Zebra MCR Software 1.0 Manual

Yin Wang
Networking Laboratory of Helsinki University of Technology
yinwang@netlab.hut.fi

Please note: operating system requirement: Linux 2.4.20-30.9 (at least)

1. Configure the Software

Zebra has an excellent configure script which automatically detects most host configurations. There are several additional configure options you can use to turn off IPv6 support, to disable the compilation of specific daemons, and to enable SNMP support.

--enable-guile
: Turn on compilation of the zebra-guile interpreter. You will need the guile library to make this. zebra-guile implementation is not yet finished. So this option is only useful for zebra-guile developers.

--disable-ipv6
: Turn off IPv6 related features and daemons. Zebra configure script automatically detects IPv6 stack. But sometimes you might want to disable IPv6 support of Zebra.

--disable-zebra
: Do not build zebra daemon.

--disable-ripd
: Do not build ripd.

--disable-ripngd
: Do not build ripngd.

--disable-ospfd
: Do not build ospfd.

--disable-ospf6d
: Do not build ospf6d.

--disable-bgpd
: Do not build bgpd.

--disable-bgp-announce
: Make bgpd which does not make bgp announcements at all. This feature is good for using bgpd as a BGP announcement listener.

--enable-netlink
: Force to enable GNU/Linux netlink interface. Zebra configure script detects netlink interface by checking a header file. When the header file does not match to the current running kernel, configure script will not turn on netlink support.

--enable-snmp
: Enable SNMP support. By default, SNMP support is disabled.

You may specify any combination of the above options to the configure script. By default, the executables are placed in /usr/local/sbin and the configuration files in /usr/local/etc. The /usr/local/ installation prefix and other directories may be changed

using the following options to the configuration script.

--prefix=*prefix*
Install architecture-independent files in *prefix* [/usr/local].
--sysconfdir=*dir*
Read-only sample configuration file in *dir* [*prefix*/etc].

```
% ./configure --disable-ipv6
```

This command will configure zebra and the routing daemons.

[Reference]:

http://www.zebra.org/zebra/Configure-the-Software.html#Configure%20the%20Software

2. Build the Software

After configuring the software, you will need to compile it for your system. Simply issue the command make in the root of the source directory and the software will be compiled. If you have *any* problems at this stage, be certain to send a bug report See Bug Reports.

```
% ./configure
.
.
.
./configure output
.
.
.
% make
```

[Reference]

http://www.zebra.org/zebra/Build-the-Software.html#Build%20the%20Software

3. Install the Software

Installing the software to your system consists of copying the compiled programs and supporting files to a standard location. After the installation process has completed, these files have been copied from your work directory to /usr/local/bin, and /usr/local/etc.

To install the Zebra suite, issue the following command at your shell prompt: make install.

```
%
% make install
```

%

Zebra daemons have their own terminal interface or VTY. After installation, you have to setup each beast's port number to connect to them. Please add the following entries to /etc/services.

```
zebrasrv      2600/tcp            # zebra service
zebra         2601/tcp            # zebra vty
ripd          2602/tcp            # RIPd vty
ripngd        2603/tcp            # RIPngd vty
ospfd         2604/tcp            # OSPFd vty
bgpd          2605/tcp            # BGPd vty
ospf6d        2606/tcp            # OSPF6d vty
```

If you use a FreeBSD newer than 2.2.8, the above entries are already added to /etc/services so there is no need to add it. If you specify a port number when starting the daemon, these entries may not be needed.

You may need to make changes to the config files in /usr/local/etc/*.conf. See Config Commands.

[Reference]

http://www.zebra.org/zebra/Install-the-Software.html#Install%20the%20Software

4. Installation Example:

./configure –-disable-ipv6 –-disable-ripd –-disable-ripngd –-disable-bgpd –-disable-ospf6d –-disable-bgp-announce –-enable-netlink –-disable-snmp –-disable-nssa
make
make install

5. Configuration files

The following configuration files must be configured before running the programs.

