



OSPF with Zebra

*Dynamic Routing on the ROSPA Meta-network
Presented Dec 11/2002 to LOSURS*

Tillman Hodgson

tillman@seekingfire.com

ROSPA





INTRODUCTION



What's this presentation about?

- I'm here to talk about routing with OSPF using Gnu Zebra, an open source routing daemon with an interface strikingly familiar to that of a Cisco router
- The drive behind this is that ROSPA is running a VPN Meta-Network project using Zebra (so I'm also here to talk about Rospa)

What's the VPN Meta-Network project?

- The VPN Meta-Network is a complex multi-subnet IP network, layered over top of the Internet
- It's created using VPN-type technologies, such as OpenVPN and IPsec
- Its purpose is to act as a sandbox for learning VPN technologies on a real live network
- It's getting large enough that the static routing is too cumbersome – we need dynamic routing!

What's the Dynamic Routing project?

- The project's purpose is to allow members to experiment with dynamic routing technologies on a live network
- The project uses Zebra, which is the replacement for the traditional open source Unix routing daemon Gated, which has “gone commercial”
- The project is “live” – this is a real network with real services that really needs dynamic routing



ROUTING BACKGROUND



Terms

- Routing is the act of passing packets from one network to another
- IP Networks are defined by network masks (layer 3), not necessarily by physical cabling (layer 2)
- Routers are computers with access to two or more networks
- Three major types of routing: static, vector and link-state

What is static routing?

- The type of routing most people are familiar with — their default gateways
- Doesn't require a special routing protocol, directly manipulates the routing table
- Simply says "To get to network X, send packets to router Y"

What is vector routing?

- Hop-count based
- Each router knows the state of only its immediate neighbours
- Each router advertises its complete routing table
- Each router builds its own routing table based on what its neighbours have told it (routing by rumour)
- RIPv2 is probably the most common vector routing protocol

What is link-state routing?

- Richer metrics for describing the cost of a link
- Each router maintains a picture of the complete network
- Each router advertises only what has changed
- Each router builds its own routing table based on that complete information
- OSPF is probably the most common link-state protocol

How does link-state work?

- Upon initialization or due to any change, a router will generate an advertisement of all its links
- Routers exchange link states by flooding:
 - Each router stores a copy in its own link-state database
 - It then forwards the advertisement to its own neighbours
- After the database of a router is updated, it calculates the Shortest Path Tree to all destinations and builds a routing table
- If no changes take place that require an advertisement, OSPF is very quiet

Why use link-state instead of vector?

- Can handle larger internetworks
- Quick convergence
- Less bandwidth overhead
- Can describe links better and make better routing choices

What is OSPF?

- “Open Shortest Path First”
- An widely implemented open standard for link-state routing
- Developed in 1988, formalized in 1991
- OSPF Version 2: RFC 2328

-
-
-



ZEBRA



-
-
-
-
-
-
-
-
-
-

What is GNU Zebra?

- GPL'ed multi-protocol routing software
- Developed by Kunihiro Ishiguro, a member of JPNIC
- Supports BGPv4, OSPFv2, RIPv1 & RIPv2
- Supports both IPv4 and IPv6
- Has a separate process (daemon) for each protocol
- <http://www.zebra.org>

Commercialism

- The need for Zebra arose when development of Gated, the traditional Unix routing daemon, went closed-source (under NextHop technologies)
- Zebra does have a commercial version available (from IP Infusion Inc)
- GPL'ed version of Zebra is usable and maintained software

The Zebra CLI

- 99% “Cisco-like”
- You can even use standard Cisco config files!
- To access it, you telnet to localhost – each process has its own port
- Has full readline support – command history, nice editing, etc.

The ROSPA Meta-Network

- Presenting only a subset of it:
 - An OpenVPN tunnel from Coyote to Scott
 - Another OpenVPN tunnel from Coyote to Scott via his second cable modem
 - An OpenVPN tunnel from Scott to another Scott network
Scott has too many networks :-)
- Note that each tunnel has its own P2P “shim” network

zebra.conf

```
hostname coyote.seekingfire.prv
password 8 dVTSWeTazpcQ
log file /var/log/zebra.log
log record-priority
service advanced-vty
service password-encryption
!
ip route 224.0.0.9/32 127.0.0.1
!
access-list term permit 127.0.0.1/32
access-list term deny any
!
line vty
  access-class term
```

ospfd.conf

```
router ospf
  ospf router-id 192.168.23.2
  auto-cost reference-bandwidth 1000
  passive-interface tunl
  passive-interface r1l
  network 10.1.2.2/32 area 0
  network 10.1.3.2/32 area 0
  network 10.1.3.6/32 area 0
  network 192.168.23.0/24 area 0
  distribute-list private out connected
!
access-list private permit 10.0.0.0/8
access-list private permit 192.168.0.0/16
access-list private deny any
!
line vty
  access-class term
```

