Linux Mobile IPv6 HOWTO

Lars Strand

<lars (at) unik no>

2004−04−20

Revision History

Revision 1.2	2004−04−20	Revised by: LKS
Updated links. Changed lisence back to GFDL		
Revision 1.1	2004−02−04	Revised by: LKS
Added "Travelling through several foregin LAN's" and "Returning home". Some cleanup and restructuring.		
Revision 1.0	2003−12−02	Revised by: TMM
Reviewed by LDP		
Revision 0.5.2	2003−11−26	Revised by: LKS
A lot of cleanup. Thanks to John Levon levon [at] movementarian.org		
Revision 0.5.1	2003−11−22	Revised by: LKS
Changed the license from GFDL to OPL due to some GFDL problems.		
Revision 0.5	2003−11−18	Revised by: LKS
Converted to XML Docbook. Some cleanup.		
Revision 0.4	2002−11−07	Revised by: LKS
Fixed some errors + update. Thanks to Henrik Petander petander (at) tcs hut fi.		
Revision 0.3.1	2003−11−03	Revised by: LKS
Updated to MIPL relase 1.0 (kernel 2.4.22).		
Revision 0.3	2003−08−05	Revised by: LKS
Initial release.		

This document describes the software and procedures to set up and use mobile IPv6 for Linux.

Table of Contents

1. Introduction **1**
- 1.1. What is Mobile IP? 1
- 1.2. Why Mobile IP? 1
- 1.3. How does it work? 1

2. IPv6 ..... **3**

3. Mobile IPv6 for Linux **4**
- 3.1. Patching the kernel 4
- 3.2. Userspace tools 5
- 3.3. MIPv6 device node 5
- 3.4. Automatic startup 5

4. Test bed **7**
- 4.1. Testcase 7
- 4.2. Step−by−step configuration 7
  - 4.2.1. Setting up a fully functional IPv6 network 7
  - 4.2.2. Configuring Mobile IPv6 8
  - 4.2.3. Configuring radvd on AR 10
  - 4.2.4. Configuring radvd on HA 11

5. Doing some tests **13**
- 5.1. Pre−test 13
- 5.2. Movement detection 13
- 5.3. ping6 15
- 5.4. Kernel IP routing table 15
- 5.5. Travelling through several foregin LAN's 16
- 5.6. Returning home 17
- 5.7. Real life testing − smooth handover 18

6. FAQ **19**

7. Useful Resources **21**

8. Copyright, acknowledgments and miscellaneous **22**
- 8.1. Copyright and License 22
- 8.2. How this document was produced 22
- 8.3. Feedback 22
- 8.4. Acknowledgments 22

A. GNU Free Documentation License **23**

A.1. PREAMBLE **24**

A.2. APPLICABILITY AND DEFINITIONS **25**

Table of Contents

A.3. VERBATIM COPYING **27**

A.4. COPYING IN QUANTITY **28**

A.5. MODIFICATIONS **29**

A.6. COMBINING DOCUMENTS **31**

A.7. COLLECTIONS OF DOCUMENTS **32**

A.8. AGGREGATION WITH INDEPENDENT WORKS **33**

A.9. TRANSLATION **34**

A.10. TERMINATION **35**

A.11. FUTURE REVISIONS OF THIS LICENSE **36**

A.12. ADDENDUM: How to use this License for your documents **37**

1. Introduction

This document describes the software and procedures to set up and use mobile IPv6 for Linux. The "Mobility Support in IPv6" draft answers the *what* and *why* of mobile IP:

1.1. What is Mobile IP?

"Each mobile node is always identified by its home address, regardless of its current point of attachment to the Internet. While situated away from its home, a mobile node is also associated with a care−of address, which provides information about the mobile node's current location. IPv6 packets addressed to a mobile node's home address are transparently routed to its care−of address via the mobile nodes Home Agent (HA). The protocol enables IPv6 nodes to cache the binding of a mobile node's home address with its care−of address, and then to send any packets destined for the mobile node directly to it at this care−of address." −−− draft−ietf−mipv6−24, page 1−2.

1.2. Why Mobile IP?

"Without specific support for mobility in IPv6, packets destined to a mobile node (host or router) would not be able to reach it while the mobile node is away from its home link (the link on which its home IPv6 subnet prefix is in use), since routing is based on the subnet prefix in a packet's destination IP address. In order to continue communication in spite of its movement, a mobile node could change its IP address each time it moves to a new link, but the mobile node would then not be able to maintain transport and higher−layer connections when it changes location. Mobility support in IPv6 is particularly important, as mobile computers are likely to account for a majority or at least a substantial fraction of the population of the Internet during the lifetime of IPv6." −−− draft−ietf−mipv6−24, page 6.