#vi /usr/local/etc/zebra.conf

#vi /usr/local/etc/ospfd.conf

6. Running Programs

/usr/local/sbin/zebra &

/usr/local/sbin/ospfd &

7. VTY commmands

For zebra router:

telnet localhost 2601

For ospfd:

telnet localhost 2604

The detailed commands are described as follow part.

MCR OSPF Commands

OSPF router

To start OSPF process you have to specify OSPF router. As of this writing ospfd does not support multiple OSPF process

Command: router ospf
Command: no router ospf
Enable or disable the OSPF process. Ospfd does not yet support multiple OSPF processes. So you can not specify an OSPF process number.

OSPF command: ospf router-id x.x.x.x
OSPF command: no ospf router-id

OSPF Command: ospf abr-type 1
OSPF Command: no ospf abr-type 1
For MCR ospfd, only standard ABR type is supported.

OSPF Command: opassive interface INTERFACE
OSPF Command: no passive interface INTERFACE

OSPF Command: timers spf <0-4294967275> < 0- 4294967295 >
OSPF Command: no timers spf

OSPF Command: timers bsp < 0-4294967275 > < 0-4294967275 >
OSPF Command: no timers bsp

Set the parameters of the second routing table calculation timers.

OSPF Command: refresh group-limit <0-10000>
OSPF Command: refresh per-slice <0-10000>
OSPF Command: refresh age-diff <0-10000>

OSPF Command: auto-cost refrence-bandwidth <1-4292967>
OSPF Command: no auto-cost refrence-bandwidth

OSPF Command: network x.x.x.x/M area x.x.x.x
OSPF Command: network x.x.x.x/M area <0-4294967275>
OSPF Command: no network x.x.x.x/M area x.x.x.x
OSPF Command: no network x.x.x.x/M area <0-4294967275>

This command specifies the OSPF enabled interfac. If the interface has an address of 10.0.0.1/8, then the command below provides network information to the ospf beasites.

router ospf
network 10.0.0.0/8 area 0

The netwrok command's mask length should be the same as the interface address's mask.

OSPF area

OSPF Command: area x.x.x.x range x.x.x.x/M
OSPF Command: area < 0-4294967275 > range x.x.x.x/M
OSPF Command: no area x.x.x.x range x.x.x.x/M
OSPF Command: no area < 0-4294967275 > range x.x.x.x/M

OSPF Command: area x.x.x.x second-range x.x.x.x/M
OSPF Command: area < 0-4294967275 > second-range x.x.x.x/M
OSPF Command: area x.x.x.x second-range x.x.x.x/M advertise
OSPF Command: area < 0-4294967275 > second-range x.x.x.x/M advertise
OSPF Command: area x.x.x.x second-range x.x.x.x/M cost <x>
OSPF Command: area < 0-4294967275 > second-range x.x.x.x/M cost <x>
OSPF Command: area x.x.x.x second-range x.x.x.x/M substitute y.y.y.y/N
OSPF Command: area < 0-4294967275 > second-range x.x.x.x/M substitute y.y.y.y/N
OSPF Command: no area x.x.x.x second-range x.x.x.x/M
OSPF Command: no area < 0-4294967275 > second-range x.x.x.x/M
OSPF Command: no area x.x.x.x second-range x.x.x.x/M cost <x>
OSPF Command: no area < 0-4294967275 > second-range x.x.x.x/M cost <x>
OSPF Command: no area x.x.x.x second-range x.x.x.x/M advertise
OSPF Command: no area < 0-4294967275 > second-range x.x.x.x/M advertise
OSPF Command: no area x.x.x.x second-range x.x.x.x/M substitute y.y.y.y/N
OSPF Command: no area < 0-4294967275 > second-range x.x.x.x/M substitute

y.y.y.y/N
Set the aggregated range of IP addresses. The range is used for the route summarization of the second routing table

OSPF Command: area x.x.x.x virtual-link x.x.x.x
OSPF Command: area < 0-4294967275 > virtual x.x.x.x
OSPF Command: no area x.x.x.x virtual-link x.x.x.x
OSPF Command: no area < 0-4294967275 > virtual x.x.x.x