Reading the Host Routing Table

Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
B - BGP, > - selected route, * - FIB route

```
K>* 0.0.0.0/0 via 24.72.10.1, r11
O>* 10.1.3.1/32 [110/110] via 10.1.3.6, tun4, 05:14:06
*
* via 10.1.3.2, tun3, 05:14:06
O 10.1.3.2/32 [110/100] is directly connected, tun3, 3d21h08m
C>* 10.1.3.2/32 is directly connected, tun3
O>* 10.1.3.5/32 [110/110] via 10.1.3.6, tun4, 05:14:06
*
* via 10.1.3.2, tun3, 05:14:06
O 10.1.3.6/32 [110/100] is directly connected, tun4, 3d21h08m
C>* 10.1.3.6/32 is directly connected, tun4
O>* 10.42.1.0/24 [110/110] via 10.1.3.6, tun4, 05:14:06
*
* via 10.1.3.2, tun3, 05:14:06
O>* 10.42.10.0/24 [110/120] via 10.1.3.6, tun4, 05:14:06
*
* via 10.1.3.2, tun3, 05:14:06
O 192.168.23.0/24 [110/100] is directly connected, r10, 3d21h08m
C>* 192.168.23.0/24 is directly connected, r10
```

Reading the Host Routing Table

Destination	Gateway	Flags	Refs	Use	Netif
default	24.72.10.1	UGSc	26	4643770	rl1
10.1.2.1	10.1.2.2	UGH1	0	0	tun2
10.1.2.2	10.1.2.1	UH	2	0	tun2
10.1.3.1	10.1.3.6	UGH1	0	0	tun4
10.1.3.2	10.1.3.1	UH	0	0	tun3
10.1.3.5	10.1.3.6	UGH1	0	0	tun4
10.1.3.6	10.1.3.5	UH	8	0	tun4
10.42.1/24	10.1.3.6	UG1c	0	0	tun4
10.42.10/24	10.1.3.6	UG1c	0	0	tun4
24.72.10/24	link#2	UC	3	0	rl1
24.72.10.1	00:07:85:08:3d:4a	UHLW	27	7	rl1
24.72.10.95	00:80:c8:d6:e2:26	UHLW	1	43574	rl1
24.72.10.212	00:50:bf:e1:f2:b7	UHLW	0	12	lo0
192.168.8	10.1.2.2	UG1c	0	6	tun2
192.168.23	link#1	UC	5	0	rl0
224.0.0.9	127.0.0.1	UGH1	0	0	lo0

Demosntrating the routing in action

- Because Scott and Tillman have redundant tunnels, one of Scott's cable modems can stop working and OSPF transparently routes around it
- Because Scott has a second network connected, Tillman can route *through* Scott's primary network to get to it
- Because OSPF handles all the routing, adding additional networks is painless and automatic

-
-
-



QUESTIONS?



-
-
-
-
-
-
-
-
-
-



ROSPA & GETTING INVOLVED



About ROSPA

- ROSPA is the Regina Open Source Professionals Association
- Formed in the summer of 2002
- A volunteer-driven peer support organization
- Focused on supporting folks who work with open source/Free software tools and open standards in a professional capacity

ROSPA Partnerships

- ROSPA and LOSURS are companion groups
 - Non-competitive – different goals
- A member of the Regina Chamber of Commerce
- Working towards being an official affiliate of USENIX & SAGE

What does ROSPA do?

- Advocacy on behalf of its members and their craft
- Meetings and technical presentations
- A forum for IT professionals working with open source tools
- Knowledge sharing & peer support
- Participating in events like ROSE
- Research & projects

About ROSPA Projects

This is where the cool hands-on stuff happens:

- Projects run on top of the Meta-Network "sandbox":
 - The Dynamic Routing project
 - The Kerberos project
 - Developing ROSPA services
- Other projects include:
 - Inter-group communications

About ROSPA Projects (cont ...)

- Ideas for new projects include:
 - Compiling a local “success stories” and case study database
 - Creating a local Certificate Authority root for use with public-key encryption
 - Writing articles for *;login:*, *Linux Journal*, *SysAdmin* and other industry magazines
 - Mentorship program (perhaps by tapping into the SAGE program)

Getting Involved ...

- Visit us at `http://www.rosipa.ca`
- Sign up for some of the mailing lists
- Join one of the projects (or start your own)
- Email `info@rosipa.ca` for more information

Colophon

This presentation was generated with $\text{\LaTeX}2\text{e}$ and the Prosper \LaTeX package

- All the normal advantages of \LaTeX
- Use any PDF viewer – cross-platform presentations!
- PDF special effects like fades, wipes and sparkles

`http://prosper.sourceforge.net/`