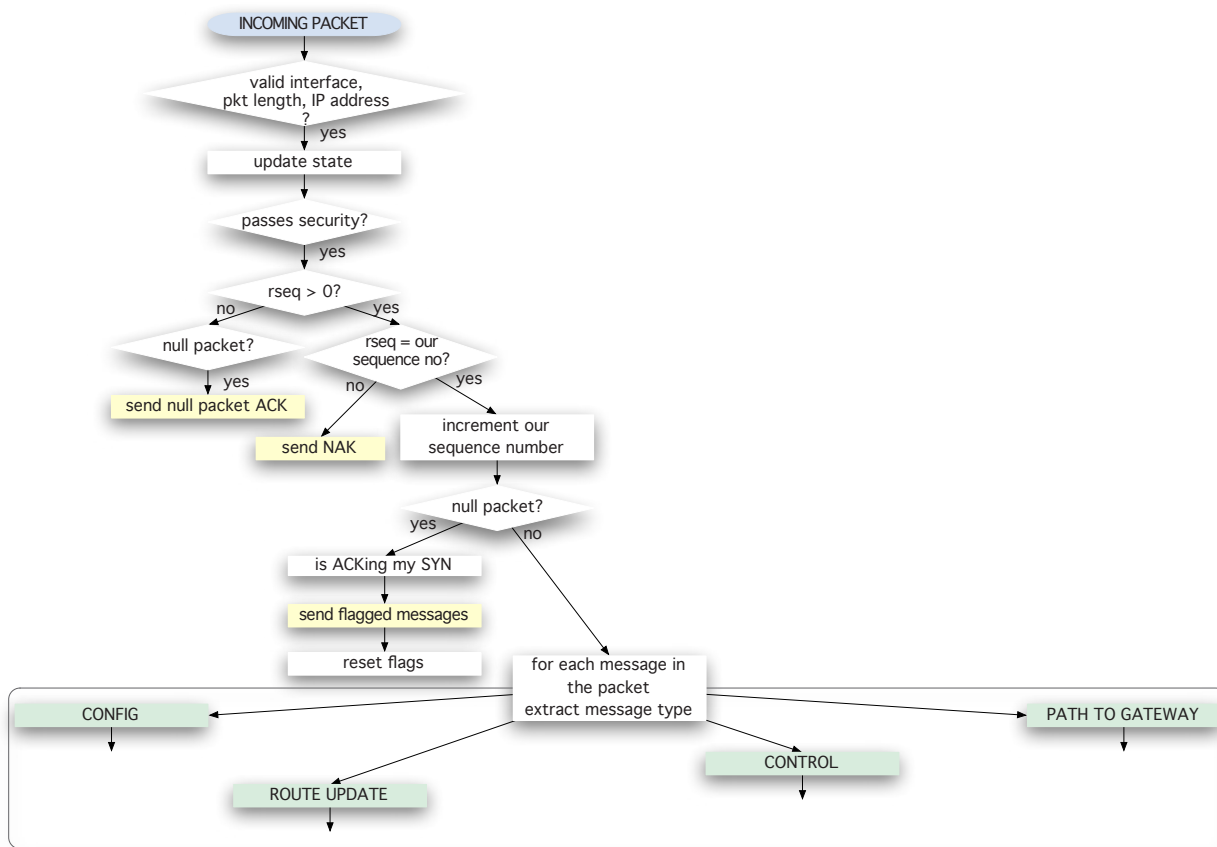


Figure 5.4  
Incoming packet flow

When an INCOMING PACKET event is triggered by the arrival of a packet, a new INCOMING PACKET thread is created and set to wait for the next INCOMING PACKET event to occur. The daemon initialises a buffer of the correct size and reads in the contents of the packet as a continuous stream. It then checks that

- the packet comes from a known interface with a legitimate address
- FRP is running on this interface and the address belongs to a known FRP peer
- the packet length fits within the minimum and maximum size range pre-defined in `frp_packet.h`

The buffer is discarded if any of the checks fail.

Using the pre-defined sizes, types and flags specified in the file `frp_packet.h`, the daemon can now decode the single sequence of data into the correctly sized pieces of information contained in each individual message.

First the packet header

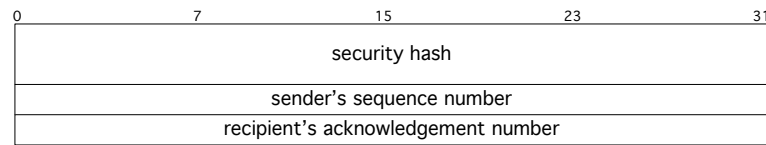


Figure 5.5  
Packet header

```
#define FRP_PKT_HDRSIZE      16          // (128 bits)
#define FRP_PKT_MINSIZE      16          // (128 bits)
#define FRP_PKT_MAXSIZE      1400       // as specified by Don

struct frp_pkt_hdr
{
    u_int8_t      hash[8];
    u_int32_t      sendSeq;
    u_int32_t      recipAck;
};
```


is extracted and the peer address confirmed. The peer `flag_alive` is turned ON and the `time_last_heard` flag time stamped. The function `checksecure`, which has been lifted directly from the original implementation of FRP, is run against the security hash using extracted peer address, the host router's local address and the secret stored in `frp_peers` for this particular peer. `checksecure` must succeed for processing of the packet to continue.

The peer's sequence number (the sender's sequence number in this instance) is checked to see whether this is the start of a new conversation or whether this packet is the next instalment in an on-going one. If it is equal to 0 and the entire packet consists only of the packet header with no messages attached (a null packet), then the peer is initiating a new conversation and the outgoing packet is a null packet ACK. Otherwise, check the peer's acknowledgement of the host's current sequence number (the recipient's acknowledgement number in this instance) against the actual sequence number stored by the host for that peer to see if they match. If they do not, then there has been a problem in the exchange of messages — a lost packet perhaps — and the host indicates this to the peer by sending a control NAK packet, essentially asking the peer to start the conversation again.

It has now been established that this is part of an on-going exchange, that the initial handshake is complete, and that the packet and the sequence numbers are valid. If the packet is a null packet, then the peer is ACKing a SYN sent by this host to start a new conversation. As communication has been established, the out-going packet will contain all flagged messages for that peer.

The only remaining possibility is that the packet contains one or messages in an on-going conversation, so the final section of the `incoming_packet` function contains a `while` loop that iterates to the end of the buffer holding the packet data extracting individual messages.

Within the `while`, the message header is extracted



*Figure 5.6*  
*Message header*

```
struct frp_msg_hdr
{ u_int8_t      length;
  u_int8_t      type;
};

// frp message types
#define FRP_MSG_CONTROL      0x01
#define FRP_MSG_IPV4CONFIG  0x41
#define FRP_MSG_IPV4GATEWAY 0x42
#define FRP_MSG_IPV4UPDATE  0x43
```

and the length checked, allowing the full message to be pulled from the buffer and stored. Pointers keep track of the start of the buffer and the current position within it. The type is then used to consult a `switch` statement which ensures the message is passed to the correct message handling function.



## Control messages




Figure 5.7  
Control message

```

struct frp_msg_control
{ struct frp_msg_hdr    msg_hdr;
  u_int8_t              type;
  u_int8_t              param;
};

// frp control types
#define FRP_CTRL_POLL  1
#define FRP_CTRL_ACK   2
#define FRP_CTRL_NAK   3

```

Control messages are used to indicate that a has reached it's natural end (type 2, ACK), that there has been a problem in the exchange of messages so the conversation is being prematurely terminated and should begin again (type 3, NAK), and to check that a peer who has not been heard from for a period of time is still alive (type 1, POLL).




Figure 5.8  
Control flow

A control ACK is different to a null ACK. The null ACK occurs when a router is responding to an initial SYN with an empty packet (header only) accepting the communication and establishing sequence numbers.



Figure 5.9  
Null ACK

The control ACK follows a successful exchange of message(s) specifying that the conversation is at an end and no further communication is expected.



Figure 5.10  
Control ACK

Consequently, the correct response to a control ACK is to do nothing except clear the `awaiting_ack` flag for that peer.

A control NAK is used when the sequence number the peer acknowledges is not the same as the one the host knows was the latest one sent. The most likely cause of this is lost or delayed messages. In response to a control NAK, a host retrieves the last message sent to that peer, strips the packet header from it. A new header is generated with a new sequence number and the security hash is re-generated before the packet is sent again.

A POLL is received because the sender wishes to know that the recipient is still alive so all that is required in response is a control ACK.

### The FRP daemon in action

The standard poll exchange illustrates the use of control messages and the difference between null packet ACKs and control ACKs.

The second packet in the trace below shows 10.0.1.20 responding to a SYN with a null packet ACK.

```

FRP PACKET HEADER
  source address: 10.0.1.50 (a.0.1.32)
  destination address: 10.0.1.20 (a.0.1.14)
  security hash: I??? {yk
  sender's seq no: 4 (0x4)
  recipient's ack no: 0 (0x0)
--SYN Packet

Waiting for incoming FRP packet

FRP PACKET HEADER
  source address: 10.0.1.20 (a.0.1.14)
  destination address: 10.0.1.50 (a.0.1.32)
  security hash: \???J??x
  sender's seq no: 4 (0x4)
  recipient's ack no: 4 (0x4)
  Null Message: 0x0
--ACK Packet

Waiting for incoming FRP packet
  
```

The third packet shows 10.0.1.50 using a control message to send a poll

and in the fourth 20 responds with a control message containing an ACK.

```
FRP PACKET HEADER
  source address: 10.0.1.50 (a.0.1.32)
 destination address: 10.0.1.20 (a.0.1.14)
  security hash: q?G??N??
 sender's seq no: 5 (0x5)
 recipient's ack no: 4 (0x4)
```

```
Control message: 0x1
--Poll, param = 0
```

Waiting for incoming FRP packet

```
FRP PACKET HEADER
  source address: 10.0.1.20 (a.0.1.14)
 destination address: 10.0.1.50 (a.0.1.32)
  security hash: ?0?)??
 sender's seq no: 5 (0x5)
 recipient's ack no: 5 (0x5)
```

```
Control message: 0x1
--Control ACK, param = 0
```

## Configuration messages




Figure 5.11  
Configuration message

```
struct frp_msg_ipv4config
{ struct frp_msg_hdr      msg_hdr;
  u_short                 cost;
  u_short                 poll;
  u_short                 retry;
  struct in_addr           id;
};
```

The easiest of the message types: copy the new cost, poll, retry and id values into the correct places in `frp_peers`.




Figure 5.12  
Configuration flow

## Path to gateway messages

Gateway path messages are used to exchange gateway routes among peers. The decision making tree is relatively straightforward but the messages have a unique complication in that it is unknown how many addresses will make up the path. Consequently the message length must be used to iterate through the addresses, storing each one in a list. The only other piece of information that needs extracting and storing is the cost of the path from the peer to the gateway. If this is infinity ( $0xffff$ ), then the peer has no gateway and is quiescent.




Figure 5.13  
Path to gateway message

```
struct frp_msg_ipv4gateway
{ struct frp_msg_hdr    msg_hdr;
  u_short               cost;
  struct in_addr        path[62];
};
```

First the message is checked to ensure that length falls within the maximum (252 bytes) and minimum (8 bytes) message size, and that the length of the path section is a multiple of 4 bytes. If it fails either of these tests, the peer's gateway cost is set to infinity ( $0xffff$ ) and the message is discarded.




Figure 5.14  
Path to gateway flow

Next the gateway cost is examined. If it is infinity (0xffff), then the peer record is updated and the message is discarded. Otherwise, the path is extracted stored in a linked list, which in turn is stored in `frp_peers`.

Finally, check to see if the host's local address is in the path. If it is not, add the cost of the link between the peer and the host is to the gateway cost just provided by the peer. This is matched against the cost of the current gateway for the host. If the new path is lower, and therefore 'better', update the router's data so that the `gateway_cost` is this peer's new gateway cost plus the link cost, `gateway_nexthop` is this peer, and `gateway_path` is this peer's new path. Iterate through the list of peers and set the `send_gateway` flag to ON and trigger an UPDATE event to send out new path to gateway messages to all peers.

### The FRP daemon in action

This trace shows the conversation between 10.0.1.20 and 10.0.1.50 once 20 is told that it is now a gateway.

Following the standard SYN/ACK exchange,

```
FRP PACKET
  source address: 10.0.1.20 (a.0.1.14)
  destination address: 10.0.1.50 (a.0.1.32)
  security hash: ??%W?J
  sender's seq no: 10 (0xa)
  recipient's ack no: 0 (0x0)
--SYN Packet
```

Waiting for incoming FRP message

```
FRP PACKET
  source address: 10.0.1.50 (a.0.1.32)
  destination address: 10.0.1.20 (a.0.1.14)
  security hash:
  sender's seq no: 10 (0xa)
  recipient's ack no: 10 (0xa)
  Null Message: 0x0
--ACK Packet
```

Waiting for incoming FRP message

```
FRP PACKET
  source address: 10.0.1.20 (a.0.1.14)
  destination address: 10.0.1.50 (a.0.1.32)
  security hash: !S<0Y?
  sender's seq no: 11 (0xb)
  recipient's ack no: 10 (0xa)
```

```
IPV4 path to gateway message: 0x42
  cost: 0
  # of nodes in path: 1
  path[0]: 10.0.1.20
```

Waiting for incoming FRP message

20 builds a path to gateway message, setting the cost to 0 as 20 is the gateway and the path length to 1 as 20's address is the only address in the path.

50 ACKs receipt.

```
FRP PACKET
  source address: 10.0.1.50 (a.0.1.32)
  destination address: 10.0.1.20 (a.0.1.14)
  security hash: ?MD?EÜ?
  sender's seq no: 11 (0xb)
  recipient's ack no: 11 (0xb)

Control Message: 0x1
--Control ACK, param = 0
```

## Route update messages

Route updates are the most intricate of the messages as they potentially trigger changes to the routing table and thus bring the FRP decision algorithm into play. They also have the added complexity of the message flag system that handles batching of multiple route updates over multiple messages, and potentially multiple packets.




Figure 5.15  
Route update message

```
struct frp_msg_ipv4update
{ struct frp_msg_hdr      msg_hdr;
  u_int8_t                flags;
  u_int8_t                length;
  u_short                 routecost;
  u_short                 gatecost;
  struct in_addr           prefix;
};

// frp update flags
#define FRP_FLAG_BEGIN      0x01
#define FRP_FLAG_COMMIT    0x02
#define FRP_FLAG_NULLRT    0x04
#define FRP_FLAG_UPDATE    0x08
#define FRP_FLAG_DELETE    0x10
#define FRP_FLAG_GATEWAY   0x80
```

The important piece of information is the update type flag, which specifies where the current route update lies in the potential chain of messages. The six flags, defined in `frp_packet.h`, are stored as single pre-specified bits in a byte of memory and a message may have more than one flag set. A bitwise AND is used to extract each of the flagged bits.

0x01		FRP_FLAG_BEGIN
0x02		FRP_FLAG_COMMIT
0x04		FRP_FLAG_NULLRT
0x08		FRP_FLAG_UPDATE
0x10		FRP_FLAG_DELETE
0x20		
0x40		
0x80		FRP_FLAG_GATEWAY

Figure 5.16  
Route update message flags

Each FRP router shares all the routes it knows about that pass the FRP algorithm test as a batch of update messages. There should not be a great many of them as FRP explicitly tries ‘to keep routing table sizes to a minimum’ [57]. The first message in the batch has the BEGIN flag set and the final message has the COMMIT flag set. It is assumed that all the routes between the BEGIN and the COMMIT belong to that batch. Any route within the batch may be flagged as a GATEWAY route.

If flags are not encountered in the correct sequence, it is assumed that a mistake has occurred and the entire batch is aborted. So if a second BEGIN appears before the first is closed by a COMMIT, the entire chain of routes arriving after the first BEGIN but before the second is decreed out-of-date. The arrival of the second BEGIN starts a new batch replacing the first one, as the peer has obviously updated it’s routing table again in the time it has taken for the messages to be delivered.

There are three alternatives to sending a batch of FRP routes. The first is to indicate that no routes are present in the message by using the NULLRT flag. The specification says “*If there are no routes in the update, send route message with length=1, flags=BEGIN + COMMIT + NULLRT*”, so when NULLRT is used, BEGIN and COMMIT flags must also be present. The “*length=1*” refers to the prefix length (as the message length must be 12) but the specification also says, “*if NULLRT is specified, fields beyond the flags field may be omitted*”. So the FRP Quagga implementation reads in the entire message but if the NULLRT flag is set, then it does not use any data beyond the 3<sup>rd</sup> byte. For compatibility, it sets the prefix length to 1 when creating a message but makes no use of it when processing.

The second and third alternatives allow a single route to be deleted or updated using the DELETE or UPDATE flag in combination with a BEGIN and a COMMIT in the same single message. The original specification is unclear about the use of BEGIN and COMMIT flags with DELETES and UPDATES and as the original version of FRP did not implement these

functions, the code does not provide additional information. Therefore, an independent decision had to be made when adding these two options to the Quagga daemon. The specification does state “*Update: single route add/change (abort batch update)*” and “*Delete: delete specified route (abort batch update)*”, and also uses BEGIN and COMMIT with the NULLRT flag. Consequently, it was decided that these two should include a BEGIN as this makes processing route updates easier; all three of the alternatives can be treated as a new batch of one. Although a COMMIT is not actually needed from a processing point of view, it seems cleaner to match the batch updates. So in the FRP Quagga daemon, BEGINs with DELETES and UPDATES are required and COMMITs are optional.

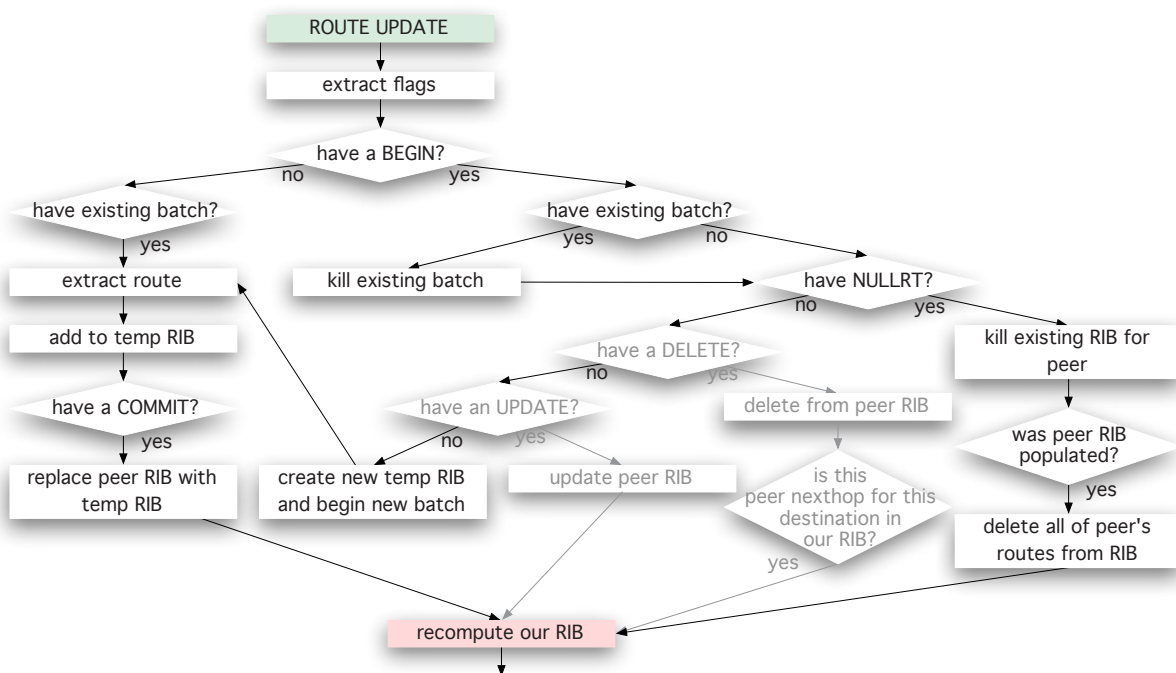


Figure 5.17  
Route update flow  
(note that DELETE and UPDATE not fully implemented in FRP Quagga daemon)

The diagram above shows the route update message decision tree. Having extracted the flags belonging to the message, the first aspect to check is the presence of a BEGIN and, no matter what the result, the next is to check for the presence of an existing batch. When distributing a batched route update, a peer sends its entire table of all acceptable FRP routes. In the FRP Quagga daemon, the recipient creates a temporary RIB (temp\_rib) for that peer when a new batch is begun and builds the new routing table for the peer in it. When a COMMIT is received, the currently stored routing table for that peer is replaced with the new one. So the presence of an existing batch can be determined by seeing if a temporary RIB has been set up for the peer.



If both a BEGIN and a temporary RIB are present, a now out-of-date batch already exists so this is aborted and a new one started. The next steps involve checking for the NULLRT, DELETE and UPDATE flags. Only one of the three can be present, so the process ends once one has been found and dealt with. If none of the three exist, then a new multi-message batch begins.

The presence of a BEGIN and a NULLRT means that the peer is sending an empty update message — that is, it has no FRP routes to share — so it is necessary to delete the RIB currently stored for that peer. Before this can happen, it is necessary to work through the host's routing table checking for routes using the peer as a next hop and deleting them. This in turn triggers the recomputation of the hosts routing table to see if other viable routes are available to replace the deleted ones.

If both a BEGIN and a DELETE appear, the host needs to first remove the route from the RIB stored for the peer and then check the host RIB to see if the route used the current peer as the nexthop for that route. If it does, delete the route from the host RIB and recompute in case another peer has provided an alternative.

In the case of receiving a BEGIN and an UPDATE, modify the affected route appropriately in the stored peer RIB, then recompute the host's RIB.

If there is no BEGIN, then a current batch should already exist and this message is an addition to the chain. Extract the route, add it to the temporary RIB for the peer, and check for a COMMIT. If one is not present, start the loop again with the next message, otherwise replace the old stored peer RIB with the new temporary RIB and recompute the host RIB.

### The FRP daemon in action

The first trace below shows 10.0.1.20 initiating a null route exchange.

```
FRP PACKET HEADER
  source address: 10.0.1.20 (a.0.1.14)
destination address: 10.0.1.50 (a.0.1.32)
  security hash: ??%W?J
  sender's seq no: 10 (0xa)
recipient's ack no: 0 (0x0)
--SYN Packet

Waiting for incoming FRP packet

FRP PACKET HEADER
  source address: 10.0.1.50 (a.0.1.32)
destination address: 10.0.1.20 (a.0.1.14)
  security hash:
  sender's seq no: 10 (0xa)
recipient's ack no: 10 (0xa)
  Null Message: 0x0
--ACK Packet

Waiting for incoming FRP packet
```

The third packet of the exchange shows the single route update message with the BEGIN, NULLRT and COMMIT flags turned on and everything else set to 0.

```
FRP PACKET HEADER
  source address: 10.0.1.20 (a.0.1.14)
  destination address: 10.0.1.50 (a.0.1.32)
  security hash: ?(:W?
  sender's seq no: 11 (0xb)
  recipient's ack no: 10 (0xa)

IPV4 Route Update Message: 0x43
  flags: 1792
  FRP_FLAG_BEGIN flag: 256
  FRP_FLAG_NULLRT flag: 1024
  FRP_FLAG_DELETE flag: 0
  FRP_FLAG_COMMIT flag: 512
  FRP_FLAG_UPDATE flag: 0
  FRP_FLAG_GATEWAY flag: 0
  length: 0
  routecost: 0
  gatecost: 0
  prefix: 0.0.0.0

Waiting for incoming FRP packet

FRP PACKET HEADER
  source address: 10.0.1.50 (a.0.1.32)
  destination address: 10.0.1.20 (a.0.1.14)
  security hash: ?MD??
  sender's seq no: 11 (0xb)
  recipient's ack no: 11 (0xb)

Control message: 0x1
--Control ACK, param = 0
```

In this second example, 10.0.1.20 is a gateway and 10.0.1.50 is not. Both hosts have the 'backbone' network, 10.0.1.0/24, set as static routes. Just before the trace begins, the VTY terminal is used to tell 20 to start advertising the network 10.0.20.0/24 via its interface address 10.0.20.1.

```
FRP PACKET HEADER
  source address: 10.0.1.20 (a.0.1.14)
  destination address: 10.0.1.50 (a.0.1.32)
  security hash: x44M?
  sender's seq no: 8 (0x8)
  recipient's ack no: 0 (0x0)
--SYN Packet

Waiting for incoming FRP packet

FRP PACKET HEADER
  source address: 10.0.1.50 (a.0.1.32)
  destination address: 10.0.1.20 (a.0.1.14)
  security hash: ?K?h??
  sender's seq no: 9 (0x9)
  recipient's ack no: 8 (0x8)
  Null Message: 0x0
--ACK Packet

Waiting for incoming FRP packet
```

Following the expected SYN and ACK, 20 builds a series of routes messages to hold the route it is now advertising to the network 10.0.20.0/24. In this case the route is only two hops long.

```
FRP PACKET HEADER
  source address: 10.0.1.20 (a.0.1.14)
  destination address: 10.0.1.50 (a.0.1.32)
  security hash: ??P??]?U
  sender's seq no: 9 (0x9)
  recipient's ack no: 9 (0x9)

IPV4 Route Update Message: 0x43
  flags: 256
  FRP_FLAG_BEGIN flag: YES
  FRP_FLAG_NULLRT flag: NO
  FRP_FLAG_DELETE flag: NO
  FRP_FLAG_COMMIT flag: NO
  FRP_FLAG_UPDATE flag: NO
  FRP_FLAG_GATEWAY flag: NO
  length: 24
  routecost: 0
  gatecost: 0
  prefix: 10.0.1.0

IPV4 Route Update Message: 0x43
  flags: 512
  FRP_FLAG_BEGIN flag: NO
  FRP_FLAG_NULLRT flag: NO
  FRP_FLAG_DELETE flag: NO
  FRP_FLAG_COMMIT flag: YES
  FRP_FLAG_UPDATE flag: NO
  FRP_FLAG_GATEWAY flag: NO
  length: 24
  routecost: 0
  gatecost: 0
  prefix: 10.0.20.0

Waiting for incoming FRP packet

FRP PACKET HEADER
  source address: 10.0.1.50 (a.0.1.32)
  destination address: 10.0.1.20 (a.0.1.14)
  security hash: ??_????Z
  sender's seq no: 10 (0xa)
  recipient's ack no: 9 (0x9)

Control message: 0x1
--Control ACK, param = 0
```

The following trace shows the FRP daemon routing table after the route update message has been received and processed on 50. the F prefix (here highlighted in red) denotes a route passed to Zebra from the FRP daemon. The first is the default route. The second is the statically set backbone route, which is inactive because the directly connected version is favoured. The third is a route that 50 has received from 20. Note that 50 has the network 10.0.50.0/24 set but as it is not a gateway, it is quiescent and therefore does not advertise the route.

```
frp50router# sh ip ro
Codes: K - kernel route, C - connected, S - static, R - RIP, F - FRP, O - OSPF,
       I - ISIS, B - BGP, > - selected route, * - FIB route

F  0.0.0.0/0 [184/3] via 10.0.1.20, en0, 00:00:39
K>* 0.0.0.0/0 via 10.0.1.1, en0
F  10.0.1.0/24 [184/3] via 10.0.1.20 inactive, 00:02:17
S  10.0.1.0/24 [1/0] is directly connected, en0
C * 10.0.1.0/24 is directly connected, en1
C>* 10.0.1.0/24 is directly connected, en0
F>* 10.0.20.0/24 [184/3] via 10.0.1.20, en0, 00:10:20
S  10.0.50.0/24 [1/0] is directly connected, en1
C>* 10.0.50.0/24 is directly connected, en1
C>* 127.0.0.0/8 is directly connected, lo0
frp50router#
```

The corresponding kernel routing table on 50. The route from 20 has been received from Quagga and inserted (highlighted in red).

```
Artemis:~ deb$ netstat -nrf inet
Routing tables

Internet:
Destination      Gateway          Flags    Refs      Use  Netif  Expire
default          10.0.1.1        UGSc     2          0    en0
10.0.1/24        link#4          UCS      4          0    en0
10.0.1.1         0:3:93:e3:fc:92 UHLW     6        191    en0   1130
10.0.1.20        0:25:4b:ca:7b:b2 UHLW     3       13223    en0   1147
10.0.1.50        127.0.0.1       UHS      0          0    lo0
10.0.1.254       34:15:9e:18:d0:28 UHLW     0         571    en0    104
10.0.1.255       ff:ff:ff:ff:ff:ff UHLWb    0          3    en0
10.0.20/24       10.0.1.20       UG1c     0          0    en0
10.0.50/24       link#5          UC       1          0    en1
10.0.50.1        0:1f:5b:c6:a4:da UHLW     0          0    lo0
127              127.0.0.1       UCS      0          0    lo0
127.0.0.1        127.0.0.1       UH       0          0    lo0
169.254          link#4          UCS      0          0    en0
Artemis:~ deb$
```

## Re-computing the RIB

Once the processing of the packet containing the route updates is complete, the `frp_recompute_rib` function is called. This steps through the routes stored for a peer, which will be the ones newly deposited in what was the `temp_rib` and is now the stored RIB for that peer. Each stored route is checked against the host's RIB. If the host does not currently have a route to that destination, the route is added with the current peer as the nexthop. If the route is already in the host's RIB, the FRP algorithm is applied and if the new route is 'better' than the old one, an exchange is made. In each case, the RIB is flagged as having changed so that once the loop through the peer's routes is complete an update can be triggered so that all the newly acquired routes can be passed on to other peers.

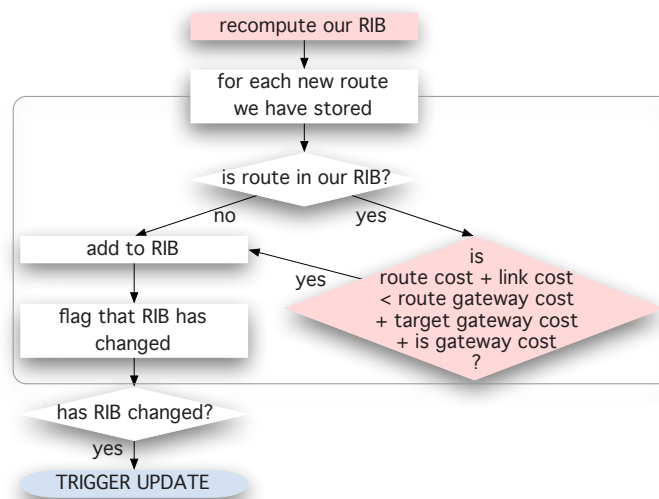


Figure 5.18  
Recomputing the routing table flow

The FRP algorithm is used when the 'new' route has the same destination as a route already in the host's RIB. A comparison is made between the existing and the new route to determine which one belongs in the table of 'best' routes.

In the case of receiving a BEGIN and an UPDATE, modify the affected route appropriately in the stored peer RIB and check to see if the current peer is the nexthop for the route in the host's RIB. If it is, check to see if the new route is 'better' than the old one, changing it in the host RIB if it is — no other peer can have a better route or it would already be in there — and forcing a recompute if it is not in case another peer has a better one. If the current peer is not the nexthop, nothing more needs to be done.

## Outgoing packets

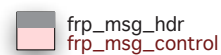
There is no one outgoing packet function but rather a series of individual ones, each handling one specific message or packet creation task. These functions are all gathered together in `frp_packet.c` and draw heavily on the packet and message structures defined in `frp_packet.h`. When building a new packet, a buffer is created and the various parts of the packet's contents are 'memcpy'd' in to the correct place as they are created. Pointers are used to keep track of the beginning and end of the buffer, as well as the current insertion point.

The exact specification for each of the different message types can be found chapter 3.

## Building a message

Each individual message within a packet begins with a message header. These are 2 bytes in length and contain the length and the type of the message. The function `make_frp_msg_header` uses a switch to create a message header with the correct values in each of the two fields.

The FRP control message is 4 bytes in length including the message header. The only useful data carried by the message is the control type — POLL, ACK or NAK — as the parameters field is not currently used. So the `make_frp_msg_control` function creates a control message header, sets the type appropriately, and sets the parameters to 0.



*Figure 5.19*  
*Control message buffer*

`make_frp_msg_ipv4_config` sets up a 12 byte message to carry the configuration settings of a router. It creates a configuration message header, then extracts the cost, poll and retry settings from the current router configuration to fill out these fields of a configuration message. The final field is filled in with the host's address of the interface that the peer communicates with the host on.



*Figure 5.20*  
*Configuration message buffer*

When it comes to gateway messages, things get a little more complex because a ‘path to gateway’ length varies. The `make_frp_msg_ipv4_gateway` function has two strands governed by the `is_gateway_flag`. If the flag is set, the host router is a gateway so the path to gateway contains one address — that of the router’s gateway interface — and the cost equals 0. Otherwise, a check is made to see if the host is quiescent. If it is, the cost is set to infinity (`0xffff`) and the path to 0. At 8 bytes, these are the smallest gateway messages.

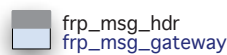


Figure 5.21  
Path to gateway message with no path buffer

If the flag is not set and the host is not quiescent, cost is set to the current gateway cost and the entire path to gateway is appended with a further 4 bytes added to accommodate each additional address.




Figure 5.22  
Path to gateway message with a path buffer

The `make_frp_msg_ipv4_update` function is more complex still. Called whenever the host’s routing table is modified, this works through the RIB putting together a routing update for each peer, excluding routes that have that peer as the nexthop. Each route update message handles a single destination so a separate message is created for each entry in the table.




Figure 5.23  
Single route update message buffer

All the messages required to send the RIB are sent as a batch using flags to provide the information necessary to interpret the entire table.




Figure 5.24  
Batch of route update messages buffer

The function starts by setting a counter to 0 to indicate that there are no messages in the batch and then enters a loop that steps through each route stored in the host RIB. The nexthop for the route is checked to make sure that the route does not originate from the current peer before the peer record for the nexthop is retrieved. This provides the necessary information to make a test of the FRP decision algorithm between the route from the RIB and the route to the same destination in the peer record. If the new route is 'better' then a message containing the route is created to send to the peer as an update.

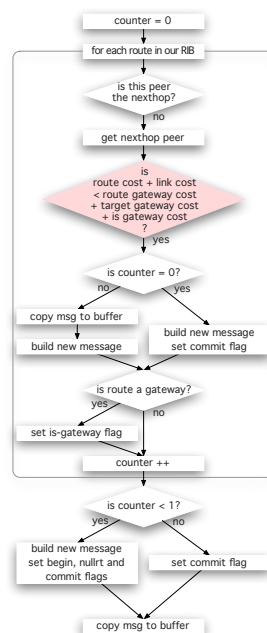


Figure 5.25  
Building a batch of route update messages flow



If the counter equals 0, then this is the first route in a new batch so a new message is created and the BEGIN update flag is set. If the counter is greater than 0, then there is a previous message to be ‘memcopied’ to the buffer before creating a new message for the current route. The BEGIN flag is obviously not used in this case.

The next check is to see if the current route is flagged by the host as a gateway route and, if so, the GATEWAY update flag is set. The message length, route cost, gateway cost, and destination prefix are then filled in appropriately and the counter is incremented. The loop ends when all routes in the RIB have been considered.

The final part of the function tidies up the batch ready to send. The counter is again checked. If it is greater than 0, then there are legitimate messages in the batch, the last one of which needs have the COMMIT update flag set before being copied to the buffer.

If the counter is still equal to 0 however, no routes are being passed to the peer so a null route packet needs to be sent. A new message is created, the message length is added appropriately, and the BEGIN, NULLRT and COMMIT update flags are set. The message is copied to the buffer and is ready to be turned into a completed packet.

### **Building a packet**

A packet is one or more messages, each one complete with its message header, chained together with a packet header attached to the beginning. Although it is the first data in the packet, the 16 byte packet header is actually the last part of the packet to be created. The function `make_frp_pkt_hdr` fills out the sender’s sequence number field by adding 1 to the last sequence number used for that peer and then replacing the stored one with the new one. The recipient’s acknowledgement number is either set to the last received sequence number stored for the peer or to 0 depending on context.

It is the security hash that requires this part of the packet to be created last as the hash uses the entire packet as a parameter. The actual hash is performed in the function `dohash` in `frp_peer.c`. This function is one of only two pieces of code lifted directly from the original FRP implementation. This is because unless the hash is performed in exactly the same way in every implementation, the reverse hash when the packet is received invariably fails. As the code is essentially a series of library calls, it seemed entirely reasonable just to use the original `dohash` and `checksecure` functions.

First a hash is formed using the contents of the packet, minus the 16 bytes of the header. This is then rehashed with the host's address, rehashed again with the peer's address, and finally again with the host's secret. The first 64 bits of the final hash form the security hash for the packet.

```

dohash (u_char* buf, int len, const char* secret, IPADDR sa, IPADDR da)
{
    static u_int8_t hash[SHA_DIGEST_LENGTH];
    SHA_CTX shctx;
    SHA_Init(&shctx);
    SHA_Update(&shctx, buf, len);
    SHA_Update(&shctx, (u_char*)&sa, sizeof(IPADDR));
    SHA_Update(&shctx, (u_char*)&da, sizeof(IPADDR));
    SHA_Update(&shctx, secret, strlen(secret));
    SHA_Final(hash, &shctx);
    return hash;
}

```




Figure 5.26  
Packet header buffer

SYN packets and the initial response ACK packets are just packet headers with no additional messages attached and so are straightforward to create. The function `build_syn_pt` finds the address of the interface the peer is attached to, creates a SYN packet (ie: a packet header), opens a socket to use and calls `frp_send_packet` to send the packet to the peer. If the packet is successfully sent, a complete copy is stored in the `packet_latest` variable in the host's peer record, along with the packet length and the sequence number used — this enables the host to resend the packet should a NAK be received. The awaiting ACK flag is also set for this peer.




Figure 5.27  
Packet header plus control message buffer

Control messages (POLL, ACK and NAK) consist of a packet header and a single control message with the correct flag set. In this implementation, there are currently two functions — `build_ack_pkt` and `build_nak_pkt` — that are virtually identical. It would be neater and more efficient to have a single `build_control_pkt` that handled all three in one function. These two functions create a buffer to build the packet in and then call `make_frp_msg_control`, setting the type field appropriately, before calling `make_frp_pkt_hdr` to create the packet header. The packet is sent, and as in the SYN packet, a copy is stored in the peer record and the awaiting ACK flag is set.




Figure 5.28  
Multiple messages in one packet buffer

SYN, ACK and control messages are sent with a single message per packet. All the other messages types can be assembled into batches and sent with multiple messages carried by a single packet. The complete flow for an outgoing packet is shown below. It can be seen that most of the complexity is inside the loop sending flagged messages triggered by an update to a router configuration, a gateway route or a routing table.

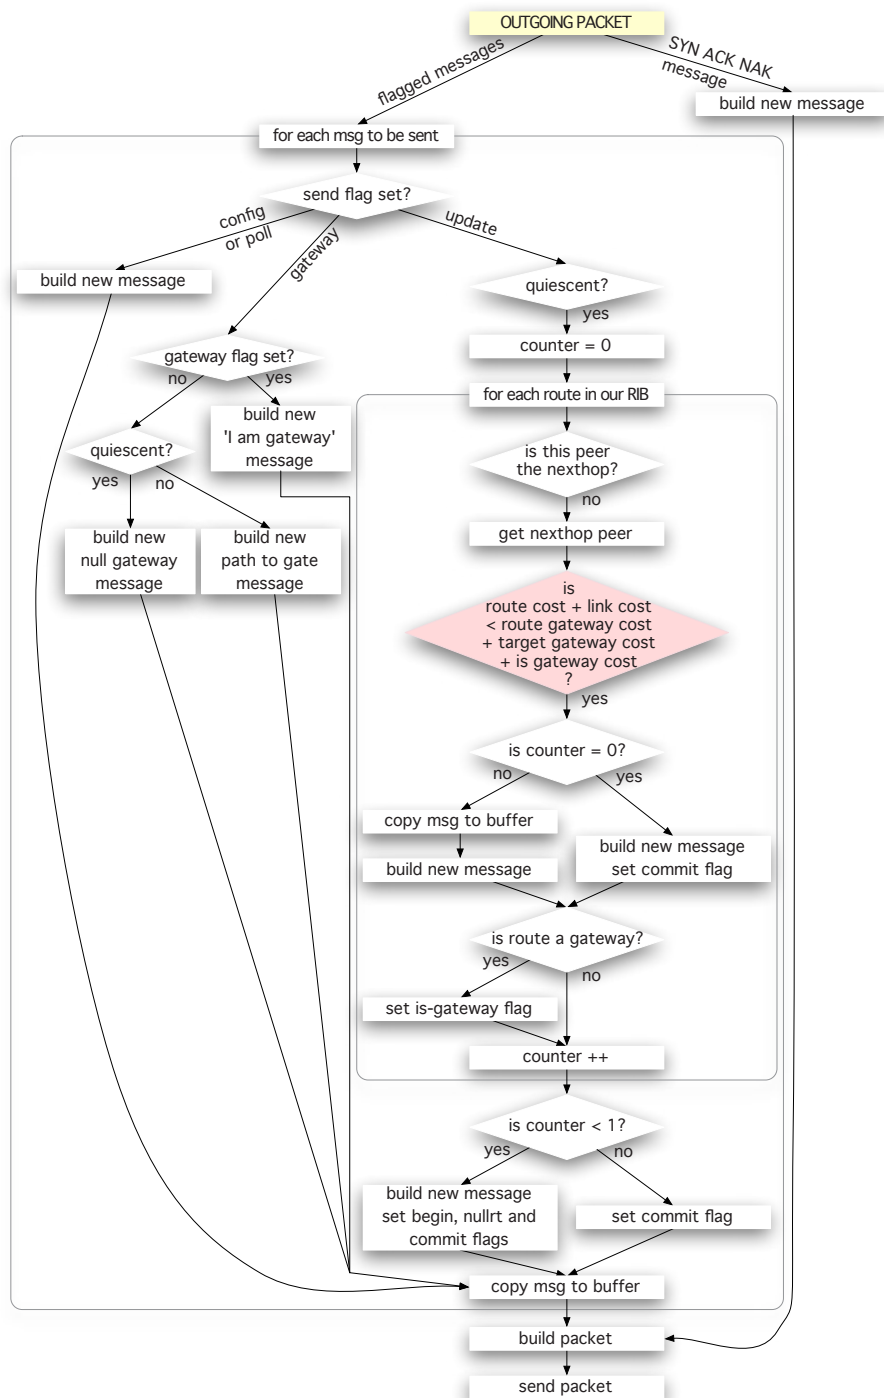


Figure 5.29  
Outgoing packet flow

The `build_batch_pkt` function handles creating packets that potentially carry more than one message — specifically, all the messages currently flagged as needing to be sent to a single peer. The first step is to create the buffer that will store the data while the packet is being built. Next a local flag is created and turned OFF, and each of the peer flags — `flag_send_poll`, `flag_send_config`, `flag_send_gateway`, `flag_send_update` — are checked to see if any have been activated. For each of these that are set, a message of the appropriate type is created and sent. A check is made to ensure that the new message, when added to the packet buffer, will not cause the packet to exceed the specified maximum size of 1400 bytes.

The poll flag, the configuration flag and the gateway flag are all handled in the same way. In each case, the code segment calls the appropriate message function, the message and message header is created and ‘memcpyed’ to the buffer, the packet length is adjusted and the local flag is switched ON. The update flag is only checked if the host is not quiescent (ie: it has a gateway) as update messages are not sent when the host has no gateway.

*If there is no path to a gateway, the routing updates enter into a quiescent state, where link state and path information is still exchanged, but no changes are made to the routing table, and no routing updates are issued. [57]*

The `make_frp_msg_ipv4_update` function must be passed a calculation of how much packet space is left as it is potentially creating a batch of more than one update message. It must also return the amount of space it has actually used to allow the packet length to be accurately tallied. Upon completion, the update segment turns ON the local flag.

Once `build_batch_pkt` has worked through the flags and created messages as necessary for each one set, the next step is to check the local flag. If this is ON, messages have indeed been created so a packet header is created and added to the front of the buffer, and the packet is sent. The peer record is updated by turning all the send flags OFF, storing a copy of the latest sent packet plus its length and sequence number, and turning on the awaiting ACK flag. If the local flag is OFF, no messages have been created so `build_batch_pkt` checks the awaiting ACK flag and if it is ON, creates and sends a control ACK message to indicate that the current conversation is over.

The final function in `frp_packet.c` is `frp_send_packet`. This function is essentially the relevant code copied and pasted from the RIP and RIPng daemons. At present, this FRP implementation does minimal checking for error conditions. The next version will need to augment this function. The send routine essentially sets up the necessary addressing, socket and port issues, opens the socket, sends the packet to the peer, and closes the socket, passing the success or failure of the send back to the calling function.

## Testing

---

Testing the *Quagga* FRP daemon did not happen as a single stage in the development but as a continuous process throughout the project. This testing was to ensure the correctness of the *Quagga* implementation — to make sure that the daemon worked as specified. Testing the design of the protocol itself to see if it resolves the problems it sets out to solve was beyond the scope of this project.

The first major round of testing occurred at the point where the shell demon was complete. It needed to successfully accomplish a number of tasks before the FRP additions were added in. The testing regime included loading the daemon from the configuration file, making sure the menu system in the VTY interface worked properly and that new commands could not only be added but also executed correctly. These are all things that could be tested directly.

It was vital that the shell daemon talked to the Zebra daemon via the *zserv* client and that routing information was passed through to Zebra and on to the kernel. This was tested by setting static routes (which also tested the VTY commands system) and then checking the Zebra and kernel routing tables to see that the routes appeared (see examples earlier in this section on page 92).

A useful testing and debugging tool is the *frp.log* file. *Quagga* provides the basic setup and support for these logs but it is possible to extend on those basics. The *frp\_debug* files handle these extensions, adding in additional, FRP specific *#ifdef* statements that appear interlaced with all the other *Quagga* debug and status output in the log and in the terminal console. The bonus of tapping in to this system is that the various levels of debugging support can be live switched on and off via the VTY system without taking the daemon down.

Implementation of the fringe routing protocol into the shell daemon created a whole new range of issues and problems to be tested. The problem is that conversations between routers are silent and invisible. A major hurdle was the exchange of packets once the security was in place. The first iteration produced a very secure protocol that wouldn't even talk to itself. Running the *Quagga* implementation against the Knossos implementation was the only way to ensure that this stage was operating correctly.

*Wireshark* was invaluable for watching traffic move between routers and for drilling down into the packet to see the raw data. *FRPsniffer* was even more useful as it showed the interpreted data in an instantly readable form. Once again, keeping an eye on the routing tables via the standard Unix tools was necessary to ensure routes were passed through

correctly. Setting FRPs administrative distance ridiculously high was another useful testing technique for ensuring route changes could be traced through the process.

One area where testing was not required was in the realm of the router, as opposed to the realm of FRP. It was simply assumed that *Quagga* was successfully handling these functions.

Finally, general protocol testing was carried out and the results of this are shown at the end of each of the earlier sections of this chapter where they illustrate that the implementation of that section does in fact work.





---

## 6. Conclusions and future work

---

This project set out to achieve two distinct goals: to understand and specify the new fringe routing protocol; and to create an independent, standalone implementation of the protocol in the *Quagga software routing suite*. Achieving the first goal required the development of a detailed knowledge of routing and routing protocols. The second generated a new challenge, that of becoming familiar with all aspects of the *Quagga* environment.

The first outcome has been achieved in chapter 3 in the expansion of the initial concepts and notes provided by Don Stokes into a detailed description of the fringe routing protocol. This has addressed a number of ambiguities and assumptions of knowledge, both implicit and explicit, in the supplied material. A future goal is to expand this description further into a formal specification to be submitted to the Internet Engineering Task Force (IETF) as an Internet Draft.

The second goal is realised as described in chapters 4 and 5. Achieving this required extending the *Quagga* suite to add FRP to the set of supported routing protocols. The new *Quagga* FRP daemon is able to bootstrap itself using a supplied configuration file and communicate with other FRP peers. It can apply the FRP routing algorithms to correctly process routing information and inject routes into the kernel. A future goal is to extend this daemon to support IPv6.

The work on this project was presented, in conjunction with Don Stokes of Knossos Networks, at the NZ Network Operators Group (NZNOG) conference in Wellington in February 2011 where it was well received by the professional networking community. Slides from the presentation are available in Appendix B.



---

## Bibliography

---

- [1] Anon, 2008, *Quagga: the easy tutorial*, viewed 18 March 2011 <<http://openmaniak.com/quagga.php>>.
- [2] Ascigil, O., Yuan, S., Griffioen, J. and Ken Calvert, K., 2008, "Deconstructing the Network Layer", *Proceedings of the 17<sup>th</sup> International Conference on Computer Communications and Networks*, 2008, ICCCN 2008, St. Thomas, US Virgin Islands, pp. 1-6.
- [3] van Beijnum, P. I. and van Beijnum, I. V., 2002, *BGP*, O'Reilly & Associates, Inc., Sebastopol, Canada.
- [4] van Beijnum, I., 2006, "IPv6 Internals", *The Internet Protocol Journal*, vol. 9, no. 3, viewed 18 March 2011, <[http://www.cisco.com/web/about/ac123/ac147/archived\\_issues/ipj\\_9-3/ipv6\\_internals.html](http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_9-3/ipv6_internals.html)>
- [5] van Beijnum, I., 2007, *Everything you need to know about IPv6*, viewed 18 March 2011 <<http://arstechnica.com/hardware/news/2007/03/IPv6.ars>>.
- [6] van Beijnum, I., 2010, *There is no Plan B: why the IPv4-to-IPv6 transition will be ugly*, viewed 18 March 2011 <<http://arstechnica.com/business/news/2010/09/there-is-no-plan-b-why-the-ipv4-to-ipv6-transition-will-be-ugly.ars/>>.
- [7] Bhargavan, K., Obradovic, D. and Gunter, C.A., 2002, "Formal Verification of Standards for Distance Vector Routing Protocols", *Journal of the ACM.*, vol. 49, no. 4, pp. 538-576.
- [8] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, Harvard University, October 1996.
- [9] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, Harvard University, March 1997.
- [10] Bush, R. and D. Meyer, "Some Internet Architectural Guidelines and Philosophy", RFC 3439, The Internet Society, December 2002.
- [11] Caesar, M. and Rexford, J., 2005, "BGP Routing Policies in ISP Networks", *IEEE Network*, vol. 19, no. 6, pp. 5-11.
- [12] Carpenter, B., Ed., "Architectural Principles of the Internet", RFC 1958, IAB, June 1996.