OSPF interface

Interface Command: ip ospf cost <0-65535>
Interface Command: no ip ospf cost

Interface Command: ip ospf bandwidth <1-10000000 kbps>
Interface Command: no ip ospf bandwidth
Set the output bandwidth of an interface. The input value will be calculated by the following equation.
72 TOS metric in Router-LSA = INPUT VALUE / 152.59 kbps

Interface Command: ip ospf dead-interval <1-65535>
Interface Command: no ip ospf dead-interval
Set number of seconds for RouterDeadInterval timer value used for Wait Timer and Inactivity Timer. This value must be the same for all routers attached to a common netwrok. The default value is 40 seconds.

Interface Command: ip ospf hello-interval <1-65535>
Interface Command: no ip ospf hello-interval
Set number of seconds for HellInterval timer value. Setting this value, Hello packet will be sent every timer value seconds on specified interface. This value must be the same for all routers attached to a common network. The default value is 10 seconds.

Interface Command: ip ospf netwrok (broadcast | non-broadcast | point-to-multipoint | point-to-point)
Interface Command: no ip ospf network
Set explicitly network type for specified interface.
Interface Command: ip ospf priority <0-255>
Interface Command: no ip ospf prioirty
Set RouterPriority integer value. Setting higher value, router will be more eligible to become Designated Router. Setting the value to 0, router is no longer eligible to Designated Router. The default value is 1.

Interface Command: ip ospf retransmit-interval <1-65535>
Interface Command: no ip ospf retransmi-interval

Set number of seconds for RxmtInterval timer value. This value is used when retransmitting Database Description and Link State Request packets. The default value is 5 seconds.

Interface Command: ip ospf transmit-delay
Interface Command: no ip ospf transmit-delay
Set number of seconds for InfTransDelay value. LSAs' age should be incremented by this value when transmitting. The default value is 1 seconds.

Redistribute routes to OSPF

The following commands are modified in order to support MCR capability.

OSPF Command: redistribute (kernel | connected | static | rip | bgp) metric <0-16777214> metric-bandwidth <1-10000000 kbps >
OSPF Command: redistribute (kernel | connected | static | rip | bgp) metric <0-16777214> metric-bandwidth <1-10000000 kbps > metric-type (1 | 2) route-map WORD
OSPF Command: redistribute (kernel | connected | static | rip | bgp) metric <0-16777214> metric-bandwidth <1-10000000 kbps > metric-type (1 | 2)
OSPF Command: redistribute (kernel | connected | static | rip | bgp) metric-type (1 | 2) metric <0-16777214> metric-bandwidth <1-10000000 kbps >
OSPF Command: redistribute (kernel | connected | static | rip | bgp) metric-type (1 | 2) metric <0-16777214> metric-bandwidth <1-10000000 kbps > route-map WORD
OSPF Command: redistribute (kernel | connected | static | rip | bgp) metric-type (1 | 2)
OSPF Command: redistribute (kernel | connected | static | rip | bgp) metric-type (1 | 2) route-map WORD
OSPF Command: redistribute (kernel | connected | static | rip | bgp)
OSPF Command: redistribute (kernel | connected | static | rip | bgp) route-map WORD
OSPF Command: redistribute (kernel | connected | static | rip | bgp) metric <0-16777214> metric-bandwidth <1-10000000 kbps > route-map WORD
OSPF Command: no redistribute (kernel | connected | static | rip | bgp)

OSPF Command: default-information originate metric <0-16777214> metric-bandwidth < 1- 10000000 kbps > metric-type (1 | 2) route-map WORD
OSPF Command: default-information originate metric <0-16777214> metric-bandwidth < 1- 10000000 kbps > metric-type (1 | 2)
OSPF Command: default-information originate metric <0-16777214> metric-bandwidth < 1- 10000000 kbps >
OSPF Command: default-information originate
OSPF Command: default-information originate metric <0-16777214> metric-bandwidth < 1- 10000000 kbps > route-map WORD

OSPF Command: default-information originate metric-type (1 | 2) metric <0-16777214> metric-bandwidth < 1- 10000000 kbps > route-map WORD
OSPF Command: default-information originate metric-type (1 | 2) metric <0-16777214> metric-bandwidth < 1- 10000000 kbps >
OSPF Command: default-information originate metric-type (1 | 2)
OSPF Command: default-information originate metric-type (1 | 2) route-map WORD
OSPF Command: default-information originate route-map WORD

OSPF Command: default-information originate always metric <0-16777214> metric-bandwidth < 1- 10000000 kbps > metric-type (1 | 2) route-map WORD
OSPF Command: default-information originate always metric <0-16777214> metric-bandwidth < 1- 10000000 kbps > metric-type (1 | 2)
OSPF Command: default-information originate always metric <0-16777214> metric-bandwidth < 1- 10000000 kbps >
OSPF Command: default-information originate always
OSPF Command: default-information originate always metric <0-16777214> metric-bandwidth < 1- 10000000 kbps > route-map WORD
OSPF Command: default-information originate always metric-type (1 | 2) metric <0-16777214> metric-bandwidth < 1- 10000000 kbps > route-map WORD
OSPF Command: default-information originate always metric-type (1 | 2) metric <0-16777214> metric-bandwidth < 1- 10000000 kbps >
OSPF Command: default-information originate always metric-type (1 | 2)
OSPF Command: default-information originate always metric-type (1 | 2) route-map WORD
OSPF Command: default-information originate always route-map WORD
OSPF Command: no default-information originate

OSPF Command: default-metric <1-16777214> default-bandwidth-metric <1-10000000 kbps >
OSPF Command: no default-metric

Showing OSPF information

The display of the following commands is changed in the terms of the characteristics of the OSPF with MCR extensions.

Command: show ip ospf
Command: show ip ospf interface [INTERFACE]
Command: show ip ospf neighbor
Command: show ip ospf neighbor detail
Command: show ip ospf database
Command: show ip ospf database (asbr-summary | external | network | router | summary | bsummary | basbr)
"bsummary" is used for Type 12 Summary-LSA (the second routing table)
"basbr" is used for Type 13 Summary-LSA (the second routing table)

Command: show ip ospf route

8. Example

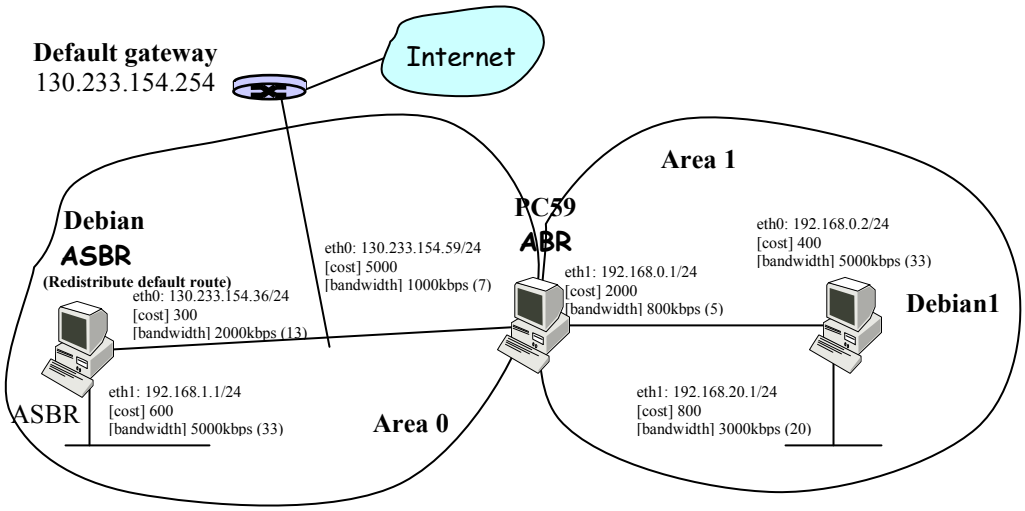

Router Debian1 configuration files

1) /usr/local/etc/zebra.conf

```
! -*- bebian1 zebra -*-
!
! zebra sample configuration file
!
! $Id: zebra.conf.sample,v 1.14 1999/02/19 17:26:38 developer Exp $
!
hostname debian1_Router
password zebra
! enable password zebra
interface eth0
  multicast
  description debian1eth0-to-PC59eth1
interface eth1
  multicast
  description debian1eth1
!
!log file zebra.log
~
```