- [13] Cisco, 2009, *Border Gateway Protocol*, viewed 18 March 2011 < [http://docwiki.cisco.com/wiki/Border\\_Gateway\\_Protocol](http://docwiki.cisco.com/wiki/Border_Gateway_Protocol)>.
- [14] Cisco, n.d., *Basic Router Configuration* viewed 18 March 2011 < <http://www.cisco.com/en/US/docs/routers/access/800/850/software/configuration/guide/routconf.pdf>>.
- [15] Cisco, n.d., *Cisco IOS IP Routing Protocols Configuration Guide*. Viewed 18 March 2011 < [http://www.cisco.com/en/US/docs/ios/12\\_3/featlist/ip\\_vcg.html](http://www.cisco.com/en/US/docs/ios/12_3/featlist/ip_vcg.html)>
- [16] Cisco, n.d., *Route Selection in Cisco Routers*, viewed 18 March 2011 < [http://www.cisco.com/en/US/tech/tk365/technologies\\_tech\\_note09186a0080094823.shtml](http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094823.shtml)>.
- [17] Cisco, n.d., *IP Addressing and Subnetting for New Users*, viewed 18 March 2011 < [http://www.cisco.com/en/US/tech/tk365/technologies\\_tech\\_note09186a00800a67f5.shtml](http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800a67f5.shtml)>.
- [18] Cisco, n.d., *Routing Basics*, viewed 18 March 2011 < [http://docwiki.cisco.com/wiki/Routing\\_Basics](http://docwiki.cisco.com/wiki/Routing_Basics)>.
- [19] Cisco, n.d., *What Is Administrative Distance?*, viewed 18 March 2011 < [http://www.cisco.com/en/US/tech/tk365/technologies\\_tech\\_note09186a0080094195.shtml](http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094195.shtml)>.
- [20] Deering, S. and Hinden, R., "Internet Protocol, Version 6 (IPv6)", RFC 2460, Cisco/Nokia, December 1998.
- [21] Feldmann, A. and Rexford, J. 2001, "IP Network Configuration for Intradomain Traffic Engineering", *IEEE Network*, vol. 15, no. 5, pp. 46-57.
- [22] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", BCP 122, RFC 4632, Cisco Systems/Tropos Networks, August 2006.
- [23] Griffin, T.G. and Sobrinho, J.L., 2005, "Metarouting", *Proceedings of the 2005 conference on Applications, technologies, architectures, and protocols for computer communications*, SIGCOMM '05, Philadelphia, Pennsylvania, USA, pp. 1-12.
- [24] Hagen, S., 2006, *IPv6 essentials*, 2nd ed., O'Reilly Media, Sebastopol, Canada.
- [25] Hall, B. J., 2009, *Beej's guide to network programming: using Internet sockets* (version 3.0.14). [Online]. Available: <http://beej.us/guide/bgnet/>
- [26] He, L., 2005, "A Verified Distance Vector Routing Protocol for Protection of Internet Infrastructure", *Lecture Notes in Computer Science*, vol. 3421, pp. 463-470.

- [27] He, J., Rexford, J. and Chiang, M., 2007, "Don't Optimize Existing Protocols, Design Optimizable Protocols", *ACM SIGCOMM Computer Communication Review*, vol. 37, no.3, pp. 53-58.
- [28] Hedrick, C., "Routing Information Protocol", RFC 1058, Rutgers University, June 1988.
- [29] Hinden, R., "Internet Engineering Task Force Internet Routing Protocol Standardization Criteria", RFC 1264, BBN, October 1991.
- [30] Hinden, R. and Deering, S., "IP Version 6 Addressing Architecture", Nokia/Cisco, February 2006.
- [31] Huitema, C., 1995, *Routing in the Internet*, Prentice-Hall, Inc., Harlow, United Kingdom.
- [32] Huitema, C., Postel, J. and Crocker, S., "Not All RFCs are Standards", RFC 1796, INRIA/ISI/CyberCash, April 1995
- [33] Hulse, R., 2005, "WIX: a distributed internet exchange", *Linux Journal*, vol. 135, no. July, pp. 2.
- [34] Hunt, C., 2002, *TCP/IP network administration*, O'Reilly Media, Inc., Sebastopol, Canada.
- [35] Ishiguro, K. et al., 2005, *Quagga: a routing software package for TCP/IP networks*, viewed 18 March 2011 <<http://wenku.baidu.com/view/fae88f74f46527d3240ce003.html>>.
- [36] Kurose, J. F. and Ross, K., 2002, *Computer Networking: a Top-Down Approach Featuring the Internet*, 2nd. Ed, Addison-Wesley Longman Publishing Co., Inc., Harlow, United Kingdom.
- [37] Le, F., Xie, G.G., Pei, D., Wang, J. and Zhang, H., 2008, "Shedding Light on the Glue Logic of the Internet Routing Architecture", *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 4, pp. 39-50.
- [38] Le, F., Xie, G.G. and Zhang, H., 2007, "Understanding route redistribution", *proceedings of the IEEE International Conference on Network Protocols*, Beijing, China, pp. 81-92.

- [39] Le, F., Xie, G.G. and Zhang, H., 2010, "Understanding route aggregation", Carnegie Mellon University Computer Science Report CS-10-106, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA.
- [40] Little, M., "Goals and functional requirements for inter-autonomous system routing", RFC 1126, SAIC, October 1989.
- [41] Malhotra, R., 2002, *IP routing*, O'Reilly Media, Inc., Sebastopol, Canada.
- [42] Malkin, G., "RIPng Protocol Applicability Statement", RFC 2081, Xylogics, January 1997.
- [43] Malkin, G., "RIP Version 2", RFC 2453, Bay Networks, November 1998.
- [44] Malkin, G. and Minnear, R., "RIPng for IPv6", RFC 2080, Xylogics/Ipsilon Networks, January 1997.
- [45] Medhi, D. and Ramasamy, K., 2007, *Network Routing: Algorithms, Protocols, and Architectures*, Morgan Kaufmann Publishers Inc., Burlington, Massachusetts, USA.
- [46] Moy, J. T., 1998, *OSPF: Anatomy of an Internet Routing Protocol*, Addison-Wesley Longman Publishing Co., Inc., Harlow, United Kingdom.
- [47] Nucci, A., Bhattacharyya, S., Taft, N. and Diot, C., 2007, "IGP Link Weight Assignment for Operational Tier-1 Backbones", *IEEE/ACM Transactions on Networking*, vol. 15, no. 4, pp. 789-802.
- [48] Parkhurst, W., 2004, *Routing First-Step*, Cisco Press, Indianapolis, Indiana, USA.
- [49] Partridge, C., Snoeren, A.C., Strayer, W.T., Schwartz, B., Condell, M. and Castiñeyra, I., 2000, "FIRE: Flexible Intra-AS Routing Environment", *ACM SIGCOMM Computer Communication Review*, vol. 30, no. 4, pp. 191-203.
- [50] Perlman, R., 2000, *Interconnections: Bridges, Routers, Switches, and Internetworking Protocols*, 2nd Ed., Addison-Wesley Longman Publishing Co., Inc., Harlow, United Kingdom.
- [51] Pilot, n.d., *Zebra hacking notes*, viewed 18 March 2011 <<http://quagga.net/faq/zebra-hacking-guide.txt>>.
- [52] Postel, J. and J. Reynolds, "Instructions to RFC Authors", RFC 2223, ISI, October 1997.

- [53] Raghavan, B., Verkaik, P. and Snoeren, A.C., 2009, "Secure and Policy-Compliant Source Routing", *IEEE/ACM Transactions on Networking*, vol. 17, no. 3, pp. 764-777.
- [54] Rekhter, Y., Li, T. and Hares, S., "A Border Gateway Protocol 4 (BGP-4)", The Internet Society, January 2006.
- [55] Stallings, W., 1993, *Data and Computer Communications*, 4<sup>th</sup> Ed., Prentice Hall, Harlow, United Kingdom.
- [56] Stevens, W. R. *et al*, *Unix network programming: the sockets networking API*, 3<sup>rd</sup> ed. Boston, MA: Addison-Wesley, 2004.
- [57] Stokes, D., 2010, *Notes on FRP*, personal communication.
- [58] Stokes, D., 2010, *FRP router implementation code*, personal communication.
- [59] Tanenbaum, A. S., 1985, *Computer Networks*, Prentice Hall, Harlow, United Kingdom.
- [60] Uriarte, Y., 2001, *Zebra for dummies, aka Zebra hacking how to*, viewed 18 March 2011 < <http://www.quagga.net/zh.html>>.
- [61] Walrand, J. and Parekh, S., 2010, "Communication networks: a concise introduction", *Synthesis Lectures on Communication Networks*, vol. 3, no. 1, pp. 1-192.
- [60] Wang, Y. and Wang, Z., 1999, "Explicit Routing Algorithms for Internet Traffic Engineering", *proceedings of the eighth international conference on Commuter Communications and Networks*, Boston, Massachusetts, USA, pp. 582-588.
- [61] White, R., 2005, "Caveats in Testing Routing Protocol Convergence", *The Internet Protocol Journal*, vol. 8, no. 4, pp. 20-27.





---

## Appendix A

---

### Initial Fringe Routing Protocol specification from Don Stokes

#### Goals

- Unicast, configured neighbor relationship
- Secure, must work on public access networks
- Intended to be subordinate to IGP; routes gatewayed at one location only
- Nodes forward routes toward their designated gateway, Nodes forward routes away from the gateway if link cost < sum of gateway cost of originated route + gateway cost of local node (i.e. it's closer to route direct than via the gateways)
- Try to keep routing table sizes to a minimum
- Keep traffic down to keepalives (5 sec?) if no updates
- Only update kernel routing table or transmit routing information if there is a path to a gateway. Activity is quiescent until gateway route is established.

#### Route forwarding

Forward route from local-node to remote-node if:

$$\text{route-cost} + \text{link-cost} < \text{route-gw-cost} + \text{target-gw-cost} + \text{is-gw-route}$$

where:

route-cost is the accumulated sum of the costs of intervening links;

link-cost is the cost of the link to remote-node;

route-gw-cost is the cost from the route's originating node to its gateway, as expressed in the routing update;

target-gw-cost is the cost from the remote-node to its gateway, as expressed in its link advertisements;

is-gw-route is 1 if the route is being passed toward the gateway, 0 otherwise. The is-gw flag is only set to 1 on a route at the originator, and set to 0 any time the route is passed to a peer rather than a gateway.

This algorithm means that routes are only passed across the network if they are "better" than going up to the backbone and back -- in many cases this will be preferable to using backup paths.

**Gateway paths**

Each host expresses to all its peers the path to its gateway. This allows a host to exclude paths that include itself from being considered as potential gateway routes, and avoids counting to infinity.

**Quiescence**

If there is no path to a gateway, the routing updates enter into a quiescent state, where link state and path information is still exchanged, but no changes are made to the routing table, and no routing updates are issued. Without this, routes quickly count to infinity when the last gateway path disappears.

**Split Horizon**

A route should not be advertised back up the link it was learned from.

**Messages**

Message header

hash (64): SHA1(secret, lseq, rseq, message ...)

lseq (32): Local (sent) sequence number

rseq (32): Receive (acknowledge) sequence number

Each packet is one message

Messages sent as UDP packets < 1400 data bytes long.

mten: 8 bits length of message in longwords including type/len

mttype: 8 bits. Note: mttype is in form <proto><type>; IPv4-specific messages are 0x4n, IPv6 messages are 0x6n.

Message type 41: IPv4 configuration

linkcost(16): link cost. Maximise received value with local configuratio.

polltime(16): Poll/keepalive frequency in tenths of seconds. Minimise received value with local configurion.

failtime(15): Timeout to failure after last acked packet, in tenths of seconds. Minimise received value with local configuration.

routid(32): Unique ID (usually IP address) of peer router, used to avoid loops.

Message type 42: Path change (path to gateway),  
 gwcost(16) Distance to gateway. 0 = is gateway, 0xffff = gateway unknown,  
 path(32 x n): Path to gateway, list of 32 bit node IDs, used to prevent counting to infinity. Router will not learn a path that contains itself. Path length defined by message length, limited to 62 entries.

Message type 0x43: Add/change/delete IPv4 route

flags(8): Flags, see below

prefixlen(8) prefix length (0-32)

rcost(16) route cost (distance vector)

rgwcost(16) distance of original route from gateway

prefix(32) IPv4 or IPv6 address

Message type 61: IPv6 configuration

linkcost(16): link cost

gw(128): Gateway address for upstream routes. Note, may be overridden by local configuration.

Message type 0x63: Add/change/delete IPv6 route

flags(8): see below

prefixlen(8) prefix length (0-32 for IPv4, 0-128 for IPv6)

rcost(16) route cost (distance vector)

rgwcost(16) distance of original route from gateway

prefix(32-128) IPv6 address, truncated to 32 bit boundary following prefix length.

## Routing flags

BEGIN (0x01): This is the first route of a batch update. Begin the update transaction, abort any already in progress.

COMMIT (0x02): Last route of batch update, commit the update.

NULLRT (0x04): Only interpret flags; no route exists in message. Note, if NULLRT is specified, fields beyond the flags field may be omitted.

UPDATE (0x08): Single route add/change. (Abort batch update.) \*

DELETE (0x10): Delete specified route. (Abort batch update.) \*

GATEWAY (0x80): Route is a gateway route

\* UPDATE and DELETE are not implemented in frpd.

## Batch updates

Batch processing is handles as follows:

If a route message is received with the BEGIN flag set, any existing batch update is flushed, and a new one started. The route should be added to the new batch.

If a route message is received with the COMMIT flag set, the route should be added to the batch, and the whole batch processed as a complete update.

If a route message with all flags (except the GATEWAY flag) clear is received, the route should be added to the inbound batch.

If any other type of message is received, the batch should be flushed and the message processed as if the batch never started.

If there are no routes in the update, send route message with length=1, flags=BEGIN + COMMIT + NULLRT.

## Initial handshake procedure

Assume isn & rsn, isn = initiator's initial sequence number, rsn = responder's initial sequence number.

Initiator sends null packet, lseq = isn, rseq = 0.

Responder replies with null packet, lseq = rsn, rseq = isn.

Initiator sends IP config packet plus path len (both messages must be present), lseq = isn+1, rseq=rsn.

Responder replies in kind, lseq = rsn+1, rseq = isn+1. Responder may send routing information in the same packet.

Initiator replies, lseq = isn+2, rseq = rsn+1, optionally with updated path and with routing information.

Responder acks, lseq = rsn + 1, rseq = isn + 2.

## Sequence numbers

Given state variables, locseq (Local sequence number), ackseq (acknowledged local sequence number) & remseq (remote sequence number):

If a packet is received with rseq = 0: Discard packet if not null. Set remseq = lseq. Send reply lsec = locseq, rseq = remseq.

All other packets, check rseq = locseq. If not, discard. Update ackseq.

If lseq != remseq, set remseq = lseq, process packet, reply lseq = locseq, rseq = remseq. (locseq may be incremented by update if data is sent that needs to be acknowledged.)

Sent packets should increment locseq prior to constructing packet. If locseq wraps to 0, set locseq to 1.

Periodically resend last packet if ackseq != locseq.

Example, n1 -> n2

```
n1 ip=10.0.0.1: lseq = 100
n2 ip=10.0.0.2: lseq = 200
10.0.0.1 -> 10.0.0.2 lseq=100 rseq=0                n2 xseq=101
10.0.0.2 -> 10.0.0.1 rseq=100 lseq=200
10.0.0.1 -> 10.0.0.2 lseq=101 rseq=200 config, path    validate xseq
10.0.0.2 -> 10.0.0.1 rseq=101 lseq=201 config, path, routes
10.0.0.1 -> 10.0.0.2 lseq=102 rseq=201 routes
10.0.0.2 -> 10.0.0.1 rseq=102 lseq=201
10.0.0.1 -> 10.0.0.2 lseq=103 rseq=201 poll
10.0.0.2 -> 10.0.0.1 rseq=103 lseq=202 response
10.0.0.1 -> 10.0.0.2 lsec=103 rseq=202
10.0.0.1 -> 10.0.0.2 lseq=103 rseq=202                Repeat packet
```



---

## Appendix B

---

Slides from the FRP presentation in conjunction with Don Stokes of Knossos Networks, at the NZ Network Operators Group (NZNOG) conference in Wellington in February 2011.

<p style="text-align: center;"><b>Developing the Fringe Routing Protocol</b></p> <p style="text-align: center;"><b>Don Stokes</b> Knossos Networks Ltd</p> <p style="text-align: center;"><b>Deb Shepherd</b> Victoria University of Wellington Catch 22</p> <p style="text-align: right;">NZNOG 2011</p>	<p style="text-align: right;">Developing the Fringe Routing Protocol</p> <hr/> <p style="text-align: center;"><b>The Fringe Routing Protocol</b></p> <p style="text-align: center;"><b>Don Stokes</b> Knossos Networks Ltd</p> <p style="text-align: right;">NZNOG 2011</p>
<p style="text-align: right;">Developing the Fringe Routing Protocol</p> <hr/> <p style="text-align: center;"><b>A new routing protocol!</b></p> <p style="text-align: center;">Don, are you insane?</p> <p style="text-align: right;">NZNOG 2011</p>	<p style="text-align: right;">Developing the Fringe Routing Protocol</p> <hr/> <p style="text-align: center;"><b>A new routing protocol!</b></p> <p style="text-align: center;">Don, are you insane?</p> <p style="text-align: center;">Probably.</p> <p style="text-align: center;">But I had some problems with existing protocols.</p> <p style="text-align: right;">NZNOG 2011</p>

Developing the Fringe Routing Protocol

## FRP design goals

- Subservient to primary IGP;
- Routing table minimisation;
- Unicast traffic;
- Asymmetric route avoidance;
- Security;
- IPv6 support in single session (v4 or v6).

NZNOG 2011

Developing the Fringe Routing Protocol

## Backbone Path Computation

Each router finds the shortest path to the backbone. This path forms the default route.

Paths can only be formed via routers that are a gateway or have a path to a gateway, thus:

- routing updates not processed unless a gateway path is established;
- paths propagate outwards from the backbone to the leaves.

Each router announces its gateway path to all its peers. This provides loop detection.

Routes announced toward the backbone retain the gateway flag (if set). Announcements to other peers clear the flag.

Only routes with the gateway flag set are announced to the backbone.

NZNOG 2011

Developing the Fringe Routing Protocol

## Routing Metrics and Propagation

Each route has two metrics:

- The path cost of the route from the origin (c)
- The path cost to the designated gateway (gc), learned at routes origin.
- Plus a flag to indicate if the route has only been announced toward the backbone.

Routes are only propagated to peers if the cost of getting to the remote peer is less than that of going via the backbone, i.e:

$$c(\text{route}) + c(\text{link}) < gc(\text{route}) + gc(\text{remote\_peer})$$

NZNOG 2011

Developing the Fringe Routing Protocol

## Protocol (1)

Unicast UDP, Multiple messages per packet. TTL set to 1.

64bit checksum, based on agreed secure hash and shared secret (currently SHA1).

Sequence numbers prevent replay attacks and detect restarts.

Separate message types for IPv4 and IPv6 actions, common link control message.

Regular Poll/ACK keeps link alive.

NZNOG 2011

Developing the Fringe Routing Protocol

## Protocol (2)

Packet header: sum(64), loc-seq(32), rem-seq(32).

Message header: length(8), type(8). Note length is count of 32 bit words; max message is 1024 octets.


Message types:

- Session control (1), Subtypes: Poll, ACK, NACK
- Configuration (0x41, 0x61): cost, poll & fail times, router-ID.
- Gateway path (0x42, 0x62): Gateway cost, path
- Route announcement (0x43, 0x63): flags, prefix length, cost, gateway cost.
  - Flags: Begin, commit, null, update, delete, is gateway

NZNOG 2011

Developing the Fringe Routing Protocol

## Example Network (1)



A announces X (c=2) & Y (c=1)

B announces Z (c=1)

D sees:

- X (via A, c=1, gw=1 gc=2)
- Y (local, gw=1, gc=1)
- Z (via E, c=1, gw=0, gc=1)


E sees:

- X (via D, c=2, gc=0 gc=2)
- Y (via D, c=1, gw=0, gc=1)
- Z (local, gw=1, gc=1)

NZNOG 2011

Developing the Fringe Routing Protocol

## Example Network (2)



Note cost of D-E link now set to 3

A announces X (c=2) & Y (c=1)

B announces Z (c=1)

D sees:

- X (via A, c=1, gw=1 gc=2)
- Y (local, gw=1, gc=1)

Routes for X & Y not announced from D to E:

$$c(X) + c(D-E) > gc(X) + gc(E) \\ 1 + 3 > 2 + 1$$


i.e. backbone path is lower cost.

Ditto announcements from E to D.

NZNOG 2011

Developing the Fringe Routing Protocol

## Example Network (3)



E now chooses D as its gateway

A announces all routes via D

D sees:

- X (via A, c=1, gw=1 gc=2)
- Y (local, gw=1, gc=1)
- Z (via E, c=3, gw=1, gc=4)

E announces routes to D.

NZNOG 2011



Developing the Fringe Routing Protocol

## Implementations

Knossos implementation:

- > FreeBSD based (also run on NetBSD)
  - > Porting would require rewrite of routing table interface.
- > IPv4 only.
- > Used in anger (i.e. paying customers) on Knossos FreeBSD routing platform.

Deb Shepherd's implementation:

- > Portable, Quagga process;
- > IPv6 capable (Deb note: eventually).

NZNOG 2011

Developing the Fringe Routing Protocol

## FRP as a Quagga daemon: edited highlights

**Deb Shepherd**  
Victoria University of Wellington  
Catch 22

NZNOG 2011

Developing the Fringe Routing Protocol

## What is Quagga?

- > software routing suite
- > open source
- > Unix / Linux platforms
- > fork of earlier GNU Zebra project




Code and documentation (of sorts)

- > <http://www.quagga.net/>
- > Manual (of user commands, a selection)
- > *Zebra for Dummies* aka *Zebra How-To*
  - > <http://www.quagga.net/zhzh.html>
- > *Zebra hacking notes*
  - > <http://quagga.net/faq/zebra-hacking-guide.txt>

NZNOG 2011

Developing the Fringe Routing Protocol

## How Quagga works



NZNOG 2011

Developing the Fringe Routing Protocol

## Developing in Quagga

User view

- > 'vt' terminal front end with Cisco style commands

Developer view

- > modular (mostly)
- > core Zebra daemon
  - > library
  - > call back functions
- > independent protocol daemons
  - > bgpd, isisd, ospfd, ospf6d, ripd, ripngd & frpd
- > various makefiles, install & config scripts, and platform specific additions

NZNOG 2011

Developing the Fringe Routing Protocol

## FRP in Quagga — getting started

- > FRP is distance vector, as are RIP and BGP
- > FRP needs to be both IPv4 and IPv6

Ripped the RIP, RIPng and BGP daemons apart

- > using RIP to look at IPv4 and RIPng to look at IPv6
- > and BGP to look at an integrated Ipv4/6 daemon

Took ripd and used find/replace to turn it into frpd

- > then started to pull out all the RIP specific code
- > the intention was to develop a daemon shell
- ...
- > then threw everything away & wrote shell from scratch

NZNOG 2011

Developing the Fringe Routing Protocol

## FRP in Quagga — the daemon

`frp_main.c`

- > set up and initialisation

`frp_zebra.c / frp_interface.c / frp_debug.c`

- > hooks into Zebra including zclient & call back functions

`frpd.h`

- > FRP daemon includes, defines, external variables, macros, prototypes

`frpd.c`

- > main FRP daemon code
- > config changes, sockets, events, incoming packets, polls

NZNOG 2011

Developing the Fringe Routing Protocol

## FRP in Quagga — the daemon

`frp_packet.h`

- > packet and message includes, defines, external variables, macros, prototypes

`frp_packet.c`

- > FRP daemon code specific to packet handling

`frp_peer.c`

- > FRP daemon code specific to peers and secrets

`frp_route.c`

- > FRP daemon code specific to routing
- > the FRP decision algorithm, the FRP RIB, injection of routes into the kernel via Zebra

NZNOG 2011

## FRP in Quagga — events & threads

```
case FRP_EVENT_INCOMING:
    frp->t_read = thread_add_read (master,
        frp_incoming_packet, NULL, sock);
case FRP_EVENT_UPDATE:
    if (frp->t_update_interval)
        frp->update_trigger = 1;
    else if (! frp->t_update)
        frp->t_update = thread_add_event (master,
            frp_update_peers, NULL, 0);
case FRP_EVENT_POLL:
    if (frp->t_poll)
    {
        thread_cancel (frp->t_poll);
        frp->t_poll=NULL;
    }
    frp->t_poll = thread_add_timer (master,
        frp_poll_peers, NULL, frp->poll);
```

NZNOG 2011

## FRP in Quagga — events & threads

```
int frp_incoming_packet (struct thread* t)
{
    variables etc ...

    /* fetch socket then register myself */
    sock = THREAD_FD (t);
    frp->t_read = NULL;
    /* add myself to the next event */
    frp_event (FRP_EVENT_INCOMING, sock);

    // read the packet from the socket
    ...
}
```

NZNOG 2011

## FRP in Quagga — Don's code

Code re-use?

- > why yes! — two whole functions worth in fact
- > Quagga is multi threaded, Don's code is not
- > so no real scope for re-use

Secrets

- > FRP peers exchange secrets
- > the diplomatic etiquette involved is exacting
- > so I lifted Don's security hash creation and checking functions and used them as black boxes

NZNOG 2011

## FRP in Quagga — configuration

Live configuration of Quagga

- > commands are executed via the 'vty' terminal
- > obviously these need to be built in to the daemon


Config file

- > loaded on startup (main.c)
- > calls & executes the same functions as the 'vty' terminal

```
router frp
network 10.0.1.0/24
secret sapphire
cost 3
poll 60
gateway yes
neighbor 10.0.1.50 secret artemis
```

NZNOG 2011

## FRP quirk — batched messages



NZNOG 2011

## Questions ...

NZNOG 2011

---

## Appendix C

---

### Changes made to the Quagga code to hook in a new routing daemon

---

#### Quagga code changes

---

##### **quagga/**

config.h (actually auto generated so will need more work)

'PATH_FRPD_PID' undeclared	647	add #define PATH_FRPD_PID "/var/run/frpd.pid"
'FRP_VTYSH_PATH' undeclared	647	add #define FRP_VTYSH_PATH "/var/run/frpd.vty"

##### **quagga/lib/**

command.h		
'FRP_NODE' undeclared	63	add FRP_NODE, to enum node_type
RIP search	275	Is FRP needed here?
command.c		
RIPNG_NODE search	2386	add case FRP_NODE: to switch (vty->node)
RIPNG_NODE search	2445	add case FRP_NODE: to switch (vty->node)
distribute.c		
RIP search	762	seems to be rip/ripng specific
if_rmap.c		
RIP search	329	seems to be rip/ripng specific
log.h		
'ZLOG_FRP' undeclared	46	add ZLOG_FRP, to enum
log.c		
RIP search	42	add "FRP", to char *zlog_proto_names[]
RIP search	298	seems to be ripng specific
note in zebra.h [449]	826	add DESC_ENTRY (ZEBRA_ROUTE_FRP, "frp", 'F' ), to struct zebra_desc_table
memory.c		
RIP search	485	add DEFUN
RIP search	498	add 3 install_element's

```

memtypes.h
    'MTYPE_FRP_ROUTE' undeclared      7      /* Auto-generated from memtypes.c by gawk. Do not
                                          edit! */
                                          memtypes.h produced from memtypes.c by memtypes.awk

memtypes.c
    note in memtypes.h                 247    add memory_list memory_list_frp[]
    note in memtypes.h                 265    add { memory_list_frp, "FRP" }, to struct mlist
                                          mlists
                                          only added structure, route info, interface and peer to
                                          begin with
                                          check this all works next time compiled

route_types.h
    RIP search                          /* Auto-generated from route_types.txt by gawk. Do
                                          not edit! */
                                          route_types.h produced from route_types.txt by
                                          route_types.awk

route_types.txt
    RIP search                          43    add ZEBRA_ROUTE_FRP, frp, frpd, 'F', 1, 1, "FRP"
    RIP search                          65    add ZEBRA_ROUTE_FRP, "Fringe Routing Protocol (FRP)"

routemap.h
    'RMAP_FRP' undeclared              42    add RMAP_FRP, to enum

vty.c
    RIPNG_NODE search                  690    add case FRP_NODE: to switch (vty->node)
    RIPNG_NODE search                  1101   add case FRP_NODE: to switch (vty->node)

zebra.h
    'ZEBRA_ROUTE_FRP' undeclared       432    add #define ZEBRA_ROUTE_FRP 11
                                          change #define ZEBRA_ROUTE_MAX 12
    'ZEBRA_FRP_DISTANCE_DEFAULT'      519    add #define ZEBRA_EBGP_DISTANCE_DEFAULT 10
    undeclared                         (used 10 because smaller than all other protocols - want FRP to have priority
                                          for now)

```

#### quagga/vtysh/

```

vtysh.c
    RIPNG_NODE search                  715    add static struct cmd_node frp_node = { FRP_NODE,
                                          "%s(config-router)# ", };
    RIPNG_NODE search                  1032   add new DEFUNSH
                                          currently based on RIPng but as other protocols have difference styles, may need more
                                          work
                                          potential for different v4/v6 setups in FRP
    RIPNG_NODE search                  1107   add case FRP_NODE: to switch (vty->node)
    RIPNG_NODE search                  2264+  variety of install_element stuff (probably needs more work)

vtysh_config.c
    RIPNG_NODE search                  160 (207) add else if (strcmp (line, "router frp", strlen ("router
                                          frp"))) == 0) config = config_get (FRP_NODE, line); to
                                          switch (c)

```

#### quagga/zebra/

```

redistribute.c
    RIP search                         248    add case ZEBRA_ROUTE_FRP: to switch (type)
    RIP search                         279    add case ZEBRA_ROUTE_FRP: to switch (type)

rib.h
    RIP search                         86    add FRP to other interior protocols * sub-queue 2: RIP, RIPng, OSPF, OSPF6,
                                          IS-IS, FRP

zebra_rib.c
    RIP search                         59    add FRP admin distance {ZEBRA_ROUTE_FRP, 10},
                                          using 10 because smaller than all other protocols - want FRP to have priority for now
    RIP search                         1222   add [ZEBRA_ROUTE_FRP] = 2, to u_char meta_queue_map

zebra_snmp.c
    RIP search                         231    not doing snmp therefore use default - no change

zebra_vty.c
    RIP search                         556    add || rib->type == ZEBRA_ROUTE_FRP to ONE_WEEK_SECOND
    RIP search                         778    add || rib->type == ZEBRA_ROUTE_FRP to if
    RIP search                         811    add R - RIP, to SHOW_ROUTE_V4_HEADER
    RIP search                         931    add "Fringe Routing Protocol (FRP)\n" to DEFUN
    RIP search                         954    add else if (strcmp (argv[0], "f", 1) == 0) type =
                                          ZEBRA_ROUTE_FRP;

```

---

## Appendix D

---

### Code

The *Quagga* FRP daemon implementation

The FRPsniffer implementation

The original implementation of FRP by Don Stokes



---

## Appendix D

---

The Quagga FRP daemon implementation code





```
1 // -----
2 // INCLUDES
3 // -----
4 #include "frpd.h"
5
6
7
8 // -----
9 // GLOBAL VARIABLES
10 // -----
11
12 /* frpd command line options */ // DEB COMMENT: see no reason to change these
13 struct option longopts[] =
14 { { "daemon", no_argument, NULL, 'd'},
15   { "config_file", required_argument, NULL, 'f'},
16   { "pid_file", required_argument, NULL, 'i'},
17   { "dryrun", no_argument, NULL, 'C'},
18   { "help", no_argument, NULL, 'h'},
19   { "vty_addr", required_argument, NULL, 'A'},
20   { "vty_port", required_argument, NULL, 'P'},
21   { "retain", no_argument, NULL, 'r'},
22   { "user", required_argument, NULL, 'u'},
23   { "group", required_argument, NULL, 'g'},
24   { "version", no_argument, NULL, 'v'},
25   { 0 }
26 };
27 /* frpd configuration file name and directory */
28 char config_default[] = SYSCONFDIR FRP_DEFAULT_CONFIG; // DEB COMMENT: FRP_DEFAULT_CONFIG frpd.h as
29 'frpd.conf'
30 char *config_file;
31 /* frpd process ID saved for use by init system */
32 const char *pid_file;
33 /* frpd VTY bind address */
34 char *vty_addr;
35 /* frpd VTY connection port */
36 int vty_port;
37 /* frpd route retain mode flag */
38 int retain_mode;
39
40 /* frpd privileges */
41 zebra_capabilities_t _caps_p [] =
42 { ZCAP_NET_RAW, // DEB COMMENT: privs.h [34]
43   ZCAP_BIND // DEB COMMENT: privs.h [31]
44 };
45 // DEB COMMENT: zebra_privs_t privs.h [57]
46 struct zebra_privs_t frpd_privs =
47 {
48   #if defined(QUAGGA_USER)
49     .user = QUAGGA_USER, // DEB COMMENT: config.h [649] as "quagga"
50   #endif
51   #if defined(QUAGGA_GROUP)
52     .group = QUAGGA_GROUP, // DEB COMMENT: config.h [643] as "quagga"
53   #endif
54   #ifdef VTY_GROUP
55     .vty_group = VTY_GROUP, // DEB COMMENT: config.h [733] (is this actually used?)
56   #endif
57     .caps_p = _caps_p, // DEB FOLLOW UP: ???
58     .cap_num_p = 2, // DEB FOLLOW UP: ???
59     .cap_num_i = 0 // DEB FOLLOW UP: ???
60 };
61
```

```
62 /* master thread */
63 struct thread_master *master;
64
65
66
67
68 // -----
69 // FUNCTIONS
70 // -----
71
72 /* frpd help information display */
73 static void
74 usage(char *progname, int status)
75 { if (status != 0)
76   fprintf(stderr, "Try `%s --help' for more information.\n", progname);
77   else
78   {
79     printf ("Usage : %s [OPTION...]\n",
80     Daemon which manages FRP.\n\n\
81     -d, --daemon Runs in daemon mode\n\
82     -f, --config_file Set configuration file name\n\
83     -i, --pid_file Set process identifier file name\n\
84     -A, --vty_addr Set vty's bind address\n\
85     -P, --vty_port Set vty's port number\n\
86     -r, --retain When program terminates, retain added route by frpd.\n\
87     -u, --user User to run as\n\
88     -g, --group Group to run as\n\
89     -v, --version Print program version\n\
90     -C, --dryrun Check configuration for validity and exit\n\
91     -h, --help Display this help and exit\n\
92     \n\
93     Report bugs to %s\n", progname, ZEBRA_BUG_ADDRESS); // DEB COMMENT: version.h [31]
94   }
95   exit (status);
96 }
97
98 /* frpd signal hang up and re-initialise handler */
99 static void
100 sighup (void)
101 { zlog_info ("SIGHUP received"); // DEB COMMENT: log.h [124]
102   frp_clean (); // DEB COMMENT: frpd.c
103   frp_reset (); // DEB COMMENT: frpd.c
104   /* reload config file */
105   vty_read_config (config_file, config_default); // DEB COMMENT: vty.c [2333]
106   /* create VTY socket */
107   vty_serv_sock (vty_addr, vty_port, FRP_VTYSH_PATH); // DEB COMMENT: vty.c [2141]
108   // DEB COMMENT: FRP_VTYSH_PATH config.h [662] as '/var/run/frpd.vty'
109 }
110 /* frpd signal interrupt handler */
111 static void
112 sigint (void)
113 { zlog_notice ("Terminating on signal");
114
115   if (!retain_mode)
116     frp_clean ();
117
118   exit (0);
119 }
120 /* frpd SIGUSR1 handler */
121 static void
122 sigusr1 (void)
```

```

123 { zlog_rotate (NULL);
124 }
125 /* frpd signal handler */
126 struct quagga_signal_t frp_signals[] =
127 { { .signal = SIGHUP,
128   .handler = &sigup,
129 },
130 { .signal = SIGUSR1,
131   .handler = &sigusr1,
132 },
133 { .signal = SIGINT,
134   .handler = &sigint,
135 },
136 { .signal = SIGTERM,
137   .handler = &sigint,
138 },
139 };
140
141 // -----
142 // FRPD MAIN
143 // -----
144
145 int
146 main (int argc, char **argv)
147 {
148     #ifdef DEB_DEBUG
149     fprintf (stderr, "DEB DEBUG: entering frp_main.c - main\n");
150     #endif //DEB_DEBUG
151
152     int daemon_mode = 0;
153     /* configuration file name and directory */
154     config_file = NULL;
155     /* process ID saved for use by init system */
156     pid_file = PATH_FRPD_PID; // DEB COMMENT: config.h [633] as '/var/run/frpd.pid'
157     int dryrun = 0;
158     /* FRP VTY bind address */
159     vty_addr = NULL;
160     /* FRP VTY connection port */
161     vty_port = FRP_VTY_PORT; // DEB COMMENT: frpd.h as '2609'
162     /* Route retain mode flag. */
163     retain_mode = 0;
164
165     char *p;
166     char *progrname;
167     struct thread thread;
168
169     /* set umask before anything for security */
170     umask (0027); // DEB COMMENT: like chmod 640
171
172     /* get program name */
173     progrname = ((p = strrchr (argv[0], '/')) ? ++p : argv[0]);
174
175     /* logging init */
176     zlog_default = openzlog(progrname, ZLOG_FRP, LOG_CONS|LOG_NDELAY|LOG_PID, LOG_DAEMON);
177
178     /* command line option parse */
179     while (1)
180     {
181         int opt;

```

```

184     opt = getopt_long (argc, argv, "dfi:hA:P:u:g:vC", longopts, 0);
185     if (opt == EOF)
186         break;
187     switch (opt)
188     {
189     case 0:
190         break;
191     case 'd':
192         daemon_mode = 1;
193         break;
194     case 'f':
195         config_file = optarg;
196         break;
197     case 'A':
198         vty_addr = optarg;
199         break;
200     case 'i':
201         pid_file = optarg;
202         break;
203     case 'P':
204         /* deal with atoi() returning 0 on failure, and frpd not listening on frpd port... */
205         if (strcmp(optarg, "0") == 0)
206             { vty_port = 0;
207               break;
208             }
209         vty_port = atoi (optarg);
210         if (vty_port <= 0 || vty_port > 0xffff)
211             vty_port = FRP_VTY_PORT;
212         break;
213     case 'r':
214         retain_mode = 1;
215         break;
216     case 'u':
217         frpd_privs.user = optarg;
218         break;
219     case 'g':
220         frpd_privs.group = optarg;
221         break;
222     case 'v':
223         print_version (progrname);
224         exit (0);
225         break;
226     case 'C':
227         dryrun = 1;
228         break;
229     case 'h':
230         usage (progrname, 0);
231         break;
232     default:
233         usage (progrname, 1);
234         break;
235     }
236 }
237
238 /* set up master thread / threads management */
239 master = thread_master_create ();
240
241 /* library initialization */
242 zprivs_init (&frpd_privs);
243 signal_init (master, Q_SIGC(frp_signals), frp_signals);
244 cmd_init (1);

```

```
245 vty_init (master);
246 memory_init ();
247
248 /* FRPd initialization */
249 frp_init ();          // DEB COMMENT: frpd.c, where most of the set up happens
250 frp_zclient_init ();  // DEB COMMENT: frp_zebra.c
251
252 /* sort installed commands */
253 sort_node ();
254
255 /* get configuration file */
256 vty_read_config (config_file, config_default);
257
258 /* start execution only if not in dry-run mode */
259 if(dryrun)
260     return(0);
261
262 /* change to the daemon program */
263 if (daemon_mode && daemon (0, 0) < 0)
264 { zlog_err("FRPd daemon failed: %s", strerror(errno));
265   exit (1);
266 }
267
268 /* create VTY socket, start tcp and unix socket listeners */
269 vty_serv_sock (vty_addr, vty_port, FRP_VTYSH_PATH);
270
271 /* create process id file */
272 pid_output (pid_file);
273
274 /* print banner */
275 zlog_notice ("FRPd %s starting: vty@%d", QUAGGA_VERSION, vty_port);
276
277 /* fetch next active thread -- main program loop */
278 while (thread_fetch (master, &thread))
279     thread_call (&thread);
280
281 /* not reached */
282 return 0;
283 }
284
```

```
1 #ifndef _QUAGGA_FRPD_H
2 #define _QUAGGA_FRPD_H
3
4
5 #define DEB_DEBUG
6 // #define DEB_DEBUG_D
7 // #define DEB_DEBUG_IF
8 // #define DEB_DEBUG_PEER
9 #define DEB_DEBUG_PKT
10
11
12 // -----
13 / * INCLUDES * /
14 // -----
15 #include <sys/types.h>
16 #include <stdlib.h>
17 #include <stdio.h>
18 #include <stdint.h>
19 #include <stddef.h>
20 #include <netinet/in.h>
21
22 #include "zebra.h"
23 #include "command.h"
24 #include "getopt.h"
25 #include "if.h"
26 #include "log.h"
27 #include "memory.h"
28 #include "privs.h"
29 #include "sigevent.h"
30 #include "thread.h"
31 #include "version.h"
32 #include "vty.h"
33 #include "prefix.h"
34 #include "zclient.h"
35 #include "table.h"
36 #include "sockopt.h"
37 #include "sockunion.h"
38 #include "stream.h"
39
40 #include "frp_packet.h"
41
42
43
44 // -----
45 / * DEFINES * /
46 // -----
47
48 / * FRP version number. * /
49 #define FRP_VERSION 1
50
51 / * Default config file name * /
52 #define FRP_DEFAULT_CONFIG "frpd.conf" // DEB COMMENT: same as all the other protocols
53
54 / * FRP ports * /
55 #define FRP_PORT_DEFAULT 343 // DEB COMMENT: port used by Don
56 #define FRP_VTY_PORT 2609 // DEB COMMENT: next available unused port in Quagga
57 // DEB COMMENT: check on quagga-dev?
58 / * frp router defaults
59 #define FRP_DEFAULT_COST 1
60 #define FRP_DEFAULT_POLL 5
61 #define FRP_DEFAULT_RETRY 1
```

```
62
63 #define FRP_METRIC
64
65 #define FRP_INFINITY 0xffff
66 / * frp route types. * /
67 #define FRP_ROUTE_RTE 0
68 #define FRP_ROUTE_STATIC 1
69
70 / * frp read/write types * /
71 #define FRP_READ_REQUEST 1
72 #define FRP_READ_RESPONSE 2
73 #define FRP_WRITE_UPDATE 3
74 #define FRP_WRITE_KEEPALIVE 4
75
76 // need a macro here that generates a random number - using 1 for debugging purposes
77 #define NEW_SEQ_NO 0
78
79 // set to determine how many polls to wait before declaring a peer 'dead'
80 #define FRP_PEER_DEAD 5
81
82 // to support code lifted from Don's implementation
83 #define IPADDR u_int32_t
84 #define FRP_HASHSIZE 8
85
86 // random number generation
87 #define RANDOM_SEED() srand(time(NULL))
88 #define RANDOM_INT(__MIN__, __MAX__) ((__MIN__) + random() % ((__MAX__+1) - (__MIN__)))
89
90 // -----
91 / * EXTERNAL VARIABLES * /
92 // -----
93
94 / * frp event. * /
95 enum frp_event
96 { FRP_EVENT_INCOMING,
97   FRP_EVENT_UPDATE,
98   FRP_EVENT_POLL,
99 };
100
101 // flags
102 enum flag
103 { OFF,
104   ON,
105 };
106
107 / * frp structure. * /
108 struct frp
109 { / * frp version * /
110   int version;
111   / * frp output buffer * /
112   struct stream* obuf;
113   / * frp socket * /
114   int sock;
115   / * frp threads * /
116   struct thread* t_read;
117   struct thread* t_poll;
118   struct thread* t_update;
119   struct thread* t_update_interval;
120   int update_trigger;
121   / * timer values. * /
122   unsigned long update_time;
```

```
123 unsigned long    timeout_time;
124 unsigned long    garbage_time;
125 /* frp routing information base (linked list) */
126 struct route_table* rib;
127 /* frp only static routing information (linked list) */
128 struct route_table* routes;
129 /* frp neighbors (linked list) */
130 struct route_table* neighbors;
131 /* frp gateway path (linked list) */
132 struct list* gateway_path;
133 // is this peer a gateway?
134 enum flag         is_gateway_flag;
135 // which peer is the current next hop to the current gateway?
136 // - use this to trigger an update if this peer changes it's path to gateway
137 struct frp_peer*  gateway_nexthop;
138 // number of hops to gateway
139 int               gateway_cost;
140 // gateway types
141 #define FRP_GATEWAY_ALWAYS 0
142 #define FRP_GATEWAY_YES   1
143 #define FRP_GATEWAY_NO    2
144 // secret
145 const char*       secret;
146 // cost of the link (16 bits)
147 u_short           cost;
148 // poll time (16 bits)
149 u_short           poll;
150 // retry time (16 bits)
151 u_short           retry;
152 };
153
154 /* FRP specific interface configuration */
155 struct frp_interface
156 { /* frp is enabled on this interface */
157     int enable_network;
158     int enable_interface;
159     /* frp is running on this interface */
160     int running;
161     /* for filter type slot */
162     #define FRP_FILTER_IN    0
163     #define FRP_FILTER_OUT   1
164     #define FRP_FILTER_MAX   2
165     /* access-list */
166     struct access_list* list[FRP_FILTER_MAX];
167     /* prefix-list */
168     struct prefix_list* prefix[FRP_FILTER_MAX];
169     /* route-map */
170     struct route_map* routemap[FRP_FILTER_MAX];
171     /* wake up thread */
172     struct thread* t_wakeup;
173     /* interface statistics */
174     int rcv_badpackets;
175     int rcv_badroutes;
176     int sent_updates;
177     /* passive interface */
178     int passive;
179 };
180
181 // frp peer
182 struct frp_peer
183 { struct frp_peer* next;
```

```
184 struct frp_peer* prev;
185 /* peer address */
186 struct in_addr    address;
187 const char*       secret;
188 u_short           cost;
189 u_short           poll;
190 u_short           retry;
191 // our current sequence number with this peer
192 u_int32_t         lseq;
193 // this peer's current sequence number
194 u_int32_t         rseq;
195 // current path to gateway for this peer
196 struct list*      gateway_path;
197 // cost of gateway, ie: number of hops in path to gateway
198 int               gateway_cost;
199 // temporary storage for new incoming routes
200 struct list*      rib;
201 struct list*      temp_rib;
202 // latest packet sent to this peer
203 u_char*           packet_latest;
204 u_int8_t           packet_latest_length;
205 u_int32_t         packet_latest_lseq;
206 // latest timer timestamps
207 time_t            time_latest_packet;
208 time_t            time_last_heard;
209 time_t            time_sent_config;
210 /* timeout thread */
211 struct thread*    t_timeout;
212 // flags
213 enum flag         flag_alive;
214 enum flag         flag_send_syn;
215 enum flag         flag_send_poll;
216 enum flag         flag_send_config;
217 enum flag         flag_send_gateway;
218 enum flag         flag_send_update;
219 enum flag         flag_awaiting_ack;
220 /* statistics */
221 int               rcv_badpackets;
222 int               rcv_badroutes;
223 };
224
225 // store route updates from a peer
226 struct frp_rte
227 { u_int8_t         flags; // update type flags (8 bits)
228   u_int8_t         length; // prefix length (8 bits)
229   u_short           routecost; // route cost (16 bits)
230   u_short           gatecost; // cost from originator to gateway (16 bits)
231   struct prefix_ipv4 prefix; // IP prefix (32 bits)
232 };
233
234 /* frp route information */
235 struct frp_info
236 { struct in_addr    nexthop;
237   u_int32_t         cost;
238   enum flag         is_gateway_flag;
239   /* this route's type */
240   int type;
241   /* sub type */
242   int sub_type;
243   // back pointer
244   struct route_node* rte_node;
```

```
245 / * which interface does this route come from * /
246 unsigned int ifindex;
247 };
248
249 / * buffer to store frp data * /
250 union frp_pkt_buf
251 { char buf[FRP_PKT_MAXSIZE];
252   struct frp_pkt_hdr frp_pkt_hdr;
253   struct frp_msg_hdr frp_msg_hdr;
254   struct frp_msg_control frp_msg_control;
255   struct frp_msg_ipv4config frp_msg_ip4config;
256   struct frp_msg_ipv4gateway frp_msg_ip4gateway;
257   struct frp_msg_ipv4update frp_msg_ip4update;
258   struct frp_msg_ipv6config frp_msg_ip6config;
259   struct frp_msg_ipv6gateway frp_msg_ip6gateway;
260   struct frp_msg_ipv6update frp_msg_ip6update;
261 };
262
263 extern struct zebra_privs_t frpd_privs;
264
265
266
267 // -----
268 // PROTOTYPES
269 // -----
270 / * there is only one frp structre * /
271 extern struct frp* frp;
272 // and only one linked list of frp peers
273 extern struct list* frp_peers;
274
275
276 / * master thread * /
277 extern struct thread_master* master;
278
279 extern void frp_init (void);
280 extern void frp_clean (void);
281 extern void frp_reset (void);
282 extern void frp_event (enum frp_event, int);
283
284 // frp_route.c
285 extern void frp_recompute_rib (struct frp_peer* peer);
286 extern void frp_delete_peer_from_rib (struct frp_peer* peer);
287
288 // frp_packet.c
289 extern int frp_send_packet (u_char * buf, int size, struct sockaddr_in *to);
290 extern struct frp_pkt_hdr make_frp_pkt_hdr (struct in_addr local, struct frp_peer* peer, int flag, int len, u_char*
291 buf);
292 extern struct frp_msg_hdr make_frp_msg_hdr (int type);
293 extern struct frp_msg_control make_frp_msg_control (int type);
294 extern struct frp_msg_ipv4config make_frp_msg_ip4config (struct in_addr peer);
295 extern int make_frp_msg_ip4gateway (struct in_addr peer, struct frp_msg_ipv4gateway* msg);
296 extern int make_frp_msg_ip4update (struct frp_peer* peer, int available_length, u_char* buf);
297 extern int build_syn_pkt (struct frp_peer* peer);
298 extern int build_ack_pkt (struct frp_peer* peer, struct sockaddr_in* from, struct in_addr local);
299 extern int build_nak_pkt (struct frp_peer* peer, struct sockaddr_in* from, struct in_addr local);
300 extern int build_batch_pkt (struct frp_peer* peer, struct sockaddr_in* from, struct in_addr local);
301
302 // frp_peer.c
303 extern void frp_peer_init (void);
304 extern int frp_neighbor_add (struct vty* vty, const char* ip_str, const char* secret);
305 extern struct frp_peer* frp_peer_lookup (struct in_addr* addr);
```

```
306 extern u_int8_t* dohash(u_char* buf, int len, const char* secret, IPADDR sa, IPADDR da);
307 extern int checksecure(u_char* pkt, int len, struct frp_peer* peer, struct in_addr local);
308 extern void secure(struct frp_peer* peer, struct in_addr local, struct frp_pkt_hdr* pkt, int len);
309 extern struct prefix* find_local_address_for_peer(struct in_addr dest);
310
311 // frp_interface.c
312 extern struct route_table *frp_enable_network;
313 extern void frp_interface_init (void);
314 extern void frp_config_write_network (struct vty* vty);
315 extern int frp_neighbor_lookup (struct sockaddr_in*);
316 extern struct in_addr frp_get_interface_address (struct in_addr peer);
317
318 // frp_zebra.c
319 extern void frp_zclient_init (void);
320 extern void frp_zebra_ipv4_add (struct prefix_ipv4 *p, struct in_addr *nexthop, u_int32_t metric, u_char distance);
321 extern void frp_zebra_ipv4_delete (struct prefix_ipv4 *p, struct in_addr *nexthop, u_int32_t metric);
322 / * zebra client API callback functions * /
323 extern int frp_interface_add (int command, struct zclient *zclient, zebra_size_t length);
324 extern int frp_interface_delete (int command, struct zclient *zclient, zebra_size_t length);
325 extern int frp_interface_up (int command, struct zclient *zclient, zebra_size_t length);
326 extern int frp_interface_down (int command, struct zclient *zclient, zebra_size_t length);
327 extern int frp_interface_address_add (int command, struct zclient *zclient, zebra_size_t length);
328 extern int frp_interface_address_delete (int command, struct zclient *zclient, zebra_size_t length);
329
330
331
332 #endif / * _QUAGGA_FRPD_H * /
333
```

```
1 // -----
2 // INCLUDES
3 // -----
4 #include "frpd.h"
5 #include "frp_debug.h"
6
7
8
9 // -----
10 // GLOBAL VARIABLES
11 // -----
12
13 /* frp node structure */
14 static struct cmd_node frp_node =
15 { FRP_NODE, "%s(config-router)# ", 1 };
16
17 /* frp structure */
18 struct frp* frp = NULL;
19
20
21
22 // -----
23 // PROTOTYPES
24 // -----
25 static int frp_config_write (struct vty* vty);
26 static int frp_create (void);
27 static int frp_create_socket (struct sockaddr_in* from);
28 static int frp_incoming_packet (struct thread* t);
29 static int16_t frp_incoming_control_msg (struct frp_msg_control* msg, struct frp_peer* peer, struct sockaddr_in*
30 from, struct in_addr local);
31 static int16_t frp_incoming_config_msg (struct frp_msg_ipv4config* msg, struct frp_peer* peer, struct sockaddr_in*
32 from, struct in_addr local);
33 static int16_t frp_incoming_gateway_msg (struct frp_msg_ipv4gateway* msg, struct frp_peer* peer, struct
34 sockaddr_in* from, struct in_addr local);
35 static int16_t frp_incoming_update_msg (struct frp_msg_ipv4update* msg, struct frp_peer* peer, struct sockaddr_in
36 from, struct in_addr local);
37
38 static int frp_update_peers (struct thread* t);
39 static int frp_update_interval (struct thread* t);
40 static int frp_poll_peers (struct thread* t);
41
42
43
44 // -----
45 // FUNCTIONS
46 // -----
47
48 /* create new frp instance and set it to global variable */
49 int
50 frp_create (void)
51 {
52     #ifdef DEB_DEBUG
53         fprintf (stderr, "DEB DEBUG: entering frpd.c - frp_create\n");
54     #endif //DEB_DEBUG
55     frp = XCALLOC (MTYPE_FRP, sizeof (struct frp));
56     /* set initial values */
57     frp->version = FRP_VERSION;
58     frp->cost = FRP_DEFAULT_COST;
59     frp->poll = FRP_DEFAULT_POLL;
60     frp->retry = FRP_DEFAULT_RETRY;
61     frp->is_gateway_flag = OFF;
```

```
62     frp->gateway_cost = FRP_INFINITY;
63     /* initialize frp routing table */ // DEB COMMENT: route_table_init table.h [56]
64     frp->rib = route_table_init ();
65     frp->routes = route_table_init ();
66     frp->neighbors = route_table_init ();
67     /* make output stream */
68     frp->obuf = stream_new (FRP_PKT_MAXSIZE);
69     /* make socket */
70     frp->sock = frp_create_socket (NULL);
71     if (frp->sock < 0)
72     {
73         #ifdef DEB_DEBUG
74             fprintf (stderr, "DEB DEBUG: -- frp_create, failed to create socket\n");
75         #endif //DEB_DEBUG
76         return frp->sock;
77     }
78     #ifdef DEB_DEBUG
79         fprintf (stderr, "DEB DEBUG: -- frp_create, created socket: %d\n", frp->sock);
80     #endif //DEB_DEBUG
81     /* create read and timer thread */
82     frp_event (FRP_EVENT_INCOMING, frp->sock);
83     frp_event (FRP_EVENT_UPDATE, 1);
84     frp_event (FRP_EVENT_POLL, 1);
85     #ifdef DEB_DEBUG
86         fprintf (stderr, "DEB DEBUG: -- frp_create, finished events\n");
87     #endif //DEB_DEBUG
88     return 0;
89 }
90
91 /* create socket for frp protocol */
92 int
93 frp_create_socket (struct sockaddr_in* from)
94 { int ret;
95   int sock;
96   struct sockaddr_in addr;
97
98   #ifdef DEB_DEBUG
99       fprintf (stderr, "DEB DEBUG: entering frpd.c - frp_create_socket\n");
100   #endif //DEB_DEBUG
101
102   // set address to '0'
103   memset (&addr, 0, sizeof (struct sockaddr_in));
104   // if address from elsewhere, set to that address
105   if (!from)
106   { addr.sin_family = AF_INET;
107     addr.sin_addr.s_addr = INADDR_ANY;
108     #ifdef HAVE_STRUCT_SOCKADDR_IN_SIN_LEN
109         addr.sin_len = sizeof (struct sockaddr_in);
110     #endif /* HAVE_STRUCT_SOCKADDR_IN_SIN_LEN */
111   } else
112   { memcpy(&addr, from, sizeof(addr));
113   }
114
115   /* sending port must always be the frp port */
116   addr.sin_port = htons (FRP_PORT_DEFAULT);
117
118   /* make datagram socket */
119   sock = socket (AF_INET, SOCK_DGRAM, 0);
120   if (sock < 0)
121   { zlog_err("Cannot create UDP socket: %s", safe_strerror(errno));
122     exit (1);
```

```
123 }
124 // set the socket options
125 sockopt_reuseaddr (sock); // DEB COMMENT: sockunion.c [455]
126 sockopt_reuseport (sock); // DEB COMMENT: sockunion.c [472]
127 setsockopt_so_rcvbuf (sock, FRP_PKT_MAXSIZE); // DEB COMMENT: sockopt.c [27]
128
129 #ifdef DEB_DEBUG
130 fprintf (stderr, "DEB DEBUG: -- frp_create_socket, after make socket, port: %d\n", (int) ntohs(addr.sin_port))
131 #endif //DEB_DEBUG
132
133 // bind address to socket
134 if (frpd_privs.change (ZPRIVS_RAISE))
135 { zlog_err ("frp_create_socket: could not raise privs");
136 }
137 if ( (ret = bind (sock, (struct sockaddr *) & addr, sizeof (addr))) < 0)
138 { int save_errno = errno;
139 if (frpd_privs.change (ZPRIVS_LOWER))
140 { zlog_err ("frp_create_socket: could not lower privs");
141 }
142 zlog_err ("%s: Can't bind socket %d to %s port %d: %s", __func__, sock, inet_ntoa(addr.sin_addr), (int)
143 ntohs(addr.sin_port), safe_strerror(save_errno));
144 close (sock);
145 #ifdef DEB_DEBUG
146 fprintf (stderr, "DEB DEBUG: -- frp_create_socket, after bind %d\n", ret);
147 #endif //DEB_DEBUG
148 return ret;
149 }
150 if (frpd_privs.change (ZPRIVS_LOWER))
151 { zlog_err ("frp_create_socket: could not lower privs");
152 }
153
154 return sock;
155 }
156
157 /* frp configuration write function */
158 // DEB COMMENT: called when FRP_NODE is installed
159 // provides the FRP part of the 'write' command from 'router frp' to the end of the config
160 // needs adding to each time a command set is added - use rip/ripng as templates
161 // commands are written out in same form as entered from command line
162 int
163 frp_config_write (struct vty* vty)
164 { int write = 0;
165 struct route_node* rn;
166 #ifdef DEB_DEBUG
167 zlog_debug ("DEB DEBUG: entering frpd.c - frp_config_write");
168 #endif //DEB_DEBUG
169 if (frp) // DEB COMMENT: ie: only in frpd(config-router)# mode
170 { /* router frp */
171 vty_out (vty, "router frp%s", VTY_NEWLINE);
172 /* network */
173 frp_config_write_network (vty);
174 // router setup
175 if (frp->is_gateway_flag == ON)
176 { vty_out (vty, "gateway yes %s", VTY_NEWLINE);
177 }
178 }
179 vty_out (vty, "secret %s %s", frp->secret, VTY_NEWLINE);
180 if (frp->cost != FRP_DEFAULT_COST)
181 { vty_out (vty, "cost %d %s", frp->cost, VTY_NEWLINE);
182 }
183 if (frp->poll != FRP_DEFAULT_POLL)
```

```
184 { vty_out (vty, "poll %d %s", frp->poll, VTY_NEWLINE);
185
186 }
187 if (frp->retry != FRP_DEFAULT_RETRY)
188 { vty_out (vty, "retry %d %s", frp->retry, VTY_NEWLINE);
189
190 }
191 /* neighbours */
192 for (rn = route_top (frp->neighbors); rn; rn = route_next (rn))
193 { if (rn->info)
194 { struct prefix* rn_p = (struct prefix*)&rn->p;
195 struct in_addr address;
196 int addr = inet_pton (AF_INET, inet_ntoa (rn_p->u.prefix4), &address);
197 #ifdef DEB_DEBUG
198 zlog_debug ("DEB DEBUG: -- frp_config_write - address=%s", inet_ntoa (address));
199 #endif //DEB_DEBUG
200 struct frp_peer* peer = frp_peer_lookup (&address);
201 #ifdef DEB_DEBUG
202 zlog_debug ("DEB DEBUG: -- frp_config_write - peer->addr=%s, peer->secret=%s", inet_ntoa
203 (peer->address), peer->secret);
204 #endif //DEB_DEBUG
205 if (peer == NULL)
206 { vty_out (vty, "neighbor %s peer is null %s", inet_ntoa (rn_p->u.prefix4), VTY_NEWLINE);
207 }
208 else
209 { vty_out (vty, "neighbor %s secret %s %s", inet_ntoa (rn_p->u.prefix4), peer->secret, VTY_NEWLINE)
210 }
211 }
212 }
213 write++;
214 }
215 return write;
216 }
217
218 void
219 frp_clean ()
220 {
221 #ifdef DEB_DEBUG
222 fprintf (stderr, "DEB DEBUG: entering frpd.c - frp_clean\n");
223 #endif //DEB_DEBUG
224 return;
225 }
226
227 void
228 frp_reset ()
229 {
230 #ifdef DEB_DEBUG
231 fprintf (stderr, "DEB DEBUG: entering frpd.c - frp_reset\n");
232 #endif //DEB_DEBUG
233 return;
234 }
235
236 // -----events-----
237
238 void
239 frp_event (enum frp_event event, int sock)
240 {
241 #ifdef DEB_DEBUG
242 zlog_debug ("DEB DEBUG: entering frpd.c - frp_event");
243 #endif //DEB_DEBUG
```



```
245 switch (event)
246 {
247     case FRP_EVENT_INCOMING:
248         #ifdef DEB_DEBUG
249             zlog_debug ("DEB DEBUG: -- frp_event, case: FRP_EVENT_INCOMING");
250         #endif //DEB_DEBUG
251         // create a new read thread and tell it to run frp_incoming_packet()
252         frp->t_read = thread_add_read (master, frp_incoming_packet, NULL, sock);
253         break;
254     case FRP_EVENT_UPDATE:
255         #ifdef DEB_DEBUG
256             zlog_debug ("DEB DEBUG: -- frp_event, case: FRP_EVENT_UPDATE");
257         #endif //DEB_DEBUG
258         if (frp->t_update_interval)
259             frp->update_trigger = 1;
260         else if (! frp->t_update)
261             frp->t_update = thread_add_event (master, frp_update_peers, NULL, 0);
262         break;
263     case FRP_EVENT_POLL:
264         #ifdef DEB_DEBUG
265             zlog_debug ("DEB DEBUG: -- frp_event, case: FRP_EVENT_POLL");
266         #endif //DEB_DEBUG
267         if (frp->t_poll)
268             { thread_cancel (frp->t_poll);
269               frp->t_poll = NULL;
270             }
271         frp->t_poll = thread_add_timer (master, frp_poll_peers, NULL, (unsigned long)frp->poll);
272         break;
273     default:
274         #ifdef DEB_DEBUG
275             fprintf (stderr, "DEB DEBUG: -- frp_event, case: default\n");
276         #endif //DEB_DEBUG
277         // complain and die
278         zlog_info ("Unknown FRP event %d received", event);
279         break;
280 }
281 }
282
283
284
285 // -----FRP_EVENT_INCOMING-----
286 int
287 frp_incoming_packet (struct thread* t)
288 {
289     #ifdef DEB_DEBUG
290         zlog_debug ("DEB DEBUG: entering frpd.c - frp_incoming_packet");
291     #endif //DEB_DEBUG
292     int sock;
293     socklen_t fromlen;
294     int pkt_length;
295     struct sockaddr_in from;
296     struct in_addr local;
297     struct interface* ifp;
298     struct connected* ifc;
299     struct frp_interface* ri;
300     // read buffer
301     union frp_pkt_buf in_pkt_buf;
302     // u_char* current_pos;
303     //peer packet is from
304     struct frp_peer* peer;
305     int secure;
```

```
306 // outgoing packet storage
307 struct frp_pkt_hdr* pkt_hdr;
308 // struct frp_msg_hdr* msg_hdr;
309 struct frp_pkt_hdr ack_pkt;
310 int sent;
311
312 // ---1: SET UP STUFF-----
313 /* fetch socket then register myself */
314 sock = THREAD_FD (t);
315 frp->t_read = NULL;
316 /* add myself to the next event */
317 frp_event (FRP_EVENT_INCOMING, sock);
318 /* frpd manages only IPv4 */
319 memset (&from, 0, sizeof (struct sockaddr_in));
320 fromlen = sizeof (struct sockaddr_in);
321 // read the packet from the socket
322 // DEB COMMENT: rip/ng have a recvmmsg/recv_packet function that does all sorts of additional checking - COME E
323 TO THIS
324 pkt_length = recvfrom (sock, (char *)&in_pkt_buf.buf, FRP_PKT_MAXSIZE, 0, (struct sockaddr *)&from, &fromlen);
325 if (pkt_length < 0)
326 { zlog_info ("recvfrom failed: %s", safe_strerror (errno));
327   return pkt_length;
328 }
329 /* which interface is this packet from */
330 ifp = if_lookup_address (from.sin_addr);
331 /* frp packet received */
332 if (IS_FRP_DEBUG_EVENT)
333     zlog_debug ("RECV packet from %s port %d on %s", inet_ntoa (from.sin_addr), ntohs (from.sin_port), ifp ?
334 ifp->name : "unknown");
335 /* if this packet comes from unknown interface, ignore it */
336 if (ifp == NULL)
337 { zlog_info ("frp_incoming_packet: cannot find interface for packet from %s port %d", inet_ntoa (from.sin_addr),
338 (from.sin_port));
339   return -1;
340 }
341 ifc = connected_lookup_address (ifp, from.sin_addr);
342 if (ifc == NULL)
343 { zlog_info ("frp_incoming_packet: cannot find connected address for packet from %s " "port %d on interface %s"
344 inet_ntoa (from.sin_addr), ntohs (from.sin_port), ifp->name);
345   return -1;
346 }
347 /* packet length check */
348 if (pkt_length < FRP_PKT_MINSIZE)
349 { zlog_warn ("packet size %d is smaller than minimum size %d", pkt_length, FRP_PKT_MINSIZE);
350   return pkt_length;
351 }
352 if (pkt_length > FRP_PKT_MAXSIZE)
353 { zlog_warn ("packet size %d is larger than max size %d", pkt_length, FRP_PKT_MAXSIZE);
354   return pkt_length;
355 }
356 /* is frp running or is this frp neighbor */
357 ri = ifp->info;
358 if (! ri->running && ! frp_neighbor_lookup (&from))
359 { if (IS_FRP_DEBUG_EVENT)
360     { zlog_debug ("FRP is not enabled on interface %s.", ifp->name);
361     }
362     return -1;
363 }
364
365 // ---2: EXTRACT PACKET HEADER AND CHECK HAVE A VALID PACKET -----
366 /* get the packet header from the buffer */
```

```
367 pkt_hdr = &in_pkt_buf.frp_pkt_hdr;
368 // which peer is the packet from?
369 peer = frp_peer_lookup (&from.sin_addr);
370 // is it from an IP address we are allowed to talk to?
371 // - need to set up access lists for this - come back to
372 // indicate this peer is still alive
373 peer->flag_alive = ON;
374 // timestamp this communication
375 peer->time_last_heard = time (NULL);
376 // does the security hash match?
377 local = frp_get_interface_address (peer->address);
378 // set secure to fail
379 secure = 0;
380 if (peer->secret != NULL)
381 { // check security hash: secret = peer's, sender = peer, dest = us
382     secure = checksecure((u_char *)&in_pkt_buf.buf, pkt_length, peer, local);
383 } else // no secret so set secure to succeed
384 { secure = 1;
385 }
386 if (!secure)
387 { zlog_info ("secure check failed");
388     return -1;
389 }
390 #ifdef DEB_DEBUG_D
391     zlog_debug ("DEB DEBUG: -- frp_incoming_packet - passed security check, SYN=%d (0x%x), ACK=%d (0x%x)"
392 ntohl(pkt_hdr->sendSeq), ntohl(pkt_hdr->sendSeq), ntohl(pkt_hdr->recipAck), ntohl(pkt_hdr->recipAck));
393 #endif //DEB_DEBUG
394
395 // ---3: DETERMINE IF NEW OR ON-GOING CONVERSATION-----
396 // store the peer's sequence number
397 peer->rseq = ntohl(pkt_hdr->sendSeq);
398 if (pkt_hdr->recipAck == 0)
399 // have a SYN packet - peer has initiated a new conversation
400 {
401     #ifdef DEB_DEBUG
402         zlog_debug ("DEB DEBUG: -- frp_incoming_packet - have a SYN packet, SYN=%d (0x%x), ACK=%d (0x%x)",
403 ntohl(pkt_hdr->sendSeq), ntohl(pkt_hdr->sendSeq), ntohl(pkt_hdr->recipAck), ntohl(pkt_hdr->recipAck));
404     #endif //DEB_DEBUG
405     // is the packet the correct length?
406     if (pkt_length != FRP_PKT_HDRSIZE)
407     { zlog_info ("packet size check failed");
408         return -1;
409     }
410     // create an ACK packet to send to new peer
411     ack_pkt = make_frp_pkt_hdr (local, peer, FRP_ACK, FRP_PKT_HDRSIZE, NULL);
412     // send ACK packet to peer
413     sent = frp_send_packet ((u_char *)&ack_pkt, FRP_PKT_HDRSIZE, &from);
414     if (sent)
415     { memcpy (peer->packet_latest, (u_char *)&ack_pkt, FRP_PKT_HDRSIZE);
416         peer->packet_latest_length = FRP_PKT_HDRSIZE;
417         peer->packet_latest_lseq = ntohl(pkt_hdr->sendSeq);
418         peer->flag_awaiting_ack = ON;
419     }
420     #ifdef DEB_DEBUG_D
421         zlog_debug ("DEB DEBUG: -- frp_incoming_packet - packet_latest_lseq=%d", peer->packet_latest_lseq);
422     #endif //DEB_DEBUG
423     } else
424     { zlog_debug ("packet number %d not sent", ack_pkt.sendSeq);
425     }
426 } else if (pkt_hdr->recipAck > 0)
427 {
428 }
```

```
428 #ifdef DEB_DEBUG
429     zlog_debug ("DEB DEBUG: -- frp_incoming_packet - have an ACK packet, SYN=%d (0x%x), ACK=%d (0x%x)"
430 ntohl(pkt_hdr->sendSeq), ntohl(pkt_hdr->sendSeq), ntohl(pkt_hdr->recipAck), ntohl(pkt_hdr->recipAck));
431 #endif //DEB_DEBUG
432 // does the packet rseq match our previous lseq?
433 if (ntohl(pkt_hdr->recipAck) == peer->lseq)
434 { // is the packet just a packet header or does it contain a payload?
435     if (pkt_length == FRP_PKT_HDRSIZE) // is acking my syn so send flagged message(s) - config, gateway, up
436     {
437         #ifdef DEB_DEBUG
438             zlog_debug ("DEB DEBUG: -- frp_incoming_packet - is acking my syn so send flagged message(s), about
439 build_batch_pkt");
440         #endif //DEB_DEBUG
441         build_batch_pkt (peer, &from, local);
442     } else // for each message in the packet
443     { char* start_pos;
444       char* current_pos;
445       struct frp_msg_hdr* current_msg;
446       int16_t msg_length;
447
448       start_pos = (char *)&in_pkt_buf.buf;
449       current_pos = start_pos + FRP_PKT_HDRSIZE;
450       current_msg = NULL;
451       msg_length = 0;
452       current_msg = (struct frp_msg_hdr*)current_pos;
453
454       // while the in buffer still has data in it
455       int nak_response_sent = 0;
456       while ((current_msg != NULL) && (current_pos < (start_pos + pkt_length)))
457       { // extract message length and message type and process message
458         msg_hdr = (struct frp_msg_hdr*)((char *)&in_pkt_buf.buf + FRP_PKT_HDRSIZE);
459         #ifdef DEB_DEBUG
460             zlog_debug ("DEB DEBUG: -- frp_incoming_packet - length=%d, type=0x%x", current_msg->length,
461 current_msg->type);
462         #endif //DEB_DEBUG
463         switch (current_msg->type)
464         { case FRP_MSG_CONTROL:
465           #ifdef DEB_DEBUG
466               zlog_debug ("DEB DEBUG: -- frp_incoming_packet - SWITCH control");
467           #endif //DEB_DEBUG
468           current_msg = (struct frp_msg_control*)current_pos;
469           msg_length = frp_incoming_control_msg ((struct frp_msg_control*)current_msg, peer, &from, lc
470           if (msg_length == -1)
471           { nak_response_sent = 1;
472             current_pos += FRP_MSG_CONTROL_SIZE;
473           } else
474           { current_pos += msg_length;
475           }
476           current_msg = (struct frp_msg_hdr*)current_pos;
477           break;
478         case FRP_MSG_IPV4CONFIG:
479           #ifdef DEB_DEBUG
480               zlog_debug ("DEB DEBUG: -- frp_incoming_packet - SWITCH config");
481           #endif //DEB_DEBUG
482           current_msg = (struct frp_msg_ipv4config*)current_pos;
483           msg_length = frp_incoming_config_msg ((struct frp_msg_ipv4config*)current_msg, peer, &from,
484           current_pos += msg_length;
485           current_msg = (struct frp_msg_hdr*)current_pos;
486           break;
487         case FRP_MSG_IPV4GATEWAY:
488           #ifdef DEB_DEBUG
```

```

489         zlog_debug ("DEB DEBUG: -- frp_incoming_packet - SWITCH gateway");
490         #endif //DEB_DEBUG
491         current_msg = (struct frp_msg_ipv4gateway*)current_pos;
492         msg_length = frp_incoming_gateway_msg ((struct frp_msg_ipv4gateway*)current_msg, peer, &f
493 local);
494         current_pos += msg_length;
495         current_msg = (struct frp_msg_hdr*)current_pos;
496         break;
497     case FRP_MSG_IPV4UPDATE:
498         #ifdef DEB_DEBUG
499             zlog_debug ("DEB DEBUG: -- frp_incoming_packet - SWITCH update");
500         #endif //DEB_DEBUG
501         current_msg = (struct frp_msg_ipv4update*)current_pos;
502         msg_length = frp_incoming_update_msg ((struct frp_msg_ipv4update*)current_msg, peer, &from,
503         current_pos += msg_length;
504         current_msg = (struct frp_msg_hdr*)current_pos;
505         break;
506     case FRP_MSG_IPV6CONFIG:
507         //call function
508         break;
509     case FRP_MSG_IPV6GATEWAY:
510         //call function
511         break;
512     case FRP_MSG_IPV6UPDATE:
513         //call function
514         break;
515     default:
516         #ifdef DEB_DEBUG
517             zlog_debug ("DEB DEBUG: -- frp_incoming_packet - SWITCH default");
518         #endif //DEB_DEBUG
519         break;
520     }
521 }
522 // send packet, unless have responded to a NAK
523 if (!nak_response_sent)
524 {
525     #ifdef DEB_DEBUG
526         zlog_debug ("DEB DEBUG: -- frp_incoming_packet - send packet, unless have responded to a NAK, ab
527 call_build_batch_pkt");
528     #endif //DEB_DEBUG
529     build_batch_pkt (peer, &from, local);
530 }
531 }
532 } else // packet rseq does not match our previous lseq
533 {
534     // SEND NAK
535     sent = build_nak_pkt (peer, &from, local);
536     #ifdef DEB_DEBUG
537         zlog_debug ("DEB DEBUG: -- 4 - pkt_hdr->recipAck != peer->lseq - STOP");
538     #endif //DEB_DEBUG
539 }
540 } else // have an invalid packet
541 {
542     zlog_info ("invalid incoming packet");
543     return -1;
544 }
545 return pkt_length;
546 }
547
548 int16_t
549 frp_incoming_control_msg (struct frp_msg_control* msg, struct frp_peer* peer, struct sockaddr_in* from, struct

```

```

550 in_addr local)
551 {
552     int sent;
553     #ifdef DEB_DEBUG
554         zlog_debug ("DEB DEBUG: entering frpd.c - frp_incoming_control_msg");
555     #endif //DEB_DEBUG
556     switch (msg->type)
557     {
558     case FRP_CTRL_POLL:
559         #ifdef DEB_DEBUG
560             zlog_debug ("DEB DEBUG: -- frp_incoming_control_msg - FRP_CTRL_POLL");
561         #endif //DEB_DEBUG
562         // send sn ACK
563         sent = build_ack_pkt (peer, from, local);
564         break;
565     case FRP_CTRL_ACK:
566         #ifdef DEB_DEBUG
567             zlog_debug ("DEB DEBUG: -- frp_incoming_control_msg - FRP_CTRL_ACK");
568         #endif //DEB_DEBUG
569         // end of conversation - do nothing except update peer record
570         peer->flag_awaiting_ack = OFF;
571         #ifdef DEB_DEBUG_D
572             zlog_debug ("DEB DEBUG: -- frp_incoming_control_msg - packet_latest_lseq=%d", peer->packet_latest_l
573         #endif //DEB_DEBUG
574         break;
575     case FRP_CTRL_NAK:
576         #ifdef DEB_DEBUG
577             zlog_debug ("DEB DEBUG: -- frp_incoming_control_msg - FRP_CTRL_NAK");
578         #endif //DEB_DEBUG
579         // update last packet header with new seqs & hash
580         struct frp_pkt_hdr new_pkt_hdr = make_frp_pkt_hdr (local, peer, FRP_ACK, peer->packet_latest_length,
581         (u_char*)peer->packet_latest);
582         memcpy((u_char*)peer->packet_latest, &new_pkt_hdr, FRP_PKT_HDRSIZE);
583         // resend last packet
584         sent = frp_send_packet ((u_char*)&peer->packet_latest, peer->packet_latest_length, from);
585         if (!sent)
586         {
587             zlog_debug ("packet not sent");
588         }
589         return -1;
590     }
591     break;
592 }
593 return msg->msg_hdr.length;
594 }
595
596 int16_t
597 frp_incoming_config_msg (struct frp_msg_ipv4config* msg, struct frp_peer* peer, struct sockaddr_in* from, struct
598 in_addr local)
599 {
600     #ifdef DEB_DEBUG
601         zlog_debug ("DEB DEBUG: entering frpd.c - frp_incoming_config_msg");
602     #endif //DEB_DEBUG
603     // extract config info and update peer record
604     peer->cost = ntohs(msg->cost);
605     peer->poll = ntohs(msg->poll);
606     peer->retry = ntohs(msg->retry);
607     #ifdef DEB_DEBUG
608         zlog_debug ("DEB DEBUG: -- frp_incoming_control_msg - msg cost=%d, poll=%d, retry=%d", ntohs(msg->cost)
609         ntohs(msg->poll), ntohs(msg->retry));
610         zlog_debug ("DEB DEBUG: -- frp_incoming_control_msg - peer cost=%d, poll=%d, retry=%d", peer->cost,
611         peer->poll, peer->retry);
612     #endif //DEB_DEBUG
613     return msg->msg_hdr.length;
614 }

```

```
611
612 int16_t
613 frp_incoming_gateway_msg (struct frp_msg_ipv4gateway* msg, struct frp_peer* peer, struct sockaddr_in* from,
614 struct in_addr local)
615 {
616     #ifdef DEB_DEBUG
617         zlog_debug ("DEB DEBUG: entering frpd.c - frp_incoming_gateway_msg");
618     #endif //DEB_DEBUG
619     // extract config info and update peer record
620     int msg_length = msg->msg_hdr.length;
621     // (msg_length - msg_hdr - cost) / length of single address in path
622     int path_length = (msg_length - 32) / 32;
623     int remainder = (msg_length - 32) % 32;
624     struct in_addr* current_pos = (struct in_addr*)((char*)msg + 32);
625
626     if ((path_length == 0) || (remainder != 0))
627     { //not a valid path
628         peer->gateway_cost = FRP_INFINITY;
629         #ifdef DEB_DEBUG
630             zlog_debug ("DEB DEBUG: -- frp_incoming_control_msg - invalid path length - msg cost=%d", ntohs(msg->cost));
631         #endif //DEB_DEBUG
632         return msg->msg_hdr.length;
633     } else
634     { // create new list
635         struct list* gate_path = list_new ();
636         // step through addresses and store in list
637         for (int i = 1; i <= path_length; i++)
638         { // collect and format address
639             struct in_addr* current_addr = XCALLOC (MTYPE_FRP_PEER, sizeof (struct in_addr));
640             memcpy (current_addr, current_pos, sizeof (struct in_addr));
641             if (current_addr->s_addr == local.s_addr)
642             { // if this is our address, don't want this path because is using us as a gateway.
643                 peer->gateway_cost = FRP_INFINITY;
644                 // get rid of list and current_addr
645                 list_free (gate_path);
646                 XFREE (MTYPE_FRP_PEER, current_addr);
647                 #ifdef DEB_DEBUG
648                     zlog_debug ("DEB DEBUG: -- frp_incoming_control_msg - we are in list - msg cost=%d", ntohs(msg->cost));
649                 #endif //DEB_DEBUG
650                 return msg->msg_hdr.length;
651             } else
652             { // add address to list
653                 listnode_add (gate_path, current_addr);
654                 #ifdef DEB_DEBUG
655                     zlog_debug ("DEB DEBUG: -- frp_incoming_control_msg - path address=%d", current_addr->s_addr);
656                 #endif //DEB_DEBUG
657             }
658             current_pos++;
659         }
660         // pass list to peer record
661         peer->gateway_path = gate_path;
662         peer->gateway_cost = ntohs(msg->cost);
663         // check if this path is better than our current gateway path
664         if ((peer->gateway_cost + peer->cost) < frp->gateway_cost)
665         { frp->gateway_cost = peer->gateway_cost + peer->cost;
666           frp->gateway_nexthop = peer;
667           frp->gateway_path = peer->gateway_path;
668           // flag each peer to send out a new gateway message
669           struct listnode* node, *nnode;
670           for (ALL_LIST_ELEMENTS (frp_peers, node, nnode, peer))
671           { peer->flag_send_gateway = ON;
```

```
672     }
673 }
674 }
675 #ifdef DEB_DEBUG
676     zlog_debug ("DEB DEBUG: -- frp_incoming_control_msg - msg cost=%d", ntohs(msg->cost));
677     zlog_debug ("DEB DEBUG: -- frp_incoming_control_msg - peer cost=%d", peer->cost);
678 #endif //DEB_DEBUG
679     return msg->msg_hdr.length;
680 }
681
682 int16_t
683 frp_incoming_update_msg (struct frp_msg_ipv4update* msg, struct frp_peer* peer, struct sockaddr_in* from, struct in_addr local)
684 {
685     #ifdef DEB_DEBUG
686         zlog_debug ("DEB DEBUG: entering frpd.c - frp_incoming_update_msg");
687     #endif //DEB_DEBUG
688
689     struct frp_rte* new_route;
690     new_route = XCALLOC (MTYPE_FRP_PEER, sizeof (struct frp_rte));
691     struct frp_rte* existing_route;
692
693     // do we have a BEGIN flag? (using bitwise &)
694     if (msg->flags & FRP_FLAG_BEGIN)
695     { // have an existing batch?
696         if (peer->temp_rib != NULL)
697         { // kill existing temp rib
698             list_delete (peer->temp_rib);
699             peer->temp_rib = NULL;
700         }
701         // have a NULLRT?
702         if (msg->flags & FRP_FLAG_NULLRT)
703         { int peer_rib_size = peer->rib->count;
704           // clear peer rib
705           list_delete_all_node (peer->rib);
706           // re-compute our routing table
707           if (peer_rib_size > 0)
708           { frp_recompute_rib (peer);
709             return msg->msg_hdr.length;
710           }
711         }
712         // extract route information
713         new_route->flags = msg->flags;
714         new_route->length = msg->length;
715         new_route->route_cost = ntohs(msg->route_cost);
716         new_route->gate_cost = ntohs(msg->gate_cost);
717         // prefix requires a little more work
718         struct prefix_ipv4 address;
719         memset (&address, 0, sizeof (address));
720         address.family = AF_INET;
721         address.prefix = msg->prefix;
722         address.prefixlen = msg->length; //or the length provided
723         apply_mask_ipv4(&address);
724         new_route->prefix = address;
725
726         // have a DELETE?
727         if (msg->flags & FRP_FLAG_DELETE)
728         { // work thru peer rib looking for destination match to delete
729             struct listnode* node;
730             struct listnode* delete_node = NULL;
731             for (ALL_LIST_ELEMENTS_RO (peer->rib, node, existing_route))
732             }
```

```
733 // compare the two prefixes to see if they are the same host and the same network
734 { if (prefix_same (&new_route->prefix, &existing_route->prefix))
735 { delete_node = node;
736 break;
737 }
738 }
739 if (delete_node != NULL)
740 { list_delete_node (peer->rib, delete_node);
741 }
742 return msg->msg_hdr.length;
743 }
744
745 // have an UPDATE?
746 if (msg->flags & FRP_FLAG_UPDATE)
747 { // work thru peer rib looking for destination match to update
748 struct listnode* node;
749 struct listnode* update_node = NULL;
750 for (ALL_LIST_ELEMENTS_RO (peer->rib, node, existing_route))
751 // compare the two prefixes to see if they are the same host and the same network
752 { if (prefix_same (&new_route->prefix, &existing_route->prefix))
753 { update_node = node;
754 break;
755 }
756 }
757 if (update_node != NULL)
758 { list_delete_node (peer->rib, update_node);
759 listnode_add (peer->rib, new_route);
760 }
761 return msg->msg_hdr.length;
762 }
763
764 // create new temp rib
765 peer->temp_rib = list_new ();
766
767 } else // no BEGIN flag
768 {
769 // have an existing batch?
770 if (peer->temp_rib == NULL)
771 { zlog_info ("invalid route update packet");
772 return msg->msg_hdr.length;
773 }
774 }
775 }
776
777 // add info to temp rib
778 listnode_add (peer->temp_rib, new_route);
779
780 // have a COMMIT?
781 if (msg->flags & FRP_FLAG_COMMIT)
782 { // replace peer rib with temp rib
783 list_delete (peer->rib);
784 peer->rib = peer->temp_rib;
785 peer->temp_rib = NULL;
786 frp_recompute_rib (peer);
787 }
788
789 return msg->msg_hdr.length;
790 }
791
792
793
```

```
794
795
796 // _____FRP_EVENT_UPDATE_____
797
798 /* Execute an event update */
799 static int
800 frp_update_peers (struct thread *t)
801 {
802 #ifdef DEB_DEBUG
803 zlog_debug ("DEB DEBUG: entering frpd.c - frp_update_peers");
804 #endif //DEB_DEBUG
805 int interval;
806 /* clear thred pointer */
807 frp->t_update = NULL;
808 /* cancel interval timer */
809 if (frp->t_update_interval)
810 { thread_cancel (frp->t_update_interval);
811 frp->t_update_interval = NULL;
812 }
813 frp->update_trigger = 0;
814 /* logging triggered update */
815 if (IS_FRP_DEBUG_EVENT)
816 { zlog_debug ("update triggered");
817 }
818
819 // for each peer, unless quiescent
820 if ((frp->gateway_cost == FRP_INFINITY) || (frp->gateway_cost == 0))
821 {
822 struct frp_peer* peer;
823 struct listnode *node, *nnode;
824 for (ALL_LIST_ELEMENTS (frp_peers, node, nnode, peer))
825 {
826 // set update flag
827 peer->flag_send_update = ON;
828
829 // if flag_awaiting_ack is off
830 if (peer->flag_awaiting_ack == OFF)
831 {
832 // send SYN packet to peer
833 int sent = build_syn_pkt (peer);
834 if (sent)
835 { zlog_debug ("sent SYN packet to %s", inet_ntoa (peer->address));
836 }
837 }
838 }
839 }
840
841 /* After a triggered update is sent, a timer should be set for a
842 random interval between 1 and 5 seconds. If other changes that
843 would trigger updates occur before the timer expires, a single
844 update is triggered when the timer expires */
845 interval = (random () % 5) + 1;
846 frp->t_update_interval = thread_add_timer (master, frp_update_interval, NULL, interval);
847 return 0;
848 }
849
850 /* Event update interval timer */
851 static int
852 frp_update_interval (struct thread *t)
853 { int frp_update_peers (struct thread *);
854 frp->t_update_interval = NULL;
```

```

855     if (frp->update_trigger)
856     { frp->update_trigger = 0;
857       frp_update_peers (t);
858     }
859     return 0;
860 }
861
862 // -----FRP_EVENT_POLL-----
863
864 /* frp's periodical timer */
865 static int
866 frp_poll_peers (struct thread *t)
867 { /* clear timer pointer */
868   frp->t_poll = NULL;
869   if (IS_FRP_DEBUG_EVENT)
870   { zlog_debug ("poll timer fired");
871     }
872   // check for peers with update flags set
873   // check for new peers and send a SYN packet
874   // check for peers with update flags set and send a batch message
875   // check for peers that haven't responded and resend last message or declare dead
876   int sent;
877   struct frp_peer* peer;
878   struct listnode *node, *nnode;
879   for (ALL_LIST_ELEMENTS (frp_peers, node, nnode, peer))
880   { // if peer hasn't responded for X polls, declare dead and move on - set flags so handle resurrection correctly
881     if ((peer->flag_alive == OFF) || (peer->flag_awaiting_ack > FRP_PEER_DEAD))
882     { peer->flag_alive = OFF;
883       peer->flag_awaiting_ack = OFF;
884       peer->flag_send_poll = OFF;
885       peer->flag_send_syn = OFF;
886       peer->flag_send_config = ON;
887       peer->flag_send_update = ON;
888       peer->flag_awaiting_ack = ON;
889       #ifdef DEB_DEBUG
890       zlog_debug ("DEB DEBUG: -- frp_poll_peers - %s is dead", inet_ntoa (peer->address));
891       #endif //DEB_DEBUG
892     } else // ie: not dead
893     { if (peer->flag_awaiting_ack == 1)
894       { peer->flag_awaiting_ack ++;
895       } else if (peer->flag_awaiting_ack > 1)
896       { // if a peer hasn't acked for more than a poll cycle, resend last message
897         struct sockaddr_in socket;
898         socket.sin_family = AF_INET;
899         socket.sin_port = htons(FRP_PORT_DEFAULT);
900         socket.sin_addr = peer-> address;
901         sent = frp_send_packet (peer->packet_latest, peer->packet_latest_length, &socket);
902         if (sent)
903         { peer->flag_awaiting_ack ++;
904           } else
905           { zlog_debug ("packet not sent");
906             }
907         }
908         #ifdef DEB_DEBUG
909         zlog_debug ("DEB DEBUG: -- frp_poll_peers - awaiting ACK %d from %s", peer->packet_latest_lseq, ir
910 (peer->address));
911         #endif //DEB_DEBUG
912       } else if ((peer->flag_send_syn) || (peer->flag_send_config) || (peer->flag_send_gateway) ||
913 (peer->flag_send_update))
914       { // if flags are set, send a batch packet
915         sent = build_syn_pkt (peer);

```

```

916     if (sent)
917     { peer->flag_send_syn = OFF;
918       zlog_debug ("sent SYN packet to %s", inet_ntoa (peer->address));
919     }
920   } else
921   { // have we heard from the peer within a poll period? if not, poll
922     time_t now = time (NULL);
923     double test = difftime (now, peer->time_last_heard);
924     if (test > frp->poll)
925     { peer->flag_send_poll = ON;
926       sent = build_syn_pkt (peer);
927       if (sent)
928       { peer->flag_send_syn = OFF;
929         zlog_debug ("sent SYN packet to %s", inet_ntoa (peer->address));
930       }
931     }
932   }
933 }
934
935 /* polls may be suppressed if a regular update is due by the time the polls would be sent */
936 if (frp->t_update_interval)
937 { thread_cancel (frp->t_update_interval);
938   frp->t_update_interval = NULL;
939 }
940 frp->update_trigger = 0;
941 /* register myself */
942 frp_event (FRP_EVENT_POLL, 0);
943 return 0;
944 }
945
946 // -----
947 // ROUTER COMMAND DEFUN
948 // -----
949
950 // Create the FRP router config command - frpd(config)#
951 DEFUN (router_frp,
952        router_frp_cmd,
953        "router frp",
954        "Enable a routing process\n"
955        "Fringe Routing Protocol (FRP)\n")
956 {
957   int ret;
958   /* if frp is not enabled before */
959   if (! frp)
960   {
961     #ifdef DEB_DEBUG
962     fprintf (stderr, "DEB DEBUG: entering frpd.c - router_frp\n");
963     #endif //DEB_DEBUG
964     ret = frp_create ();
965     if (ret < 0)
966     { zlog_info ("Can't create FRP");
967       return CMD_WARNING;
968     }
969   }
970 }