2) /usr/local/etc/ospfd.conf

```
! -*- debian1 ospf -*-
!
```

```
! OSPFd sample configuration file
!
hostname debian1_ospfd
password zebra
!enable password please-set-at-here
!
interface eth0
  ip ospf cost 400
  ip ospf bandwidth 5000
interface eth1
  ip ospf cost 800
  ip ospf bandwidth 3000
router ospf
    router-id 192.168.20.1
    network 192.168.0.0/24 area 1
    network 192.168.20.0/24 area 1
!
log stdout
```

PC59 Configuration files

1) /usr/local/etc/zebra.conf

```
! -*- zebra -*-
!
! zebra sample configuration file
!
! $Id: zebra.conf.sample,v 1.14 1999/02/19 17:26:38 developer Exp $
!
hostname PC59_Router
password zebra
!enable password zebra
interface eth0
 multicast
 description My interface
interface eth1
 multicast
 description My interface 2
!log file zebra.log
```

2) /usr/local/etc/ospfd.conf

```
! -*- ospf -*-
!
! OSPFd sample configuration file
!
hostname PC59_ospfd
```

```
password zebra
!enable password please-set-at-here
!
interface eth0
   ip ospf cost 5000
   ip ospf bandwidth 1000
interface eth1
   ip ospf cost 2000
   ip ospf bandwidth 800
router ospf
  ospf router-id 192.168.0.1
  ospf abr-type standard
  network 130.233.154.0/24 area 0
  network 192.168.0.0/24 area 1
!  area 0 range 130.233.0.0/16 cost 150
!  area 1 second-range 192.168.0.0/16 cost 1234
!
log stdout
```

Debian Configuration files

1) /usr/local/etc/zebra.conf

```
! -*- debian zebra -*-
!
! zebra sample configuration file
!
! $Id: zebra.conf.sample,v 1.14 1999/02/19 17:26:38 developer Exp $
!
hostname debian_Router
password zebra
! enable password zebra
interface eth0
  multicast
  description debian_eth0-to-pc59_eth0
interface eth1
  multicast
  description debian_eth1
!
!log file zebra.log
~
```

2) /usr/local/etc/ospfd.conf! -*- debian ospf -*-

```
!
! OSPFd sample configuration file
!
hostname debian_ospfd
```

```
password zebra
!enable password please-set-at-here
!
interface eth0
  ip ospf cost 300
  ip ospf bandwidth 2000
interface eth1
  ip ospf cost 600
  ip ospf bandwidth 5000
router ospf
  ospf router-id 172.16.1.1
  network 130.233.154.0/24 area 0
  network 172.16.1.0/16 area 0
  redistribute static
  default-information originate metric 300 metric-bandwidth 100000 metric-type
2
!log stdout
```

Routing tables

- PC59' routing table:

SP routing table

```
PC59_ospfd# show ip ospf route
============ OSPF network routing table ============
N    130.233.154.0/24      [5000] area: 0.0.0.0
                           directly attached to eth0
N    192.168.0.0/24        [2000] area: 0.0.0.1
                           directly attached to eth1
N    192.168.1.0/24        [5600] area: 0.0.0.0
                           via 130.233.154.36, eth0
N    192.168.20.0/24       [2800] area: 0.0.0.1
                           via 192.168.0.2, eth1

============ OSPF router routing table =============
R    192.168.1.1           [5000] area: 0.0.0.0, ASBR
                           via 130.233.154.36, eth0

============ OSPF external routing table ===========
N E2 0.0.0.0/0             [5000/300] tag: 0
                           via 130.233.154.254, eth0

============ OSPF BSP network routing table ============
N    130.233.154.0/24      [9362] area: 0.0.0.0
                           directly attached to eth0
N    192.168.0.0/24        [13107] area: 0.0.0.1
                           directly attached to eth1
N    192.168.1.0/24        [11348] area: 0.0.0.0
                           via 130.233.154.36, eth0
N    192.168.20.0/24       [16384] area: 0.0.0.1
                           via 192.168.0.2, eth1

============ OSPF BSP router routing table =============
R    192.168.1.1           [9362] area: 0.0.0.0, ASBR
                           via 130.233.154.36, eth0

============ OSPF BSP external routing table ===========
N E2 0.0.0.0/0             [9362/168] tag: 0
                           via 130.233.154.254, eth0
```

BSP routing table

Inter-area routes do not exist in the PC59's SP routing table and BSP routing table because the PC59 is a unique ABR in the test environment.

- Debian's routing table

```
debian_ospfd# show ip ospf route
============ OSPF network routing table ============
N    130.233.154.0/24      [300] area: 0.0.0.0
                           directly attached to eth0
N IA 192.168.0.0/24        [2300] area: 0.0.0.0
                           via 130.233.154.59, eth0
N    192.168.1.0/24        [600] area: 0.0.0.0
                           directly attached to eth1
N IA 192.168.20.0/24       [3100] area: 0.0.0.0
                           via 130.233.154.59, eth0

============ OSPF router routing table =============
R    192.168.0.1           [300] area: 0.0.0.0, ABR
                           via 130.233.154.59, eth0

============ OSPF external routing table ===========

============ OSPF BSP network routing table ============
N    130.233.154.0/24      [5041] area: 0.0.0.0
                           directly attached to eth0
N IA 192.168.0.0/24        [18148] area: 0.0.0.0
                           via 130.233.154.59, eth0
N    192.168.1.0/24        [1986] area: 0.0.0.0
                           directly attached to eth1
N IA 192.168.20.0/24       [21425] area: 0.0.0.0
                           via 130.233.154.59, eth0

============ OSPF BSP router routing table =============
R    192.168.0.1           [5041] area: 0.0.0.0, ABR
                           via 130.233.154.59, eth0

============ OSPF BSP external routing table ===========
debian_ospfd#
```

SP routing table

BSP routing table

In the Debian's BSP routing table and SP routing table, there are no external routes because the external routing information is redistributed by the Debian itself.

- Debian1's routing table

```
debian1_ospfd# show ip ospf route
============ OSPF network routing table ============
N IA 130.233.154.0/24      [5400] area: 0.0.0.1
                           via 192.168.0.1, eth0
N    192.168.0.0/24        [400] area: 0.0.0.1
                           directly attached to eth0
N IA 192.168.1.0/24        [6000] area: 0.0.0.1
                           via 192.168.0.1, eth0
N    192.168.20.0/24       [800] area: 0.0.0.1
                           directly attached to eth1

============ OSPF router routing table =============
R    192.168.0.1           [400] area: 0.0.0.1, ABR
                           via 192.168.0.1, eth0
R    192.168.1.1        IA [5400] area: 0.0.0.1, ASBR
                           via 192.168.0.1, eth0

============ OSPF external routing table ===========
N E2 0.0.0.0/0             [5400/300] tag: 0
                           via 192.168.0.1, eth0

============ OSPF BSP network routing table ============
N IA 130.233.154.0/24      [11348] area: 0.0.0.1
                           via 192.168.0.1, eth0
```

SP routing table

BSP routing table

```
N    192.168.0.0/24        [1986] area: 0.0.0.1
                           directly attached to eth0
N IA 192.168.1.0/24        [13334] area: 0.0.0.1
                           via 192.168.0.1, eth0
N    192.168.20.0/24       [3277] area: 0.0.0.1
                           directly attached to eth1

============ OSPF BSP router routing table =============
R    192.168.0.1           [1986] area: 0.0.0.1, ABR
                           via 192.168.0.1, eth0
R    192.168.1.1        IA [11348] area: 0.0.0.1, ASBR
                           via 192.168.0.1, eth0

============ OSPF BSP external routing table ===========
N E2 0.0.0.0/0             [11348/168] tag: 0
                           via 192.168.0.1, eth0
```