```

```
977 vty->node = FRP_NODE;
978 vty->index = frp;
979 return CMD_SUCCESS;
980 }
981 DEFUN (no_router_frp,
982 no_router_frp_cmd,
983 "no router frp",
984 NO_STR
985 "Enable a routing process\n"
986 "Routing Information Protocol (FRP)\n")
987 {
988 if (frp)
989 frp_clean ();
990 return CMD_SUCCESS;
991 }
992
993 // is this router a gateway>
994 DEFUN (frp_gateway,
995 frp_gateway_cmd,
996 "gateway (yes|no)",
997 "Is this router a FRP gateway? (yes|no)\n"
998 "\n"
999 "\n")
1000 {
1001 const char* test1 = "yes";
1002 const char* test2 = argv[0];
1003 int result = strcmp(test1, test2);
1004 if (result == 0 || (*argv[0] == 'y'))
1005 { frp->is_gateway_flag = ON;
1006 frp->gateway_cost = 0;
1007 } else
1008 { frp->is_gateway_flag = OFF;
1009 frp->gateway_cost = FRP_INFINITY;
1010 }
1011 // flag each peer to send out a new gateway message
1012 struct frp_peer* peer;
1013 struct listnode *node, *nnode;
1014 for (ALL_LIST_ELEMENTS (frp_peers, node, nnode, peer))
1015 { frp_recompute_rib (peer);
1016 peer->flag_send_gateway = ON;
1017 peer->flag_send_update = ON;
1018 }
1019 #ifdef DEB_DEBUG
1020 vty_out (vty, "DEB DEBUG: -- frp_gateway_cmd - result=%d, frp->is_gateway_flag=%d %s", result,
1021 frp->is_gateway_flag, VTY_NEWLINE);
1022 #endif //DEB_DEBUG
1023 return CMD_SUCCESS;
1024 }
1025
1026 // set our secret
1027 // - this is only set in the conf file, not on the command line
1028 // - so the command is not installed
1029 DEFUN (frp_secret,
1030 frp_secret_cmd,
1031 "secret WORD",
1032 "\n"
1033 "\n")
1034 {
1035 if (!strcmp(argv[0], ""))
1036 { vty_out (vty, "Please specify address and secret A.B.C.D secret WORD%s", VTY_NEWLINE);
1037 return CMD_WARNING;
```

```
1038 }
1039 char* secret = (char*)argv[0];
1040 frp->secret = (char*) XCALLOC (MTYPE_FRP, strlen(secret) + 1);
1041 memcpy (frp->secret, secret, strlen(secret));
1042 #ifdef DEB_DEBUG
1043 vty_out (vty, "DEB DEBUG: -- frp_secret_cmd - frp->secret = %s %s", frp->secret, VTY_NEWLINE);
1044 #endif //DEB_DEBUG
1045 return CMD_SUCCESS;
1046 }
1047
1048 // set the cost of the link
1049 DEFUN (frp_cost,
1050 frp_cost_cmd,
1051 "cost INT",
1052 "Cost of link\n"
1053 "Set cost of link as an integer\n")
1054 {
1055 // set cost
1056 frp->cost = atoi(argv[0]);
1057 // flag each peer to send out a new config message
1058 struct frp_peer* peer;
1059 struct listnode *node, *nnode;
1060 for (ALL_LIST_ELEMENTS (frp_peers, node, nnode, peer))
1061 { peer->flag_send_config = ON;
1062 }
1063 #ifdef DEB_DEBUG
1064 vty_out (vty, "DEB DEBUG: -- frp_cost_cmd - frp->cost = %d %s", frp->cost, VTY_NEWLINE);
1065 #endif //DEB_DEBUG
1066 return CMD_SUCCESS;
1067 }
1068
1069 // set the poll time
1070 DEFUN (frp_poll,
1071 frp_poll_cmd,
1072 "poll INT",
1073 "Poll / keepalive frequency\n"
1074 "Set the poll time in seconds as an integer\n")
1075 {
1076 // set poll
1077 frp->poll = atoi(argv[0]);
1078 // flag each peer to send out a new config message
1079 struct frp_peer* peer;
1080 struct listnode *node, *nnode;
1081 for (ALL_LIST_ELEMENTS (frp_peers, node, nnode, peer))
1082 { peer->flag_send_poll = ON;
1083 }
1084 #ifdef DEB_DEBUG
1085 vty_out (vty, "DEB DEBUG: -- frp_poll_cmd - frp->poll = %d %s", frp->poll, VTY_NEWLINE);
1086 #endif //DEB_DEBUG
1087 return CMD_SUCCESS;
1088 }
1089
1090 // set the retry time
1091 DEFUN (frp_retry,
1092 frp_retry_cmd,
1093 "retry INT",
1094 "Timeout to failure after acked packet\n"
1095 "Set the retry time in seconds as an integer\n")
1096 {
1097 // set retry
1098 frp->retry = atoi(argv[0]);
```



```
1099 // flag each peer to send out a new config message
1100 struct frp_peer* peer;
1101 struct listnode *node, *nnode;
1102 for (ALL_LIST_ELEMENTS (frp_peers, node, nnode, peer))
1103 { peer->flag_send_config = ON;
1104 }
1105 #ifdef DEB_DEBUG
1106 vty_out (vty, "DEB DEBUG: -- frp_retry_cmd - frp->retry = %d %s", frp->retry, VTY_NEWLINE);
1107 #endif //DEB_DEBUG
1108 return CMD_SUCCESS;
1109 }
1110
1111
1112
1113
1114 // -----
1115 // INIT FUNCTION
1116 // -----
1117 // initialise frp structure and set commands
1118 void
1119 frp_init (void)
1120 {
1121 #ifdef DEB_DEBUG
1122 fprintf (stderr, "DEB DEBUG: entering frpd.c - frp_init\n");
1123 #endif //DEB_DEBUG
1124 /* Install FRP_NODE */
1125 // DEB COMMENT: seems to be the initial install point, used in all other daemons (although not necessarily in the
1126 daemon.c)
1127 install_node (&frp_node, frp_config_write); // DEB COMMENT: install_node command.h [33:
1128
1129 // DEB COMMENT: seems to be the place to install frp specific commands)
1130 install_default (FRP_NODE); // DEB COMMENT: install_default command.h [3
1131 // DEB COMMENT: FRP_NODE needs to be define
1132
1133 // create the frp router config command - frpd(config)#
1134 install_element (CONFIG_NODE, &router_frp_cmd);
1135 install_element (CONFIG_NODE, &no_router_frp_cmd);
1136
1137 // install the frp router commands
1138 install_element (FRP_NODE, &frp_secret_cmd);
1139 install_element (FRP_NODE, &frp_cost_cmd);
1140 install_element (FRP_NODE, &frp_poll_cmd);
1141 install_element (FRP_NODE, &frp_retry_cmd);
1142 install_element (FRP_NODE, &frp_gateway_cmd);
1143
1144 /* initialise frp debugging commands and functions */
1145 frp_debug_init ();
1146
1147 /* initialise frp interface related commands and functions */
1148 frp_interface_init ();
1149
1150 /* initialise frp neighbour related commands and functions */
1151 frp_peer_init ();
1152
1153 #ifdef DEB_DEBUG
1154 fprintf (stderr, "DEB DEBUG: -- leaving frp_init\n");
1155 #endif //DEB_DEBUG
1156
1157 return;
1158 }
1159 }
```



```
1 // -----
2 // INCLUDES
3 // -----
4 #include "frpd.h"
5 #include "frp_debug.h"
6 #include "linklist.h"
7
8 // to support code lifted from Don's implementation
9 #include <openssl/sha.h>
10
11 #define IPADDR u_int32_t
12
13
14
15
16 // -----
17 // GLOBAL VARIABLES
18 // -----
19
20 /* linked list of frp peers */
21 struct list* frp_peers = NULL;
22
23 /* static prototypes */
24 static int frp_neighbor_delete (struct prefix_ipv4 *p);
25
26 static void frp_peer_free (struct frp_peer *peer);
27 static int frp_frps_cmp (struct frp_peer *p1, struct frp_peer *p2);
28 static struct frp_peer* frp_peer_add (struct in_addr* addr, const char* secret);
29 static struct frp_peer* frp_peer_new (void);
30
31
32 /* prototypes */
33 struct frp_peer* frp_peer_lookup (struct in_addr *addr);
34 struct frp_peer* frp_peer_lookup_next (struct in_addr *addr);
35
36
37
38 // -----
39 // FUNCTIONS
40 // -----
41
42 // ----SUPPORT-----
43
44
45 struct frp_peer*
46 frp_peer_new (void)
47 { return XCALLOC (MTYPE_FRP_PEER, sizeof (struct frp_peer)); }
48
49
50 void
51 frp_peer_free (struct frp_peer *peer)
52 { XFREE (MTYPE_FRP_PEER, peer); }
53
54
55 struct frp_peer*
56 frp_peer_lookup (struct in_addr *addr)
57 { struct frp_peer *peer;
58   struct listnode *node, *nnode;
59   for (ALL_LIST_ELEMENTS (frp_peers, node, nnode, peer))
60     { if (IPV4_ADDR_SAME (&peer->address, addr))
61       return peer;
62     }
```

```
62 }
63 return NULL;
64 }
65
66 struct frp_peer*
67 frp_peer_lookup_next (struct in_addr *addr)
68 { struct frp_peer *peer;
69   struct listnode *node, *nnode;
70   for (ALL_LIST_ELEMENTS (frp_peers, node, nnode, peer))
71     { if (htonl (peer->address.s_addr) > htonl (addr->s_addr))
72       return peer;
73     }
74 return NULL;
75 }
76
77 static struct frp_peer*
78 frp_peer_add (struct in_addr* addr, const char* secret)
79 { struct frp_peer *peer;
80   peer = frp_peer_lookup (addr);
81   if (peer)
82     { peer->secret = secret;
83     } else
84     { peer = frp_peer_new ();
85       peer->address = *addr;
86       peer->secret = (char*) XCALLOC (MTYPE_FRP_PEER, strlen(secret) + 1);
87       memcpy (peer->secret, secret, strlen(secret));
88       peer->packet_latest = (u_char*) XCALLOC (MTYPE_FRP_PEER, FRP_PKT_MAXSIZE);
89       listnode_add_sort (frp_peers, peer);
90       peer->rib = list_new ();
91       peer->temp_rib = NULL;
92     }
93 return peer;
94 }
95
96 int
97 frp_frps_cmp (struct frp_peer *p1, struct frp_peer *p2)
98 { return htonl (p1->address.s_addr) > htonl (p2->address.s_addr); }
99
100
101 struct prefix*
102 find_local_address_for_peer(struct in_addr dest)
103 { struct interface* ifp;
104   struct listnode* node;
105   if (iflist == NULL)
106     {
107       #ifdef DEB_DEBUG
108       zlog_debug ("DEB DEBUG: -- find_local_address_for_peer - iflist=null");
109       #endif //DEB_DEBUG
110     }
111   /* Check each interface. */
112   for (ALL_LIST_ELEMENTS_RO (iflist, node, ifp))
113     { struct connected* conn = connected_lookup_address (ifp, dest);
114       if (conn != NULL)
115         { return conn->address;
116         } else
117         {
118           #ifdef DEB_DEBUG
119           zlog_debug ("DEB DEBUG: -- find_local_address_for_peer - conn=null");
120           #endif //DEB_DEBUG
121         }
122     }
```

```
123     return NULL;
124 }
125
126 // ----SECURITY-----
127 // LIFTED STRAIGHT FROM DON'S CODE
128
129 /*
130 Security stuff
131 Hash calculation: SHA1 over packet, IP addresses & secret
132 Note that only the first 64 bits (FRP_HASHSIZE) of the hash are used.
133 */
134 // #include <sha.h>
135 u_int8_t*
136 dohash(u_char* buf, int len, const char* secret, IPADDR sa, IPADDR da)
137 { static u_int8_t hash[SHA_DIGEST_LENGTH];
138   SHA_CTX shctx;
139   SHA_Init(&shctx);
140   SHA_Update(&shctx, buf, len);
141   SHA_Update(&shctx, (u_char*)&sa, sizeof(IPADDR));
142   SHA_Update(&shctx, (u_char*)&da, sizeof(IPADDR));
143   SHA_Update(&shctx, secret, strlen(secret));
144   SHA_Final(hash, &shctx);
145   #ifdef DEB_DEBUG_PEER
146     zlog_debug ("DEB DEBUG: -- dohash - hash: %d%d%d%d%d%d%d%d%d%d%d%d%d%d",
147               hash[0],hash[1],hash[2],hash[3],hash[4],hash[5],hash[6],hash[7],hash[8],hash[9],
148               hash[10],hash[11],hash[12],hash[13],hash[14],hash[15],hash[16],hash[17],hash[18],hash[19]);
149   #endif //DEB_DEBUG
150   return hash;
151 }
152
153 /*
154 Check that hash matches packet hash.
155 Do not call if peer->secret is null.
156 */
157 int
158 checksecure(u_char* pkt, int len, struct frp_peer* peer, struct in_addr local)
159 { u_int8_t *hash;
160   hash = dohash(pkt + FRP_HASHSIZE, len - FRP_HASHSIZE, peer->secret, peer->address.s_addr, local.s_addr);
161   #ifdef DEB_DEBUG_PEER
162     zlog_debug ("DEB DEBUG: -- checksecure - packet len: %d", len);
163     zlog_debug ("DEB DEBUG: -- checksecure - checking from: %s secret: %s", inet_ntoa(peer->address),
164 peer->secret);
165     zlog_debug ("DEB DEBUG: -- checksecure - checking to: %s secret: %s", inet_ntoa(local), peer->secret);
166   #endif //DEB_DEBUG
167   if(!memcmp(hash, pkt, FRP_HASHSIZE))
168   { return 1;
169   }
170   zlog_debug("packet security failure");
171   return 0;
172 }
173
174 /*
175 Compute packet hash
176 Or zero hash field if secret is null
177 */
178 // NOT ACTUALLY USED
179 void
180 secure(struct frp_peer* peer, struct in_addr local, struct frp_pkt_hdr* pkt, int len)
181 { u_int8_t *hash;
182   if(!peer->secret)
183   { memset(pkt->hash, 0, FRP_HASHSIZE);
```

```
184     return;
185   }
186   hash = dohash(&pkt->hash[FRP_HASHSIZE], len-FRP_HASHSIZE, peer->secret, local.s_addr, peer->address.s_addr);
187   memcpy(pkt->hash, hash, FRP_HASHSIZE);
188 }
189
190 // ----NEIGHBOURS-----
191 /* add new frp neighbor to struct route_table *neighbors */
192 int
193 frp_neighbor_add (struct vty *vty, const char* ip_str, const char* secret)
194 {
195   #ifdef DEB_DEBUG
196     vty_out (vty, "DEB DEBUG: entering frp_peer.c - frp_neighbor_add %s %s %s", ip_str, secret, VTY_NEWLINE);
197     zlog_debug ("DEB DEBUG: entering frp_peer.c - frp_neighbor_add %s secret %s", ip_str, secret);
198   #endif //DEB_DEBUG
199
200   /* convert incoming peer address to a struct prefix */
201   struct prefix_ipv4 p;
202   int ret = str2prefix_ipv4 (ip_str, &p); // DEB COMMENT: prefix.c [217] (also v6 & combiner
203   // check have a valid address
204   if (ret <= 0)
205   {
206     #ifdef DEB_DEBUG
207       zlog_debug ("DEB DEBUG: -- frp_neighbor_add - invalid address");
208     #endif //DEB_DEBUG
209     vty_out (vty, "Please specify address as A.B.C.D%s", VTY_NEWLINE);
210     return CMD_WARNING;
211   }
212   // check have a secret
213   if (!strcmp(secret, ""))
214   {
215     #ifdef DEB_DEBUG
216       zlog_debug ("DEB DEBUG: -- frp_neighbor_add - invalid secret");
217     #endif //DEB_DEBUG
218     vty_out (vty, "Please specify address and secret as A.B.C.D secret WORD%s", VTY_NEWLINE);
219     return CMD_WARNING;
220   }
221
222   // store peer in struct frp.neighbors
223   struct route_node* node;
224   node = route_node_get (frp->neighbors, (struct prefix *) &p); // DEB COMMENT: table.c [272]
225   if (node->info) // if peer already exists
226   {
227     #ifdef DEB_DEBUG
228       zlog_debug ("DEB DEBUG: -- frp_neighbor_add - peer already exists");
229     #endif //DEB_DEBUG
230     route_unlock_node (node);
231     return -1;
232   }
233   node->info = frp->neighbors;
234
235   // convert the struct prefix to an in_addr
236   struct prefix* node_p;
237   struct in_addr address;
238   node_p = (struct prefix*)&node->p;
239   #ifdef DEB_DEBUG
240     zlog_debug ("DEB DEBUG: -- frp_neighbor_add - node_p=%s", inet_ntoa (node_p->u.prefix4));
241   #endif //DEB_DEBUG
242   int addr_test = inet_pton (AF_INET, inet_ntoa (node_p->u.prefix4), &address);
243
244   // store address and secret of peer in list frp_peers
```

```
245 struct frp_peer* peer;
246 if (addr_test)
247 { peer = frp_peer_add (&address, secret);
248 #ifdef DEB_DEBUG
249 zlog_debug ("DEB DEBUG: -- frp_neighbor_add, %s, %s added to frp_peers", inet_ntoa (peer->address),
250 peer->secret);
251 #endif //DEB_DEBUG
252 } else
253 {
254 #ifdef DEB_DEBUG
255 zlog_debug ("DEB DEBUG: -- frp_neighbor_add, add to frp_peers failed");
256 #endif //DEB_DEBUG
257 route_unlock_node (node);
258 return -1;
259 }
260
261 // set peer defaults
262 //peer->lseq = NEW_SEQ_NO;
263 peer->lseq = RANDOM_INT(0,INT_MAX);
264 peer->rseq = 0;
265 peer->gateway_path = NULL;
266 peer->gateway_cost = FRP_INFINITY;
267 if ((frp->gateway_cost == FRP_INFINITY) || (frp->gateway_cost == 0))
268 { peer->flag_send_gateway = OFF;
269 peer->flag_send_update = OFF;
270 } else
271 { peer->flag_send_gateway = ON;
272 peer->flag_send_update = ON;
273 }
274 peer->flag_alive = ON;
275 peer->flag_send_syn = ON;
276 peer->flag_send_config = ON;
277 // route_unlock_node (node);
278 return 0;
279 }
280
281 /* delete a frp neighbor from struct route_table *neighbors */
282 int
283 frp_neighbor_delete (struct prefix_ipv4 *p)
284 { struct route_node *node;
285 /* lock for look up */
286 node = route_node_lookup (frp->neighbors, (struct prefix *) p);
287 if (! node)
288 { return -1;
289 }
290 node->info = NULL;
291 /* unlock lookup lock */
292 route_unlock_node (node);
293 /* unlock real neighbor information lookup */
294 route_unlock_node (node);
295 return 0;
296 }
297
298
299
300
301
302
303 // -----
304 // ROUTER COMMAND DEFUNs
305 // -----
```

```
306
307 DEFUN (frp_neighbor,
308 frp_neighbor_cmd,
309 "neighbor A.B.C.D secret WORD",
310 "Specify a neighbor router\n"
311 "Neighbor address\n"
312 "Neighbor secret\n"
313 "the secret\n")
314 {
315 int ret = frp_neighbor_add (vty, argv[0], argv[1]);
316 if (ret == 0)
317 {
318 #ifdef DEB_DEBUG
319 vty_out (vty, "DEB DEBUG: -- successfully called frp_neighbor_add%s", VTY_NEWLINE);
320 #endif //DEB_DEBUG
321 return CMD_SUCCESS;
322 } else if (ret == -1)
323 {
324 #ifdef DEB_DEBUG
325 vty_out (vty, "DEB DEBUG: -- frp_neighbor_add - address already exists%s", VTY_NEWLINE);
326 #endif //DEB_DEBUG
327 return ret;
328 }
329 return ret;
330 }
331
332 DEFUN (no_frp_neighbor,
333 no_frp_neighbor_cmd,
334 "no neighbor A.B.C.D",
335 NO_STR
336 "Specify a neighbor router\n"
337 "Neighbor address\n")
338 { int ret;
339 struct prefix_ipv4 p;
340 ret = str2prefix_ipv4 (argv[0], &p);
341 if (ret <= 0)
342 { vty_out (vty, "Please specify address by A.B.C.D%s", VTY_NEWLINE);
343 return CMD_WARNING;
344 }
345 frp_neighbor_delete (&p);
346 return CMD_SUCCESS;
347 }
348
349
350
351 // -----
352 // INIT FUNCTION
353 // -----
354 void
355 frp_peer_init (void)
356 {
357 #ifdef DEB_DEBUG
358 fprintf (stderr, "DEB DEBUG: entering frp_peer.c - frp_peer_init\n");
359 #endif //DEB_DEBUG
360
361 //initialise the random number generation
362 RANDOM_SEED();
363
364 // initialise peer secret structure
365 frp_peers = list_new ();
366 frp_peers->cmp = (int (*)(void *, void *)) frp_frp_peers_cmp;
```

```
367
368 // create the FRP neighbor command
369 install_element (FRP_NODE, &frp_neighbor_cmd);
370 install_element (FRP_NODE, &no_frp_neighbor_cmd);
371
372 #ifdef DEB_DEBUG
373     fprintf (stderr, "DEB DEBUG: -- leaving frp_peer_init\n");
374 #endif //DEB_DEBUG
375
376 return;
377 }
378
```

```
1 #define FRP_PKT_HDRSIZE 16 // (128 bits)
2 #define FRP_PKT_MINSIZE 16 // (128 bits)
3 #define FRP_PKT_MAXSIZE 1400 // as specified by Don
4 #define FRP_MSG_CONTROL_SIZE 4 // (32 bits)
5 #define FRP_MSG_IPV4CONFIG_SIZE 12 // (96 bits)
6 #define FRP_MSG_IPV4GATEWAY_MINSIZE 8 // (64 bits)
7 #define FRP_MSG_IPV4UPDATE_SIZE 12 // (96 bits)
8 #define FRP_MSG_IPV6CONFIG_SIZE 24 // (192 bits)
9 #define FRP_MSG_IPV6GATEWAY_MINSIZE 20 // (160 bits)
10 #define FRP_MSG_IPV6UPDATE_SIZE 24 // (192 bits)
11
12 // frp message types
13 #define FRP_MSG_CONTROL 0x01
14 #define FRP_MSG_IPV4CONFIG 0x41
15 #define FRP_MSG_IPV4GATEWAY 0x42
16 #define FRP_MSG_IPV4UPDATE 0x43
17 #define FRP_MSG_IPV6CONFIG 0x61
18 #define FRP_MSG_IPV6GATEWAY 0x62
19 #define FRP_MSG_IPV6UPDATE 0x63
20
21 // frp control types
22 #define FRP_CTRL_POLL 1
23 #define FRP_CTRL_ACK 2
24 #define FRP_CTRL_NAK 3
25
26 // frp update flags
27 #define FRP_FLAG_BEGIN 0x01
28 #define FRP_FLAG_COMMIT 0x02
29 #define FRP_FLAG_NULLRT 0x04
30 #define FRP_FLAG_UPDATE 0x08
31 #define FRP_FLAG_DELETE 0x10
32 #define FRP_FLAG_GATEWAY 0x80
33
34 // frp packet header (128 bits)
35 struct frp_pkt_hdr
36 {
37     u_int8_t hash[8]; // security hash (64 bits)
38     u_int32_t sendSeq; // sender's sequence number (32 bits)
39     u_int32_t recipAck; // recipient's acknowledgement number (32 bits)
40 };
41
42 // frp header (16 bits)
43 struct frp_msg_hdr
44 {
45     u_int8_t length; // message length (8 bits)
46     u_int8_t type; // message type (8 bits)
47 };
48
49 // frp control (32 bits)
50 struct frp_msg_control
51 {
52     struct frp_msg_hdr msg_hdr; // message header (16 bits)
53     u_int8_t type; // control type (8 bits)
54     u_int8_t param; // control parameters (8 bits)
55 };
56
57 // frp IPv4 configuration (96 bits)
58 struct frp_msg_ipv4config
59 {
60     struct frp_msg_hdr msg_hdr; // message header (16 bits)
61     u_short cost; // cost of the link (16 bits)
62     u_short poll; // poll time (16 bits)
63     u_short retry; // retry time (16 bits)
64     struct in_addr id; // router-ID of peer (32 bits)
```

```
62 };
63
64 // frp IPv4 path to gateway (64-2016 bits)
65 struct frp_msg_ipv4gateway
66 {
67     struct frp_msg_hdr msg_hdr; // message header (16 bits)
68     u_short cost; // cost from peer to gateway (16 bits)
69     struct in_addr path[62]; // path from peer to gateway (32xn bits, 1 <= n <= 62)
70 };
71
72 // frp IPv4 route update (96 bits)
73 struct frp_msg_ipv4update
74 {
75     struct frp_msg_hdr msg_hdr; // message header (16 bits)
76     u_int8_t flags; // update type flags (8 bits)
77     u_int8_t length; // prefix length (8 bits)
78     u_short routecost; // route cost (16 bits)
79     u_short gatecost; // cost from originator to gateway (16 bits)
80     struct in_addr prefix; // IP prefix (32 bits)
81 };
82
83 // frp IPv6 configuration (192 bits)
84 struct frp_msg_ipv6config
85 {
86     struct frp_msg_hdr msg_hdr; // message header (16 bits)
87     u_short cost; // cost of the link (16 bits)
88     u_short poll; // poll time (16 bits)
89     u_short retry; // retry time (16 bits)
90     struct in6_addr id; // router-ID of peer (128 bits)
91 };
92
93 // frp IPv6 path to gateway (160-7968 bits)
94 struct frp_msg_ipv6gateway
95 {
96     struct frp_msg_hdr msg_hdr; // message header (16 bits)
97     u_short cost; // cost from peer to gateway (16 bits)
98     struct in6_addr path[1]; // path from peer to gateway (32xn bits, 1 <= n <= 62)
99 };
100
101 // frp IPv6 route update (192 bits)
102 struct frp_msg_ipv6update
103 {
104     struct frp_msg_hdr msg_hdr; // message header (16 bits)
105     u_int8_t flags; // update type flags (8 bits)
106     u_int8_t length; // prefix length (8 bits)
107     u_short routecost; // route cost (16 bits)
108     u_short gatecost; // cost from originator to gateway (16 bits)
109     struct in6_addr prefix; // IP prefix (128 bits)
110 };
111
112 // -----
113 // GLOBAL VARIABLES AND DEFINES
114 // -----
115
116 struct frp_pkt_hdr pkt_hdr;
117 struct frp_pkt_hdr* pkt_hdr_ptr;
118
119 /* frp event. */
120 enum frp_makepkthdr_flag
121 {
122     FRP_ACK,
123     FRP_SYN,
124 };
```

```
1 // -----
2 // INCLUDES
3 // -----
4 #include "frpd.h"
5 #include "frp_debug.h"
6
7
8
9 // -----
10 // FUNCTIONS
11 // -----
12
13 /* frp packet send to destination address, on interface denoted by
14 * by connected argument. NULL to argument denotes destination should be
15 * should be frp multicast group
16 */
17 // DEB COMMENT: no multicast in frp
18 int
19 //frp_send_packet (u_char* buf, int size, struct sockaddr_in *to, struct connected *ifc)
20 frp_send_packet (u_char* buf, int size, struct sockaddr_in *to)
21 { int ret, send_sock;
22   struct sockaddr_in sin;
23   send_sock = frp->sock;
24   sin.sin_port = htons (FRP_PORT_DEFAULT);
25   sin.sin_addr.s_addr = htonl (0);
26   // assert (ifc != NULL);
27   if (IS_FRP_DEBUG_PACKET)
28   {
29 #define ADDRESS_SIZE 20
30     char dst[ADDRESS_SIZE];
31     dst[ADDRESS_SIZE - 1] = '\0';
32     if (to)
33     { strncpy (dst, inet_ntoa(to->sin_addr), ADDRESS_SIZE - 1);
34     }
35 #undef ADDRESS_SIZE
36 //     zlog_debug("frp_send_packet %s > %s (%s)", inet_ntoa(ifc->address->u.prefix4), dst, ifc->ifp->name);
37     zlog_debug("frp_send_packet %s", dst);
38   }
39   /* Make destination address. */
40   memset (&sin, 0, sizeof (struct sockaddr_in));
41   sin.sin_family = AF_INET;
42 #ifdef HAVE_STRUCT_SOCKADDR_IN_SIN_LEN
43   sin.sin_len = sizeof (struct sockaddr_in);
44 #endif /* HAVE_STRUCT_SOCKADDR_IN_SIN_LEN */
45   /* When destination is specified, use it's port and address. */
46   if (to)
47   { sin.sin_port = to->sin_port;
48     #ifdef DEB_DEBUG
49     zlog_debug ("DEB DEBUG: -- frp_send_packet, port = %d", ntohs (sin.sin_port));
50     #endif //DEB_DEBUG
51     sin.sin_addr = to->sin_addr;
52     send_sock = frp->sock;
53   }
54   ret = sendto (send_sock, buf, size, 0, (struct sockaddr *)&sin,
55     sizeof (struct sockaddr_in));
56   if (IS_FRP_DEBUG_EVENT)
57     zlog_debug ("SEND to %s.%d", inet_ntoa(sin.sin_addr), ntohs (sin.sin_port));
58   if (ret < 0)
59     zlog_warn ("can't send packet : %s", safe_strerror (errno));
60   if (!to)
61     close(send_sock);
```

```
62   return ret;
63 }
64
65 struct frp_pkt_hdr
66 make_frp_pkt_hdr (struct in_addr local, struct frp_peer* peer, int flag, int len, u_char* buf)
67 { struct frp_pkt_hdr pkt_hdr;
68   u_int8_t* hash;
69   u_int32_t seq_no;
70   u_int32_t ack_no;
71   #ifdef DEB_DEBUG
72     zlog_debug ("DEB DEBUG: entering frp_packet.c - make_frp_pkt_hdr - flag=%d", flag);
73   #endif //DEB_DEBUG
74   // create and fill SYN and ACK fields - store current lseq for this peer
75   #ifdef DEB_DEBUG_PEER
76     zlog_debug ("DEB DEBUG: -- make_frp_pkt_hdr - peer->lseq=%u, peer->rseq=%u", peer->lseq, peer->rseq);
77   #endif //DEB_DEBUG
78   seq_no = peer->lseq + 1;
79   pkt_hdr.sendSeq = htonl(seq_no);
80   peer->lseq = seq_no;
81   #ifdef DEB_DEBUG_PEER
82     zlog_debug ("DEB DEBUG: -- C - peer->lseq=%d, peer->rseq=%d", peer->lseq, peer->rseq);
83     zlog_debug ("DEB DEBUG: -- C - sendSeq=%d", ntohl(pkt_hdr.sendSeq));
84   #endif //DEB_DEBUG
85   if (flag) // ie: starting a new conversation
86   { // create packet header fields - security hash: secret = ours, sender = us, dest = peer
87     ack_no = 0;
88     pkt_hdr.recipAck = htonl(ack_no);
89     hash = dohash((u_char*)&pkt_hdr + FRP_HASHSIZE, FRP_PKT_HDRSIZE - FRP_HASHSIZE, frp->secret, local.s_a
90 peer->address.s_addr);
91     // hash = dohash((u_char*)&pkt_hdr + FRP_HASHSIZE, len, frp->secret, local.s_addr, peer->address.s_addr);
92     memcpy(pkt_hdr.hash, hash, FRP_HASHSIZE);
93     #ifdef DEB_DEBUG_PEER
94     zlog_debug ("DEB DEBUG: -- make_frp_pkt_hdr - packet len: %d", len);
95     zlog_debug ("DEB DEBUG: -- make_frp_pkt_hdr - packet length: %d", FRP_PKT_HDRSIZE - FRP_HASHSIZE);
96     zlog_debug ("DEB DEBUG: -- B - peer->lseq=%d, peer->rseq=%d", peer->lseq, peer->rseq);
97     zlog_debug ("DEB DEBUG: -- B - sendSeq=%d, recipAck=%d", ntohl(pkt_hdr.sendSeq), ntohl(pkt_hdr.recipAck));
98     #endif //DEB_DEBUG
99     #ifdef DEB_DEBUG
100     zlog_debug ("DEB DEBUG: -- make_frp_pkt_hdr - SYN from: %s secret: %s", inet_ntoa (local), frp->secret);
101     zlog_debug ("DEB DEBUG: -- make_frp_pkt_hdr - SYN to: %s secret: %s", inet_ntoa (peer->address),
102 frp->secret);
103     zlog_debug ("DEB DEBUG: -- make_frp_pkt_hdr - SYN sendSeq=%d, recipAck=%d", ntohl(pkt_hdr.sendSeq),
104 ntohl(pkt_hdr.recipAck));
105     #endif //DEB_DEBUG
106   } else // ie: continuing a conversation
107   { // create packet header fields - security hash: secret = ours, sender = us, dest = peer
108     ack_no = peer->rseq;
109     pkt_hdr.recipAck = htonl(ack_no);
110     if (buf != NULL)
111     { memcpy(buf, &pkt_hdr, FRP_PKT_HDRSIZE);
112       hash = dohash(buf + FRP_HASHSIZE, len - FRP_HASHSIZE, frp->secret, local.s_addr, peer->address.s_addr);
113     } else
114     { hash = dohash((u_char*)&pkt_hdr + FRP_HASHSIZE, FRP_PKT_HDRSIZE - FRP_HASHSIZE, frp->secret,
115 local.s_addr, peer->address.s_addr);
116     }
117     memcpy(pkt_hdr.hash, hash, FRP_HASHSIZE);
118     #ifdef DEB_DEBUG
119     zlog_debug ("DEB DEBUG: -- make_frp_pkt_hdr - ACK from: %s secret: %s", inet_ntoa (local), frp->secret);
120     zlog_debug ("DEB DEBUG: -- make_frp_pkt_hdr - ACK to: %s secret: %s", inet_ntoa (peer->address),
121 frp->secret);
122     zlog_debug ("DEB DEBUG: -- make_frp_pkt_hdr - ACK sendSeq=%d, recipAck=%d", ntohl(pkt_hdr.sendSeq),
```

```
123 ntohl(pkt_hdr.recipAck));
124 #endif //DEB_DEBUG
125 #ifdef DEB_DEBUG_PEER
126     zlog_debug ("DEB DEBUG: -- D - peer->lseq=%d, peer->rseq=%d", peer->lseq, peer->rseq);
127     zlog_debug ("DEB DEBUG: -- D - sendSeq=%d, recipAck=%d", ntohl(pkt_hdr.sendSeq), ntohl(pkt_hdr.recipAck));
128 #endif //DEB_DEBUG
129 }
130 return pkt_hdr;
131 }
132
133 struct frp_msg_hdr
134 make_frp_msg_hdr (int type)
135 { struct frp_msg_hdr msg_hdr;
136   memset (&msg_hdr, 0, sizeof (msg_hdr));
137   // length field
138   switch (type)
139   { case FRP_MSG_CONTROL:
140     msg_hdr.length = FRP_MSG_CONTROL_SIZE;
141     break;
142     case FRP_MSG_IPV4CONFIG:
143     msg_hdr.length = FRP_MSG_IPV4CONFIG_SIZE;
144     break;
145     case FRP_MSG_IPV4GATEWAY:
146     msg_hdr.length = FRP_MSG_IPV4GATEWAY_MINSIZE;
147     break;
148     case FRP_MSG_IPV4UPDATE:
149     msg_hdr.length = FRP_MSG_IPV4UPDATE_SIZE;
150     break;
151     default:
152     break;
153   }
154   // type field
155   msg_hdr.type = type;
156   #ifdef DEB_DEBUG
157     zlog_debug ("DEB DEBUG: -- make_frp_msg_hdr - type=0x%x", msg_hdr.type);
158   #endif //DEB_DEBUG
159   return msg_hdr;
160 }
161
162
163 struct frp_msg_control
164 make_frp_msg_control (int type)
165 { struct frp_msg_control msg;
166   msg.msg_hdr = make_frp_msg_hdr (FRP_MSG_CONTROL);
167   msg.type = type;
168   msg.param = 0;
169   return msg;
170 }
171
172 struct frp_msg_ipv4config
173 make_frp_msg_ipv4config (struct in_addr peer)
174 { struct frp_msg_ipv4config msg;
175   msg.msg_hdr = make_frp_msg_hdr (FRP_MSG_IPV4CONFIG);
176   msg.cost = htons (frp->cost);
177   msg.poll = htons (frp->poll);
178   msg.retry = htons (frp->retry);
179   msg.id = frp_get_interface_address (peer);
180   #ifdef DEB_DEBUG_PEER
181     zlog_debug ("DEB DEBUG: -- make_frp_msg_ipv4config - cost=%d, poll=%d, retry=%d, id=%s", ntohs(msg.cost), ntohs(msg.poll), ntohs(msg.retry), inet_ntoa (peer));
182   #endif //DEB_DEBUG
183 }
```

```
184 return msg;
185 }
186
187 int
188 make_frp_msg_ipv4gateway (struct in_addr peer, struct frp_msg_ipv4gateway* msg)
189 { msg->msg_hdr = make_frp_msg_hdr (FRP_MSG_IPV4GATEWAY);
190   switch (frp->is_gateway_flag)
191   { // we are the gateway
192     case FRP_GATEWAY_ALWAYS:
193     case FRP_GATEWAY_YES:
194     msg->cost = 0;
195     msg->path[0] = frp_get_interface_address (peer);
196     return FRP_MSG_IPV4GATEWAY_MINSIZE;
197     break;
198     // we are not the gateway
199     case FRP_GATEWAY_NO:
200     if (frp->gateway_cost == 0) // have no gateway
201     { msg->cost = htons (FRP_INFINITY);
202       msg->path[0].s_addr = 0;
203       return FRP_MSG_IPV4GATEWAY_MINSIZE;
204     } else // someone else is the gateway
205     { msg->cost = htons (frp->gateway_cost);
206       struct listnode *node;
207       struct in_addr *path_step_addr;
208       msg->path[0] = frp_get_interface_address (peer);
209       int i = 1;
210       for (ALL_LIST_ELEMENTS_RO (frp->gateway_path, node, path_step_addr))
211       { msg->path[i] = *path_step_addr;
212         i++;
213       }
214       return FRP_MSG_IPV4GATEWAY_MINSIZE + ((i - 1) * sizeof (struct in_addr));
215     }
216     break;
217     default:
218     break;
219   }
220   return FRP_MSG_IPV4GATEWAY_MINSIZE;
221 }
222
223 int
224 make_frp_msg_ipv4update (struct frp_peer* peer, int available_length, u_char* buf)
225 {
226   int counter = 0;
227   int length = 0;
228   struct route_node* rn;
229   struct frp_info* route_info;
230   struct frp_peer* nexthop;
231   struct frp_msg_ipv4update msg;
232   msg.flags = 0;
233   msg.length = 0;
234   msg.routecost = 0;
235   msg.gatecost = 0;
236   struct in_addr nullPrefix;
237   nullPrefix.s_addr = 0;
238   msg.prefix = nullPrefix;
239
240   // for each route in our RIB
241   for (rn = route_top (frp->rib); rn; rn = route_next (rn))
242   {
243     if (rn->info)
244     {
```

```
245     route_info = (struct frp_info*)(rn->info);
246     // if this peer is not the nexthop
247     if (peer->address.s_addr != route_info->nexthop.s_addr)
248     {
249         // get nexthop peer
250         nexthop = frp_peer_lookup (&route_info->nexthop);
251         if (nexthop != NULL)
252         { // FRP decision algorithm
253             // if route cost + link cost < route gateway cost + target gateway cost + is gateway route
254             if ((route_info->cost + nexthop->cost) < (nexthop->gateway_cost + peer->gateway_cost +
255 route_info->is_gateway_flag))
256             {
257                 if (counter == 0) // if counter = 0
258                 {
259                     // make new route message with BEGIN flag (& GATE flag if appropriate)
260                     msg.msg_hdr = make_frp_msg_hdr (FRP_MSG_IPV4UPDATE);
261                     if (route_info->is_gateway_flag)
262                     { msg.flags = FRP_FLAG_BEGIN + FRP_FLAG_GATEWAY;
263                     } else
264                     { msg.flags = FRP_FLAG_BEGIN;
265                     }
266                     msg.length = htons(rn->p.prefixlen);
267                     // msg.length = 24;
268                     msg.routecost = htons(route_info->cost + nexthop->cost);
269                     msg.gatecost = htons(nexthop->gateway_cost);
270                     msg.prefix = rn->p.u.prefix4;
271                     // increment counter
272                     counter++;
273                 } else // counter != 0
274                 {
275                     // memcpy message to buffer
276                     if ((length + FRP_MSG_IPV4UPDATE_SIZE) <= available_length)
277                     { memcpy(buf + length, &msg, FRP_MSG_IPV4UPDATE_SIZE);
278                     length += FRP_MSG_IPV4UPDATE_SIZE;
279                     // make new route message (with & GATE flag if appropriate)
280                     msg.msg_hdr = make_frp_msg_hdr (FRP_MSG_IPV4UPDATE);
281                     if (route_info->is_gateway_flag)
282                     { msg.flags = FRP_FLAG_GATEWAY;
283                     } else
284                     { msg.flags = 0;
285                     }
286                     msg.length = htons(rn->p.prefixlen);
287                     msg.routecost = htons(route_info->cost + nexthop->cost);
288                     msg.gatecost = htons(nexthop->gateway_cost);
289                     msg.prefix = rn->p.u.prefix4;
290                     // increment counter
291                     counter++;
292                 }
293             }
294         }
295     } else // nexthop is null
296     {
297         // complain
298     }
299 }
300 }
301 }
302 // add in static routes (enabled networks)
303 for (rn = route_top (frp_enable_network); rn; rn = route_next (rn))
304 { if (rn->info)
305 {
```

```
306     #ifdef DEB_DEBUG
307         zlog_debug ("DEB DEBUG: -- make_frp_msg_ipv4update %s/%d", inet_ntoa (rn->p.u.prefix4), rn->p.prefixlen);
308     #endif //DEB_DEBUG
309     if (counter == 0) // if counter = 0
310     {
311         // make new route message with BEGIN flag (& GATE flag if appropriate)
312         msg.msg_hdr = make_frp_msg_hdr (FRP_MSG_IPV4UPDATE);
313         msg.flags = FRP_FLAG_BEGIN;
314         msg.length = htons(rn->p.prefixlen);
315         msg.routecost = 0;
316         msg.gatecost = htons(frp->gateway_cost);
317         msg.prefix = rn->p.u.prefix4;
318         // increment counter
319         counter++;
320     } else // counter != 0
321     {
322         // memcpy message to buffer
323         if ((length + FRP_MSG_IPV4UPDATE_SIZE) <= available_length)
324         { memcpy(buf + length, &msg, FRP_MSG_IPV4UPDATE_SIZE);
325         length += FRP_MSG_IPV4UPDATE_SIZE;
326         // make new route message (with & GATE flag if appropriate)
327         msg.msg_hdr = make_frp_msg_hdr (FRP_MSG_IPV4UPDATE);
328         msg.flags = 0;
329         msg.length = htons(rn->p.prefixlen);
330         msg.routecost = 0;
331         msg.gatecost = htons(frp->gateway_cost);
332         msg.prefix = rn->p.u.prefix4;
333         // increment counter
334         counter++;
335     }
336 }
337 }
338 }
339 }
340
341 if (counter > 0) //sending routes to this peer
342 {
343     // set COMMIT flag, memcpy message to buffer
344     msg.flags += FRP_FLAG_COMMIT;
345 } else // sending nullrt to this peer
346 {
347     // make new route message with BEGIN & COMMIT & NULLRT flags
348     #ifdef DEB_DEBUG
349         zlog_debug ("DEB DEBUG: -- make_frp_msg_ipv4update, sending nullrt to this peer");
350     #endif //DEB_DEBUG
351     msg.msg_hdr = make_frp_msg_hdr (FRP_MSG_IPV4UPDATE);
352     msg.flags = FRP_FLAG_BEGIN + FRP_FLAG_COMMIT + FRP_FLAG_NULLRT;
353 }
354 }
355 // memcpy final message to buffer
356 if ((length + FRP_MSG_IPV4UPDATE_SIZE) <= available_length)
357 { memcpy(buf + length, &msg, FRP_MSG_IPV4UPDATE_SIZE);
358 length += FRP_MSG_IPV4UPDATE_SIZE;
359 }
360 return length;
361 }
362
363 // build and send a SYN packet to a peer
364 int
365 build_syn_pkt (struct frp_peer* peer)
366 {
```



```
367 { // find local address of interface peer is attached to
368     struct prefix* local_p;
369     struct in_addr local_a;
370     int sent;
371     local_p = find_local_address_for_peer(peer->address);
372     if (local_p != NULL)
373     { int local_p_test = inet_pton(AF_INET, inet_ntoa(local_p->u.prefix4), &local_a);
374       else
375     {
376         #ifdef DEB_DEBUG
377         zlog_debug ("DEB DEBUG: -- build_syn_pkt, failed to get local interface address");
378         #endif //DEB_DEBUG
379         return 0;
380     }
381     // create a SYN packet to send
382     struct frp_pkt_hdr syn_pkt;
383     syn_pkt = make_frp_pkt_hdr (local_a, peer, FRP_SYN, FRP_PKT_HDRSIZE, NULL);
384     // create socket
385     struct sockaddr_in socket;
386     socket.sin_family = AF_INET;
387     socket.sin_port = htons(FRP_PORT_DEFAULT);
388     socket.sin_addr = peer-> address;
389     #ifdef DEB_DEBUG
390     zlog_debug ("DEB DEBUG: -- build_syn_pkt - SYN=%u (0x%x), ACK=%u (0x%x), addr=%s, port=%u",
391     ntohl(syn_pkt.sendSeq), ntohl(syn_pkt.sendSeq), syn_pkt.recipAck, syn_pkt.recipAck, inet_ntoa(socket.sin_addr),
392     ntohs(socket.sin_port));
393     #endif //DEB_DEBUG
394     // send SYN packet to peer
395     sent = frp_send_packet ((u_char*)&syn_pkt, FRP_PKT_HDRSIZE, &socket);
396     if (sent)
397     { memcpy (peer->packet_latest, (u_char*)&syn_pkt, FRP_PKT_HDRSIZE);
398       peer->packet_latest_length = FRP_PKT_HDRSIZE;
399       peer->packet_latest_lseq = ntohl(syn_pkt.sendSeq);
400       peer->flag_awaiting_ack = ON;
401       #ifdef DEB_DEBUG_PKT
402       zlog_debug ("DEB DEBUG: -- build_syn_pkt - packet_latest_lseq=%d", peer->packet_latest_lseq);
403       #endif //DEB_DEBUG
404     } else
405     { zlog_debug ("packet number %d not sent", syn_pkt.sendSeq);
406     }
407     return sent;
408 }
409
410 // build and send an ACK packet to a peer
411 int
412 build_ack_pkt (struct frp_peer* peer, struct sockaddr_in* from, struct in_addr local)
413 { u_char* out_pkt_buf;
414   int out_pkt_len;
415   int sent;
416   struct frp_pkt_hdr out_pkt_hdr;
417   struct frp_msg_control out_pkt;
418   // create the control message
419   out_pkt_len = FRP_PKT_HDRSIZE + FRP_MSG_CONTROL_SIZE;
420   out_pkt_buf = XCALLOC (MTYPE_FRP, out_pkt_len);
421   out_pkt = make_frp_msg_control (FRP_CTRL_ACK);
422   memcpy(out_pkt_buf + FRP_PKT_HDRSIZE, &out_pkt, FRP_MSG_CONTROL_SIZE);
423   // create the packet header - has to be last because of the hash
424   out_pkt_hdr = make_frp_pkt_hdr (local, peer, FRP_ACK, out_pkt_len, out_pkt_buf);
425   memcpy(out_pkt_buf, &out_pkt_hdr, FRP_PKT_HDRSIZE);
426   #ifdef DEB_DEBUG_PEER
427   zlog_debug ("DEB DEBUG: -- 7 - out_pkt_hdr size =%d", sizeof(out_pkt_hdr));
```

```
428     zlog_debug ("DEB DEBUG: -- 7 - out_pkt_hdr sendSeq=%d", ntohl(out_pkt_hdr.sendSeq));
429     zlog_debug ("DEB DEBUG: -- 7 - out_pkt_hdr recipAck=%d", ntohl(out_pkt_hdr.recipAck));
430     zlog_debug ("DEB DEBUG: -- 7 - out_pkt_hdr length=%d, type=0x%x, type=%d, param=%d",
431     out_pkt.msg_hdr.length, out_pkt.msg_hdr.type, out_pkt.type, out_pkt.param);
432     #endif //DEB_DEBUG
433     // send control ack packet
434     sent = frp_send_packet (out_pkt_buf, out_pkt_len, from);
435     if (sent)
436     { memcpy (peer->packet_latest, (u_char*)out_pkt_buf, out_pkt_len);
437       peer->packet_latest_length = out_pkt_len;
438       peer->packet_latest_lseq = ntohl(out_pkt_hdr.sendSeq);
439       peer->flag_awaiting_ack = ON;
440       #ifdef DEB_DEBUG_PKT
441       zlog_debug ("DEB DEBUG: -- build_ack_pkt - packet_latest_lseq=%d", peer->packet_latest_lseq);
442       #endif //DEB_DEBUG
443     } else
444     { zlog_debug ("packet number %d not sent", out_pkt_hdr.sendSeq);
445     }
446     return sent;
447 }
448
449 // build and send a NAK packet to a peer
450 int
451 build_nak_pkt (struct frp_peer* peer, struct sockaddr_in* from, struct in_addr local)
452 { u_char* out_pkt_buf;
453   int out_pkt_len;
454   int sent;
455   struct frp_pkt_hdr out_pkt_hdr;
456   struct frp_msg_control out_pkt;
457   // create the control message
458   out_pkt_len = FRP_PKT_HDRSIZE + FRP_MSG_CONTROL_SIZE;
459   out_pkt_buf = XCALLOC (MTYPE_FRP, out_pkt_len);
460   out_pkt = make_frp_msg_control (FRP_CTRL_NAK);
461   memcpy(out_pkt_buf + FRP_PKT_HDRSIZE, &out_pkt, FRP_MSG_CONTROL_SIZE);
462   // create the packet header - has to be last because of the hash
463   out_pkt_hdr = make_frp_pkt_hdr (local, peer, FRP_ACK, out_pkt_len, out_pkt_buf);
464   memcpy(out_pkt_buf, &out_pkt_hdr, FRP_PKT_HDRSIZE);
465   // send control nak packet
466   sent = frp_send_packet (out_pkt_buf, out_pkt_len, from);
467   if (sent)
468   { memcpy (peer->packet_latest, (u_char*)out_pkt_buf, out_pkt_len);
469     peer->packet_latest_length = out_pkt_len;
470     peer->packet_latest_lseq = ntohl(out_pkt_hdr.sendSeq);
471     peer->flag_awaiting_ack = ON;
472     #ifdef DEB_DEBUG_PKT
473     zlog_debug ("DEB DEBUG: -- build_nak_pkt - packet_latest_lseq=%d", peer->packet_latest_lseq);
474     #endif //DEB_DEBUG
475   } else
476   { zlog_debug ("packet number %d not sent", out_pkt_hdr.sendSeq);
477   }
478   return sent;
479 }
480
481 // build and send a batch packet to a peer
482 int
483 build_batch_pkt (struct frp_peer* peer, struct sockaddr_in* from, struct in_addr local)
484 {
485     #ifdef DEB_DEBUG
486     zlog_debug ("DEB DEBUG: entering frp_packet.c - build_batch_pkt");
487     #endif //DEB_DEBUG
488     // create an ACK packet to send to peer
```

```
489 u_char*          out_pkt_buf;
490 int              out_pkt_len;
491 enum flag        localflag;
492 int              sent;
493 struct frp_pkt_hdr out_pkt_hdr;
494 struct frp_msg_control out_pkt_ctrl;
495 struct frp_msg_ipv4config out_pkt_conf;
496 struct frp_msg_ipv4gateway out_pkt_gate;
497 // create buffer for outgoing packet
498 localflag = OFF;
499 out_pkt_buf = XCALLOC (MTYPE_FRP, FRP_PKT_MAXSIZE);
500 out_pkt_len = FRP_PKT_HDRSIZE;
501 // if required, create a control packet, send a poll and add to outgoing packet
502 if ((peer->flag_send_poll) && ((out_pkt_len + FRP_MSG_CONTROL_SIZE) < FRP_PKT_MAXSIZE))
503 { out_pkt_ctrl = make_frp_msg_control (FRP_CTRL_POLL);
504   memcpy(out_pkt_buf + out_pkt_len, &out_pkt_ctrl, FRP_MSG_CONTROL_SIZE);
505   out_pkt_len += FRP_MSG_CONTROL_SIZE;
506   localflag = ON;
507 }
508 // if required, create a config packet and add to outgoing packet
509 if ((peer->flag_send_config) && ((out_pkt_len + FRP_MSG_IPV4CONFIG_SIZE) < FRP_PKT_MAXSIZE))
510 { out_pkt_conf = make_frp_msg_ipv4config (peer->address);
511   memcpy(out_pkt_buf + out_pkt_len, &out_pkt_conf, FRP_MSG_IPV4CONFIG_SIZE);
512   out_pkt_len += FRP_MSG_IPV4CONFIG_SIZE;
513   localflag = ON;
514 }
515 // if required, create a path to gateway packet and add to outgoing packet
516 if (peer->flag_send_gateway)
517 { struct frp_msg_ipv4gateway out_pkt_gate;
518   int out_pkt_gate_length;
519   out_pkt_gate_length = make_frp_msg_ipv4gateway (peer->address, &out_pkt_gate);
520   if ((out_pkt_len + out_pkt_gate_length) < FRP_PKT_MAXSIZE)
521   { memcpy(out_pkt_buf + out_pkt_len, &out_pkt_gate, out_pkt_gate_length);
522     out_pkt_len += out_pkt_gate_length;
523     localflag = ON;
524   }
525 }
526 // if not quiescent
527 if (frp->is_gateway_flag == ON)
528 { // if required, create a route update packet and add to outgoing packet
529   if ((peer->flag_send_update) && ((out_pkt_len + FRP_MSG_IPV4UPDATE_SIZE) < FRP_PKT_MAXSIZE))
530   { int available_length = 0;
531     int used_length = 0;
532     // calculate available space left
533     available_length = FRP_PKT_MAXSIZE - out_pkt_len;
534     // return how much was added
535     used_length = make_frp_msg_ipv4update (peer, available_length, out_pkt_buf + out_pkt_len);
536     out_pkt_len += used_length;
537     localflag = ON;
538   }
539 }
540
541 if (localflag == ON)
542 { // create the packet header - has to be last because of the hash
543   out_pkt_hdr = make_frp_pkt_hdr (local, peer, FRP_ACK, out_pkt_len, out_pkt_buf);
544   #ifdef DEB_DEBUG_PEER
545     zlog_debug ("DEB DEBUG: -- 7 - out_pkt_hdr.sendSeq=%d", ntohl(out_pkt_hdr.sendSeq));
546     zlog_debug ("DEB DEBUG: -- 7 - out_pkt_hdr.recipAck=%d", ntohl(out_pkt_hdr.recipAck));
547     zlog_debug ("DEB DEBUG: -- 7 - out_pkt_hdr size=%d", sizeof(out_pkt_hdr));
548   #endif //DEB_DEBUG
549   memcpy(out_pkt_buf, &out_pkt_hdr, FRP_PKT_HDRSIZE);
```

```
550 // send packet
551 #ifdef DEB_DEBUG_PEER
552   zlog_debug ("DEB DEBUG: -- 6 - out_pkt_len=%d", out_pkt_len);
553   struct frp_pkt_hdr* test_hdr = (struct frp_pkt_hdr*)out_pkt_buf;
554   zlog_debug ("DEB DEBUG: -- 6 - out_pkt_buf.sendSeq=%d", ntohl(test_hdr->sendSeq));
555   zlog_debug ("DEB DEBUG: -- 6 - out_pkt_buf.recipAck=%d", ntohl(test_hdr->recipAck));
556 #endif //DEB_DEBUG
557 sent = frp_send_packet (out_pkt_buf, out_pkt_len, from);
558 // reset peer flags
559 if (sent)
560 { peer->flag_send_poll = OFF;
561   peer->flag_send_config = OFF;
562   peer->flag_send_gateway = OFF;
563   peer->flag_send_update = OFF;
564   memcpy (peer->packet_latest, (u_char*)out_pkt_buf, out_pkt_len);
565   peer->packet_latest_length = out_pkt_len;
566   peer->packet_latest_lseq = ntohl(out_pkt_hdr.sendSeq);
567   peer->flag_awaiting_ack = ON;
568   #ifdef DEB_DEBUG_PKT
569     zlog_debug ("DEB DEBUG: -- build_batch_pkt - packet_latest_lseq=%d", peer->packet_latest_lseq);
570   #endif //DEB_DEBUG
571 } else
572 { zlog_debug ("packet number %d not sent", out_pkt_hdr.sendSeq);
573 }
574 } else //if no messages created because localflag == OFF
575 { // if peer awaiting ACK, send ACK packet to finish conversation
576   if (peer->flag_awaiting_ack == ON)
577   {
578     #ifdef DEB_DEBUG_PKT
579       zlog_debug ("DEB DEBUG: -- build_batch_pkt - IF peer->flag_awaiting_ack == ON");
580     #endif //DEB_DEBUG
581     sent = build_ack_pkt (peer, from, local);
582     peer->flag_awaiting_ack = OFF;
583   }
584 }
585 return sent;
586 }
587 }
```

```

1 // -----
2 // INCLUDES
3 // -----
4 #include "frpd.h"
5
6
7
8 // -----
9 // FUNCTIONS
10 // -----
11
12 // works thru a list of routes given to us by a peer
13 // checks to see if each route is in our RIB
14 // where appropriate, replaces old routes with 'better' new routes
15 // deletes old routes from and passes new routes to Zebra
16 void
17 frp_recompute_rib (struct frp_peer* peer)
18 {
19     #ifdef DEB_DEBUG
20         zlog_debug ("DEB_DEBUG: entering route.c - frp_recompute_rib");
21     #endif //DEB_DEBUG
22     int flag = 0;
23     // step thru each new route in peer rib
24     struct listnode* node;
25     struct frp_rte* rte;
26     struct interface* peer_if = if_lookup_address (peer->address);
27     for (ALL_LIST_ELEMENTS_RO (peer->rib, node, rte))
28     {
29         struct route_node* rib_rte_ptr;
30         rib_rte_ptr = route_node_get (frp->rib, (struct prefix*)&rte->prefix);
31         // is the route in our rib?
32         if (rib_rte_ptr->info != NULL)
33         {
34             // calculate existing route cost
35             struct frp_info* rte_info = (struct frp_info*)rib_rte_ptr->info;
36             u_int32_t old_cost = rte_info->cost;
37             // calculate new route cost
38             u_int32_t new_cost = rte->routecost + peer->cost;
39             // is it a 'better' route using the frp algorithm?
40             if (new_cost < old_cost)
41             {
42                 // delete old from zebra
43                 frp_zebra_ipv4_delete ((struct prefix_ipv4*)&rte->prefix, &rte_info->nexthop, old_cost);
44                 // new route info for rib
45                 struct frp_info new_route_info;
46                 new_route_info.nexthop = peer->address;
47                 new_route_info.cost = new_cost;
48                 new_route_info.rte_node = rib_rte_ptr;
49                 new_route_info.type = ZEBRA_ROUTE_FRP;
50                 new_route_info.sub_type = FRP_ROUTE_RTE;
51                 new_route_info.ifindex = peer_if->ifindex;
52                 // replace old route with new in our rib
53                 memcpy (rte_info, &new_route_info, sizeof(struct frp_info));
54                 // add new route to zebra
55                 frp_zebra_ipv4_add ((struct prefix_ipv4*)&rte->prefix, &peer->address, new_cost,
56 ZEBRA_FRP_DISTANCE_DEFAULT);
57                 // set update flag
58                 flag = 1;
59             }
60             // no else because ignore route
61         } else // not in our rib

```

```

62     {
63         // calculate new route cost
64         u_int32_t new_cost = rte->routecost + peer->cost;
65         // new route info for rib
66         struct frp_info new_route_info;
67         new_route_info.nexthop = peer->address;
68         new_route_info.cost = new_cost;
69         new_route_info.rte_node = rib_rte_ptr;
70         new_route_info.type = ZEBRA_ROUTE_FRP;
71         new_route_info.sub_type = FRP_ROUTE_RTE;
72         new_route_info.ifindex = peer_if->ifindex;
73         // add to our rib
74         struct frp_info* rte_info = XCALLOC (MTYPE_FRP, sizeof (struct frp_info));
75         memcpy (rte_info, &new_route_info, sizeof(struct frp_info));
76         rib_rte_ptr->info = rte_info;
77         // add to zebra rib
78         frp_zebra_ipv4_add ((struct prefix_ipv4*)&rte->prefix, &peer->address, new_cost,
79 ZEBRA_FRP_DISTANCE_DEFAULT);
80         // set update flag
81         flag = 1;
82     }
83     //route_unlock_node (rib_rte_ptr);
84 }
85 if (flag)
86 {
87     // frp_event (FRP_EVENT_UPDATE, 1);
88 }
89 }
90
91
92 // deletes all routes from a particular peer - in our RIB and in Zebra
93 // checks each other peer for the best route to replace a deleted one
94 // passes new route to Zebra
95 void
96 frp_delete_peer_from_rib (struct frp_peer* deleted_peer)
97 {
98     #ifdef DEB_DEBUG
99         zlog_debug ("DEB_DEBUG: entering route.c - frp_delete_peer_from_rib");
100     #endif //DEB_DEBUG
101     struct route_node* rib_rte_ptr;
102     // for each entry in our rib
103     for (rib_rte_ptr = route_top (frp->rib); rib_rte_ptr; rib_rte_ptr = route_next (rib_rte_ptr))
104     {
105         struct frp_info* rte_info = (struct frp_info*)rib_rte_ptr->info;
106         // nexthop matches the deleted peer address
107         if ((rte_info != NULL) && (rte_info->nexthop.s_addr == deleted_peer->address.s_addr))
108         {
109             int flag = 0;
110             struct prefix_ipv4 deleted_dest;
111             memset (&deleted_dest, 0, sizeof (struct prefix_ipv4));
112             deleted_dest.family = AF_INET;
113             deleted_dest.prefix = rib_rte_ptr->p.u.prefix4;
114             deleted_dest.prefixlen = IPV4_MAX_BITLEN;
115             // delete from zebra
116             frp_zebra_ipv4_delete (&deleted_dest, &rte_info->nexthop, rte_info->cost);
117             // delete from rib (actually just sets the node info in the list to null)
118             XFREE (MTYPE_FRP, rte_info);
119             rib_rte_ptr->info = NULL;
120
121             struct listnode* node1;
122             struct frp_peer* list_peer;

```

```

123 struct route_node* new_route_node = NULL;
124 // for each peer
125 for (ALL_LIST_ELEMENTS_RO (frp_peers, node1, list_peer))
126 {
127     // this is not the deleted peer
128     if (list_peer->address.s_addr != deleted_peer->address.s_addr)
129     {
130         struct listnode* node2;
131         struct frp_rte* rte;
132         //for each route in the peer rib
133         for (ALL_LIST_ELEMENTS_RO (list_peer->rib, node2, rte))
134         {
135             // this route is to the deleted destination
136             // compare the two prefixes to see if they are the same host and the same network
137             if (prefix_same (&rte->prefix, &deleted_dest))
138             {
139                 struct route_node* new_rte_ptr;
140                 new_rte_ptr = route_node_get(frp->rib, (struct prefix*)&rte->prefix);
141                 new_route_node = new_rte_ptr;
142                 // this route is not in the rib
143                 if (new_rte_ptr->info == NULL)
144                 {
145                     // add route to rib
146                     struct frp_info new_route_info;
147                     struct interface* peer_if = if_lookup_address (list_peer->address);
148                     new_route_info.nexthop = list_peer->address;
149                     new_route_info.cost = rte->route_cost;
150                     new_route_info.is_gateway_flag = (rte->flags & FRP_FLAG_GATEWAY);
151                     new_route_info.rte_node = new_rte_ptr;
152                     new_route_info.type = ZEBRA_ROUTE_FRP;
153                     new_route_info.sub_type = FRP_ROUTE_RTE;
154                     new_route_info.ifindex = peer_if->ifindex;
155                     struct frp_info* rte_info = XCALLOC (MTYPE_FRP, sizeof (struct frp_info));
156                     memcpy (rte_info, &new_route_info, sizeof (struct frp_info));
157                     new_rte_ptr->info = rte_info;
158                     flag = 1;
159                 } else // is already in the rib
160                 {
161                     // calculate existing route cost
162                     struct frp_info* rte_info = (struct frp_info*)new_rte_ptr->info;
163                     u_int32_t old_cost = rte_info->cost;
164                     // calculate new route cost
165                     u_int32_t new_cost = rte->route_cost + list_peer->cost;
166                     // this route is better than the existing route
167                     if (new_cost < old_cost)
168                     {
169                         // replace existing route with this route
170                         struct frp_info new_route_info;
171                         struct interface* peer_if = if_lookup_address (list_peer->address);
172                         new_route_info.nexthop = list_peer->address;
173                         new_route_info.cost = rte->route_cost;
174                         new_route_info.is_gateway_flag = (rte->flags & FRP_FLAG_GATEWAY);
175                         new_route_info.rte_node = new_rte_ptr;
176                         new_route_info.type = ZEBRA_ROUTE_FRP;
177                         new_route_info.sub_type = FRP_ROUTE_RTE;
178                         new_route_info.ifindex = peer_if->ifindex;
179                         memcpy (rte_info, &new_route_info, sizeof (struct frp_info));
180                         flag = 1;
181                     }
182                 }
183             }
184             route_unlock_node (new_rte_ptr);

```

```

184     }
185 }
186 }
187 }
188 if ((flag)&&(new_route_node != NULL))
189 { struct frp_info* rte_info = (struct frp_info*)new_route_node->info;
190     // add new route to zebra
191     frp_zebra_ipv4_add (&deleted_dest, &rte_info->nexthop, rte_info->cost, ZEBRA_FRP_DISTANCE_DEFAULT)
192     flag = 0;
193 }
194 }
195 }
196 }
197 }

```

```
1 // -----
2 // INCLUDES
3 // -----
4 #include "frpd.h"
5
6
7
8 // -----
9 // GLOBAL VARIABLES
10 // -----
11 /* create an instance of the zebra client */
12 struct zclient *zclient = NULL; // DEB COMMENT: zclient.h [35]
13
14 /* static prototypes */
15 static int frp_zebra_read_ipv4 (int command, struct zclient *zclient, zebra_size_t length);
16
17
18
19 // -----
20 // ZEBRA CALLBACK FUNCTIONS
21 // -----
22 // DEB COMMENT: need to supply functions to implement each of the zebra client API callback functions
23 // -- taken from zclient.h [70]
24 // DEB COMMENT: prototypes defined in frpd.h
25 // int (*router_id_update) (int, struct zclient *, uint16_t); // not needed
26 int (*interface_add) (int, struct zclient *, uint16_t); // defined in frp_interface.c [44]
27 int (*interface_delete) (int, struct zclient *, uint16_t); // defined in frp_interface.c [67]
28 int (*interface_up) (int, struct zclient *, uint16_t); // defined in frp_interface.c [91]
29 int (*interface_down) (int, struct zclient *, uint16_t); // defined in frp_interface.c [115]
30 int (*interface_address_add) (int, struct zclient *, uint16_t); // defined in frp_interface.c [136]
31 int (*interface_address_delete) (int, struct zclient *, uint16_t); // defined in frp_interface.c [163]
32 int (*ipv4_route_add) (int, struct zclient *, uint16_t); // defined in frp_zebra.c [39]
33 int (*ipv4_route_delete) (int, struct zclient *, uint16_t); // defined in frp_zebra.c [39]
34 // int (*ipv6_route_add) (int, struct zclient *, uint16_t); // defined in frp_zebra.c
35 // int (*ipv6_route_delete) (int, struct zclient *, uint16_t); // defined in frp_zebra.c
36
37
38 /* zebra route add and delete */
39 int
40 frp_zebra_read_ipv4 (int command, struct zclient *zclient, zebra_size_t length)
41 { struct stream *s;
42   struct zapi_ipv4 api;
43   unsigned long ifindex;
44   struct in_addr nexthop;
45   struct prefix_ipv4 p;
46   #ifdef DEB_DEBUG
47     zlog_debug ("DEB_DEBUG: entering frp_interface.c - frp_zebra_read_ipv4");
48   #endif //DEB_DEBUG
49   s = zclient->ibuf;
50   ifindex = 0;
51   nexthop.s_addr = 0;
52   /* type, flags, message */
53   api.type = stream_getc (s);
54   api.flags = stream_getc (s);
55   api.message = stream_getc (s);
56   /* IPv4 prefix */
57   memset (&p, 0, sizeof (struct prefix_ipv4));
58   p.family = AF_INET;
59   p.prefixlen = stream_getc (s);
60   stream_get (&p.prefix, s, PSIZE (p.prefixlen));
61   /* nexthop, ifindex, distance, metric */
```

```
62   if (CHECK_FLAG (api.message, ZAPI_MESSAGE_NEXTHOP))
63   { api.nexthop_num = stream_getc (s);
64     nexthop.s_addr = stream_get_ipv4 (s);
65   }
66   if (CHECK_FLAG (api.message, ZAPI_MESSAGE_IFINDEX))
67   { api.ifindex_num = stream_getc (s);
68     ifindex = stream_getl (s);
69   }
70   if (CHECK_FLAG (api.message, ZAPI_MESSAGE_DISTANCE))
71   { api.distance = stream_getc (s);
72   } else
73   { api.distance = 255;
74   }
75   if (CHECK_FLAG (api.message, ZAPI_MESSAGE_METRIC))
76   { api.metric = stream_getl (s);
77   } else
78   { api.metric = 0;
79   }
80   return 0;
81 }
82
83
84
85 // -----
86 // FUNCTIONS
87 // -----
88
89 /* frpd to zebra command interface. */
90 void
91 frp_zebra_ipv4_add (struct prefix_ipv4 *p, struct in_addr *nexthop, u_int32_t metric, u_char distance)
92 { struct zapi_ipv4 api;
93   if (zclient->redist[ZEBRA_ROUTE_FRP])
94   { api.type = ZEBRA_ROUTE_FRP;
95     api.flags = 0;
96     api.message = 0;
97     SET_FLAG (api.message, ZAPI_MESSAGE_NEXTHOP);
98     api.nexthop_num = 1;
99     api.nexthop = *nexthop;
100    api.ifindex_num = 0;
101    SET_FLAG (api.message, ZAPI_MESSAGE_METRIC);
102    api.metric = metric;
103    if (distance && distance != ZEBRA_FRP_DISTANCE_DEFAULT)
104    { SET_FLAG (api.message, ZAPI_MESSAGE_DISTANCE);
105      api.distance = distance;
106    }
107    zapi_ipv4_route (ZEBRA_IPV4_ROUTE_ADD, zclient, p, &api);
108    // DEB COMMENT: frp is not doing SNMP
109    // frp_global_route_changes++;
110  }
111 }
112
113 void
114 frp_zebra_ipv4_delete (struct prefix_ipv4 *p, struct in_addr *nexthop, u_int32_t metric)
115 { struct zapi_ipv4 api;
116   if (zclient->redist[ZEBRA_ROUTE_FRP])
117   { api.type = ZEBRA_ROUTE_FRP;
118     api.flags = 0;
119     api.message = 0;
120     SET_FLAG (api.message, ZAPI_MESSAGE_NEXTHOP);
121     api.nexthop_num = 1;
122     api.nexthop = *nexthop;
```

```
123     api.ifindex_num = 0;
124     SET_FLAG (api.message, ZAPI_MESSAGE_METRIC);
125     api.metric = metric;
126     zapi_ipv4_route (ZEBRA_IPV4_ROUTE_DELETE, zclient, p, &api);
127 }
128 }
129
130
131 // -----
132 // INIT FUNCTION
133 // -----
134 // initialise the zebra client structure and it's commands
135 // DEB COMMENT: zclient etc zclient.h
136
137 void
138 frp_zclient_init (void)
139 {
140     #ifdef DEB_DEBUG
141         fprintf (stderr, "DEB DEBUG: entering frp_zebra.c - frp_zclient_init\n");
142     #endif //DEB_DEBUG
143
144     /* allocate and initialise zebra structure */
145     zclient = zclient_new ();
146     zclient_init (zclient, ZEBRA_ROUTE_FRP);          // DEB COMMENT: ZEBRA_ROUTE_FRP zebra.h [4
147
148     /* set up callback functions */
149     zclient->interface_add = frp_interface_add;
150     zclient->interface_delete = frp_interface_delete;
151     zclient->interface_up = frp_interface_up;
152     zclient->interface_down = frp_interface_down;
153     zclient->interface_address_add = frp_interface_address_add;
154     zclient->interface_address_delete = frp_interface_address_delete;
155     zclient->ipv4_route_add = frp_zebra_read_ipv4;
156     zclient->ipv4_route_delete = frp_zebra_read_ipv4;
157     // zclient->ipv6_route_add = frp_ipv6_route_add;
158     // zclient->ipv6_route_delete = frp_ipv6_route_delete;
159
160     #ifdef DEB_DEBUG
161         fprintf (stderr, "DEB DEBUG: -- leaving frp_zclient_init\n");
162     #endif //DEB_DEBUG
163 }
164
```

```

1 // -----
2 // INCLUDES
3 // -----
4 #include "frpd.h"
5 #include "frp_debug.h"
6
7
8
9 // -----
10 // GLOBAL VARIABLES
11 // -----
12
13 /* frp enabled interface table */
14 struct route_table *frp_enable_network;
15
16 /* frp enabled network vector */
17 vector frp_enable_interface;
18
19 /* static prototypes */
20 static void frp_apply_address_del (struct connected *ifc);
21 static void frp_connect_set (struct interface *ifp, int set);
22 static void frp_enable_apply (struct interface *);
23 static void frp_enable_apply_all (void);
24 static int frp_enable_if_add (const char *ifname);
25 static int frp_enable_if_lookup (const char *ifname);
26 static int frp_enable_network_add (struct prefix *p);
27 static int frp_enable_network_lookup_if (struct interface *ifp);
28 static int frp_if_down(struct interface *ifp);
29 static int frp_if_ipv4_address_check (struct interface *ifp);
30 static struct frp_interface * frp_interface_new (void);
31 static int frp_interface_delete_hook (struct interface *ifp);
32 static int frp_interface_new_hook (struct interface *ifp);
33 static int frp_interface_wakeup (struct thread *t);
34
35
36
37
38 // -----
39 // ZEBRA CALLBACK FUNCTIONS
40 // -----
41 // DEB COMMENT: need to supply a function to implement each of the zebra client API callback functions
42 // -- taken from zclient.h [70]
43 // DEB COMMENT: prototypes defined in frpd.h
44
45 int
46 frp_interface_add (int command, struct zclient *zclient, zebra_size_t length)
47 { struct interface *ifp;
48   #ifdef DEB_DEBUG_IF
49     zlog_debug ("DEB DEBUG: entering frp_interface.c - frp_interface_add");
50   #endif //DEB_DEBUG
51   ifp = zebra_interface_add_read (zclient->ibuf);
52   if (IS_FRP_DEBUG_ZEBRA)
53   { zlog_debug ("interface add %s index %d flags %llx metric %d mtu %d", ifp->name, ifp->ifindex, (unsigned long
54 long) ifp->flags, ifp->metric, ifp->mtu);
55   }
56   /* check if this interface is frp enabled or not */
57   frp_enable_apply (ifp);
58   return 0;
59 }
60
61 int

```

```

62 frp_interface_delete (int command, struct zclient *zclient, zebra_size_t length)
63 { struct interface *ifp;
64   struct stream *s;
65   #ifdef DEB_DEBUG_IF
66     zlog_debug ("DEB DEBUG: entering frp_interface.c - frp_interface_delete");
67   #endif //DEB_DEBUG
68   s = zclient->ibuf;
69   /* zebra_interface_state_read() updates interface structure in iflist */
70   ifp = zebra_interface_state_read(s);
71   if (ifp == NULL)
72   { return 0;
73   }
74   if (if_is_up (ifp))
75   { frp_if_down(ifp);
76   }
77   zlog_info("interface delete %s index %d flags %llx metric %d mtu %d", ifp->name, ifp->ifindex, (unsigned long l
78 ifp->flags, ifp->metric, ifp->mtu);
79   /* To support pseudo interface do not free interface structure. */
80   /* if_delete(ifp); */
81   ifp->ifindex = IFINDEX_INTERNAL;
82   return 0;
83 }
84
85 int
86 frp_interface_up (int command, struct zclient *zclient, zebra_size_t length)
87 { struct interface *ifp;
88   #ifdef DEB_DEBUG_IF
89     zlog_debug ("DEB DEBUG: entering frp_interface.c - frp_interface_up");
90   #endif //DEB_DEBUG
91   /* zebra_interface_state_read () updates interface structure in iflist */
92   ifp = zebra_interface_state_read (zclient->ibuf);
93   if (ifp == NULL)
94   { return 0;
95   }
96   if (IS_FRP_DEBUG_ZEBRA)
97   { zlog_debug ("interface %s index %d flags %llx metric %d mtu %d is up", ifp->name, ifp->ifindex, (unsigned lo
98 long) ifp->flags, ifp->metric, ifp->mtu);
99   }
100   /* check if this interface is frp enabled or not */
101   frp_enable_apply (ifp);
102   return 0;
103 }
104
105 int
106 frp_interface_down (int command, struct zclient *zclient, zebra_size_t length)
107 { struct interface *ifp;
108   struct stream *s;
109   #ifdef DEB_DEBUG_IF
110     zlog_debug ("DEB DEBUG: entering frp_interface.c - frp_interface_up");
111   #endif //DEB_DEBUG
112   s = zclient->ibuf;
113   /* zebra_interface_state_read() updates interface structure in iflist. */
114   ifp = zebra_interface_state_read(s);
115   if (ifp == NULL)
116   { return 0;
117   }
118   frp_if_down(ifp);
119   if (IS_FRP_DEBUG_ZEBRA)
120   { zlog_debug ("interface %s index %d flags %llx metric %d mtu %d is down", ifp->name, ifp->ifindex, (unsigned l
121 long) ifp->flags, ifp->metric, ifp->mtu);
122   }

```

```

123     return 0;
124 }
125
126 int
127 frp_interface_address_add (int command, struct zclient *zclient, zebra_size_t length)
128 { struct connected *ifc;
129   struct prefix *p;
130   #ifdef DEB_DEBUG_IF
131     zlog_debug ("DEB DEBUG: entering frp_interface.c - frp_interface_address_add");
132   #endif //DEB_DEBUG
133   ifc = zebra_interface_address_read (ZEBRA_INTERFACE_ADDRESS_ADD, zclient->ibuf);
134   if (ifc == NULL)
135   { return 0;
136   }
137   p = ifc->address;
138   if (p->family == AF_INET)
139   { if (IS_FRP_DEBUG_ZEBRA)
140     { zlog_debug ("connected address %s/%d is added", inet_ntoa (p->u.prefix4), p->prefixlen);
141     }
142     frp_enable_apply(ifc->ifp);
143   }
144   return 0;
145 }
146
147 int
148 frp_interface_address_delete (int command, struct zclient *zclient, zebra_size_t length)
149 { struct connected *ifc;
150   struct prefix *p;
151   #ifdef DEB_DEBUG_IF
152     zlog_debug ("DEB DEBUG: entering frp_interface.c - frp_interface_address_delete");
153   #endif //DEB_DEBUG
154   ifc = zebra_interface_address_read (ZEBRA_INTERFACE_ADDRESS_DELETE, zclient->ibuf);
155   if (ifc)
156   { p = ifc->address;
157     if (p->family == AF_INET)
158     { if (IS_FRP_DEBUG_ZEBRA)
159       { zlog_debug ("connected address %s/%d is deleted", inet_ntoa (p->u.prefix4), p->prefixlen);
160       }
161       /* Check whether this prefix needs to be removed */
162       frp_apply_address_del(ifc);
163     }
164     connected_free (ifc);
165   }
166   return 0;
167 }
168
169
170
171 // -----
172 // FUNCTIONS
173 // -----
174
175 /* Allocate new frp's interface configuration. */
176 struct frp_interface*
177 frp_interface_new (void)
178 { struct frp_interface *ri;
179   ri = XCALLOC (MTYPE_FRP_INTERFACE, sizeof (struct frp_interface));
180   return ri;
181 }
182
183 /* add frp enable network */

```

```

184 int
185 frp_enable_network_add (struct prefix *p)
186 {
187   #ifdef DEB_DEBUG_IF
188     zlog_debug ("DEB DEBUG: entering frp_interface.c - frp_enable_network_add");
189   #endif //DEB_DEBUG
190   struct route_node *node;
191   node = route_node_get (frp_enable_network, p);
192   #ifdef DEB_DEBUG_IF
193     zlog_debug ("DEB DEBUG: -- frp_enable_network_add, after node =");
194   #endif //DEB_DEBUG
195   if (node->info)
196   {
197     #ifdef DEB_DEBUG_IF
198       zlog_debug ("DEB DEBUG: -- frp_enable_network_add, IF");
199     #endif //DEB_DEBUG
200     route_unlock_node (node);
201     return -1;
202   } else
203   {
204     #ifdef DEB_DEBUG_IF
205       zlog_debug ("DEB DEBUG: -- frp_enable_network_add, ELSE");
206     #endif //DEB_DEBUG
207     node->info = (char *) "enabled";
208
209     struct frp_peer* peer;
210     struct listnode *node, *nnode;
211     for (ALL_LIST_ELEMENTS (frp_peers, node, nnode, peer))
212     {
213       frp_recompute_rib (peer);
214     }
215   }
216   #ifdef DEB_DEBUG_IF
217     zlog_debug ("DEB DEBUG: -- frp_enable_network_add, before frp_enable_apply_all");
218   #endif //DEB_DEBUG
219   frp_enable_apply_all();
220   #ifdef DEB_DEBUG_IF
221     zlog_debug ("DEB DEBUG: -- frp_enable_network_add, after frp_enable_apply_all");
222   #endif //DEB_DEBUG
223   // route_unlock_node (node);
224   #ifdef DEB_DEBUG_IF
225     zlog_debug ("DEB DEBUG: -- leaving frp_enable_network_add");
226   #endif //DEB_DEBUG
227   return 1;
228 }
229
230 /* add interface to frp_enable_if. */
231 int
232 frp_enable_if_add (const char *ifname)
233 { int ret;
234   ret = frp_enable_if_lookup (ifname);
235   if (ret >= 0)
236   { return -1;
237   }
238   vector_set (frp_enable_interface, strdup (ifname));
239   frp_enable_apply_all();
240   return 1;
241 }
242
243 /* apply network configuration to all interface */
244 void

```



```

245 frp_enable_apply_all ()
246 { struct interface *ifp;
247   struct listnode *node, *nnode;
248   /* Check each interface. */
249   for (ALL_LIST_ELEMENTS (iflist, node, nnode, ifp))
250   { frp_enable_apply (ifp);
251   }
252 }
253
254 /* update interface status */
255 void
256 frp_enable_apply (struct interface *ifp)
257 { int ret;
258   struct frp_interface *ri = NULL;
259   #ifdef DEB_DEBUG_IF
260     zlog_debug ("DEB DEBUG: entering frp_interface.c - frp_enable_apply");
261   #endif //DEB_DEBUG
262   /* check interface */
263   if (!if_is_operative (ifp))
264   { return;
265   }
266   ri = ifp->info;
267   #ifdef DEB_DEBUG_IF
268     if (ri == NULL)
269       zlog_debug ("DEB DEBUG: -- frp_enable_apply - IF ifp->info = null");
270     else
271       zlog_debug ("DEB DEBUG: -- frp_enable_apply - ELSE ifp->info = %p", ifp->info);
272   #endif //DEB_DEBUG
273   /* check network configuration */
274   ret = frp_enable_network_lookup_if (ifp);
275   #ifdef DEB_DEBUG_IF
276     zlog_debug ("DEB DEBUG: -- frp_enable_apply - ret = %d", ret);
277   #endif //DEB_DEBUG
278   /* if the interface is matched */
279   if (ret > 0)
280   { ri->enable_network = 1;
281   } else
282   { ri->enable_network = 0;
283   }
284   #ifdef DEB_DEBUG_IF
285     zlog_debug ("DEB DEBUG: -- frp_enable_apply - ri->enable_network = %d", ri->enable_network);
286   #endif //DEB_DEBUG
287   /* check interface name configuration */
288   ret = frp_enable_if_lookup (ifp->name);
289   #ifdef DEB_DEBUG_IF
290     zlog_debug ("DEB DEBUG: -- frp_enable_apply - ret = %d", ret);
291   #endif //DEB_DEBUG
292   if (ret >= 0)
293   { ri->enable_interface = 1;
294   } else
295   { ri->enable_interface = 0;
296   }
297   #ifdef DEB_DEBUG_IF
298     zlog_debug ("DEB DEBUG: -- frp_enable_apply - ri->enable_interface = %d", ri->enable_interface);
299   #endif //DEB_DEBUG
300   /* any interface MUST have an IPv4 address */
301   if (! frp_if_ipv4_address_check (ifp) )
302   { ri->enable_network = 0;
303     ri->enable_interface = 0;
304   }
305   /* update running status of the interface */

```

```

306   if (ri->enable_network || ri->enable_interface)
307   { if (IS_FRP_DEBUG_EVENT)
308     { zlog_debug ("turn on %s", ifp->name);
309     }
310     /* add interface wake up thread */
311     if (!ri->t_wakeup)
312     { ri->t_wakeup = thread_add_timer (master, frp_interface_wakeup, ifp, 1);
313     }
314   } else
315   { if (ri->running)
316     { /* might as well clean up the route table as well frp_if_down sets to 0 ri->running, and displays "turn off %s"
317       frp_if_down(ifp);
318       frp_connect_set (ifp, 0);
319     }
320   }
321 }
322
323
324 /* check interface is enabled by ifname statement */
325 int
326 frp_enable_if_lookup (const char *ifname)
327 { unsigned int i;
328   char *str;
329   for (i = 0; i < vector_active (frp_enable_interface); i++)
330   { if ((str = vector_slot (frp_enable_interface, i)) != NULL)
331     { if (strcmp (str, ifname) == 0)
332       { return i;
333       }
334     }
335   }
336   return -1;
337 }
338
339 /* check interface is enabled by network statement */
340 /* check whether the interface has at least a connected prefix that is within the frp_enable_network table */
341 int
342 frp_enable_network_lookup_if (struct interface *ifp)
343 { struct listnode *node, *nnode;
344   struct connected *connected;
345   struct prefix_ipv4 address;
346   #ifdef DEB_DEBUG_IF
347     zlog_debug ("DEB DEBUG: entering frp_interface.c - frp_enable_network_lookup_if");
348     zlog_debug ("DEB DEBUG: -- frp_enable_network_lookup_if - ifp->connected = %p", ifp->connected);
349   #endif //DEB_DEBUG
350   if (ifp->connected == NULL)
351     return -1;
352   for (ALL_LIST_ELEMENTS (ifp->connected, node, nnode, connected))
353   {
354     struct prefix *p;
355     struct route_node *node;
356     #ifdef DEB_DEBUG_IF
357       if (connected == NULL)
358         zlog_debug ("DEB DEBUG: -- frp_enable_network_lookup_if - IF connected = null");
359     else
360       zlog_debug ("DEB DEBUG: -- frp_enable_network_lookup_if - ELSE connected->address = %s",
361 inet_ntoa(connected->address->u.prefix4));
362     #endif //DEB_DEBUG
363     p = connected->address;
364     if (p->family == AF_INET)
365     { address.family = AF_INET;
366       address.prefix = p->u.prefix4;

```

```

367     address.prefixlen = IPV4_MAX_BITLEN;
368     node = route_node_match (frp_enable_network, (struct prefix *)&address);
369     if (node)
370     { route_unlock_node (node);
371       return 1;
372     }
373   }
374 }
375 return -1;
376 }
377
378 /* is there an address on interface that can be used */
379 int
380 frp_if_ipv4_address_check (struct interface *ifp)
381 { struct listnode *nn;
382   struct connected *connected;
383   int count = 0;
384   if (ifp->connected == NULL)
385     return 0;
386   for (ALL_LIST_ELEMENTS_RO (ifp->connected, nn, connected))
387   { struct prefix *p;
388     p = connected->address;
389     if (p->family == AF_INET)
390     { count++;
391     }
392   }
393   return count;
394 }
395
396 /* join to multicast group and send request to the interface */
397 int
398 frp_interface_wakeup (struct thread *t)
399 { struct interface *ifp;
400   struct frp_interface *ri;
401   /* get interface */
402   ifp = THREAD_ARG (t);
403   ri = ifp->info;
404   ri->t_wakeup = NULL;
405   /* set running flag */
406   ri->running = 1;
407   return 0;
408 }
409
410 void
411 frp_connect_set (struct interface *ifp, int set)
412 { struct listnode *node, *nnode;
413   struct connected *connected;
414   struct prefix_ipv4 address;
415   for (ALL_LIST_ELEMENTS (ifp->connected, node, nnode, connected))
416   { struct prefix *p;
417     p = connected->address;
418     if (p->family != AF_INET)
419     { continue;
420     }
421     address.family = AF_INET;
422     address.prefix = p->u.prefix4;
423     address.prefixlen = p->prefixlen;
424     apply_mask_ipv4 (&address);
425   }
426 }
427

```

```

428 int
429 frp_if_down(struct interface *ifp)
430 { struct route_node *rp;
431   struct frp_info *rinfo;
432   struct frp_interface *ri = NULL;
433   if (frp)
434   { for (rp = route_top (frp->rib); rp = route_next (rp))
435     { if ((rinfo = rp->info) != NULL)
436       { /* routes got through this interface. */
437         if (rinfo->ifindex == ifp->ifindex && rinfo->type == ZEBRA_ROUTE_FRP && rinfo->sub_type ==
438             FRP_ROUTE_RTE)
439         { frp_zebra_ipv4_delete ((struct prefix_ipv4 *) &rp->p, &rinfo->nexthop, rinfo->cost);
440           } //else
441         }
442       }
443   }
444   ri = ifp->info;
445   if (ri->running)
446   { if (IS_FRP_DEBUG_EVENT)
447     { zlog_debug ("turn off %s", ifp->name);
448       ri->running = 0;
449     }
450   }
451   return 0;
452 }
453
454 void
455 frp_apply_address_del (struct connected *ifc) {
456   struct prefix_ipv4 address;
457   struct prefix *p;
458   if (!frp)
459     return;
460
461   if (!if_is_up(ifc->ifp))
462     return;
463
464   p = ifc->address;
465
466   memset (&address, 0, sizeof (address));
467   address.family = p->family;
468   address.prefix = p->u.prefix4;
469   address.prefixlen = p->prefixlen;
470   apply_mask_ipv4 (&address);
471 }
472
473 // find local address of interface a peer is attached to
474 struct in_addr
475 frp_get_interface_address (struct in_addr peer)
476 { struct prefix* local_p;
477   struct in_addr local_a;
478   local_p = find_local_address_for_peer(peer);
479   if (local_p != NULL)
480   {
481     #ifdef DEB_DEBUG_IF
482     zlog_debug ("DEB DEBUG: -- frp_get_interface_address - got a valid interface address");
483     #endif //DEB_DEBUG
484     int local_p_test = inet_pton (AF_INET, inet_ntoa (local_p->u.prefix4), &local_a);
485     } else
486   {
487     #ifdef DEB_DEBUG_IF
488     zlog_debug ("DEB DEBUG: -- frp_get_interface_address - failed to get local interface address");
489     }
490 }

```

```

489     #endif //DEB_DEBUG
490 }
491 return local_a;
492 }
493
494 int
495 frp_neighbor_lookup (struct sockaddr_in* from)
496 { struct prefix_ipv4 p;
497   struct route_node *node;
498   memset (&p, 0, sizeof (struct prefix_ipv4));
499   p.family = AF_INET;
500   p.prefix = from->sin_addr;
501   p.prefixlen = IPV4_MAX_BITLEN;
502   node = route_node_lookup (frp->neighbors, (struct prefix *) &p);
503   if (node)
504   { route_unlock_node (node);
505     return 1;
506   }
507   return 0;
508 }
509
510 void
511 frp_config_write_network (struct vty *vty)
512 { unsigned int i;
513   char *ifname;
514   struct route_node* rn;
515   /* network type frp enable interface statement */
516   for (rn = route_top (frp_enable_network); rn; rn = route_next (rn))
517   { if (rn->info)
518     { vty_out (vty, "%s%s/%d%s", " network", inet_ntoa (rn->p.u.prefix4), rn->p.prefixlen, VTY_NEWLINE);
519     }
520   }
521   /* interface name frp enable statement */
522   for (i = 0; i < vector_active (frp_enable_interface); i++)
523   { if ((ifname = vector_slot (frp_enable_interface, i)) != NULL)
524     { vty_out (vty, "%s%s%s", " network", ifname, VTY_NEWLINE);
525     }
526   }
527 }
528
529 /* called when interface structure allocated */
530 int
531 frp_interface_new_hook (struct interface *ifp)
532 {
533   #ifdef DEB_DEBUG_IF
534   zlog_debug ("DEB DEBUG: entering frp_interface_new_hook - interface->name = %s, interface->desc = %s, ",
535             ifp->name, ifp->desc);
536   #endif //DEB_DEBUG
537   ifp->info = frp_interface_new ();
538   return 0;
539 }
540
541 /* called when interface structure deleted */
542 int
543 frp_interface_delete_hook (struct interface *ifp)
544 { XFREE (MTYPE_FRP_INTERFACE, ifp->info);

```

```

550   ifp->info = NULL;
551   return 0;
552 }
553
554 // -----
555 // ROUTER COMMAND DEFUNS
556 // -----
557
558 /* frp enable a network or interface configuration */
559 DEFUN (frp_network,
560 frp_network_cmd,
561 "network (A.B.C.D/MIWORD)",
562 "Enable routing on an IP network\n"
563 "IP prefix <network>/<length>, e.g., 35.0.0.0/8\n"
564 "Interface name\n")
565 { int ret;
566   struct prefix_ipv4 p;
567   #ifdef DEB_DEBUG_IF
568   vty_out (vty, "DEB DEBUG: entering frp_interface.c - frp_network_cmd %s", VTY_NEWLINE);
569   #endif //DEB_DEBUG
570   ret = str2prefix_ipv4 (argv[0], &p);
571   if (ret)
572   {
573     #ifdef DEB_DEBUG_IF
574     vty_out (vty, "DEB DEBUG: -- frp_network_cmd - inside IF %s", VTY_NEWLINE);
575     #endif //DEB_DEBUG
576     ret = frp_enable_network_add ((struct prefix *) &p);
577     #ifdef DEB_DEBUG_IF
578     vty_out (vty, "DEB DEBUG: -- frp_network_cmd - successfully called frp_enable_network_add %s",
579             VTY_NEWLINE);
580     #endif //DEB_DEBUG
581   } else
582   {
583     #ifdef DEB_DEBUG_IF
584     vty_out (vty, "DEB DEBUG: -- frp_network_cmd - inside ELSE %s", VTY_NEWLINE);
585     #endif //DEB_DEBUG
586     ret = frp_enable_if_add (argv[0]);
587     #ifdef DEB_DEBUG_IF
588     vty_out (vty, "DEB DEBUG: -- frp_network_cmd - successfully called frp_enable_if_add %s", VTY_NEWLINE);
589     #endif //DEB_DEBUG
590   }
591   if (ret < 0)
592   { vty_out (vty, "There is a same network configuration %s%s", argv[0], VTY_NEWLINE);
593     return CMD_WARNING;
594   } else //send an update event
595   { frp_event (FRP_EVENT_UPDATE, 1);
596   }
597   return CMD_SUCCESS;
598 }
599
600 // -----
601 // INIT FUNCTION
602 // -----
603 void
604 frp_interface_init (void)
605 {
606   #ifdef DEB_DEBUG_IF

```

```
611     fprintf (stderr, "DEB DEBUG: entering frp_interface.c - frp_interface_init\n");
612     #endif //DEB_DEBUG
613
614     /* default initial size of interface vector */
615     if_init();
616     if_add_hook (IF_NEW_HOOK, frp_interface_new_hook);
617     if_add_hook (IF_DELETE_HOOK, frp_interface_delete_hook);
618
619     /* frp network init. */
620     frp_enable_interface = vector_init (1);
621     frp_enable_network = route_table_init ();
622
623     // create the FRP network command
624     install_element (FRP_NODE, &frp_network_cmd);
625     // install_element (FRP_NODE, &no_frp_neighbor_cmd);
626
627     return;
628 }
629
630
631
```

```
1  #ifndef _ZEBRA_FRP_DEBUG_H
2  #define _ZEBRA_FRP_DEBUG_H
3
4
5  // -----
6  /* INCLUDES */
7  // -----
8  #include "frpd.h"
9
10
11 // -----
12 /* DEFINES */
13 // -----
14 /* frp debug zebra flags */
15 #define FRP_DEBUG_ZEBRA    0x01
16 /* frp debug event flags */
17 #define FRP_DEBUG_EVENT    0x01
18 /* frp debug packet flags */
19 #define FRP_DEBUG_PACKET    0x01
20 #define FRP_DEBUG_SEND      0x20
21 #define FRP_DEBUG_RECV      0x40
22 #define FRP_DEBUG_DETAIL    0x80
23
24
25 // -----
26 /* EXTERNAL VARIABLES */
27 // -----
28 extern unsigned long frp_debug_zebra;
29 extern unsigned long frp_debug_event;
30 extern unsigned long frp_debug_packet;
31
32
33 // -----
34 /* MACROS */
35 // -----
36 #define IS_FRP_DEBUG_ZEBRA  (frp_debug_zebra & FRP_DEBUG_ZEBRA)
37
38 #define IS_FRP_DEBUG_EVENT  (frp_debug_event & FRP_DEBUG_EVENT)
39
40 #define IS_FRP_DEBUG_PACKET (frp_debug_packet & FRP_DEBUG_PACKET)
41 #define IS_FRP_DEBUG_SEND   (frp_debug_packet & FRP_DEBUG_SEND)
42 #define IS_FRP_DEBUG_RECV   (frp_debug_packet & FRP_DEBUG_RECV)
43 #define IS_FRP_DEBUG_DETAIL (frp_debug_packet & FRP_DEBUG_DETAIL)
44
45
46 // -----
47 // PROTOTYPES
48 // -----
49 extern void frp_debug_init (void);
50 extern void frp_debug_reset (void);
51
52
53 #endif /* _ZEBRA_FRP_DEBUG_H */
54
```

```
1 // -----
2 // INCLUDES
3 // -----
4 #include <zebra.h>
5 #include "command.h"
6
7 #include "frpd.h"
8 #include "frp_debug.h"
9
10 // -----
11 // GLOBAL VARIABLES
12 // -----
13
14 unsigned long frp_debug_zebra = 0;
15 unsigned long frp_debug_event = 0;
16 unsigned long frp_debug_packet = 0;
17
18 /* frp debug node */
19 static struct cmd_node debug_node =
20 { DEBUG_NODE,
21   "", /* Debug node has no interface. */
22   1
23 };
24
25 // -----
26 // FUNCTIONS
27 // -----
28
29 static int
30 frp_config_write_debug (struct vty *vty)
31 { int write = 0;
32   if (IS_FRP_DEBUG_ZEBRA)
33   { vty_out (vty, "debug frp zebra%s", VTY_NEWLINE);
34     write++;
35   }
36   if (IS_FRP_DEBUG_EVENT)
37   { vty_out (vty, "debug frp events%s", VTY_NEWLINE);
38     write++;
39   }
40   if (IS_FRP_DEBUG_PACKET)
41   { if (IS_FRP_DEBUG_SEND && IS_FRP_DEBUG_RECV)
42     { vty_out (vty, "debug frp packet%s%s", IS_FRP_DEBUG_DETAIL ? " detail" : "", VTY_NEWLINE);
43       write++;
44     } else
45     { if (IS_FRP_DEBUG_SEND)
46       { vty_out (vty, "debug frp packet send%s%s", IS_FRP_DEBUG_DETAIL ? " detail" : "", VTY_NEWLINE);
47     } else
48     { vty_out (vty, "debug frp packet rcv%s%s", IS_FRP_DEBUG_DETAIL ? " detail" : "", VTY_NEWLINE);
49     }
50     write++;
51   }
52   }
53   return write;
54 }
55
56 void
57 frp_debug_reset (void)
58 { frp_debug_zebra = 0;
59   frp_debug_event = 0;
60 }
61
```

```
62   frp_debug_packet = 0;
63 }
64
65 // -----
66 // ROUTER COMMAND DEFUNs
67 // -----
68
69 DEFUN (show_debugging_frp,
70       show_debugging_frp_cmd,
71       "show debugging frp",
72       SHOW_STR
73       DEBUG_STR
74       FRP_STR) /* DEB COMMENT: these 3 _STR defined in command.h
75 { vty_out (vty, "FRP debugging status:%s", VTY_NEWLINE);
76   if (IS_FRP_DEBUG_ZEBRA)
77   { vty_out (vty, " FRP zebra debugging is on%s", VTY_NEWLINE);
78   }
79   if (IS_FRP_DEBUG_EVENT)
80   { vty_out (vty, " FRP event debugging is on%s", VTY_NEWLINE);
81   }
82   if (IS_FRP_DEBUG_PACKET)
83   { if (IS_FRP_DEBUG_SEND && IS_FRP_DEBUG_RECV)
84     { vty_out (vty, " FRP packet%s debugging is on%s", IS_FRP_DEBUG_DETAIL ? " detail" : "", VTY_NEWLINE);
85     } else
86     { if (IS_FRP_DEBUG_SEND)
87       { vty_out (vty, " FRP packet send%s debugging is on%s", IS_FRP_DEBUG_DETAIL ? " detail" : "", VTY_NEW
88       } else
89       { vty_out (vty, " FRP packet receive%s debugging is on%s", IS_FRP_DEBUG_DETAIL ? " detail" : "",
90       VTY_NEWLINE);
91     }
92   }
93   }
94   return CMD_SUCCESS;
95 }
96
97 DEFUN (debug_frp_zebra,
98       debug_frp_zebra_cmd,
99       "debug frp zebra",
100       DEBUG_STR
101       FRP_STR
102       "FRP and ZEBRA communication\n")
103 { frp_debug_zebra = FRP_DEBUG_ZEBRA;
104   return CMD_WARNING;
105 }
106
107 DEFUN (no_debug_frp_zebra,
108       no_debug_frp_zebra_cmd,
109       "no debug frp zebra",
110       NO_STR
111       DEBUG_STR
112       FRP_STR
113       "FRP and ZEBRA communication\n")
114 { frp_debug_zebra = 0;
115   return CMD_WARNING;
116 }
117
118 DEFUN (debug_frp_events,
119       debug_frp_events_cmd,
120       "debug frp events",
121       DEBUG_STR
122
```

```
123     FRP_STR
124     "FRP events\n")
125 {
126     frp_debug_event = FRP_DEBUG_EVENT;
127     return CMD_WARNING;
128 }
129 DEFUN (no_debug_frp_events,
130     no_debug_frp_events_cmd,
131     "no debug frp events",
132     NO_STR
133     DEBUG_STR
134     FRP_STR
135     "FRP events\n")
136 {
137     frp_debug_event = 0;
138     return CMD_SUCCESS;
139 }
140
141 DEFUN (debug_frp_packet,
142     debug_frp_packet_cmd,
143     "debug frp packet",
144     DEBUG_STR
145     FRP_STR
146     "FRP packet\n")
147 {
148     frp_debug_packet = FRP_DEBUG_PACKET;
149     frp_debug_packet != FRP_DEBUG_SEND;
150     frp_debug_packet != FRP_DEBUG_RECV;
151     return CMD_SUCCESS;
152 }
153 DEFUN (no_debug_frp_packet,
154     no_debug_frp_packet_cmd,
155     "no debug frp packet",
156     NO_STR
157     DEBUG_STR
158     FRP_STR
159     "FRP packet\n")
160 {
161     frp_debug_packet = 0;
162     return CMD_SUCCESS;
163 }
164 DEFUN (debug_frp_packet_direct,
165     debug_frp_packet_direct_cmd,
166     "debug frp packet (recv|send)",
167     DEBUG_STR
168     FRP_STR
169     "FRP receive packet\n"
170     "FRP send packet\n")
171 {
172     frp_debug_packet != FRP_DEBUG_PACKET;
173     if (strcmp ("send", argv[0], strlen (argv[0])) == 0)
174     {
175         frp_debug_packet != FRP_DEBUG_SEND;
176     }
177     if (strcmp ("recv", argv[0], strlen (argv[0])) == 0)
178     {
179         frp_debug_packet != FRP_DEBUG_RECV;
180     }
181     frp_debug_packet &= ~FRP_DEBUG_DETAIL;
182     return CMD_SUCCESS;
183 }
184 DEFUN (no_debug_frp_packet_direct,
185     no_debug_frp_packet_direct_cmd,
186     "no debug frp packet (recv|send)",
```

```
184     NO_STR
185     DEBUG_STR
186     FRP_STR
187     "FRP packet\n"
188     "FRP option set for receive packet\n"
189     "FRP option set for send packet\n")
190 {
191     if (strcmp ("send", argv[0], strlen (argv[0])) == 0)
192     {
193         if (IS_FRP_DEBUG_RECV)
194         {
195             frp_debug_packet &= ~FRP_DEBUG_SEND;
196         }
197         else
198         {
199             frp_debug_packet = 0;
200         }
201     }
202     else if (strcmp ("recv", argv[0], strlen (argv[0])) == 0)
203     {
204         if (IS_FRP_DEBUG_SEND)
205         {
206             frp_debug_packet &= ~FRP_DEBUG_RECV;
207         }
208         else
209         {
210             frp_debug_packet = 0;
211         }
212     }
213     return CMD_SUCCESS;
214 }
215 DEFUN (debug_frp_packet_detail,
216     debug_frp_packet_detail_cmd,
217     "debug frp packet (recv|send) detail",
218     DEBUG_STR
219     FRP_STR
220     "FRP packet\n"
221     "FRP receive packet\n"
222     "FRP send packet\n"
223     "Detailed information display\n")
224 {
225     frp_debug_packet != FRP_DEBUG_PACKET;
226     if (strcmp ("send", argv[0], strlen (argv[0])) == 0)
227     {
228         frp_debug_packet != FRP_DEBUG_SEND;
229     }
230     if (strcmp ("recv", argv[0], strlen (argv[0])) == 0)
231     {
232         frp_debug_packet != FRP_DEBUG_RECV;
233     }
234     frp_debug_packet != FRP_DEBUG_DETAIL;
235     return CMD_SUCCESS;
236 }
237 DEFUN (debug_write_peers,
238     debug_write_peers_cmd,
239     "debug peers",
240     "Display current peers\n"
241     "Detailed peer information display\n")
242 {
243     #ifdef DEB_DEBUG
244     printf (stderr, "DEB DEBUG: entering frp_debug.c - debug_write_peers\n");
245     #endif //DEB_DEBUG
246     struct frp_peer* peer;
247     struct listnode *node, *nnode;
248     for (ALL_LIST_ELEMENTS (frp_peers, node, nnode, peer))
249     {
250         vty_out (vty, "peer information %s", VTY_NEWLINE);
251         vty_out (vty, "    address: %s %s", inet_ntoa (peer->address), VTY_NEWLINE);
252         vty_out (vty, "    secret: %s %s", peer->secret, VTY_NEWLINE);
253         vty_out (vty, "    cost: %d %s", peer->cost, VTY_NEWLINE);
254         vty_out (vty, "    poll: %d %s", peer->poll, VTY_NEWLINE);
255         vty_out (vty, "    retry: %d %s", peer->retry, VTY_NEWLINE);
256         vty_out (vty, "    lseq: %d %s", peer->lseq, VTY_NEWLINE);
257     }
258 }
```

```

245     vty_out (vty, "   rseq:  %d %s", peer->rseq, VTY_NEWLINE);
246     vty_out (vty, "   packet_latest_lseq:  %d %s", peer->packet_latest_lseq, VTY_NEWLINE);
247     vty_out (vty, "   gw cost:%d %s", peer->gateway_cost, VTY_NEWLINE);
248     vty_out (vty, "   gateway path %s", VTY_NEWLINE);
249     struct in_addr* gw_path_addr;
250     struct listnode *gw_node, *gw_nnode;
251     for (ALL_LIST_ELEMENTS (peer->gateway_path, gw_node, gw_nnode, gw_path_addr))
252     { vty_out (vty, "       address:  %s %s", inet_ntoa(*gw_path_addr), VTY_NEWLINE);
253     }
254     vty_out (vty, "   rib %s", VTY_NEWLINE);
255     struct frp_rte* rib_route;
256     struct listnode* rib_node;
257     for (ALL_LIST_ELEMENTS_RO (peer->rib, rib_node, rib_route))
258     { vty_out (vty, "       route %s", VTY_NEWLINE);
259       vty_out (vty, "         length:  %d %s", rib_route->length, VTY_NEWLINE);
260       vty_out (vty, "         routecost: %d %s", rib_route->routecost, VTY_NEWLINE);
261       vty_out (vty, "         gatecost:  %d %s", rib_route->routecost, VTY_NEWLINE);
262       vty_out (vty, "         prefix:   %s %s", inet_ntoa(rib_route->prefix.prefix), VTY_NEWLINE);
263     }
264 }
265 }
266 return CMD_SUCCESS;
267 }
268
269 DEFUN (debug_write_frp_rib,
270       debug_write_frp_rib_cmd,
271       "debug rib",
272       "Display the current rib\n"
273       "Detailed rib information display\n")
274 {
275     #ifdef DEB_DEBUG
276         fprintf (stderr, "DEB DEBUG: entering frp_debug.c - debug_write_frp_rib\n");
277     #endif //DEB_DEBUG
278     vty_out (vty, "frp rib %s", VTY_NEWLINE);
279     struct route_node* rib_rte_ptr;
280     for (rib_rte_ptr = route_top (frp->rib); rib_rte_ptr; rib_rte_ptr = route_next (rib_rte_ptr))
281     { //only display entries with frp information
282         struct frp_info* rte_info = (struct frp_info*)rib_rte_ptr->info;
283         if (rte_info != NULL)
284         { vty_out (vty, "   rib entry %s", VTY_NEWLINE);
285           vty_out (vty, "     prefix:  %s %s", inet_ntoa(rib_rte_ptr->p.u.prefix4), VTY_NEWLINE);
286           vty_out (vty, "     nexthop:  %s %s", inet_ntoa(rte_info->nexthop), VTY_NEWLINE);
287           vty_out (vty, "     cost:    %d %s", rte_info->cost, VTY_NEWLINE);
288           vty_out (vty, "     is_gateway:  %d %s", rte_info->is_gateway_flag, VTY_NEWLINE);
289         }
290         route_unlock_node (rib_rte_ptr);
291     }
292     return CMD_SUCCESS;
293 }
294
295 // -----
296 // INIT FUNCTION
297 // -----
298 void
299 frp_debug_init (void)
300 {
301     #ifdef DEB_DEBUG
302         fprintf (stderr, "DEB DEBUG: entering frp_debug.c - frp_debug_init\n");
303     #endif //DEB_DEBUG

```

```

306
307     frp_debug_zebra = 0;
308     frp_debug_event = 0;
309     frp_debug_packet = 0;
310
311     install_node (&debug_node, frp_config_write_debug);
312
313     install_element (ENABLE_NODE, &show_debugging_frp_cmd);
314
315     install_element (ENABLE_NODE, &debug_frp_zebra_cmd);
316     install_element (CONFIG_NODE, &debug_frp_zebra_cmd);
317     install_element (ENABLE_NODE, &no_debug_frp_zebra_cmd);
318     install_element (CONFIG_NODE, &no_debug_frp_zebra_cmd);
319
320     install_element (ENABLE_NODE, &debug_frp_events_cmd);
321     install_element (CONFIG_NODE, &debug_frp_events_cmd);
322     install_element (ENABLE_NODE, &no_debug_frp_events_cmd);
323     install_element (CONFIG_NODE, &no_debug_frp_events_cmd);
324
325     install_element (ENABLE_NODE, &debug_frp_packet_cmd);
326     install_element (CONFIG_NODE, &debug_frp_packet_cmd);
327     install_element (ENABLE_NODE, &no_debug_frp_packet_cmd);
328     install_element (CONFIG_NODE, &no_debug_frp_packet_cmd);
329     install_element (ENABLE_NODE, &debug_frp_packet_direct_cmd);
330     install_element (CONFIG_NODE, &debug_frp_packet_direct_cmd);
331     install_element (ENABLE_NODE, &no_debug_frp_packet_direct_cmd);
332     install_element (CONFIG_NODE, &no_debug_frp_packet_direct_cmd);
333     install_element (ENABLE_NODE, &debug_frp_packet_detail_cmd);
334     install_element (CONFIG_NODE, &debug_frp_packet_detail_cmd);
335
336     install_element (CONFIG_NODE, &debug_write_peers_cmd);
337
338     install_element (CONFIG_NODE, &debug_write_frp_rib_cmd);
339
340     #ifdef DEB_DEBUG
341         fprintf (stderr, "DEB DEBUG: -- leaving frp_debug_init\n");
342     #endif //DEB_DEBUG
343 }
344

```



---

## Appendix D

---

### The FRPsniffer implementation



```

1 // -----
2 // FRPsniffer
3 // a dedicated packet sniffer for the FRP protocol
4 // written by Deb Shepherd
5 // -----
6
7
8 #include <stdio.h>
9 #include <netinet/in.h>
10 #include <sys/types.h>
11 #include <stdlib.h>
12 #include <stdint.h>
13 #include <stddef.h>
14 #include <arpa/inet.h>
15
16 // PACKET CAPTURE LIBRARY
17 #include <pcap.h>
18 // integer, defines max bytes to be captured
19 #define SNAPLEN 65535
20 // when true brings into promiscuous mode
21 #define PROMISC 1
22 //read timeout in milliseconds (0 means no time out)
23 #define TO_MS 10000
24
25 // contains definition for 'in_addr' used in ip4 header
26 #include <netinet/in.h>
27
28 // ethernet headers are 14 bytes
29 #define EN_HEADER_LEN 14
30 // ethernet addresses are 6 bytes
31 #define EN_ADDRESS_LEN 6
32 // Ethernet Header (14 bytes)
33 struct ethernet
34 { u_char enDestination[EN_ADDRESS_LEN]; // destination address (8 bits x 6 = 6 bytes)
35   u_char enSource[EN_ADDRESS_LEN]; // source address (8 bits x 6 = 6 bytes)
36   u_short enType; // IP, ARP, etc (16 bits = 2 bytes)
37 };
38
39 // IPv4 Header (20+ bytes)
40 struct ip4
41 { u_char ip4vhl; // version and header length (4+4 bits)
42   // extract version from ip4vhl by bit shifting 4 (4 bits)
43   #define IP4VERSION(ipv) (((ip4vhl)->ip4vhl) >> 4)
44   // extract header length from ip4vhl by masking with 0000 1111 (4 bits)
45   #define IP4HEADER_LEN(ip4hl) (((ip4hl)->ip4vhl) & 0x0f)
46   u_char ip4service; // type of service (8 bits)
47   u_short ip4length; // total datagram length (16 bits)
48   u_short ip4id; // identifier (16 bits)
49   u_short ip4offset; // flags and fragmentation offset (3+13 bits)
50   // reserved, fragment flag, mask (1 bit) 1000 0000 0000 0000
51   #define IP4RESERVED 0x8000
52   // don't fragment flag, mask (1 bit) 0100 0000 0000 0000
53   #define IP4DONT 0x4000
54   // more fragments flag, mask (1 bit) 0010 0000 0000 0000
55   #define IP4MORE 0x2000
56   // mask for fragmenting bits (13 bits) 0001 1111 1111 1111
57   #define IP4MASK 0x1fff
58   u_char ip4ttl; // time to live (8 bits)
59   u_char ip4protocol; // protocol (8 bits)
60   u_short ip4checksum; // checksum (16 bits)
61   struct in_addr ip4source; // source address (32 bits) (needs include <netinet/in.h>)

```

```

62 struct in_addr ip4destination; // destination address (32 bits) (needs include <netinet/in.h>)
63 };
64
65 // udp headers are 8 bytes
66 #define UDP_HEADER_LEN 8
67 // UDP Header (8 bytes)
68 struct udp
69 { u_short udpSource; // source port (16 bits)
70   u_short udpDestination; // destination port (16 bits)
71   u_short udpLength; // message length (16 bits)
72   u_short udpChecksum; // checksum (16 bits)
73 };
74

```

```
1 // -----
2 // FRPsniffer
3 // a dedicated packet sniffer for the FRP protocol
4 // written by Deb Shepherd
5 // -----
6
7
8 #define FRP_PKT_HDRSIZE 16 // (128 bits)
9 #define FRP_PKT_MINSIZE 16 // (128 bits)
10 #define FRP_PKT_MAXSIZE 1400 // as specified by Don
11 #define FRP_MSG_CONTROL_SIZE 4 // (32 bits)
12 #define FRP_MSG_IPV4CONFIG_SIZE 12 // (96 bits)
13 #define FRP_MSG_IPV4GATEWAY_MINSIZE 8 // (64 bits)
14 #define FRP_MSG_IPV4UPDATE_SIZE 12 // (96 bits)
15 #define FRP_MSG_IPV6CONFIG_SIZE 24 // (192 bits)
16 #define FRP_MSG_IPV6GATEWAY_MINSIZE 20 // (160 bits)
17 #define FRP_MSG_IPV6UPDATE_SIZE 24 // (192 bits)
18
19 // frp message types
20 #define FRP_MSG_NULL 0x00
21 #define FRP_MSG_CONTROL 0x01
22 #define FRP_MSG_IPV4CONFIG 0x41
23 #define FRP_MSG_IPV4GATEWAY 0x42
24 #define FRP_MSG_IPV4UPDATE 0x43
25 #define FRP_MSG_IPV6CONFIG 0x61
26 #define FRP_MSG_IPV6GATEWAY 0x62
27 #define FRP_MSG_IPV6UPDATE 0x63
28
29 // frp control types
30 #define FRP_CTRL_POLL 1
31 #define FRP_CTRL_ACK 2
32 #define FRP_CTRL_NAK 3
33
34 // frp update flags
35 #define FRP_FLAG_BEGIN 0x01
36 #define FRP_FLAG_COMMIT 0x02
37 #define FRP_FLAG_NULLRT 0x04
38 #define FRP_FLAG_UPDATE 0x08
39 #define FRP_FLAG_DELETE 0x10
40 #define FRP_FLAG_GATEWAY 0x80
41
42
43 // frp packet header (128 bits)
44 struct frp_pkt_hdr
45 { u_int8_t hash[8]; // security hash (64 bits)
46   u_int32_t sendSeq; // sender's sequence number (32 bits)
47   u_int32_t recipAck; // recipient's acknowledgement number (32 bits)
48 };
49
50 // frp header (16 bits)
51 struct frp_msg_hdr
52 { u_int8_t length; // message length (8 bits)
53   u_int8_t type; // message type (8 bits)
54 };
55
56 // frp control (32 bits)
57 struct frp_msg_control
58 { struct frp_msg_hdr msg_hdr; // message header (16 bits)
59   u_int8_t type; // control type (8 bits)
60   u_int8_t param; // control parameters (8 bits)
61 };
```

```
62
63 // frp IPv4 configuration (96 bits)
64 struct frp_msg_ipv4config
65 { struct frp_msg_hdr msg_hdr; // message header (16 bits)
66   u_short cost; // cost of the link (16 bits)
67   u_short poll; // poll time (16 bits)
68   u_short retry; // retry time (16 bits)
69   struct in_addr id; // router-ID of peer (32 bits)
70 };
71
72 // frp IPv4 path to gateway (64-2016 bits)
73 struct frp_msg_ipv4gateway
74 { struct frp_msg_hdr msg_hdr; // message header (16 bits)
75   u_short cost; // cost from peer to gateway (16 bits)
76   struct in_addr path[62]; // path from peer to gateway (32xn bits, 1 <= n <= 62)
77 };
78
79 // frp IPv4 route update (96 bits)
80 struct frp_msg_ipv4update
81 { struct frp_msg_hdr msg_hdr; // message header (16 bits)
82   u_int8_t flags; // update type flags (8 bits)
83   u_int8_t length; // prefix length (8 bits)
84   u_short routecost; // route cost (16 bits)
85   u_short gatecost; // cost from originator to gateway (16 bits)
86   struct in_addr prefix; // IP prefix (32 bits)
87 };
88
89 // frp IPv6 configuration (192 bits)
90 struct frp_msg_ipv6config
91 { struct frp_msg_hdr msg_hdr; // message header (16 bits)
92   u_short cost; // cost of the link (16 bits)
93   u_short poll; // poll time (16 bits)
94   u_short retry; // retry time (16 bits)
95   struct in6_addr id; // router-ID of peer (128 bits)
96 };
97
98 // frp IPv6 path to gateway (160-7968 bits)
99 struct frp_msg_ipv6gateway
100 { struct frp_msg_hdr msg_hdr; // message header (16 bits)
101   u_short cost; // cost from peer to gateway (16 bits)
102   struct in6_addr path[1]; // path from peer to gateway (32xn bits, 1 <= n <= 62)
103 };
104
105 // frp IPv6 route update (192 bits)
106 struct frp_msg_ipv6update
107 { struct frp_msg_hdr msg_hdr; // message header (16 bits)
108   u_int8_t flags; // update type flags (8 bits)
109   u_int8_t length; // prefix length (8 bits)
110   u_short routecost; // route cost (16 bits)
111   u_short gatecost; // cost from originator to gateway (16 bits)
112   struct in6_addr prefix; // IP prefix (128 bits)
113 };
114
115
116
117 // -----
118 // GLOBAL VARIABLES AND DEFINES
119 // -----
120
121 struct frp_pkt_hdr pkt_hdr;
122 struct frp_pkt_hdr* pkt_hdr_ptr;
```

```
123
124 /* frp event. */
125 enum frp_makepkthdr_flag
126 { FRP_ACK,
127   FRP_SYN,
128 };
129
```

```
1 // -----
2 // FRPsniffer
3 // a dedicated packet sniffer for the FRP protocol
4 // written by Deb Shepherd
5 // -----
6
7
8 #include "sniffer.h"
9 #include "frp_packet.h"
10
11 u_short displayControlMessage (const struct frp_msg_control* frpControlMessage)
12 { switch (frpControlMessage->type) {
13     case FRP_CTRL_POLL:
14         printf("--Poll, param = %d \n", frpControlMessage->param);
15         break;
16     case FRP_CTRL_ACK:
17         printf("--Control ACK, param = %d \n", frpControlMessage->param);
18         break;
19     case FRP_CTRL_NAK:
20         printf("--Control NAK, param = %d \n", frpControlMessage->param);
21         break;
22     default:
23         break;
24 }
25 return frpControlMessage->msg_hdr.length;
26 }
27
28 u_short displayIP4ConfigMessage (const struct frp_msg_ip4config* frpConfigMessage)
29 { printf("cost: %d\n", ntohs(frpConfigMessage->cost));
30   printf("poll: %d\n", ntohs(frpConfigMessage->poll));
31   printf("retry: %d\n", ntohs(frpConfigMessage->retry));
32   printf("router-id: %s\n", inet_ntoa(frpConfigMessage->id));
33   return frpConfigMessage->msg_hdr.length;
34 }
35
36 u_short displayIP4GatewayMessage(const struct frp_msg_ip4gateway* frpGatewayMessage)
37 {
38     printf("cost: %d\n", ntohs(frpGatewayMessage->cost));
39     u_short numNodes = frpGatewayMessage->msg_hdr.length / 8;
40     printf(" # of nodes in path: %d\n", numNodes);
41     for (int i = 0; i < numNodes; i++)
42     {
43         printf("path[%d]: %s\n", i, inet_ntoa(frpGatewayMessage->path[i]));
44     }
45     return frpGatewayMessage->msg_hdr.length;
46 }
47
48 u_short displayIP4UpdateMessage(const struct frp_msg_ip4update* frpUpdateMessage)
49 { printf("flags: %d\n", ntohs(frpUpdateMessage->flags));
50   if (ntohs(frpUpdateMessage->flags & FRP_FLAG_BEGIN)){
51       printf("FRP_FLAG_BEGIN flag: YES\n");
52   }
53   else {
54       printf("FRP_FLAG_BEGIN flag: NO\n");
55   }
56   if (ntohs(frpUpdateMessage->flags & FRP_FLAG_NULLRT)){
57       printf("FRP_FLAG_NULLRT flag: YES\n");
58   }
59   else {
60       printf("FRP_FLAG_NULLRT flag: NO\n");
61   }
62 }
```

```
62 if (ntohs(frpUpdateMessage->flags & FRP_FLAG_DELETE)){
63     printf("FRP_FLAG_DELETE flag: YES\n");
64 }
65 else {
66     printf("FRP_FLAG_DELETE flag: NO\n");
67 }
68 if (ntohs(frpUpdateMessage->flags & FRP_FLAG_COMMIT)){
69     printf("FRP_FLAG_COMMIT flag: YES\n");
70 }
71 else {
72     printf("FRP_FLAG_COMMIT flag: NO\n");
73 }
74 if (ntohs(frpUpdateMessage->flags & FRP_FLAG_UPDATE)){
75     printf("FRP_FLAG_UPDATE flag: YES\n");
76 }
77 else {
78     printf("FRP_FLAG_UPDATE flag: NO\n");
79 }
80 if (ntohs(frpUpdateMessage->flags & FRP_FLAG_GATEWAY)){
81     printf("FRP_FLAG_GATEWAY flag: YES\n");
82 }
83 else {
84     printf("FRP_FLAG_GATEWAY flag: NO\n");
85 }
86 printf("length: %d\n", frpUpdateMessage->length);
87 printf("routecost: %d\n", ntohs(frpUpdateMessage->routecost));
88 printf("gatecost: %d\n", ntohs(frpUpdateMessage->gatecost));
89 printf("prefix: %s\n", inet_ntoa(frpUpdateMessage->prefix));
90 return frpUpdateMessage->msg_hdr.length;
91 }
92
93 int main()
94 { // INTERFACE VARIABLES
95     char* interface = "en0\0"; // a pointer to the interface to sniff - hard code for now ...
96     bpf_u_int32 netaddress; // interface net address
97     bpf_u_int32 netmask; // interface netmask
98     // PCAP SESSION VARIABLES
99     char errorbuffer[PCAP_ERRBUF_SIZE]; // pcap error string
100    pcap_t* handle = NULL; // a pointer to the pcap session handle
101    // PCAP filter variables
102    char filterexpression[] = "udp port 343"; // filter expression
103    struct bpf_program filter; // compiled filter expression
104    // packet variables
105    struct pcap_pkthdr packetHeader; // the packet header from pcap
106    const u_char* packet; // a pointer to the actual packet
107    const struct ethernet* ethernetHeader; // a pointer to the ethernet header
108    const struct ip4* ip4Header; // a pointer to the IPv4 header
109    u_int ip4HeaderSize; // actual size of the IPv4 header
110    const struct udp* udpHeader; // a pointer to the UDP header
111    u_int udpHeaderSize; // actual size of the UDP header
112    const struct frp_pkt_hdr* frpHeader; // a pointer to the FRP header
113    u_int frpHeaderSize; // actual size of the FRP header
114    //const char* payload; // a pointer to the packet payload
115
116    // find out about the interface (pcap_lookupnet)
117    printf("Interface: %s\n", interface);
118    if (pcap_lookupnet(interface, &netaddress, &netmask, errorbuffer) == -1)
119    { fprintf(stderr, "Can't get netmask for interface %s\n", interface);
120      netaddress = 0;
121      netmask = 0;
122    }
123 }
```

```
123
124 // set up a session handle and open session (pcap_open_live)
125 handle = pcap_open_live(interface, SNAPLEN, PROMISC, TO_MS, errorbuffer);
126 if (handle == NULL)
127 { fprintf(stderr, "Couldn't open interface %s: %s\n", interface, errorbuffer);
128   return(1);
129 }
130
131 // compile and apply 'udp port 343' filter
132 if (pcap_compile(handle, &filter, filterexpression, 0, netmask) == -1)
133 { fprintf(stderr, "Couldn't parse filter %s: %s\n", filterexpression, pcap_geterr(handle));
134   return(2);
135 }
136 if (pcap_setfilter(handle, &filter) == -1)
137 { fprintf(stderr, "Couldn't install filter %s: %s\n", filterexpression, pcap_geterr(handle));
138   return(2);
139 }
140
141 while (1)
142 { printf("Waiting for incoming FRP packet\n");
143   // collect a single packet from the interface (pcap_next)
144   packet = pcap_next(handle, &packetHeader);
145   if (packet == NULL)
146   { //printf("No packet collected");
147     } else
148     {
149       printf("\n");
150       printf("\nFRP PACKET HEADER\n");
151       // printf("Packet timestamp: [%d]\n", packetHeader.ts);
152       // printf("Packet available bytes: [%d]\n", packetHeader.caplen);
153       // printf("Packet length: [%d]\n", packetHeader.len);
154
155       // read packet headers
156       // ethernet
157       ethernetHeader = (struct ethernet*)(packet); // read from point 0 as an ethernet header
158       // printf("ETHERNET\n");
159       u_short byte0 = ethernetHeader->enDestination[0];
160       u_short byte1 = ethernetHeader->enDestination[1];
161       u_short byte2 = ethernetHeader->enDestination[2];
162       u_short byte3 = ethernetHeader->enDestination[3];
163       u_short byte4 = ethernetHeader->enDestination[4];
164       u_short byte5 = ethernetHeader->enDestination[5];
165       // printf("destination address: %x:%x:%x:%x:%x:%x\n", byte0, byte1, byte2, byte3, byte4, byte5);
166       byte0 = ethernetHeader->enSource[0];
167       byte1 = ethernetHeader->enSource[1];
168       byte2 = ethernetHeader->enSource[2];
169       byte3 = ethernetHeader->enSource[3];
170       byte4 = ethernetHeader->enSource[4];
171       byte5 = ethernetHeader->enSource[5];
172       // printf("source address: %x:%x:%x:%x:%x:%x\n", byte0, byte1, byte2, byte3, byte4, byte5);
173       // printf("ethernet type: %x\n", ntohs(ethernetHeader->enType));
174       // IPv4
175       ip4Header = (struct ip4*)(packet + EN_HEADER_LEN); // read from point 14 as an IPv4 header
176       ip4HeaderSize = IP4HEADER_LEN(ip4Header)*4; // calculate actual IPv4 header length
177       // printf("IPv4\n");
178       if (ip4HeaderSize < 20) // check IPv4 header is at least the min 20 bytes
179       {
180         // printf(" * Invalid IP header length: %u (%u) bytes\n", ip4HeaderSize, IP4HEADER_LEN(ip4Header));
181         return (2);
182       } else
183       {
```

```
184 // printf("version: %u (0x%x)\n", IP4VERSION(ip4Header), IP4VERSION(ip4Header));
185 // printf("header length: %u bytes (0x%x)\n", ip4HeaderSize, ip4HeaderSize);
186 // printf("type of service: 0x%x\n", ip4Header->ip4service);
187 // printf("datagram length: %u bytes (0x%x)\n", ntohs(ip4Header->ip4length), ntohs(ip4Header->ip4len));
188 // printf("identifier: %u (0x%x)\n", ntohs(ip4Header->ip4id), ntohs(ip4Header->ip4id));
189 byte0 = (ip4Header->ip4offset & IP4RESERVED);
190 byte1 = (ip4Header->ip4offset & IP4DONT);
191 byte2 = (ip4Header->ip4offset & IP4MORE);
192 byte3 = (ip4Header->ip4offset & IP4MASK);
193 // printf("flags: reserved(%x) don't(%x) more(%x) offset(%x)\n", byte0, byte1, byte2, byte3);
194 // printf("time to live: %u (0x%x)\n", ip4Header->ip4ttl, ip4Header->ip4ttl);
195 // printf("protocol: %u (0x%x)\n", ip4Header->ip4protocol, ip4Header->ip4protocol);
196 // printf("checksum: %u (0x%x)\n", ntohs(ip4Header->ip4checksum), ntohs(ip4Header->ip4checksum));
197 byte0 = (ip4Header->ip4source.s_addr & 0xFF);
198 byte1 = (ip4Header->ip4source.s_addr & 0xFF00) >> 8;
199 byte2 = (ip4Header->ip4source.s_addr & 0xFF0000) >> 16;
200 byte3 = (ip4Header->ip4source.s_addr & 0xFF000000) >> 24;
201 printf("source address: %u.%u.%u.%u (%x.%x.%x.%x)\n", byte0, byte1, byte2, byte3, byte0, byte1,
202 byte2, byte3);
203 byte0 = (ip4Header->ip4destination.s_addr & 0xFF);
204 byte1 = (ip4Header->ip4destination.s_addr & 0xFF00) >> 8;
205 byte2 = (ip4Header->ip4destination.s_addr & 0xFF0000) >> 16;
206 byte3 = (ip4Header->ip4destination.s_addr & 0xFF000000) >> 24;
207 printf("destination address: %u.%u.%u.%u (%x.%x.%x.%x)\n", byte0, byte1, byte2, byte3, byte0, byte1,
208 byte2, byte3);
209 }
210 // UDP
211 udpHeader = (struct udp*)(packet + EN_HEADER_LEN + ip4HeaderSize); // read from point 14+20+? as a U
212 header
213 udpHeaderSize = UDP_HEADER_LEN; // UDP header length
214 // printf("UDP\n");
215 // printf("source port: %u (0x%x)\n", ntohs(udpHeader->udpSource), ntohs(udpHeader->udpSource));
216 // printf("destination port: %u (0x%x)\n", ntohs(udpHeader->udpDestination),
217 ntohs(udpHeader->udpDestination));
218 // printf("header length: %u bytes (0x%x)\n", ntohs(udpHeader->udpLength), ntohs(udpHeader->udpLength));
219 // printf("checksum: %u (0x%x)\n", ntohs(udpHeader->udpChecksum), ntohs(udpHeader->udpChecksum));
220
221 // FRP
222 frpHeader = (struct frp_pkt_hdr*)(packet + EN_HEADER_LEN + ip4HeaderSize + udpHeaderSize); // read f
223 point 14+20+7+8 as a FRP header
224 frpHeaderSize = FRP_PKT_HDRSIZE; // FRP header length
225 // printf("FRP\n");
226
227 //Display the FRP packet header
228 // printf("\nFRP MESSAGE\n");
229
230
231 printf("security hash: %s\n", frpHeader->hash);
232 printf("sender's seq no: %d (0x%x)\n", ntohl(frpHeader->sendSeq), ntohl(frpHeader->sendSeq));
233 printf("recipient's ack no: %d (0x%x)\n", ntohl(frpHeader->recipAck), ntohl(frpHeader->recipAck));
234 if (ntohl(frpHeader->recipAck) == 0)
235 { printf("--SYN Packet\n\n");
236   continue;
237 }
238 //if (ntohl(frpHeader->recipAck) != 0)
239 //{ printf("message length: %u (0x%x)\n", frpHeader->msgLength, frpHeader->msgLength);
240 // printf("message type: 0x%x\n", frpHeader->msgType);
241 //}
242 u_long currentPos = EN_HEADER_LEN + ip4HeaderSize + udpHeaderSize + frpHeaderSize;
243 struct frp_msg_hdr* messageHeader = (struct frp_msg_hdr*)(packet + currentPos);
244
```

```
245 //Work through the packet extracting and displaying the messages until we get to the end of the packet
246 while ((messageHeader != 0)&&(currentPos < packetHeader.len))
247 { switch (messageHeader->type) {
248     case FRP_MSG_CONTROL:
249         printf("\nControl message: 0x%x\n", messageHeader->type);
250         currentPos += displayControlMessage((struct frp_msg_control*)messageHeader);
251         break;
252     case FRP_MSG_IPV4CONFIG:
253         printf("\nIPv4 Config message: 0x%x\n", messageHeader->type);
254         currentPos += displayIP4ConfigMessage((struct frp_msg_ipv4config*)messageHeader);
255         break;
256     case FRP_MSG_IPV4GATEWAY:
257         printf("\nIPv4 path to gateway message: 0x%x\n", messageHeader->type);
258         currentPos += displayIP4GatewayMessage((struct frp_msg_ipv4gateway*)messageHeader);
259         break;
260     case FRP_MSG_IPV4UPDATE:
261         printf("\nIPv4 route update message: 0x%x\n", messageHeader->type);
262         currentPos += displayIP4UpdateMessage((struct frp_msg_ipv4update*)messageHeader);
263         break;
264     case FRP_MSG_IPV6CONFIG:
265         printf("    IPv6 Config Message: 0x%x\n", messageHeader->type);
266         printf("    Not currently implemented\n");
267         //Not currently implemented
268         break;
269     case FRP_MSG_IPV6GATEWAY:
270         printf("    IPv6 Path To Gateway Message: 0x%x\n", messageHeader->type);
271         printf("    Not currently implemented\n");
272         //Not currently implemented
273         break;
274     case FRP_MSG_IPV6UPDATE:
275         printf("    IPv6 Route Update Message: 0x%x\n", messageHeader->type);
276         printf("    Not currently implemented\n");
277         //Not currently implemented
278         break;
279     case FRP_MSG_NULL:
280         printf("    Null Message: 0x%x\n", messageHeader->type);
281         if ((ntohl(frpHeader->recipAck) != 0)&&(ntohl(frpHeader->sendSeq) != 0))
282         { printf("--ACK Packet\n\n");
283         }
284         messageHeader = 0;
285         continue;
286         break;
287     default:
288         printf("    INVALID MESSAGE TYPE: 0x%x\n", messageHeader->type);
289         printf("    CORRUPT MESSAGE STOPPING PARSE\n");
290         messageHeader = 0;
291         continue;
292         break;
293 }
294 printf("\n");
295 //Advance the message pointer
296 messageHeader = (struct frp_msg_hdr*)(packet + currentPos);
297
298 // the remainder is actual data
299 //payload = (char *) (packet + EN_HEADER_LEN + ip4HeaderSize + udpHeaderSize + frpHeaderSize);
300 // printf("\nEN_HEADER_LEN: %u, ip4HeaderSize: %u, udpHeaderSize: %u, frpHeaderSize: %u", EN_HEAD
301 ip4HeaderSize, udpHeaderSize, frpHeaderSize);
302 // payload = (char *) (packet + EN_HEADER_LEN + ip4HeaderSize + tcpHeaderSize);
303 // printf("\nPAYLOAD: %s\n\n", payload);
304 }
305 }
```

```
306 }
307
308 // close session
309 pcap_close(handle);
310 return(0);
311 }
312
```



---

## Appendix D

---

**The original implementation of FRP by Don Stokes**



```

1  /*
2  Packet header
3  */
4  #define FRP_HASHSIZE 8
5  struct frp_hdr {
6      u_int8_t hash[FRP_HASHSIZE]; /* Security hash of packet/secret */
7      u_int32_t lseq; /* Local sequence number */
8      u_int32_t rseq; /* Remote sequence number */
9  };
10
11
12 /*
13 Message header
14 */
15 struct frp_mhdr {
16     u_int8_t len; /* Length of message, in 32 bit words */
17     u_int8_t type; /* Message type */
18 };
19
20 /*
21 Control message
22 */
23 #define FRP_CTRL 0x01
24 struct frp_ctrl {
25     struct frp_mhdr mh;
26     u_int8_t ctrl;
27     u_int8_t param;
28 };
29 #define FRP_CTRL_POLL 1
30 #define FRP_CTRL_ACK 2
31 #define FRP_CTRL_NAK 3
32
33 /*
34 IPv4 protocol messages
35 */
36 #define FRP_CONFIG 0x41 /* IPv4 static configuration */
37 struct frp_config {
38     struct frp_mhdr mh;
39     u_int16_t cost; /* Cost of link */
40     u_int16_t poll; /* Poll time */
41     u_int16_t fail; /* Retry time */
42     IPADDR id; /* Router-ID of peer */
43 };
44
45 #define FRP_PATH 0x42 /* IPv4 path to gateway */
46 struct frp_path {
47     struct frp_mhdr mh;
48     u_int16_t gwcost; /* Cost from peer to gateway */
49     IPADDR path[]; /* Path from peer to gateway */
50 };
51 #define FRP_NOGATEWAY 0xffff /* gwcost value if gateway not reachable */
52
53 #define FRP_ROUTE 0x43 /* IPv4 route update */
54 struct frp_route {
55     struct frp_mhdr mh;
56     u_int8_t flags; /* See below */
57     u_int8_t bits; /* Prefix length */
58     u_int16_t cost; /* Route cost */
59     u_int16_t gwcost; /* Cost from originator to gateway */
60     IPADDR ip; /* IP prefix */
61 };

```

```

62 #define FRP_FLAG_BEGIN 1 /* First route of batch update */
63 #define FRP_FLAG_COMMIT 2 /* Last route of batch update, commit */
64 #define FRP_FLAG_NULL 4 /* Null route (do not add) */
65 // #define FRP_FLAG_UPDATE 8 /* Add/change single route */
66 // #define FRP_FLAG_DELETE 16 /* Delete route */
67 #define FRP_FLAG_GATEWAY 128 /* Route is a gateway route */
68 #define FRP_NULLRT_SIZE 4 /* Size of null route message */
69
70
71 /*
72 IPv6 message types
73 */
74 #define FRP_CONFIG6 0x61 /* IPv6 configuration */
75 struct frp_config6 {
76     struct frp_mhdr mh;
77     u_int16_t cost; /* Cost of link */
78     struct in6_addr gw; /* Gateway address to use */
79 };
80
81 #define FRP_PATH6 0x62 /* IPv6 path to gateway */
82 struct frp_path6 {
83     struct frp_mhdr mh;
84     u_int16_t gwcost; /* Cost from peer to gateway */
85     struct in6_addr path[]; /* Path from peer to gateway */
86 };
87
88 #define FRP_ROUTE6 0x63 /* IPv6 route update */
89 struct frp_route6 {
90     struct frp_mhdr mh;
91     u_int8_t flags; /* See above */
92     u_int8_t bits; /* Prefix length */
93     u_int16_t cost; /* Route cost */
94     u_int16_t gwcost; /* Cost from originator to gateway */
95     struct in6_addr ip6; /* IPv6 prefix, truncate to nearest 32 bits */
96 };
97

```

```
1 #include <stdio.h>
2 #include <string.h>
3 #include <stdlib.h>
4 #include <stdarg.h>
5 #include <unistd.h>
6 #include <ctype.h>
7 #include <errno.h>
8 #include <sys/types.h>
9 #include <sys/socket.h>
10 #include <sys/select.h>
11 #include <net/if.h>
12 #include <net/if_dl.h>
13 #include <net/if_types.h>
14 #include <net/route.h>
15 #include <netinet/in.h>
16 #include <netinet/if_ether.h>
17 #include <arpa/inet.h>
18 #include <sys/sysctl.h>
19 #include <fcntl.h>
20 #include <signal.h>
21 #include <ifaddrs.h>
22 #include <syslog.h>
23
24 #define IPADDR u_int32_t
25 #define TIMETEN u_int32_t
26
27 #define WHITESPACE (" \t\n\r")
28
29 struct ace {
30     struct ace *next;
31     int permit;
32     IPADDR ip;
33     int bits;
34     int maxbits;
35     int rtype;
36     struct acl *acl;
37     int ifindex;
38 };
39 #define RTYPE_ANY 0
40 #define RTYPE_LAYER2 1
41 #define RTYPE_STATIC 2
42 #define RTYPE_PROTOCOL 4
43 #define RTYPE_LOCAL 8
44 #define RTYPE_REMOTE 16
45
46 #define MAX_PACKET 1360
47 #define MAX_PATH 63
48 #define DEFAULT_POLL 50 /* 5 seconds */
49 #define DEFAULT_FAIL 150 /* 15 seconds */
50 #define DEFAULT_RETRY 20 /* 2 seconds */
51
52 struct acl {
53     struct acl *next;
54     char *name;
55     struct ace *head;
56     struct ace *tail;
57 };
58
59 struct peer {
60     struct peer *next;
61     int cloned; /* If 1, object was cloned */
```

```
62 int state;
63 int fd;
64
65 struct sockaddr_in lsa; /* Local & remote peer sockaddrs */
66 struct sockaddr_in rsa;
67
68 IPADDR routerid; /* Router ID of peer */
69 IPADDR nexthop; /* IP address for next hop advertisement */
70 int ttl; /* TTL of packets */
71 char *secret; /* Shared secret */
72 int confcost; /* Configured link cost */
73 TIMETEN confpoll; /* Configured poll interval */
74 TIMETEN conffail; /* Configured fail timeout */
75 TIMETEN confretry; /* Configured retry timeout */
76 int cost; /* Link cost (maximised with peer) */
77 TIMETEN poll; /* Poll interval (minimised with peer) */
78 TIMETEN fail; /* Timeout to fail (minimised with peer) */
79 TIMETEN retry; /* Timeout to retry packet */
80 int ifindex; /* Interface index */
81 IPADDR localroute; /* Network base address of containing subnet */
82 int localbits; /* Prefix length of containing network */
83 int localannounce; /* Announce local net, 0 = if winner, */
84 /* 1 = always, 2 = never */
85
86 struct acl *remote; /* Permit connections from (if listen) */
87 struct acl *announce; /* Announce routes to peer */
88 struct acl *accept; /* Accept routes from peer */
89
90 u_int32_t lseq; /* Send sequence number */
91 u_int32_t aseq; /* Last acknowledged sequence number */
92 u_int32_t rseq; /* Expected sequence number */
93 u_int32_t xseq; /* Re-sync candidate sequence number */
94 int pathlen; /* Length of path to gateway via this link */
95 IPADDR path[MAX_PATH]; /* Path to gateway via this link */
96 int gwcost; /* Cost to gateway via this link */
97
98 struct iproute *routes; /* Inbound peer routing table */
99 struct iproute *newrts; /* Assemble new routing table here */
100 struct iproute *annrts; /* Routing announcement */
101
102 u_char *lastpacket; /* Packet in flight */
103 int lastpktlen; /* Length of packet in flight */
104 u_char *nextpacket; /* Packet buffer */
105 int nextpktlen; /* Pointer into packet buffer */
106 int ackreq; /* Acknowledgement required by peer */
107
108 TIMETEN lasttime; /* Time last message received */
109 TIMETEN nexttime; /* Time to transmit next retry */
110 int quickstart; /* Quick retries on startup */
111
112 int synreq; /* Need to synchronise */
113 int configreq; /* Need to send config */
114 int pathreq; /* Need to send path */
115 int pollreq; /* Need to send poll request */
116 int respreq; /* Need to send poll response */
117 int sendack; /* Inbound message requires ACK */
118 int shutdown; /* Need to shut down peer */
119 struct iproute *nexttrt; /* Next route to add to outbound packet */
120 int nullrt; /* Send null update */
121 };
122 #define LOCIP(peer) ((peer)->lsa.sin_addr.s_addr)
```

```
123 #define REMIP(peer) ((peer)->rsa.sin_addr.s_addr)
124
125
126 struct iproute {
127     struct iproute *next;
128     IPADDR ip; /* Base IP address of route */
129     int bits; /* Mask length */
130     int isgw; /* 1 if reverse path is via default gateways */
131     int cost; /* Cost to originator via this link */
132     int gwcost; /* Cost from originator to gateway */
133     int rtype; /* Route type */
134     struct peer *peer; /* Link associated with route */
135     int ifindex; /* Interface index of local route */
136     int inuse; /* Flag used to detect unused entries */
137 };
138
139
140 struct acl *acls; /* List header of named ACLs */
141
142 IPADDR maskbits[33]; /* Map prefix length to mask */
143
144 struct iproute *freelist; /* Free list of route objects */
145 struct iproute *localroutes; /* Local routing table header */
146
147 struct peer *peers; /* Running peers */
148 struct peer *listens; /* "Meta" peers */
149 struct peer *gwpeer; /* Gateway peer (if not a gateway) */
150 int gwcost; /* 0=gw, -1=no gw, >0 = gw cost */
151
152 int checkroutes; /* Routing update required */
153 int rtsocket; /* Routing socket fd */
154 int udpport; /* Default port number */
155 int isgateway; /* Set if router is a gateway */
156 int gwalways; /* Advertise gateway route if no default */
157 TIMETEN now; /* Current time */
158 int routeflag; /* Kernel route flag (RTF_PROTO<n>) */
159 int otherflag; /* Additional route flags */
160 int overrflag; /* Flags to override */
161 IPADDR defaultroute; /* Default route to apply if none available */
162 IPADDR defgateway; /* Current default gateway */
163
164 IPADDR routerid; /* Router ID for paths */
165
166 char *statusfile; /* File to dump status information to */
167 char *pidfile; /* File to write PID into */
168 pid_t pid; /* Our PID */
169
170 int debug; /* Debug level: 0 = no messages */
171 #define LOG0 msg /* unconditional */
172 #define LOG1 if(debug) msg /* Topology changes, major events */
173 #define LOG2 if(debug >= 2) msg /* Routing changes */
174 #define LOG3 if(debug >= 3) msg /* Exceptions */
175 #define LOG4 if(debug >= 4) dbg /* Routing announcements */
176 #define LOG5 if(debug >= 5) dbg /* Kernel routes added/deleted */
177 #define LOG6 if(debug >= 6) dbg /* Message processing */
178 #define LOG7 if(debug >= 7) dbg /* Packets sent/received */
179 #define LOG8 if(debug >= 8) dbg /* Internal decisions */
180
181
182
183
```

```
184 #define FOREACH(item, head) for((item) = (head); (item); (item) = (item)->next)
185 #define FREEROUTE(ipr) { (ipr)->next = freelist; freelist = (ipr); }
186 #define NEWROUTE(ipr) { if(((ipr) = freelist)) freelist = (ipr)->next; \
187     else (ipr) = malloc(sizeof(struct iproute)); }
188 #define ADDROUTEAFTER(ipr, table, after) { NEWROUTE(ipr) \
189     if(after) { (ipr)->next = (after)->next; \
190         (after)->next = (ipr); } \
191     else { (ipr)->next = (table); \
192         (table) = (ipr); } }
193 #define KILLROUTES(table) { struct iproute *ipr, *nipr; \
194     for(ipr= (table); ipr; ipr = nipr) { \
195         nipr = ipr->next; FREEROUTE(ipr); } \
196     (table) = 0; }
197
198 #define NEXTSEQ(r) ((r) + 1) == 0 ? 1 : (r) + 1
199
200 char *formatip(IPADDR ip);
201 char *formattime(TIMETEN t);
202
203 struct iproute *findroute(struct iproute *ipr, IPADDR ip, int bits,
204     struct iproute **prev);
205 struct acl *findacl(char *name, int create);
206 int checkacl(struct acl *acl, IPADDR ip, int bits, int rtype, int ifindex);
207 void dumproutes(struct iproute *ipr);
208
209 int getroutes();
210
211 int parse_config(char *file);
212 void parse_acl_std(void);
213 char *parseip(char *s, IPADDR *ip_p, int *bits_p, int *maxbits_p);
214 char *getlocaladdr(struct peer *peer, IPADDR peerip, int listen);
215
216 void msg(char *fmt, ...);
217 void dbg(char *fmt, ...);
218
219 /* EOF */
220
```

```
1 #include "frpd.h"
2 #include "frp.h"
3
4 /*
5  Globals
6  */
7 struct acl *acls = 0; /* List header of named ACLs */
8
9 IPADDR maskbits[33]; /* Map prefix length to mask */
10
11 struct iproute *freelist = 0; /* Free list of route objects */
12 struct iproute *localroutes = 0; /* Local routing table header */
13
14 struct peer *peers = 0; /* Running peers */
15 struct peer *listens = 0; /* "Meta" peers */
16
17 int gwcost = -1; /* 0 = we are a gateway; */
18 /* -1 = no gateway, 1-65534 = cost */
19 int rtsocket = -1; /* Routing socket */
20
21 int udpport = 343; /* Default port number */
22 int isgateway = 0; /* Is a gateway */
23 int gwalways = 0; /* Always a gateway */
24 IPADDR routerid = 0; /* Router ID for paths */
25 struct peer *gwpeer = 0; /* Gateway peer */
26 char *statusfile = 0; /* Dump status info to here */
27 int checkroutes = 0; /* Routing update required */
28 TIMETEN now; /* Current time, 100 ms increments */
29 int debug = 1; /* Debug messages */
30 char *pidfile = 0; /* Write PID here */
31
32 void
33 msg(char *fmt, ...) {
34     va_list ap;
35     va_start(ap, fmt);
36     vsyslog(LOG_INFO, fmt, ap);
37     va_end(ap);
38 }
39 void
40 dbg(char *fmt, ...) {
41     va_list ap;
42     va_start(ap, fmt);
43     vsyslog(LOG_DEBUG, fmt, ap);
44     va_end(ap);
45 }
46
47 /*
48 Find a route in a routing table, sorted by ip, bits
49 Update prev (if provided) to be previous routing entry, useful for adding
50 a route if not found.
51 */
52
53 struct iproute *
54 findroute(struct iproute *ipr, IPADDR ip, int bits, struct iproute **prev) {
55     if(prev) *prev = 0;
56     IPADDR ri;
57
58     ip = ntohl(ip);
59     FOREACH(ipr, ipr) {
60         ri = ntohl(ipr->ip);
61         if(ri > ip)
```

```
62     return 0;
63     if(ip == ri) {
64         if(bits == ipr->bits)
65             return ipr;
66         if(bits > ipr->bits)
67             return 0;
68     }
69     if(prev) *prev = ipr;
70 }
71 return 0;
72 }
73
74
75 static void
76 dumproute(FILE *file, struct iproute *ipr) {
77     char ib[IFNAMSIZ];
78     fprintf(file, "%15s/%-2d %1s%1s %-15s %-8s rc %-5d gc %d\n",
79             formatip(ipr->ip), ipr->bits,
80             (ipr->rtype & RTYPE_LAYER2) ? "2" :
81             (ipr->rtype & RTYPE_STATIC) ? "S" :
82             (ipr->rtype & RTYPE_PROTOCOL) ? "P" :
83             (ipr->rtype & RTYPE_LOCAL) ? "L" : "R",
84             ipr->isgw ? "G" : "",
85             ipr->peer ? formatip(REMIP(ipr->peer)) : "local",
86             ipr->ifindex ? if_indextoname(ipr->ifindex, ib) : "-",
87             ipr->cost, ipr->gwcost);
88 }
89 void
90 dumproutes(struct iproute *ipr) {
91     char ib[IFNAMSIZ];
92     FOREACH(ipr, ipr)
93         dbg("%s/%d %s%s -> %s iface %s rc %d gc %d",
94             formatip(ipr->ip), ipr->bits,
95             (ipr->rtype & RTYPE_LAYER2) ? "2" :
96             (ipr->rtype & RTYPE_STATIC) ? "S" :
97             (ipr->rtype & RTYPE_PROTOCOL) ? "P" :
98             (ipr->rtype & RTYPE_LOCAL) ? "L" : "R",
99             ipr->isgw ? "G" : "",
100             ipr->peer ? formatip(REMIP(ipr->peer)) : "local",
101             ipr->ifindex ? if_indextoname(ipr->ifindex, ib) : "-",
102             ipr->cost, ipr->gwcost);
103 }
104
105
106 /*
107 Dump internal state to a dump file
108 */
109 static void
110 braindump(void) {
111     FILE *file;
112     struct peer *peer;
113     struct iproute *ipr;
114     int i;
115     static char *statustemp = 0;
116
117     /*
118     Open status file. If it didn't work say so and don't try again
119     */
120     if(!statusfile) return;
121     if(!statustemp) {
122         statustemp = malloc(strlen(statusfile) + 5);
```

```
123     sprintf(statusfile, "%s.tmp", statusfile);
124 }
125 file = fopen(statusfile, "w");
126 if(!file) {
127     LOG1("Could not write %s: %m", statusfile);
128     statusfile = 0;
129     return;
130 }
131
132 /*
133 Local status
134 */
135 fprintf(file, "RouterID %s Gateway %s GWCost %d%s\nPath",
136         formatip(routerid),
137         defgateway ? formatip(defgateway) : "unset",
138         gwcost, isgateway ? " gateway" : "");
139 if(isgateway)
140     fprintf(file, " local\n");
141 else if(gwcost != -1 && gwpeer) {
142     for(i = 0; i < gwpeer->pathlen; i++)
143         fprintf(file, " %s", formatip(gwpeer->path[i]));
144     putc('\n', file);
145 }
146 else fprintf(file, " none\n");
147
148 /*
149 Listens
150 */
151 FOREACH(peer, listens)
152     fprintf(file, "Listening on %s\n", formatip(LOCIP(peer)));
153
154 /*
155 Main routing table
156 */
157 fprintf(file, "\nRouting table\n");
158 FOREACH(ipr, localroutes)
159     dumproute(file, ipr);
160
161 /*
162 Peers, including received and announced routes
163 */
164 FOREACH(peer, peers) {
165     fprintf(file, "\nPeer %s Source %s NextHop %s\n",
166             formatip(REMIP(peer)), formatip(LOCIP(peer)),
167             formatip(peer->nexthop));
168     fprintf(file, "Cost %d GWCost %d Poll %s Retry %s Fail %s"
169             " Type %s Status %s\n",
170             peer->cost, peer->gwcost,
171             formattime(peer->poll),
172             formattime(peer->retry),
173             formattime(peer->fail),
174             (peer->cloned) ? "dynamic" : "static",
175             (peer == gwpeer) ? "up-gw" :
176             (peer->synreq) ? "down" : "up");
177     fprintf(file, "Path");
178     if(peer->gwcost == -1)
179         fprintf(file, " none\n");
180     else {
181         for(i = 0; i < peer->pathlen; i++)
182             fprintf(file, " %s", formatip(peer->path[i]));
183         putc('\n', file);
184     }
```

```
184 }
185 fprintf(file, "Received routes\n");
186 FOREACH(ipr, peer->routes)
187     dumproute(file, ipr);
188 fprintf(file, "Announced routes\n");
189 FOREACH(ipr, peer->annrts)
190     dumproute(file, ipr);
191 }
192 fclose(file);
193
194 /*
195 Done
196 */
197 if(rename(statusfile, statusfile) == -1) {
198     LOG1("Could not write %s: %m", statusfile);
199     statusfile = 0;
200     unlink(statusfile);
201 }
202 }
203
204 /*
205 Dump a packet
206 */
207
208 #define DFP_S(pkt,len,peer) if(debug >= 6) dp(1, pkt, len, peer)
209 #define DFP_R(pkt,len,peer) if(debug >= 6) dp(0, pkt, len, peer)
210 static void
211 dp(int outbound, void *p, int len, struct peer *peer) {
212     u_char *ptr = p;
213     u_char *end;
214     struct frp_hdr *hdr;
215     struct frp_mhdr *mh;
216     struct frp_ctrl *cml;
217     struct frp_config *cmc;
218     struct frp_path *cmp;
219     struct frp_route *cmr;
220     int i;
221     char *t;
222     char buf[16 * MAX_PATH];
223
224     if(len == -1) {
225         dbg("From: %s error: %m", formatip(REMIP(peer)));
226         return;
227     }
228
229     end = ptr + len;
230     hdr = (struct frp_hdr *) ptr;
231     ptr += sizeof(struct frp_hdr);
232
233     if(outbound)
234         dbg("S %s -> %s lseq=%08x rseq=%08x len=%d",
235             formatip(LOCIP(peer)), formatip(REMIP(peer)),
236             ntohl(hdr->lseq), ntohl(hdr->rseq), len);
237     else dbg("R %s <- %s rseq=%08x lseq=%08x",
238             formatip(LOCIP(peer)), formatip(REMIP(peer)),
239             ntohl(hdr->rseq), ntohl(hdr->lseq), len);
240
241     while(ptr < end) {
242         mh = (struct frp_mhdr *) ptr;
243         ptr += mh->len * sizeof(u_int32_t);
244         if(!mh->len) return;
```

```
245     if(ptr > end) return;
246
247     switch(mh->type) {
248     case FRP_CTRL:
249         cml = (struct frp_ctrl *) mh;
250         switch(cml->ctrl) {
251             case FRP_CTRL_POLL: t = "poll"; break;
252             case FRP_CTRL_ACK: t = "ack"; break;
253             case FRP_CTRL_NAK: t = "nak"; break;
254             default: t = "unknown"; break;
255         }
256         dbg("CTRL C=%d (%s) P=%x", cml->ctrl, t, cml->param);
257         break;
258     case FRP_CONFIG:
259         cmc = (struct frp_config *) mh;
260         dbg("CONFIG id=%s c=%d poll=%d fail=%d",
261             formatip(cmc->id), ntohs(cmc->cost),
262             ntohs(cmc->poll), ntohs(cmc->fail));
263         break;
264     case FRP_PATH:
265         cmp = (struct frp_path *) mh;
266         t = buf;
267         for(i = 0; i < mh->len - 1; i++) {
268             sprintf(t, "%s%s", i ? ":" : "",
269                 formatip(cmp->path[i]));
270             t = strchr(t,0);
271         }
272         dbg("PATH gd=%d p=%s", ntohs(cmp->gwcost), buf);
273         break;
274     case FRP_ROUTE:
275         cmr = (struct frp_route *) mh;
276         if(cmr->flags & FRP_FLAG_NULL)
277             dbg("ROUTE NULL flags=%x", cmr->flags);
278         else dbg("ROUTE %s/%u c=%d gc=%d flags=%x (%s%s%s)",
279             formatip(cmr->ip), cmr->bits,
280             ntohs(cmr->cost), ntohs(cmr->gwcost),
281             cmr->flags,
282             cmr->flags & FRP_FLAG_GATEWAY ? "G" : "",
283             cmr->flags & FRP_FLAG_BEGIN ? "B" : "",
284             cmr->flags & FRP_FLAG_COMMIT ? "C" : "");
285         break;
286     case FRP_CONFIG6:
287         dbg("CONFIG6 unsupported");
288         break;
289     case FRP_PATH6:
290         dbg("PATH6 unsupported");
291         break;
292     case FRP_ROUTE6:
293         dbg("ROUTE6 unsupported");
294         break;
295     default:
296         dbg("UNKNOWN T=%02x L=%d", mh->type, mh->len * 4);
297     }
298 }
299 }
300
301 /*
302 Process an ACL
303 */
304 int
```

```
306 checkacl(struct acl *acl, IPADDR ip, int bits, int rtype, int ifindex) {
307     struct ace *ace;
308
309     FOREACH(ace, acl->head) {
310         if(ace->bits > 0 && (ip & maskbits[ace->bits]) != ace->ip)
311             continue;
312         if(bits != -1 && (bits < ace->bits || bits > ace->maxbits))
313             continue;
314         if(rtype && ace->rtype && !(rtype & ace->rtype))
315             continue;
316         if(ifindex && ace->ifindex && ace->ifindex != ifindex)
317             continue;
318         if(ace->acl && !checkacl(ace->acl, ip, bits, rtype, ifindex))
319             continue;
320         return ace->permit;
321     }
322     return 0;
323 }
324
325 /*
326 Security stuff
327 Hash calculation: SHA1 over packet, IP addresses & secret
328 Note that only the first 64 bits (FRP_HASHZIE) of the hash are used.
329 */
330 #include <sha.h>
331 static u_int8_t *
332 dohash(u_char *buf, int len, char *secret, IPADDR sa, IPADDR da) {
333     static u_int8_t hash[SHA_DIGEST_LENGTH];
334     SHA_CTX shctx;
335
336     SHA_Init(&shctx);
337     SHA_Update(&shctx, buf, len);
338     SHA_Update(&shctx, (u_char *)&sa, sizeof(IPADDR));
339     SHA_Update(&shctx, (u_char *)&da, sizeof(IPADDR));
340     SHA_Update(&shctx, secret, strlen(secret));
341     SHA_Final(hash, &shctx);
342     return hash;
343 }
344
345 /*
346 Check that hash matches packet hash.
347 Do not call if peer->secret is null.
348 */
349 static int
350 checksecure(u_char *pkt, int len, struct peer *peer) {
351     u_int8_t *hash;
352
353     hash = dohash(pkt + FRP_HASHSIZE, len - FRP_HASHSIZE, peer->secret,
354         REMIP(peer), LOCIP(peer));
355     if(!memcmp(hash, pkt, FRP_HASHSIZE))
356         return 1;
357     LOG2("packet security failure");
358     return 0;
359 }
360
361 /*
362 Compute packet hash
363 Or zero hash field if secret is null
364 */
365 static void
```



```
367 secure(struct peer *peer, struct frp_hdr *pkt, int len) {
368     u_int8_t *hash;
369
370     if(!peer->secret) {
371         memset(pkt->hash, 0, FRP_HASHSIZE);
372         return;
373     }
374     hash = dohash(&pkt->hash[FRP_HASHSIZE], len-FRP_HASHSIZE, peer->secret,
375                 LOCIP(peer), REMIP(peer));
376     memcpy(pkt->hash, hash, FRP_HASHSIZE);
377 }
378
379 /*
380  * Return a random sequence number (that isn't 0)
381  */
382 static u_int32_t
383 initseq(void) {
384     static u_int32_t r = 0xbabefee1;
385     int fd;
386     struct timeval tv;
387
388     do {
389         fd = open("/dev/urandom", O_RDONLY);
390         if(fd != -1) {
391             read(fd, &r, sizeof(r));
392             close(fd);
393         }
394         gettimeofday(&tv, 0);
395         r = r ^ routerid ^ tv.tv_sec ^ tv.tv_usec ^ (getpid() << 16)
396             ^ getpid();
397     }
398     while(r == 0);
399     return r;
400 }
401
402 /*
403  * Send a SYN packet
404  */
405 static void
406 sendsyn(struct peer *peer) {
407     struct frp_hdr syn;
408
409     syn.lseq = htonl(peer->lseq);
410     syn.rseq = 0;
411     secure(peer, &syn, sizeof(syn));
412     DFP_S(&syn, sizeof(syn), peer);
413     write(peer->fd, &syn, sizeof(syn));
414     return;
415 }
416
417 /*
418  * Send a NAK
419  */
420 static void
421 sendnak(struct peer *peer, int fd, u_int32_t lseq, u_int32_t rseq) {
422     struct {
423         struct frp_hdr hdr;
424         struct frp_ctrl ctrl;
425     };
```

```
428     } nak;
429     static TIMETEN nextnak = 0;
430
431     /*
432      * Rate limit NAKs, stop NAK wars.
433      */
434     if(now < nextnak)
435         return;
436     nextnak = now + 10;
437
438     /*
439      * Format & send a NAK
440      */
441     nak.hdr.lseq = htonl(htonl(rseq) + 1);
442     if(nak.hdr.lseq == 0)
443         nak.hdr.lseq = htonl(1);
444     nak.hdr.rseq = lseq;
445     nak.ctrl.mh.type = FRP_CTRL;
446     nak.ctrl.mh.len = 1;
447     nak.ctrl.ctrl = FRP_CTRL_NAK;
448     nak.ctrl.param = 0;
449     secure(peer, &nak.hdr, sizeof(nak));
450     DFP_S(&nak, sizeof(nak), peer);
451     write(fd, &nak, sizeof(nak));
452 }
453
454 /*
455  * Turn everything off on a peer
456  * If "full" is 1, completely shut down the peer. If zero, just shut the
457  * session down, but don't reset routing state.
458  * Note that clearing out the tables of routes received and announced is
459  * done in the mainline (as it's common to deletion of both static and
460  * dynamic peers).
461  */
462 static void
463 resetpeer(struct peer *peer, int full) {
464     /*
465      * Reset stuff
466      */
467     peer->lastpktlen = 0;
468     peer->nextpktlen = 0;
469     peer->lseq = initseq();
470     peer->aseq = 0;
471     peer->rseq = 0;
472     peer->nexttrt = 0;
473     peer->>nullrt = 0;
474     peer->ackreq = 0;
475     peer->synreq = 1;
476     peer->configreq = 0;
477     peer->pathreq = 0;
478     peer->pollreq = 0;
479     peer->respreq = 0;
480     peer->sendack = 0;
481     peer->nexttime = 0;
482     peer->poll = peer->confpoll;
483     peer->fail = peer->conffail;
484     peer->cost = peer->confcost;
485     peer->retry = peer->confretry;
486     peer->quickstart = 0;
487
488 }
```

```
489  /*
490  If peer is being shut down, kill off the path and gw info, and
491  reset all the timers.
492  */
493  if(full) {
494      peer->lasttime = 0;
495      peer->shutdown = 0;
496      peer->routerid = 0;
497      peer->gwcost = -1;
498      peer->pathlen = 0;
499      checkroutes = 1;
500  }
501 }
502
503
504 static void *
505 addmessage(struct peer *peer, int size, int type) {
506     u_char *ptr;
507     struct frp_mhdr *hdr;
508
509     if(!peer->nextpacket) {
510         peer->nextpacket = malloc(MAX_PACKET);
511         peer->nextpktlen = 0;
512     }
513     if(!peer->nextpktlen)
514         peer->nextpktlen = sizeof(struct frp_hdr);
515     if(!size || peer->nextpktlen + size > MAX_PACKET)
516         return 0;
517     ptr = peer->nextpacket + peer->nextpktlen;
518     peer->nextpktlen += size;
519     peer->ackreq = 1;
520     hdr = (struct frp_mhdr *) ptr;
521     hdr->len = size / sizeof(u_int32_t);
522     hdr->type = type;
523     return ptr;
524 }
525
526
527
528 #define DROP(msg) { LOG3("dropping packet: %s", msg); return; }
529
530 static void
531 process_packet(u_char *ptr, int len, struct peer *peer) {
532     u_char *end;
533     struct frp_hdr *hdr;
534     struct frp_mhdr *mh;
535     struct frp_ctrl *cml;
536     struct frp_config *cmc;
537     struct frp_path *cmp;
538     struct frp_route *cmr;
539     int c;
540     int i;
541     struct frp_hdr ack;
542     u_int32_t lseq, rseq, xseq;
543     struct iproute *ipr;
544
545     end = ptr + len;
546     hdr = (struct frp_hdr *) ptr;
547     lseq = ntohl(hdr->lseq);
548     rseq = ntohl(hdr->rseq);
549     ptr += sizeof(struct frp_hdr);
```

```
550
551 LOG8("from=%s:%d to=%s:%d syn=%d conf=%d path=%d"
552      " lseq=%x aseq=%x rseq=%x xseq=%x\n",
553      formatip(REMIP(peer)), ntohs(peer->rsa.sin_port),
554      formatip(LOCIP(peer)), ntohs(peer->lsa.sin_port),
555      peer->synreq, peer->configreq, peer->pathreq,
556      peer->lseq, peer->aseq, peer->rseq, peer->xseq);
557
558 if(!hdr->lseq)
559     DROP("zero in lseq")
560
561 /*
562 If this is a synchronisation packet, note the remote sequence
563 number in xseq. Drop any such packets with payload, as these are
564 not to be trusted.
565 If it's a kosher-looking sync, send an ACK and return. Note that
566 state won't get updated until we get a data packet.
567 */
568 if(!hdr->rseq) {
569     if(len != sizeof(struct frp_hdr))
570         DROP("null rseq with payload")
571     peer->xseq = NEXTSEQ(lseq);
572     ack.lseq = htonl(peer->lseq);
573     ack.rseq = htonl(lseq);
574     secure(peer, &ack, sizeof(ack));
575     DFP_S(&ack, sizeof(ack), peer);
576     write(peer->fd, &ack, sizeof(ack));
577     return;
578 }
579
580 /*
581 If this isn't a SYN packet (rseq != 0), check that both the
582 local and remote sequence numbers are in the expected ranges
583 */
584 if(rseq != peer->lseq &&
585    rseq != peer->aseq) {
586     sendnak(peer, peer->fd, hdr->lseq, hdr->rseq);
587     DROP("unexpected packet rseq")
588 }
589 xseq = peer->rseq;
590 if(xseq) {
591     for(i = 2; i > 0; i--) {
592         if(lseq == xseq)
593             break;
594         xseq = NEXTSEQ(xseq);
595     }
596     if(!i && (!peer->xseq || peer->xseq != lseq))
597         DROP("unexpected packet lseq")
598 }
599
600 /*
601 If the packet acknowledges one of ours, update the acknowledged
602 sequence.
603 If it was in response to one of our syn requests, we need to send
604 path and config details.
605 */
606 if(rseq == peer->lseq && peer->aseq != peer->lseq) {
607     LOG8("ack received");
608     peer->aseq = peer->lseq;
609     peer->lasttime = now;
610     peer->lastpktlen = 0;
```

```
611     if(peer->synreq) {
612         LOG6("%s: SYN acknowledged", formatip(REMIP(peer)));
613         peer->synreq = 0;
614         peer->configreq = 1;
615         peer->pathreq = 1;
616         peer->nexttrt = 0;
617         peer->>nullrt = 0;
618     }
619 }
620
621 /*
622  If the remote sequence number is different from last time,
623  force a response.
624  If this is the result of a new connection (xseq != 0), send config
625  and path info
626  */
627 xseq = peer->xseq;
628 peer->xseq = 0;
629 if(lseq != peer->rseq) {
630     peer->sendack = 1;
631     peer->rseq = lseq;
632     LOG8("ACK flag set");
633     if(xseq) {
634         peer->synreq = 0;
635         peer->configreq = 1;
636         peer->pathreq = 1;
637         peer->pollreq = 0;
638         peer->nexttrt = 0;
639         peer->>nullrt = 0;
640     }
641 }
642
643 /*
644  If it isn't, then we have a duplicate or a simple ack.
645  If there is data, re-ack the packet; if there isn't, just drop it.
646  */
647 else {
648     if(len > sizeof(struct frp_hdr))
649         peer->sendack = 1;
650     return;
651 }
652
653 while(ptr < end) {
654     mh = (struct frp_mhdr *) ptr;
655
656     ptr += mh->len * sizeof(u_int32_t);
657     if(lmh->len) DROP("bad mh->len") /* Paranoia */
658     if(ptr > end) DROP("mh->len past end") /* More paranoia */
659
660     switch(mh->type) {
661     case FRP_CTRL:
662         cml = (struct frp_ctrl *) mh;
663         if(cml->ctrl == FRP_CTRL_POLL)
664             peer->respreq = 1;
665         else if(cml->ctrl == FRP_CTRL_NAK) {
666             LOG3("%s: NAK received for peer",
667                 formatip(REMIP(peer)));
668             resetpeer(peer, 0);
669             break;
670         }
671         break;
```

```
672
673     /*
674     Configuration message -- agree on configuration
675     */
676     case FRP_CONFIG:
677         cmc = (struct frp_config *) mh;
678         LOG6("%s: router ID = %s cost = %d",
679             formatip(REMIP(peer)),
680             formatip(cmc->id), ntohs(cmc->cost));
681
682     /*
683     If the router-ID has changed, signal a routing re-run
684     But if it's ours, drop this connection on the floor
685     */
686     if(peer->routerid != cmc->id) {
687         if(cmc->id == routerid) {
688             peer->shutdown = 1;
689             DROP("duplicate router ID")
690         }
691         peer->routerid = cmc->id;
692         checkroutes = 1;
693     }
694
695     /*
696     Maximise the cost. If it changes, bump the routing
697     */
698     c = ntohs(cmc->cost);
699     if(c < peer->confcost)
700         c = peer->confcost;
701     if(c != peer->cost) {
702         peer->cost = c;
703         checkroutes = 1;
704     }
705
706     /*
707     Minimise the fail time. Minimum of one second to
708     allow for poll and retry
709     */
710     c = ntohs(cmc->fail);
711     if(c >= 10 && c < peer->conffail)
712         peer->fail = c;
713     else peer->fail = peer->conffail;
714
715     /*
716     Now the poll time
717     */
718     c = ntohs(cmc->poll);
719     if(c >= 1 && c < peer->confpoll)
720         peer->poll = c;
721     else peer->poll = peer->confpoll;
722
723     /*
724     Compute the retry time to require three retries
725     before giving up (min 1 second). Can have more by
726     specifying a smaller retry time.
727     */
728     c = (peer->fail - peer->poll - 1) / 3;
729     if(c < 1) c = 1;
730     if(peer->confretry > c)
731         peer->retry = c;
732     else peer->retry = peer->confretry;
```

```
733
734 /*
735 If we are the smaller routerid, bump the poll
736 time by one to avoid duelling polls
737 */
738 if(ntohl(routerid) < ntohl(peer->routerid))
739     peer->poll++;
740
741 LOG1("%s: Configured: id=%s cost=%d poll=%d"
742       " retry=%d fail=%d",
743       formatip(REMIP(peer)),
744       formatip(peer->routerid), peer->cost,
745       peer->poll, peer->retry, peer->fail);
746 break;
747
748 /*
749 Path message
750 Copy the path and length
751 */
752 case FRP_PATH:
753     cmp = (struct frp_path *) mh;
754     c = mh->len - 1;
755     if(c >= MAX_PATH) {
756         peer->shutdown = 1;
757         break;
758     }
759     peer->pathlen = c;
760     for(i = 0; i < c; i++)
761         peer->path[i] = cmp->path[i];
762
763     c = ntohs(cmp->gwcost);
764     if(c == FRP_NOGATEWAY)
765         c = -1;
766     if(peer->gwcost != c)
767         peer->gwcost = c;
768     checkroutes = 1;
769     if(peer->annrts) {
770         peer->nullrt = 0;
771         peer->nextrt = peer->annrts;
772     }
773     else peer->nullrt = 1;
774     break;
775
776 /*
777 Route message
778 */
779 case FRP_ROUTE:
780     cmr = (struct frp_route *) mh;
781     /*
782     BEGIN indicates that we should scrub anything we
783     have already
784     */
785     if(cmr->flags & FRP_FLAG_BEGIN)
786         KILLROUTES(peer->newrts)
787
788     /*
789     Add the route, unless it's a null route
790     */
791     if(cmr->flags & FRP_FLAG_NULL) {
792         LOG6("%s: Null route",formatip(REMIP(peer)));
793     }
```

```
794     else {
795         NEWROUTE(ipr)
796         ipr->ip = cmr->ip;
797         ipr->bits = cmr->bits;
798         ipr->isgw = !(cmr->flags & FRP_FLAG_GATEWAY);
799         ipr->cost = ntohs(cmr->cost);
800         ipr->gwcost = ntohs(cmr->gwcost);
801         ipr->peer = peer;
802         ipr->ifindex = 0;
803         ipr->rtype = 0;
804         ipr->next = peer->newrts;
805         peer->newrts = ipr;
806         LOG6("%s: Route %s/%d flags=%x gw=%d c=%d"
807               " gc=%d",
808               formatip(REMIP(peer)),
809               formatip(cmr->ip), ipr->bits,
810               ipr->bits, ipr->isgw, ipr->cost,
811               ipr->gwcost);
812     }
813
814     /*
815     COMMIT says the update is done, so delete the old
816     current routes list, and move the new list to the
817     current routes.
818     Signal the change.
819     */
820     if(cmr->flags & FRP_FLAG_COMMIT) {
821         KILLROUTES(peer->routes)
822         peer->routes = peer->newrts;
823         peer->newrts = 0;
824         checkroutes = 1;
825     }
826     break;
827
828     /*
829     IPv6 messages are not supported
830     */
831 case FRP_CONFIG6:
832 case FRP_PATH6:
833 case FRP_ROUTE6:
834     DROP("IPv6 not supported")
835
836     /*
837     Nor are unknown messages ...
838     */
839 default:
840     DROP("bad msg type")
841 }
842 }
843 }
844
845
846
847 /*
848 Connect or listen to a peer/listen
849 Return fd on success, -1 on fail
850 */
851 static int
852 connectpeer(struct peer *peer, int doconnect) {
853     int fd;
854     int i;
```

```
855 fd = socket(PF_INET, SOCK_DGRAM, 0);
856 if(fd < 0)
857     goto oops;
858 i = 1;
859 if(setsockopt(fd, SOL_SOCKET, SO_REUSEPORT, &i, sizeof(int)))
860     goto oops;
861 if(setsockopt(fd, IPPROTO_IP, IP_TTL, &peer->ttl, sizeof(int)))
862     goto oops;
863 LOG8("bind %s:%d", formatip(LOCIP(peer)), ntohs(peer->lsa.sin_port));
864 if(bind(fd, (struct sockaddr *)&peer->lsa, sizeof(struct sockaddr_in)))
865     goto oops;
866 if(doconnect) {
867     LOG8("connect %s:%d", formatip(REMIP(peer)),
868         ntohs(peer->rsa.sin_port));
869     if(connect(fd, (struct sockaddr *)&peer->rsa,
870         sizeof(struct sockaddr_in)))
871         goto oops;
872 }
873 return fd;
874
875 oops: LOG1("Could not %s to %s: %m", doconnect ? "connect to" : "listen on",
876     formatip(doconnect ? REMIP(peer) : LOCIP(peer)));
877 if(fd != -1)
878     close(fd);
879 return -1;
880 }
881
882
883
884
885 static int restart = 0;
886 static int shutdown = 0;
887 static int exitnow = 0;
888
889 static void
890 ouch(int sig) {
891     switch(sig) {
892     case SIGUSR1: if(debug < 8)
893         LOG0("Debug level raised to %d", ++debug);
894         break;
895     case SIGUSR2: if(debug > 0)
896         LOG1("Debug level lowered to %d", --debug);
897         break;
898     case SIGHUP: exitnow = restart = 1; /* Restart program */
899         LOG1("SIGHUP received, restarting ...");
900         break;
901     case SIGTERM: exitnow = shutdown = 1;
902         LOG1("SIGTERM received, full shutdown");
903         break;
904     case SIGINT: exitnow = 1;
905         LOG1("SIGINT received, shutting down");
906         break;
907     }
908 }
909
910 int
911 main(int argc, char **argv) {
912     int i, j;
913     socklen_t sal;
914     fd_set fds;
915     int fdc;
```

```
916 struct timeval tv;
917 u_char buf[65536];
918 struct rt_msghdr *rtm;
919 struct frp_hdr *hdr;
920 TIMETEN lasttime;
921 int len;
922 struct peer *peer, *lastpeer, *nextpeer;
923 struct peer *listen;
924 struct frp_ctrl *cml;
925 struct frp_config *cmc;
926 struct frp_path *cmp;
927 struct frp_route *cmr;
928 struct iproute *ipr, *pipr, *nipr;
929 u_int32_t rseq;
930 int pfd;
931 FILE *pf;
932 extern char *optarg;
933 char *configfile;
934 int foreground;
935 int debuglvl;
936 int forceexit;
937 time_t basetime;
938
939 /*
940  * Parse arguments
941  */
942 configfile = "frpd.conf";
943 foreground = 0;
944 debuglvl = -1;
945 forceexit = 0;
946 while((i = getopt(argc, argv, "c:d:fp:xX")) != -1) switch(i) {
947     case 'c': configfile = optarg; break;
948     case 'd': debuglvl = atoi(optarg); break;
949     case 'f': foreground = 1; break;
950     case 'p': pidfile = optarg; break;
951     case 'x': forceexit = SIGINT; break;
952     case 'X': forceexit = SIGTERM; break;
953     default: return 1;
954 }
955
956 /*
957  * Compute mask bits array
958  */
959 for(i = 0; i <= 32; i++)
960     maskbits[i] = htonl(0xffffffff << (32 - i));
961 parse_acl_std();
962
963 /*
964  * If we were asked to kill the running process, do so now.
965  */
966 if(forceexit) {
967     if(!pidfile)
968         parse_config(configfile);
969     if(!pidfile) {
970         fprintf(stderr, "No PID file defined\n");
971         return 1;
972     }
973     pf = fopen(pidfile, "r");
974     if(!pf) {
975         fprintf(stderr, "Could not open PID file %s: %s\n",
976             pidfile, strerror(errno));
977     }
```

```
977     return 0;
978 }
979 i = -1;
980 fscanf(pf, "%d", &i);
981
982 if(i > 1 && flock(fileno(pf), LOCK_EX | LOCK_NB) == -1
983     && errno == EWOULDBLOCK) {
984     fclose(pf);
985     if(kill(i, forceexit) == 0)
986         return 0;
987     fprintf(stderr, "Could not kill PID %d: %s\n",
988         i, strerror(errno));
989 }
990 flock(fileno(pf), LOCK_UN);
991 fclose(pf);
992 fprintf(stderr, "Existing daemon not running\n");
993 return 1;
994 }
995
996 /*
997 Parse the configuration
998 */
999 if(!parse_config(configfile))
1000     return 1;
1001 if(!peers && !listeners) {
1002     fprintf(stderr, "No peers defined\n");
1003     return 1;
1004 }
1005 if(debuglvl >= 0 && debuglvl <= 8)
1006     debug = debuglvl;
1007
1008 /*
1009 Go into the background
1010 */
1011 if(!foreground)
1012     daemon(1, 1);
1013 pid = getpid();
1014
1015 /*
1016 Open the routing socket
1017 Just listen to IPv4 routing changes
1018 Don't bother with our own messages (like, we probably know)
1019 */
1020 rtsocket = socket(PF_ROUTE, SOCK_RAW, AF_INET);
1021 if(rtsocket == -1) {
1022     perror("routing socket");
1023     return 1;
1024 }
1025 i = 0;
1026 if(setsockopt(rtsocket, SOL_SOCKET, SO_USELOOPBACK, &i, sizeof(i)) < 0) {
1027     perror("setsockopt");
1028     return 1;
1029 }
1030
1031 /*
1032 Set up syslog
1033 */
1034 openlog("frpd", (foreground ? LOG_PERROR : 0), LOG_DAEMON);
1035
1036 /*
1037 Set up signal handler
```

```
1038 */
1039 signal(SIGHUP, ouch);
1040 signal(SIGINT, ouch);
1041 signal(SIGTERM, ouch);
1042 signal(SIGUSR1, ouch);
1043 signal(SIGUSR2, ouch);
1044
1045 /*
1046 Write the PID file
1047 Open it read/write
1048 If we could, see if it's locked.
1049 If it is, read the pid, and knock the existing process on the head
1050 If any of that failed, bail, otherwise try the whole procedure again.
1051 If we got the lock, just overwrite the PID with ours (keeping the lock)
1052 If the file was not there, lock it and write our PID.
1053 Note that when we exit, we always unlink the file before dropping
1054 the lock.
1055 */
1056 pf = 0;
1057 if(pidfile) {
1058     while((pf = fopen(pidfile, "r+")) {
1059         if(flock(fileno(pf), LOCK_EX | LOCK_NB) == -1
1060             && errno == EWOULDBLOCK) {
1061             i = -1;
1062             fscanf(pf, "%d", &i);
1063             if(i > 1) {
1064                 LOG1("Killing existing process %d", i);
1065                 if(kill(i, SIGINT) == 0) {
1066                     flock(fileno(pf), LOCK_EX);
1067                     flock(fileno(pf), LOCK_UN);
1068                     fclose(pf);
1069                     sleep(1);
1070                     pf = 0;
1071                     continue;
1072                 }
1073             }
1074             LOG0("Could not lock PID file");
1075             return 1;
1076         }
1077         rewind(pf);
1078         fprintf(pf, "%d\n", (int) pid);
1079         fflush(pf);
1080         break;
1081     }
1082     if(!pf) {
1083         pf = fopen(pidfile, "w");
1084         if(pf) {
1085             flock(fileno(pf), LOCK_EX);
1086             fprintf(pf, "%d\n", (int) pid);
1087             fflush(pf);
1088         }
1089         else LOG1("Could not write PID file %s: %m",
1090             pidfile);
1091     }
1092 }
1093
1094 /*
1095 Initialise time.
1096 */
1097 gettimeofday(&tv, 0);
1098 basetime = tv.tv_sec;
```

```
1099 now = tv.tv_usec / 100000;
1100
1101
1102 /*
1103  Initialise the peers & listens. If anything failed, drop them
1104  on the floor. And stamp on them.
1105  */
1106 lastpeer = nextpeer = 0;
1107 for(peer = peers; peer; peer = nextpeer) {
1108     nextpeer = peer->next;
1109     resetpeer(peer, 1);
1110     peer->fd = connectpeer(peer, 1);
1111     if(peer->fd == -1) {
1112         if(lastpeer)
1113             lastpeer->next = nextpeer;
1114         else peers = nextpeer;
1115         free(peer);
1116     }
1117     else lastpeer = peer;
1118 }
1119 for(peer = listens; peer; peer = nextpeer) {
1120     nextpeer = peer->next;
1121     peer->fd = connectpeer(peer, 0);
1122     if(peer->fd == -1) {
1123         if(lastpeer)
1124             lastpeer->next = nextpeer;
1125         else listens = nextpeer;
1126         free(peer);
1127     }
1128     else lastpeer = peer;
1129 }
1130 if(!peers && !listens) {
1131     LOG1("No active peers/listens");
1132     exitnow = 1;
1133 }
1134
1135 /*
1136  Set the router ID, if it hasn't been manually set
1137  */
1138 if(!routerid) {
1139     if(listens)
1140         routerid = LOCIP(listens);
1141     else if(peers)
1142         routerid = LOCIP(peers);
1143 }
1144 LOG1("Starting, routerid=%s, PID=%d", formatip(routerid), pid);
1145
1146 checkroutes = 1;
1147 lasttime = 0;
1148 while(!exitnow) {
1149     /*
1150      Rate-limit this
1151      */
1152     if(now != lasttime) {
1153         lasttime = now;
1154     }
1155     /*
1156      Time related protocol stuff with each peer:
1157      - Timeouts;
1158      - Startup polls;
1159      - Retransmission;
```

```
1160     - Regular keepalives.
1161     */
1162     FOREACH(peer, peers) {
1163         /*
1164          If the peer has timed out, shut it down
1165          */
1166         if(peer->lasttime &&
1167            now >= peer->lasttime + peer->fail) {
1168             peer->shutdown = 1;
1169             continue;
1170         }
1171
1172         /*
1173          Only do the protocol retries if nexttime
1174          has been reached
1175          */
1176         if(now < peer->nexttime)
1177             continue;
1178
1179         /*
1180          If we need to send a SYN packet, do so
1181          now.
1182          */
1183         if(peer->synreq) {
1184             sendsyn(peer);
1185             if(peer->quickstart < peer->poll)
1186                 peer->nexttime = now
1187                     + (+peer->quickstart);
1188             else peer->nexttime = now
1189                 + peer->poll;
1190         }
1191
1192         /*
1193          Otherwise, if we need to retransmit an
1194          un-acknowledged packet, do so now.
1195          Note that if rseq has advanced since our
1196          last retry, we need to update (and re-secure)
1197          the packet accordingly to acknowledge the
1198          received data.
1199          */
1200         else if(peer->lseq != peer->aseq
1201                && peer->lastpktlen) {
1202             DFP_S(peer->lastpacket, peer->lastpktlen,
1203                 peer);
1204             hdr = (struct frp_hdr *)
1205                 peer->lastpacket;
1206             rseq = ntohl(peer->rseq);
1207             if(hdr->rseq != rseq) {
1208                 hdr->rseq = rseq;
1209                 secure(peer, hdr,
1210                     peer->lastpktlen);
1211             }
1212             write(peer->fd, peer->lastpacket,
1213                 peer->lastpktlen);
1214             peer->nexttime = now + peer->retry;
1215         }
1216
1217         /*
1218          If we've reached time to poll, do so
1219          */
1220         else if(now >= peer->lasttime + peer->poll)
```

```
1221     peer->pollreq = 1;
1222 }
1223
1224 /*
1225  Check the routing table for changes. Only do this
1226  once per second to allow things to settle.
1227  Run this again if the gateway status changed and
1228  we are a gateway.
1229  If it failed prematurely, getroutes() returns
1230  1, else 0, which sets checkroutes appropriately.
1231  */
1232 if(checkroutes) {
1233     i = gwcost;
1234     checkroutes = getroutes();
1235     if(isgateway && i != gwcost)
1236         checkroutes = getroutes();
1237     if(statusfile)
1238         braindump();
1239 }
1240 }
1241
1242 /*
1243  Shut down requested peers
1244  Return to pristine state; if cloned, delete completely
1245  */
1246 lastpeer = nextpeer = 0;
1247 for(peer = peers; peer; peer = nextpeer) {
1248     nextpeer = peer->next;
1249     if(peer->shutdown) {
1250         LOG1("Peer %s shut down",
1251             formatip(REMIP(peer)));
1252     }
1253     /*
1254     Free route records
1255     */
1256     KILLROUTES(peer->routes);
1257     KILLROUTES(peer->newrts);
1258     KILLROUTES(peer->annrts);
1259     checkroutes = 2;
1260
1261     /*
1262     If it's been shut down, and it's a cloned
1263     peer, kill it off completely.
1264     Make sure we don't have any dangling pointers
1265     either in gwpeer or in the routing table
1266     */
1267     if(peer->cloned) {
1268         LOG8("Cloned peer deleted");
1269         close(peer->fd);
1270         if(peer == gwpeer)
1271             gwpeer = 0;
1272         pipr = nipr = 0;
1273         for(ipr=localroutes; ipr; ipr = nipr) {
1274             nipr = ipr->next;
1275             if(ipr->peer == peer) {
1276                 if(pipr)
1277                     pipr->next = nipr;
1278                 else localroutes = nipr;
1279                 FREEROUTE(ipr)
1280             }
1281             else pipr = ipr;
```

```
1282         if(lastpeer)
1283             lastpeer->next = nextpeer;
1284         else peers = nextpeer;
1285         if(peer->lastpacket)
1286             free(peer->lastpacket);
1287         if(peer->nextpacket)
1288             free(peer->nextpacket);
1289         free(peer);
1290         continue;
1291     }
1292
1293     /*
1294     Otherwise just clean things out
1295     */
1296     resetpeer(peer, 1);
1297 }
1298 lastpeer = peer;
1299 }
1300
1301 /*
1302  Process outbound packets
1303  */
1304 FOREACH(peer, peers) {
1305     /*
1306     Don't bother with any of this if we're still waiting
1307     for our last packet
1308     */
1309     if(peer->lseq != peer->aseq || !peer->rseq
1310        || peer->synreq)
1311         continue;
1312
1313     /*
1314     If we need to send a config packet, do so now;
1315     */
1316     if(peer->configreq) {
1317         cmc = addmessage(peer, sizeof(*cmc),
1318             FRP_CONFIG);
1319         if(cmc) {
1320             cmc->cost = htons(peer->cost);
1321             cmc->id = routerid;
1322             cmc->poll = htons(peer->confpoll);
1323             cmc->fail = htons(peer->conffail);
1324             peer->configreq = 0;
1325             LOG6("%s: Added config",
1326                 formatip(REMIP(peer)));
1327         }
1328     }
1329
1330     /*
1331     Ditto the path packet
1332     */
1333     if(peer->pathreq) {
1334         if(gwpeer)
1335             j = gwpeer->pathlen;
1336         else j = 0;
1337
1338         cmp = addmessage(peer, sizeof(*cmp)
1339             + sizeof(IPADDR) * (j+1),
1340             FRP_PATH);
1341         if(cmp) {
1342             cmp->gwcost = htons(gwcost);
```



```
1343     cmp->path[0] = routerid;
1344     for(i = 0; i < j; i++)
1345         cmp->path[i+1] = gwpeer->path[i];
1346     peer->pathreq = 0;
1347     LOG6("%s: Added path gwcost=%d",
1348         formatip(REMIP(peer)),
1349         gwcost);
1350 }
1351 }
1352
1353 /*
1354  And if we're to send some routing info, do that
1355  now too
1356  */
1357 while(peer->nexttrt) {
1358     cmr = addmessage(peer, sizeof(*cmr), FRP_ROUTE);
1359     if(!cmr) break;
1360     ipr = peer->nexttrt;
1361     cmr->flags = ((ipr == peer->annrts) ?
1362         FRP_FLAG_BEGIN : 0)
1363         | ((ipr->next == 0) ?
1364         FRP_FLAG_COMMIT : 0)
1365         | ((ipr->isgw) ?
1366         FRP_FLAG_GATEWAY : 0);
1367     cmr->bits = ipr->bits;
1368     cmr->cost = hton(ipr->cost);
1369     cmr->gwcost = hton(ipr->gwcost);
1370     cmr->ip = ipr->ip;
1371     peer->nexttrt = ipr->next;
1372     LOG6("%s: Added route %s/%d",
1373         formatip(REMIP(peer)),
1374         formatip(ipr->ip), ipr->bits);
1375 }
1376
1377 /*
1378  If we have to send a null route list, do so now
1379  */
1380 if(peer->nullrt) {
1381     cmr = addmessage(peer, FRP_NULLRT_SIZE,
1382         FRP_ROUTE);
1383     if(cmr) {
1384         cmr->flags = FRP_FLAG_NULL
1385             | FRP_FLAG_BEGIN
1386             | FRP_FLAG_COMMIT;
1387         cmr->bits = 0;
1388         peer->nullrt = 0;
1389         LOG6("%s: Added null route list",
1390             formatip(REMIP(peer)));
1391     }
1392 }
1393
1394 /*
1395  Send a poll request if required
1396  */
1397 if(peer->pollreq) {
1398     cml = addmessage(peer, sizeof(struct frp_ctrl),
1399         FRP_CTRL);
1400     if(cml) {
1401         cml->ctrl = FRP_CTRL_POLL;
1402         cml->param = 0;
1403         peer->pollreq = 0;
1404     }
1405 }
```

```
1404     }
1405 }
1406
1407 /*
1408  If we need to respond to a poll request, do so.
1409  Don't worry if there's no room in the packet.
1410  */
1411 if(peer->respreq) {
1412     cml = addmessage(peer, sizeof(struct frp_ctrl),
1413         FRP_CTRL);
1414     if(cml) {
1415         cml->ctrl = FRP_CTRL_ACK;
1416         cml->param = 0;
1417     }
1418     peer->respreq = 0;
1419 }
1420
1421 /*
1422  Finally, if we haven't already ACKed an inbound
1423  packet, do so now
1424  */
1425 if(peer->sendack) {
1426     addmessage(peer, 0, 0);
1427     peer->sendack = 0;
1428 }
1429
1430 /*
1431  If we have an outbound packet, and we're not still
1432  waiting for the last one to be acknowledged, send
1433  the packet now
1434  */
1435 if(peer->nextpktlen) {
1436     if(peer->ackreq) {
1437         peer->lseq = NEXTSEQ(peer->lseq);
1438         peer->ackreq = 0;
1439     }
1440     hdr = (struct frp_hdr *)peer->nextpacket;
1441     hdr->lseq = htonl(peer->lseq);
1442     hdr->rseq = htonl(peer->rseq);
1443     secure(peer, hdr, peer->nextpktlen);
1444     DFP_S(hdr, peer->nextpktlen, peer);
1445     write(peer->fd, hdr, peer->nextpktlen);
1446     peer->nextpacket = peer->lastpacket;
1447     peer->lastpacket = (u_char *)hdr;
1448     peer->lastpktlen = peer->nextpktlen;
1449     peer->nextpktlen = 0;
1450     peer->nexttime = now + peer->retry;
1451 }
1452 }
1453
1454
1455 fdc = rtsocket + 1;
1456 FD_ZERO(&fds);
1457 FD_SET(rtsocket, &fds);
1458
1459 FOREACH(peer, listens) {
1460     if(peer->fd >= fdc)
1461         fdc = peer->fd + 1;
1462     FD_SET(peer->fd, &fds);
1463 }
1464 i = 10;
```

```
1465     if(checkroutes)
1466         i = 1;
1467     FOREACH(peer, peers) {
1468         if(peer->fd >= fdc)
1469             fdc = peer->fd + 1;
1470         FD_SET(peer->fd, &fds);
1471         if(peer->lseq != peer->aseq) {
1472             j = peer->nexttime - now;
1473             if(j >= 0 && j < i)
1474                 i = j;
1475         }
1476     } else {
1477         j = peer->lasttime + peer->poll - now;
1478         if(j >= 0 && j < i)
1479             i = j;
1480     }
1481     LOG8("now=%d nexttime=%d lasttime=%d poll=%d"
1482          " retry=%d fail=%d i=%d",
1483          now, peer->nexttime, peer->lasttime,
1484          peer->poll, peer->retry, peer->fail, i);
1485 }
1486 LOG8("Waiting %s seconds", formattime(i));
1487 tv.tv_sec = i / 10;
1488 tv.tv_usec = (i % 10) * 100000;
1489 if(!i) tv.tv_usec = 20000;
1490 i = select(fdc, &fds, 0, 0, &tv);
1491
1492 gettimeofday(&tv, 0);
1493 now = tv.tv_sec - basetime;
1494 now = now * 10 + tv.tv_usec / 100000;
1495
1496 if(i <= 0) {
1497     if(i == 0 || errno == EINTR)
1498         continue;
1499     perror("select");
1500     return 1;
1501 }
1502
1503 /*
1504 If we got a routing socket message, read it and see if it's
1505 interesting. Adds, changes & deletes might interest us,
1506 but only if they don't involved cloned routes (ARP
1507 activity), broadcasts, multicasts, ICMP
1508 redirects/unreachables etc
1509 */
1510 if(FD_ISSET(rtsocket, &fds)) {
1511     len = read(rtsocket, buf, sizeof(buf));
1512     rtm = (struct rt_msghdr *) buf;
1513     if(!checkroutes && len >= sizeof(struct rt_msghdr)
1514         && ( rtm->rtm_type == RTM_ADD
1515             || rtm->rtm_type == RTM_DELETE
1516             || rtm->rtm_type == RTM_CHANGE )
1517         && !( rtm->rtm_flags & ( RTF_REJECT
1518                               | RTF_DYNAMIC
1519                               | RTF_MODIFIED
1520                               | RTF_WASCLONED
1521                               | RTF_BROADCAST
1522                               | RTF_MULTICAST )))
1523         checkroutes = 1;
1524     LOG8("Route message type %d flags %x check=%d",
1525          rtm->rtm_type, rtm->rtm_flags, checkroutes);
```

```
1526     }
1527
1528     /*
1529     See if we got any unsolicited requests on any of the listen
1530     sockets.
1531     */
1532     FOREACH(listen, listens) if(FD_ISSET(listen->fd, &fds)) {
1533         sal = sizeof(struct sockaddr_in);
1534         len = recvfrom(listen->fd, buf, sizeof(buf), 0,
1535             (struct sockaddr *) &listen->rsa, &sal);
1536         DFP_R(buf, len, listen);
1537
1538         /*
1539         Check the remote ACL (if there is one). Just drop
1540         the packet if so
1541         */
1542         if(listen->remote) {
1543             if(!checkacl(listen->remote,
1544                 REMIP(listen), -1, 0, 0)) {
1545                 LOG2("rejected packet from %s\n",
1546                     formatip(REMIP(listen)));
1547                 continue;
1548             }
1549         }
1550
1551         /*
1552         Check if the packet has the right secret
1553         */
1554         if(listen->secret && !checksecure(buf, len, listen))
1555             continue;
1556
1557         /*
1558         Having got this far, establish a new connection.
1559         */
1560         pfd = connectpeer(listen, 1);
1561         if(pfd == -1)
1562             continue;
1563
1564         /*
1565         If this is not a SYN packet (rseq == 0), NAK it
1566         and drop the connection. Typically, this happens
1567         if we have restarted, and the other end thinks it
1568         has a session open. By giving it a sensible looking
1569         NAK, we tell it that it should drop the connection and
1570         restart from scratch.
1571         */
1572         hdr = (struct frp_hdr *) buf;
1573         if(len != sizeof(struct frp_hdr) || hdr->rseq != 0) {
1574             LOG6("NAK %s", formatip(REMIP(listen)));
1575             sendnak(listen, pfd, hdr->lseq, hdr->rseq);
1576             close(pfd);
1577             continue;
1578         }
1579
1580         /*
1581         We're happy at this point. Create a new peer
1582         object, copy the template listen object onto
1583         it, mark it as cloned (so we know to get rid of
1584         it on shutdown), do the initialisation and link
1585         it to the list of peers.
1586         */
```

```
1587     peer = malloc(sizeof(struct peer));
1588     *peer = *listen;
1589     peer->next = peers;
1590     peers = peer;
1591
1592     getlocaladdr(peer, REMIP(listen), 0);
1593     peer->lsa = listen->lsa;
1594     peer->rsa = listen->rsa;
1595     if(!peer->nexthop)
1596         peer->nexthop = REMIP(peer);
1597     peer->cloned = 1;
1598     peer->fd = pfd;
1599     peer->shutdown = 0;
1600     resetpeer(peer, 1);
1601
1602     /*
1603     Process the received packet
1604     */
1605     LOG1("Connect from %s established",
1606         formatip(REMIP(peer)));
1607     process_packet(buf, len, peer);
1608 }
1609
1610 /*
1611 See if we got any packets on established peers ...
1612 Check basic sanity. If it's sane, and it is encoded with
1613 the right secret, process it.
1614 */
1615 FOREACH(peer, peers) if(FD_ISSET(peer->fd, &fds)) {
1616     len = read(peer->fd, buf, sizeof(buf));
1617     DFP_R(buf, len, peer);
1618
1619     if(len < (int) sizeof(struct frp_hdr)) {
1620         if(len != -1)
1621             LOG3("%s: short packet",
1622                 formatip(REMIP(peer)));
1623         continue;
1624     }
1625     if(!peer->secret || checksecure(buf, len, peer))
1626         process_packet(buf, len, peer);
1627 }
1628 }
1629
1630 /*
1631 If shutdown request, do brutal shutdown stuff
1632 Basically, stuff up the peers list etc so that the routing code
1633 thinks there are no peers and therefore no default route, so
1634 all routes get deleted.
1635 */
1636 if(shutdwn) {
1637     peers = 0;
1638     getroutes();
1639 }
1640
1641 /*
1642 Last gasp packets
1643 Send a NAK packet to every active peer. This should cause the
1644 peer to restart.
1645 */
1646 FOREACH(peer, peers) if(!peer->synreq)
1647     sendnak(peer, peer->fd, htonl(peer->rseq), htonl(peer->lseq));
```

```
1648
1649 /*
1650 If there's a PID file, drop it now
1651 Ditto the status file
1652 */
1653 if(pf) {
1654     unlink(pidfile);
1655     flock(fileno(pf), LOCK_UN);
1656     fclose(pf);
1657 }
1658 if(statusfile)
1659     unlink(statusfile);
1660
1661 /*
1662 If we got asked to restart, have a go at staring ourselevs
1663 */
1664 if(restart)
1665     execv(argv[0], argv);
1666
1667 return 0;
1668 }
1669
1670 /* EOF */
1671
```

```
1 #include "frpd.h"
2
3
4 static char *secret = 0;
5
6
7
8 static char *
9 gettmp(int size) {
10     static char buf[64];
11     static char *ptr = buf;
12     char *s;
13     if(ptr + size >= buf + sizeof(buf))
14         ptr = buf;
15     s = ptr;
16     ptr += size;
17     return s;
18 }
19
20
21 char *
22 formatip(IPADDR ip) {
23     union { u_int8_t b[4]; IPADDR a; } ipa;
24     char *s;
25
26     s = gettmp(16);
27     ipa.a = ip;
28     sprintf(s, "%u.%u.%u.%u", ipa.b[0], ipa.b[1], ipa.b[2], ipa.b[3]);
29     return s;
30 }
31
32 char *
33 formattime(TIMETEN t) {
34     char *s = gettmp(16);
35     int d = t % 10;
36     sprintf(s, "%u%c%c", t / 10, d ? ':' : 0, d + '0');
37     return s;
38 }
39
40 static char *
41 parsetime(char *s, TIMETEN *tp) {
42     TIMETEN t;
43     char *p;
44
45     t = strtol(s, &p, 10);
46     if(t > 1000000)
47         return "Invalid interval";
48     t *= 10;
49     if(*p == '.' && isdigit(p[1]) && p[2] == 0)
50         t += p[1] - '0';
51     else if(p == s || *p != 0)
52         return "Invalid interval";
53     *tp = t;
54     return 0;
55 }
56
57
58 char *
59 parseip(char *s, IPADDR *ip_p, int *bits_p, int *maxbits_p) {
60     int bits;
61     int maxbits;
```

```
62     union { u_int8_t b[4]; IPADDR a; } ipa;
63     int i, j;
64     char *t;
65
66     ipa.a = 0;
67     for(i = 0; i < 4; i++) {
68         j = strtol(s, &t, 10);
69         if(t == s) return "Invalid IP address";
70         s = t + 1;
71         ipa.b[i] = (unsigned) j;
72         if(*t != '.') break;
73     }
74
75     if(*t == '/') {
76         if(!bits_p) return "Invalid IP address";
77         bits = strtol(s, &t, 10);
78         if(s == t || bits < 0 || bits > 32)
79             return "Invalid prefix length";
80         if(*t == '-') {
81             if(!maxbits_p) return "Invalid prefix length";
82             s = t + 1;
83             maxbits = strtol(s, &t, 10);
84             if(s == t || maxbits < bits || maxbits > 32)
85                 return "Invalid prefix range";
86         }
87         else maxbits = bits;
88     }
89     else maxbits = bits = 32;
90     if(*t) return "Invalid IP address";
91     if((ipa.a & maskbits[bits]) != ipa.a)
92         return "IP address mask mismatch";
93     *ip_p = ipa.a;
94     if(bits_p) *bits_p = bits;
95     if(maxbits_p) *maxbits_p = maxbits;
96     return 0;
97 }
98
99
100 struct acl *
101 findacl(char *name, int create) {
102     struct acl *acl;
103     FOREACH(acl, acs)
104         if(!strcmp(acl->name, name))
105             return acl;
106     if(!create) return 0;
107     acl = calloc(1, sizeof(struct acl));
108     acl->next = acs;
109     acs = acl;
110     acl->name = strdup(name);
111     return acl;
112 }
113
114
115 static char *
116 parse_acl(char *name, char *expression) {
117     struct acl *acl;
118     struct ace *ace, proto;
119     char *s;
120     char buf[4096];
121
122     if(expression) {
```

```

123     strcpy(buf, expression);
124     s = strtok(buf, WHITESPACE);
125 }
126 else s = strtok(0, WHITESPACE);
127
128 if(!s) return "Expected 'permit' or 'deny'";
129
130 memset(&proto, 0, sizeof(proto));
131 if(!strcmp(s, "permit"))
132     proto.permit = 1;
133 else if(!strcmp(s, "deny"))
134     return "Expected 'permit' or 'deny'";
135 proto.maxbits = 32;
136
137 while((s = strtok(0, WHITESPACE))) {
138     if(!strcmp(s, "ip")) { // n.n.n.n[/bits[-mbits]]
139         s = strtok(0, WHITESPACE);
140         if(!s) return "Expected IP prefix";
141         s = parseip(s, &proto.ip, &proto.bits, &proto.maxbits);
142         if(s) return s;
143     }
144     else if(!strcmp(s, "default")) {
145         proto.ip = 0;
146         proto.maxbits = proto.bits = 0;
147     }
148     else if(!strcmp(s, "interface")) {
149         s = strtok(0, WHITESPACE);
150         if(!s) return "Expected interface name";
151         proto.ifindex = if_nametoindex(s);
152         if(!proto.ifindex)
153             return "Unknown interface";
154     }
155     else if(!strcmp(s, "acl")) {
156         s = strtok(0, WHITESPACE);
157         if(!s) return "Expected ACL name";
158         proto.acl = findacl(s, 0);
159         if(!proto.acl) return "ACL not found";
160     }
161     else if(!strcmp(s, "layer2"))
162         proto.rtype |= RTYPE_LAYER2;
163     else if(!strcmp(s, "static"))
164         proto.rtype |= RTYPE_STATIC;
165     else if(!strcmp(s, "local"))
166         proto.rtype |= RTYPE_LOCAL;
167     else if(!strcmp(s, "protocol"))
168         proto.rtype |= RTYPE_PROTOCOL;
169     else if(!strcmp(s, "remote"))
170         proto.rtype |= RTYPE_REMOTE;
171
172     else return "Unrecognised keyword";
173 }
174
175 acl = findacl(name, 1);
176 ace = malloc(sizeof(struct ace));
177 *ace = proto;
178 if(!acl->head)
179     acl->head = ace;
180 else acl->tail->next = ace;
181 acl->tail = ace;
182
183 return 0;

```

```

184 }
185
186
187 /*
188 static void
189 dumpacl(char *name) {
190     struct acl *acl;
191     struct ace *ace;
192     union { u_int8_t b[4]; IPADDR a; } ipa;
193
194     acl = findacl(name, 0);
195     if(!acl) {
196         printf("No such ACL %s\n", name);
197         return;
198     }
199     FOREACH(ace, acl->head) {
200         ipa.a = ace->ip;
201         printf("%-6s %u.%u.%u.%u/%d-%d rtype=%x, acl=%s\n",
202             ace->permit ? "permit" : "deny",
203             ipa.b[0], ipa.b[1], ipa.b[2], ipa.b[3],
204             ace->bits, ace->maxbits, ace->rtype,
205             ace->acl ? ace->acl->name : "(none)");
206     }
207 }
208 */
209
210
211 /*
212 Find interface information associated with IP address
213 */
214 char *
215 getlocaladdr(struct peer *peer, IPADDR peerip, int listen) {
216     static struct ifaddrs *ifphdr = 0;
217     struct ifaddrs *ifp;
218     IPADDR a, m;
219     int b;
220     IPADDR localaddr;
221     static time_t lasttime = 0;
222     time_t t;
223
224     localaddr = 0;
225     peer->localbits = 0;
226
227     t = time(0);
228     if(ifphdr && t != lasttime) {
229         freeifaddrs(ifphdr);
230         ifphdr = 0;
231     }
232     if(!ifphdr) if(getifaddrs(&ifphdr)) {
233         LOG1("Error getting interface addresses: %m");
234         return "Could not obtain interface addresses";
235     }
236     lasttime = t;
237
238     for(ifp = ifphdr; ifp; ifp = ifp->ifa_next)
239         if(ifp->ifa_addr->sa_family == AF_INET) {
240             a = ((struct sockaddr_in *) (ifp->ifa_addr))->sin_addr.s_addr;
241             m = ((struct sockaddr_in *) (ifp->ifa_netmask))->sin_addr.s_addr;
242             b = (m & 0x55555555) + ((m >> 1) & 0x55555555);
243             b = (b & 0x33333333) + ((b >> 2) & 0x33333333);
244             b = (b & 0x0f0f0f0f) + ((b >> 4) & 0x0f0f0f0f);

```

```

245     b = (b & 0x0fff0fff) + ((b >> 8) & 0x0fff0fff);
246     b = (b & 0x0000ffff) + ((b >> 16) & 0x0000ffff);
247     if((!listen && a == peerip) ||
248         (!listen && (peerip & m) == (a & m)
249             && b > peer->localbits) ||
250         (!listen && peerip ==
251             ((struct sockaddr_in *) (ifp->ifa_dstaddr))->
252                 sin_addr.s_addr)) {
253         localaddr = a;
254         peer->localroute = peerip & m;
255         peer->localbits = b;
256         peer->ifindex = if_nametoindex(ifp->ifa_name);
257     }
258 }
259 if(!localaddr && listen)
260     return "Interface address not local";
261 else if(!localaddr)
262     return "IP address is not directly connected";
263
264 if(listen)
265     LOCIP(peer) = peerip;
266 else {
267     REMIP(peer) = peerip;
268     LOCIP(peer) = localaddr;
269 }
270 return 0;
271 }
272
273
274
275 static char *
276 addpeer(IPADDR peerip, int listen) {
277     char *s;
278     struct peer *peer;
279     int p;
280     int new;
281
282     if(listen) {
283         FOREACH(peer, listens)
284             if(LOCIP(peer) == peerip)
285                 break;
286     }
287     else {
288         FOREACH(peer, peers)
289             if(REMIP(peer) == peerip)
290                 break;
291     }
292     if(!peer) {
293         new = 1;
294         peer = calloc(1, sizeof(struct peer));
295         peer->lsa.sin_family = AF_INET;
296         peer->lsa.sin_port = htons(udpport);
297         peer->rsa.sin_family = AF_INET;
298         peer->rsa.sin_port = htons(udpport);
299
300         if(listen)
301             peer->nexthop = 0;
302         else peer->nexthop = peerip;
303         peer->tvl = 1;
304         peer->secret = secret;
305         peer->confcost = 1;

```

```

306     peer->confpoll = DEFAULT_POLL;
307     peer->conffail = DEFAULT_FAIL;
308     peer->confretry = DEFAULT_RETRY;
309     peer->localannounce = 2;
310
311     s = getlocaladdr(peer, peerip, listen);
312     if(s) return s;
313 }
314 else new = 0;
315
316 while((s = strtok(0, WHITESPACE))) {
317     if(!strcmp(s, "port")) {
318         if(!new) return "Port must be on first line";
319         s = strtok(0, WHITESPACE);
320         if(!s) return "Expected port number";
321         p = atoi(s);
322         if(p < 1 || p > 65535)
323             return "Invalid port number";
324         if(listen)
325             peer->lsa.sin_port = htons(p);
326         else peer->rsa.sin_port = htons(p);
327     }
328     else if(!strcmp(s, "source-port") && !listen) {
329         if(!new) return "Source port must be on first line";
330         s = strtok(0, WHITESPACE);
331         if(!s) return "Expected port number";
332         p = atoi(s);
333         if(p < 1 || p > 65535)
334             return "Invalid port number";
335         peer->lsa.sin_port = htons(p);
336     }
337     else if(!strcmp(s, "cost")) {
338         s = strtok(0, WHITESPACE);
339         if(!s) return "Expected cost";
340         peer->confcost = atoi(s);
341         if(peer->confcost < 1 || peer->confcost > 256)
342             return "Invalid cost";
343         peer->cost = peer->confcost;
344     }
345     else if(!strcmp(s, "ttl")) {
346         if(!new) return "TTL must be on first line";
347         s = strtok(0, WHITESPACE);
348         if(!s) return "Expected TTL";
349         peer->tvl = atoi(s);
350         if(peer->tvl < 1 || peer->tvl > 254)
351             return "Invalid TTL";
352     }
353     else if(!strcmp(s, "accept")) {
354         s = strtok(0, WHITESPACE);
355         if(!s) return "Expected ACL name";
356         peer->accept = findacl(s, 0);
357         if(!peer->accept)
358             return "Unknown ACL";
359     }
360     else if(!strcmp(s, "announce")) {
361         s = strtok(0, WHITESPACE);
362         if(!s) return "Expected ACL name";
363         peer->announce = findacl(s, 0);
364         if(!peer->announce)
365             return "Unknown ACL";
366     }

```

```

367     else if(!strcmp(s, "accept-connect") && listen) {
368         s = strtok(0, WHITESPACE);
369         if(!s) return "Expected ACL name";
370         peer->remote = findacl(s, 0);
371         if(!peer->remote)
372             return "Unknown ACL";
373     }
374     else if(!strcmp(s, "nexthop") && !listen) {
375         s = strtok(0, WHITESPACE);
376         if(!s) return "Expected IP address";
377         s = parseip(s, &peer->nexthop, 0, 0);
378         if(s) return s;
379     }
380     else if(!strcmp(s, "source-ip") && !listen) {
381         if(!new) return "Source IP must be on first line";
382         s = strtok(0, WHITESPACE);
383         if(!s) return "Expected IP address";
384         s = parseip(s, &LOCIP(peer), 0, 0);
385         if(s) return s;
386     }
387     else if(!strcmp(s, "secret")) {
388         s = strtok(0, WHITESPACE);
389         if(!s) return "Expected secret";
390         peer->secret = strdup(s);
391     }
392     else if(!strcmp(s, "poll")) {
393         s = strtok(0, WHITESPACE);
394         if(!s) return "Expected poll interval";
395         s = parsetime(s, &peer->confpoll);
396         if(s) return s;
397         peer->poll = peer->confpoll;
398     }
399     else if(!strcmp(s, "fail")) {
400         s = strtok(0, WHITESPACE);
401         if(!s) return "Expected timeout";
402         s = parsetime(s, &peer->conffail);
403         if(s) return s;
404         peer->fail = peer->conffail;
405     }
406     else if(!strcmp(s, "retry")) {
407         s = strtok(0, WHITESPACE);
408         if(!s) return "Expected retry interval";
409         s = parsetime(s, &peer->confretry);
410         if(s) return s;
411         peer->retry = peer->confretry;
412     }
413     else if(!strcmp(s, "announce-local")) {
414         s = strtok(0, WHITESPACE);
415         if(!s) return "Expected yes, no or choose";
416         if(!strcmp(s, "choose"))
417             peer->localannounce = 0;
418         else if(!strcmp(s, "yes"))
419             peer->localannounce = 1;
420         else if(!strcmp(s, "no"))
421             peer->localannounce = 2;
422         else return "Expected yes, no or choose";
423     }
424     else return "Unknown keyword";
425 }
426
427 if(peer->confpoll + 3 * peer->confretry > peer->conffail)

```

```

428     peer->conffail = peer->confpoll + 3 * peer->confretry;
429
430     if(new) {
431         if(listen) {
432             peer->next = listens;
433             listens = peer;
434         }
435         else {
436             peer->next = peers;
437             peers = peer;
438         }
439     }
440     return 0;
441 }
442
443
444 void
445 parse_acl_std(void) {
446     parse_acl("rfc1918", "permit ip 10.0.0.0/8-32");
447     parse_acl("rfc1918", "permit ip 172.16.0.0/12-32");
448     parse_acl("rfc1918", "permit ip 192.168.0.0/16-32");
449
450     parse_acl("bogons", "permit ip 0.0.0.0/8-32"); /* Local host */
451     parse_acl("bogons", "permit ip 127.0.0.0/8-32"); /* Loopback */
452     parse_acl("bogons", "permit ip 169.254.0.0/16-32"); /* Link-local */
453     parse_acl("bogons", "permit ip 224.0.0.0/3-32"); /* Non-unicast */
454
455     parse_acl("public", "deny acl bogons");
456     parse_acl("public", "deny acl rfc1918");
457     parse_acl("public", "permit");
458
459     parse_acl("any", "permit");
460     parse_acl("none", "deny");
461 }
462
463
464 int
465 parse_config(char *file) {
466     FILE *cf;
467     char buf[4096];
468     char *s, *t;
469     int i;
470     IPADDR peerip;
471     int lc, lcc;
472     int ok;
473
474     if(file) {
475         cf = fopen(file, "r");
476         if(!cf) {
477             perror(file);
478             return 0;
479         }
480     }
481     else cf = stdin;
482
483 #define CFOOPS(msg) { printf("%s:%d: Error: %s\n", file, lc, msg); \
484     ok = 0; goto next;}
485     lcc = 1;
486     lc = 0;
487     ok = 1;
488     while(fgets(buf, sizeof(buf), cf)) {

```

```
489     lc += lcc;
490     lcc = 1;
491     t = 0;
492     for(s = buf; *s; ) {
493         if(*s == '#' || *s == '\n' || !*s) {
494             *s = 0;
495             if(t) {
496                 i = sizeof(buf) - (t - buf);
497                 if(i >= sizeof(buf) || !fgets(t, i, cf))
498                     break;
499             }
500             else break;
501             lcc++;
502         }
503         if(*s == '\\') t = s;
504         else if(lisspace(*s))
505             t = 0;
506         s++;
507     }
508
509     s = strtok(buf, WHITESPACE);
510     if(!s) continue;
511
512     if(!strcmp(s, "acl")) {
513         s = strtok(0, WHITESPACE);
514         if(!s) CFOOPS("Expected ACL name");
515         s = parse_acl(s, 0);
516         if(s) CFOOPS(s)
517     }
518
519     else if(!strcmp(s, "listen")) {
520         s = strtok(0, WHITESPACE);
521         if(!s) CFOOPS("Expected peer address");
522         s = parseip(s, &peerip, 0, 0);
523         if(s || !peerip)
524             CFOOPS(s);
525         s = addpeer(peerip, 1);
526         if(s) CFOOPS(s);
527     }
528
529     else if(!strcmp(s, "peer")) {
530         s = strtok(0, WHITESPACE);
531         if(!s) CFOOPS("Expected peer address");
532         s = parseip(s, &peerip, 0, 0);
533         if(s || !peerip)
534             CFOOPS(s);
535         s = addpeer(peerip, 0);
536         if(s) CFOOPS(s);
537     }
538
539     else if(!strcmp(s, "id")) {
540         s = strtok(0, WHITESPACE);
541         if(!s) CFOOPS("Expected router ID");
542         s = parseip(s, &routerid, 0, 0);
543         if(s) CFOOPS(s)
544     }
545
546     else if(!strcmp(s, "port")) {
547         s = strtok(0, WHITESPACE);
548         if(!s) CFOOPS("Expected port number");
549         udpport = atoi(s);
550         if(udpport < 1 || udpport >= 65535)
551             CFOOPS("Invalid port number")
552     }
553 }
```

```
550     else if(!strcmp(s, "gateway")) {
551         s = strtok(0, WHITESPACE);
552         if(!s) CFOOPS("Expected yes, no or always")
553     }
554     else if(!strcmp(s, "always"))
555         gwalways = isgateway = 1;
556     else if(!strcmp(s, "yes"))
557         isgateway = 1;
558     else if(!strcmp(s, "no"))
559         isgateway = 1;
560     else CFOOPS("Expected yes, no or always")
561 }
562
563 else if(!strcmp(s, "statusfile")) {
564     s = strtok(0, WHITESPACE);
565     if(!s) CFOOPS("Expected status file");
566     statusfile = strdup(s);
567 }
568
569 else if(!strcmp(s, "debug")) {
570     s = strtok(0, WHITESPACE);
571     if(!s) CFOOPS("Expected debug level");
572     debug = atoi(s);
573     if(debug < 0 || debug > 8)
574         CFOOPS("Bad debug level");
575 }
576
577 else if(!strcmp(s, "pidfile")) {
578     s = strtok(0, WHITESPACE);
579     if(!s) CFOOPS("Expected PID file");
580     pidfile = strdup(s);
581 }
582
583 else if(!strcmp(s, "secret")) {
584     s = strtok(0, WHITESPACE);
585     if(!s) CFOOPS("Expected secret");
586     secret = strdup(s);
587 }
588
589 else if(!strcmp(s, "route-flag")) {
590     s = strtok(0, WHITESPACE);
591     if(!s) CFOOPS("expected routing flag")
592     if(!strcmp(s, "proto1"))
593         routeflag = RTF_PROTO1;
594     else if(!strcmp(s, "proto2"))
595         routeflag = RTF_PROTO2;
596     else if(!strcmp(s, "proto3"))
597         routeflag = RTF_PROTO3;
598     else if(!strcmp(s, "static"))
599         otherflag |= RTF_STATIC;
600     else CFOOPS("invalid routing keyword")
601 }
602
603 else if(!strcmp(s, "default-gateway")) {
604     s = strtok(0, WHITESPACE);
605     if(!s) CFOOPS("Expected default gateway address");
606     s = parseip(s, &defaulttroute, 0, 0);
607     if(s) CFOOPS(s)
608 }
609
610 else if(!strcmp(s, "override")) {
611     s = strtok(0, WHITESPACE);
612     if(!s) CFOOPS("expected routing override flag")
613     if(!strcmp(s, "proto1"))
614         overrflag |= RTF_PROTO1;
615     else if(!strcmp(s, "proto2"))
616         overrflag |= RTF_PROTO2;
617     else if(!strcmp(s, "proto3"))
618         overrflag |= RTF_PROTO3;
619 }
```



```
611     else if(!strcmp(s, "static"))
612         overrflag |= RTF_STATIC;
613     else CFOOPS("invalid routing override flag")
614 }
615
616
617     else CFOOPS("Unrecognised keyword")
618
619 next:  continue; /* FORTRAN IV, anyone? */
620 }
621 if(cf != stdin) fclose(cf);
622 return ok;
623 }
624
```

```

1 #include "frpd.h"
2
3 int routeflag = RTF_PROTO2; /* Route flag RTF_PROTO<n> */
4 int otherflag = 0; /* Other routing flags */
5 int overrflag = 0; /* Override routes with this flag */
6
7 IPADDR defaultroute = 0; /* Default default route */
8 IPADDR defgateway = 0; /* Current default route */
9
10 /*
11  Add an IP route to the kernel routing table
12  */
13 static void
14 addroute(IPADDR ip, int bits, IPADDR gw) {
15     u_char buf[256];
16     u_char *saptr;
17     struct rt_msghdr *rtm;
18     struct sockaddr_in *ipsa;
19     struct sockaddr_in *gwsa;
20     struct sockaddr_in *mask;
21     int size;
22
23     /*
24      Don't do this if we're quiescent
25      */
26     LOG5("add route %s/%d -> %s routeflag=%x",
27          formatip(ip), bits, formatip(gw), routeflag);
28     if(gwcost == -1 && ip)
29         return;
30
31     /*
32      Zap the buffer.
33      Assemble message header
34      Note that the RTF_HOST flag and the RTA_NETMASK address flag
35      are added conditionally further down.
36      */
37     memset(buf, 0, sizeof(buf));
38     rtm = (struct rt_msghdr *)buf;
39     rtm->rtm_version = RTM_VERSION;
40     rtm->rtm_type = RTM_ADD;
41     rtm->rtm_pid = pid;
42     rtm->rtm_flags = RTF_GATEWAY | RTF_UP | routeflag | otherflag;
43     rtm->rtm_addrs = RTA_DST | RTA_GATEWAY;
44     saptr = buf + sizeof(struct rt_msghdr);
45
46     /*
47      Now the IP address
48      */
49     ipsa = (struct sockaddr_in *) saptr;
50     ipsa->sin_len = sizeof(struct sockaddr_in);
51     ipsa->sin_family = AF_INET;
52     ipsa->sin_addr.s_addr = ip;
53     saptr += SA_SIZE(saptr);
54
55     /*
56      And the gateway address
57      */
58     gwsa = (struct sockaddr_in *) saptr;
59     gwsa->sin_len = sizeof(struct sockaddr_in);
60     gwsa->sin_family = AF_INET;
61     gwsa->sin_addr.s_addr = gw;

```

```

62     saptr += SA_SIZE(saptr);
63
64     /*
65      And finally the mask
66      Otherwise add a mask address.
67      The following is based on observation. (Ugh!)
68
69      Masks are a sockaddr_in, but the length is oddball, and the family
70      is unset.
71
72      Default masks have the length byte set to 0.
73      Otherwise, the length (mask->sin_len) byte is the mask length
74      divided by 8 plus 4, which leaves just the "interesting" bytes
75      of mask->sin_addr.
76
77      Note that the SA_SIZE() macro rounds up to 4 byte increments, so the
78      actual size of the mask sockaddr is 4 for a default mask, and 8
79      for other mask lengths. Note that a sockaddr (of any description)
80      is usually padded to 16 bytes, as it is in the case of the address
81      and gateway sockaddrs.
82
83      So, we get:
84      len family port addr
85      /0 00 00 00 00 Default
86      /1 05 xx xx xx 80 xx xx xx xx = anything
87      /2 05 xx xx xx C0 xx xx xx (we leave it as 0)
88      /8 05 xx xx xx FF xx xx xx
89      /9 06 xx xx xx FF 80 xx xx
90      /16 06 xx xx xx FF FF xx xx
91      /24 07 xx xx xx FF FF FF xx
92      /28 08 xx xx xx FF FF FF F0
93      /32 08 xx xx xx FF FF FF FF Or omit and set HOST
94      */
95     if(bits > 0) { /* Leave as all zeroes if bits == 0 */
96         mask = (struct sockaddr_in *) saptr;
97         mask->sin_len = (bits + 7) / 8 + 4;
98         mask->sin_addr.s_addr = maskbits[bits];
99     }
100     saptr += SA_SIZE(saptr);
101     rtm->rtm_addrs |= RTA_NETMASK;
102
103     /*
104      Update the message size and send the routing message to the kernel.
105      */
106     size = saptr - buf;
107     rtm->rtm_msglen = size;
108     write(rtsocket, buf, size);
109 }
110
111 /*
112  Main routing engine
113  The steps are as follows:
114  (1) Compute the best path to the gateway; this sets the new gateway peer
115      and gateway cost values.
116  (2) Construct the local routing table from the set of routes learned
117      from our peers.
118  (3) Parse the kernel routing table, inserting (or replacing) locally
119      originated routes. At this point, delete stale routes from the
120      kernel table.
121  (4) Insert any routes learned from peers into the routing table (if it

```

```

123     wasn't there already).
124 (5) Check the gateway status, trigger a path update if necessary.
125 (6) Deal with local routes associated with inactive peer interfaces.
126 (7) Generate route announcements to all active peers; if the announcement
127     to a peer has changed, send the update.
128 */
129 int
130 getroutes() {
131     int mib[6] = { CTL_NET, PF_ROUTE, 0, 0, NET_RT_DUMP, 0 };
132     static size_t oldbufsiz = 0;
133     static u_char *buf = 0;
134     static IPADDR oldpath[MAX_PATH];
135     static int oldpathlen = 0;
136     size_t bufsiz;
137     u_char *ptr, *end, *saptr;
138     struct rt_msghdr *rtm;
139     struct sockaddr *sa;
140     struct sockaddr_in *sin;
141     struct sockaddr_dl *sdl;
142     u_int u;
143     int i, j;
144     int valid;
145     IPADDR ip;
146     IPADDR gw;
147     int bits;
148     int ifindex;
149     int rtype;
150     struct iproute *ipr;
151     struct iproute *r;
152     struct iproute *pipr;
153     struct iproute *annh;
154     struct iproute *annt;
155     struct iproute *oldr;
156     struct peer *peer;
157     struct peer *gp;
158     int gc;
159     IPADDR defaultgw;
160     IPADDR oldgw;
161     int doann;
162     int newpath;
163
164     /*
165      * Compute gateway path on non-gateway nodes
166      */
167     newpath = 0;
168     if(!isgateway) {
169         /*
170          * Search peers for the best route to the gateway
171          */
172         gp = 0;
173         gc = -1;
174         FOREACH(peer, peers) {
175             /*
176              * Ignore peers that haven't figured out their gateway,
177              * and peers that include our router-id in their path to
178              * the gateway.
179              * The latter avoids counting to infinity.
180              */
181             if(peer->gwcost == -1)
182                 continue;
183             for(i = 0; i < peer->pathlen; i++)

```

```

184         if(peer->path[i] == routerid)
185             break;
186         if(i < peer->pathlen)
187             continue;
188
189     /*
190      * Gateway selection:
191      * Pick this one if it's the first valid one we found;
192      * or if it's better than the best we've found so far;
193      * or if it's as good as the best so far and it's the
194      * existing gateway peer (this buys us stability).
195      */
196     i = peer->gwcost + peer->cost;
197     if(gc == -1 || i < gc || (i == gc && peer == gwpeer)) {
198         gp = peer;
199         gc = i;
200     }
201
202     /*
203      * If the route has changed, log the change and set
204      * gwpeer & gwcost accordingly
205      */
206     if(gp != gwpeer || gc != gwcost) {
207         LOG1("Gateway changed from %s (cost %d)"
208             " to %s (cost %d)",
209             gwpeer ? formatip(gwpeer->nexthop) :
210             defaultroute ? formatip(defaultroute) : "none",
211             gwcost,
212             gp ? formatip(gp->nexthop) :
213             defaultroute ? formatip(defaultroute) : "none",
214             gc);
215         gwpeer = gp;
216         gwcost = gc;
217         newpath = 1;
218     }
219
220     /*
221      * Find the default route
222      * Check if the path has changed
223      */
224     if(gwpeer) {
225         if(oldpathlen != gwpeer->pathlen)
226             newpath = 1;
227         else if(!newpath) {
228             for(i = 0; i < gwpeer->pathlen; i++)
229                 if(oldpath[i] != gwpeer->path[i])
230                     newpath = 1;
231         }
232         defaultgw = gwpeer->nexthop;
233     }
234     else defaultgw = 0;
235
236     /*
237      * If the path, gateway cost etc has changed, tell all the
238      * other nodes of this fact
239      */
240     if(newpath) {
241         FOREACH(peer, peers)
242             if(!peer->synreq)
243                 peer->pathreq = 1;
244     }

```

```
245     if(gwpeer) {
246         for(i = 0; i <= gwpeer->pathlen; i++)
247             oldpath[i] = gwpeer->path[i];
248         oldpathlen = gwpeer->pathlen;
249     }
250     else oldpathlen = 0;
251 }
252 }
253 else {
254     gwpeer = 0;
255     defaultgw = 0;
256     oldpathlen = 0;
257 }
258
259 /*
260 Blast the local routing table
261 */
262 KILLROUTES(localroutes)
263
264 /*
265 Add all the peer routes to the local routing table
266 Mark them all as new; the subsequent pass through the routing table
267 will mark the existing ones as such.
268 */
269 FOREACH(peer, peers) {
270     LOG8("Searching for routes from peer %s",
271         formatip(REMIP(peer)));
272     FOREACH(ipr, peer->routes) {
273         if(peer->accept && !checkacl(peer->accept,
274             ipr->ip, ipr->bits, ipr->rtype, 0))
275             continue;
276         r = findroute(localroutes, ipr->ip, ipr->bits, &pipr);
277         /*
278         If we found the route, don't update it unless it's
279         better than the old one
280         */
281         if(r) {
282             if(ipr->cost > r->cost ||
283                 (ipr->cost == r->cost &&
284                     ipr->isgw < r->isgw))
285                 continue;
286             LOG8("Updated %s/%d -> %s c=%d",
287                 formatip(ipr->ip), ipr->bits,
288                 formatip(peer->nexthop), ipr->cost);
289         }
290     }
291     /*
292     If the route doesn't exist, add it
293     */
294     else {
295         ADDROUTEAFTER(r, localroutes, pipr)
296         LOG8("Added %s/%d -> %s c=%d",
297             formatip(ipr->ip), ipr->bits,
298             formatip(peer->nexthop), ipr->cost);
299         r->ip = ipr->ip;
300         r->bits = ipr->bits;
301     }
302     r->isgw = ipr->isgw;
303     r->cost = ipr->cost;
304     r->gwcost = ipr->gwcost;
305     r->rtype = RTYPE_REMOTE;
```

```
306     r->peer = peer;
307     r->ifindex = peer->ifindex;
308     r->inuse = 0;
309 }
310 }
311
312 /*
313 Dump the routing table. First find out how big it is, make sure
314 the buffer is big enough, then actually go and get it.
315 */
316 LOG8("Retrieving routing table");
317 if(sysctl(mib, 6, NULL, &bufsiz, NULL, 0) < 0) {
318     LOG1("error obtaining routing table size: %m");
319     return 1;
320 }
321 if((oldbufsiz || bufsiz > oldbufsiz) {
322     if(buf) free(buf);
323     buf = malloc(bufsiz);
324     oldbufsiz = bufsiz;
325 }
326 i = sysctl(mib, 6, buf, &bufsiz, NULL, 0);
327 if(i == -1 && errno == ENOMEM) {
328     LOG2("Routing table size error");
329     return 1;
330 }
331 else if(i == -1) {
332     LOG1("Error retrieving routing table: %m");
333     return 1;
334 }
335
336 /*
337 For each route in the returned bunch of routes ...
338 */
339 oldgw = defgateway;
340 defgateway = 0;
341 ptr = buf;
342 for(end = buf + bufsiz; ptr < end; ptr += rtm->rtm_msglen) {
343
344     /*
345     Peel apart the route header
346     */
347     rtm = (struct rt_msghdr *)ptr;
348     saptr = ptr + sizeof(struct rt_msghdr);
349     valid = 0;
350     ifindex = 0;
351     ip = 0;
352     gw = 0;
353     if(rtm->rtm_flags & RTF_HOST)
354         bits = 32;
355     else bits = -1;
356
357     /*
358     Ignore down, dynamic, cloned, broadcast & multicast routes
359     */
360     rtype = 0;
361     if(!((rtm->rtm_flags & RTF_UP) ||
362         (rtm->rtm_flags & (RTF_DYNAMIC |
363             RTF_WASCLONED |
364             RTF_BROADCAST |
365             RTF_MULTICAST))))
366         continue;
```

```
367
368 /*
369 Set the route type flags based on the route flags for ACL
370 matching
371 */
372 if(rtm->rtm_flags & RTF_LLINFO)
373     rtype = RTYPE_LAYER2;
374 else if(rtm->rtm_flags & RTF_STATIC)
375     rtype = RTYPE_STATIC;
376 else if(rtm->rtm_flags & RTF_PROTO1)
377     rtype = RTYPE_PROTOCOL;
378 else rtype = RTYPE_LOCAL;
379
380 /*
381 Now parse each of the sockaddr objects from the route
382 */
383 for(i = 0; i < RTAX_MAX; i++) if(rtm->rtm_addrs & (1 << i)) {
384     sa = (struct sockaddr *) saptr;
385
386     /*
387     Destination IP address. Check that it is in fact
388     an IP address; just ignore it if it isn't
389     Have a first cut at determining what the prefix length
390     is from the route class. (CIDR? What's CIDR?)
391     */
392     if(i == RTAX_DST) {
393         if(sa->sa_family != AF_INET)
394             break;
395         sin = (struct sockaddr_in *) saptr;
396         ip = sin->sin_addr.s_addr;
397         if(ip == ntohl(INADDR_LOOPBACK))
398             break;
399         j = *(u_char *) &sin->sin_addr;
400         if(bits == -1) {
401             if(j >= 128) bits = 24;
402             else if(j >= 64) bits = 16;
403             else bits = 8;
404         }
405         valid = 1;
406     }
407
408     /*
409     Gateway IP address, if specified and it's IPv4.
410     */
411     else if(i == RTAX_GATEWAY && sa->sa_family == AF_INET) {
412         sin = (struct sockaddr_in *) saptr;
413         gw = sin->sin_addr.s_addr;
414     }
415
416     /*
417     The netmask. Why they couldn't just use a sockaddr_in
418     for this I don't know. See comments in addroute()
419     above for more info on this ugly thing.
420     */
421     else if(i == RTAX_NETMASK) {
422         bits = 0;
423         for(j = 4; j < sa->sa_len; j++) {
424             u = saptr[j];
425             u = (u & 0x55) + ((u >> 1) & 0x55);
426             u = (u & 0x33) + ((u >> 2) & 0x33);
427             u = (u & 0x0f) + ((u >> 4) & 0x0f);
```

```
428         bits += u;
429     }
430 }
431
432 /*
433 Interface specification. We're just interested in
434 the ifindex for ACL matching
435 */
436 else if(i == RTAX_IFP) {
437     sdl = (struct sockaddr_dl *) saptr;
438     ifindex = sdl->sdl_index;
439 }
440 saptr += SA_SIZE(saptr);
441 }
442 if(!valid) continue;
443
444 /*
445 Check the general sanity of the route
446 */
447 LOG8("Kernel route %s/%d -> %s flags=%x ifindex=%d",
448     formatip(ip), bits, formatip(gw), rtm->rtm_flags,
449     ifindex);
450 if(ip != (ip & maskbits[bits]))
451     continue;
452
453 /*
454 If we found a default route ...
455 If it's a learned route, and it doesn't point to the correct
456 default gateway, then delete it.
457 Otherwise note that we have a default route.
458 */
459 if(bits == 0) {
460     if(gw != defaultgw && (rtm->rtm_flags & routeflag)) {
461         LOG2("delete stale default route -> %s",
462             formatip(gw));
463         rtm->rtm_type = RTM_DELETE;
464         write(rtsocket, ptr, rtm->rtm_msglen);
465     }
466     else defgateway = gw;
467     continue;
468 }
469
470 /*
471 See if this route is one we (or a previous incarnation of we)
472 inserted; if so, see if it's in our routing table.
473 If it isn't, or the route should be a local one, then it's a
474 stray, so put it down humanely.
475 */
476 if(rtm->rtm_flags & routeflag) {
477     ipr = findroute(localroutes, ip, bits, 0);
478     if(!ipr || ipr->peer || gw != ipr->peer->nexthop) {
479         if(gwcost == -1)
480             continue;
481         LOG2("delete stale route %s/%d -> %s",
482             formatip(ip), bits, formatip(gw));
483         rtm->rtm_type = RTM_DELETE;
484         write(rtsocket, ptr, rtm->rtm_msglen);
485     }
486     else ipr->inuse = 1;
487     continue;
488 }
```

```
489
490 /*
491 Find the route in our routing table
492 If it's not there, add it
493 But, if it's to be overridden, delete it
494 */
495 ipr = findroute(localroutes, ip, bits, &pipr);
496 LOG8("Route %s/%d type=%d ifindex=%d %s",
497      formatip(ip), bits, rtype, ifindex,
498      ipr ? "found" : "new");
499 if(!ipr) {
500     ADDROUTEAFTER(ipr, localroutes, pipr)
501     LOG2("Added %s/%d local", formatip(ip), bits);
502     ipr->ip = ip;
503     ipr->bits = bits;
504 }
505 else if(rtm->rtm_flags & overrflag) {
506     if(gwcost == -1)
507         continue;
508     LOG2("Overriding %s/%d", formatip(ip), bits);
509     rtm->rtm_type = RTM_DELETE;
510     write(rtsocket, ptr, rtm->rtm_msglen);
511     continue;
512 }
513 /*
514 Here is where we fill out a route as an originated route
515 Prime the cost to 0, isGW flag to true, and GWcost to the
516 local gateway cost.
517 */
518 ipr->isgw = 1;
519 ipr->cost = 0;
520 ipr->gwcost = gwcost;
521 ipr->rtype = rtype;
522 ipr->peer = 0;
523 ipr->ifindex = ifindex;
524 ipr->inuse = 1;
525 }
526
527 /*
528 For each route not found in the kernel, add it to the kernel table
529 */
530 FOREACH(ipr, localroutes) if(!ipr->inuse)
531     addroute(ipr->ip, ipr->bits, ipr->peer->nexthop);
532
533 /*
534 If this is a gateway, set the gateway flag
535 Do this if we actually have a default gateway, or if the "always
536 gateway" flag is set.
537 */
538 if(isgateway) {
539     if(defgateway || gwalways)
540         gwcost = 0;
541     else gwcost = -1;
542 }
543
544 /*
545 Otherwise, if no default route was found (or one was found that
546 didn't point to the correct peer, and was therefore deleted),
547 set the gateway to the address of the gateway peer.
548 If there is no gateway peer, set the default to the configured
549 default (if any).
```

```
550 If we set a gateway, put it into the routing table.
551 If the result of this is a change, force a PATH update to be
552 sent to all peers.
553 */
554 else if(!defgateway) {
555     if(gwpeer)
556         defgateway = gwpeer->nexthop;
557     else if(defaultroute)
558         defgateway = defaultroute;
559     if(defgateway)
560         addroute(0, 0, defgateway);
561     if(defgateway != oldgw)
562         FOREACH(peer, peers)
563             peer->pathreq = 1;
564 }
565
566 LOG4("Local table:");
567 if(debug >= 4) dumproutes(localroutes);
568
569 /*
570 If there are no routes, purge the announced routes and we're done
571 */
572 if(gwcost == -1) {
573     LOG4("Gateway not reachable, not announcing routes");
574     FOREACH(peer, peers) {
575         KILLROUTES(peer->annrts);
576         peer->nextrt = 0;
577         peer->>nullrt = 0;
578     }
579     return 0;
580 }
581
582 /*
583 Search the routing table for local routes which are the
584 subject of a peer relationship.
585 For these routes, we need to pick winners & losers; winners get
586 to announce the route to their peers.
587 The winner is the peer with the lowest gwcost; failing that the
588 lower IP address.
589 Note that if the remote peer is down, we count ourselves as
590 a loser. This deals with peer/listen relationships where the
591 peer disappears when it goes down.
592 Reminder: peer->localannounce: 0 = pick winner; 1 = always announce;
593          2 = never announce.
594 */
595 FOREACH(ipr, localroutes) {
596     ipr->inuse = 1;
597     FOREACH(peer, peers) {
598         if(ipr->ip == peer->localroute &&
599            ipr->bits == peer->localbits &&
600            peer->localannounce != 1 &&
601            (peer->localannounce == 2 ||
602             peer->gwcost < gwcost ||
603             (peer->gwcost == gwcost &&
604              ntohl(LOCIP(peer)) > ntohl(REMIP(peer)))) {
605             LOG8("Route %s/%d is local loser",
606                  formatip(ipr->ip), ipr->bits);
607             ipr->inuse = 0;
608         }
609     }
610 }
```

```
611
612 /*
613 Prepare the peer announcements
614 Ignore peers that are down
615 */
616 FOREACH(peer, peers) {
617     if(peer->gwcost == -1)
618         continue;
619     LOG4("Routing for %s", formatip(REMIP(peer)));
620
621     annh = annt = 0;
622     doann = 0;
623     if(peer->nexttrt || peer->>nullrt)
624         doann = 1;
625     oldr = peer->annrts;
626     FOREACH(ipr, localroutes) {
627         /*
628         Don't announce losing routes
629         */
630         if(!ipr->inuse)
631             continue;
632
633         /*
634         Split horizon, allowing for multiple links
635         Don't advertise routes back to where we got them from
636         Don't advertise local routes up the interface
637         they point at.
638         ??? May need to restrict this to just iface routes
639         */
640         if(ipr->peer && ipr->peer->routerid == peer->routerid)
641             continue;
642         if(!ipr->peer && ipr->ifindex == peer->ifindex)
643             continue;
644
645         /*
646         Check announcement ACL, skip if not permitted
647         */
648         if(peer->announce && !checkacl(peer->announce,
649             ipr->ip, ipr->bits,
650             ipr->rtype, ipr->ifindex))
651             continue;
652
653         /*
654         See if route would reach "over the horizon"
655         That is, if it would be shorter to route via
656         our gateway and back via the route originator's
657         gateway, then don't bother advertising the route
658         */
659         LOG8("%s/%d rc=%d + lc=%d"
660             " < rgc=%d + pgc=%d + isgw=%d",
661             formatip(ipr->ip), ipr->bits,
662             ipr->cost, peer->cost, ipr->gwcost,
663             peer->gwcost, (ipr->isgw && peer == gwpeer));
664         if(ipr->cost + peer->cost >=
665             ipr->gwcost + peer->gwcost
666             + (ipr->isgw && peer == gwpeer))
667             continue;
668
669         /*
670         Add the route to the announcement
671         */
```

```
672     ADDROUTEAFTER(r, annh, annt)
673     r->ip = ipr->ip;
674     r->bits = ipr->bits;
675     r->isgw = (ipr->isgw && (peer == gwpeer));
676     r->cost = ipr->cost + peer->cost;
677     r->gwcost = ipr->gwcost;
678     r->rtype = 0;
679     r->ifindex = 0;
680     r->peer = 0;
681     annt = r;
682
683     LOG8("new: %s/%d gw=%d cost=%d gwcost=%d rtype=%d",
684         formatip(r->ip), r->bits, r->isgw,
685         r->cost, r->gwcost, r->rtype);
686
687     /*
688     If this route is different to the corresponding
689     previous announcement, set the "do announcement"
690     flag.
691     */
692     if(!doann) {
693         if(!oldr || oldr->ip != r->ip
694             || oldr->bits != r->bits
695             || oldr->isgw != r->isgw
696             || oldr->cost != r->cost
697             || oldr->gwcost != r->gwcost
698             || oldr->rtype != r->rtype)
699             doann = 1;
700         if(oldr) oldr = oldr->next;
701     }
702 }
703
704 /*
705 If there are any unchecked previous announcement records
706 left, then the announcement has changed and we need to
707 send it out.
708 */
709 if(oldr) doann = 1;
710
711 /*
712 If the routing table has changed, force an update
713 If there are no routes to be announced to this peer,
714 set the nullrt flag to force a null route announcement
715 If the announcement hasn't changed, just blow the newly
716 constructed announcement away.
717 */
718 LOG4("Announcements for %s: (update %srequired)",
719     formatip(REMIP(peer)), doann ? "" : "not ");
720 if(debug >= 4) dumproutes(annh);
721 if(doann) {
722     KILLROUTES(peer->annrts)
723     peer->annrts = annh;
724     peer->nexttrt = annh;
725     if(annh)
726         peer->>nullrt = 0;
727     else peer->>nullrt = 1;
728 }
729 else KILLROUTES(annh)
730 }
731 return 0;
732 }
```

---

733  
734 /\* EOF \*/  
735