For all the details, read the "Mobility Support in IPv6" draft

1.3. How does it work?

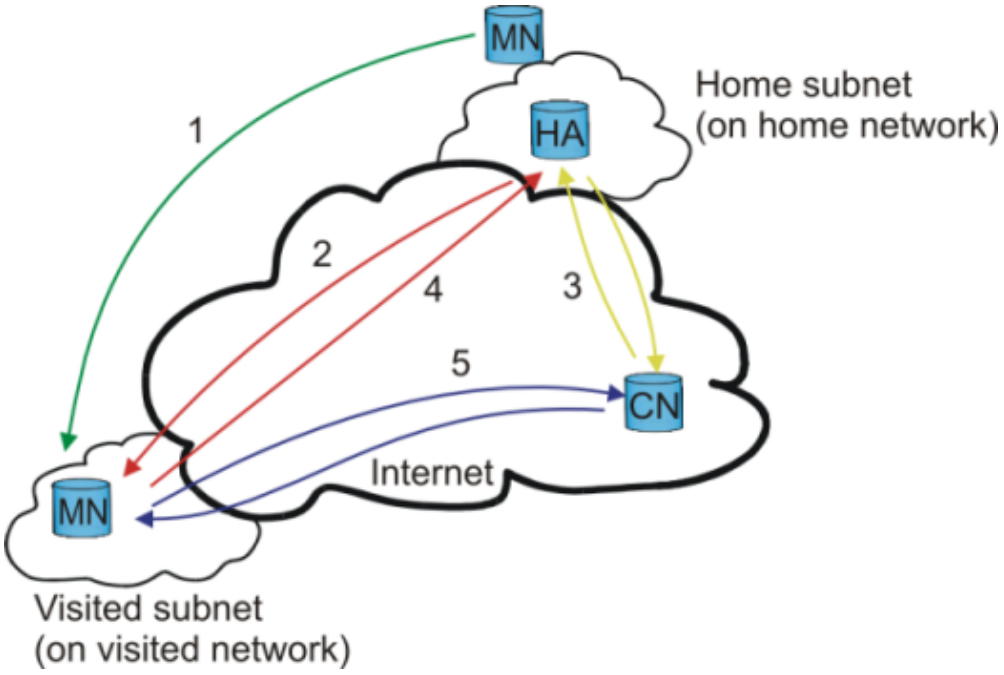

Mobile IP

1. The Mobile Node (MN) travels to a foreign network and gets a new care−of−address.
2. The MN performs a binding update to its Home Agent (HA) (the new care−of−address gets registered at HA). HA sends a binding acknowledgement to MN.
3. A Correspondent Node (CN) wants to contact the MN. The HA intercepts packets destined to the MN.
4. The HA then tunnels all packets to the MN from the CN using MN's care−of−address.
5. When the MN answers the CN, it may use its current care−of−address (and perform a binding to the CN) and communicate with the CN directly (optimized routing) or it can tunnel all its packets through the HA.

See figure "Mobile IP" for an explanation.

---

2. IPv6

IP version 6 (IPv6) is a new version of the Internet Protocol, designed as the successor to IP version 4 (IPv4) [RFC−791]. The changes from IPv4 to IPv6 fall primarily into the following categories:

- Expanded addressing capabilities
- Header format simplification
- Improved support for extensions and options
- Flow labeling capability
- Authentication and privacy capabilities

You should have basic knowledge of IPv6 stateless auto−configuring to fully understand how 'mobile IPv6' (MIPv6) works. You can read up on IPv6 Stateless Address Autoconfiguration in [RFC2462].

For more information on IPv6 in general, visit the IETF's IPv6 Working Group.

---

3. Mobile IPv6 for Linux

There are currently two Mobile IPv6 Linux implementations available. The Lancaster University in the UK has the oldest(?) implementation (http://www.cs−ipv6.lancs.ac.uk/MobileIP/). The latest kernel supported is 2.1.90, and is compatible with IETF mobile IPv6 draft−v5 (the current revision is v24). The code and website has not been updated since 1998, so it is considered obsolete.

The other implementation, which is up−to−date, is Helsinki University of Technology's MIPL project. The latest supported kernel is 2.4.22, and they have patches for the upcoming 2.6 kernel (see the FAQ). Visit http://www.mobile−ipv6.org/ for papers, software or to browse the mail archive.

3.1. Patching the kernel

The MIPL MIPv6 implementation requires a kernel patch. The implementation modifies the IPv6 kernel stack, so a kernel recompile is necessary. The installation process is well documented, but I will give a brief step−by−step howto.

Please note! The need for two different kernels, one for MN and one for HA, is obsolete. Just compile support for MN and HA in the same kernel. It is not possible to run as both an MN and an HA at the same time; which mode is chosen depends on which of the modules are loaded.

1. Download the latest Linux MIPv6 source code from http://www.mobile−ipv6.org/. The latest release today is: *mipv6−1.0−v2.4.22*. The last four numbers corresponds to the Linux kernel the patch should be applied to:
   ```
   # cd /usr/local/src
   # wget http://www.mobile−ipv6.org/download/mipv6−1.0−v2.4.22.tar.gz
   # tar zxfv mipv6−1.0−v2.4.22.tar.gz
   ```
2. Download and unpack the correspondent Linux kernel version from ftp.kernel.org:
   ```
   # cd /usr/src
   # wget ftp://ftp.kernel.org/pub/linux/kernel/v2.4/linux−2.4.22.tar.bz2
   # tar jxvf linux−2.4.22.tar.bz2
   # ln −s linux−2.4.22 linux
   # cd linux
   ```
3. Apply the MIPv6 patch:
   ```
   # patch −p1 −−dry−run < /usr/local/src/mipv6−1.0−v2.4.22/mipv6−1.0−v2.4.22.patch
   ```
   The −−dry−run option checks that the patch will apply correctly. If you get any failed hunks, you should *not* proceed. If everything went fine do:
   ```
   # patch −p1 < /usr/local/src/mipv6−1.0−v2.4.22/mipv6−1.0−v2.4.22.patch
   ```
4. Now your kernel tree is ready for configuration. Run your favorite **make *config**. The MIPv6 options are under "Networking Options". The following options should be present in *".config"*:
   ```
   CONFIG_EXPERIMENTAL=y
   CONFIG_SYSCTL=y
   CONFIG_PROC_FS=y
   CONFIG_MODULES=y
   CONFIG_NET=y
   CONFIG_NETFILTER=y
   CONFIG_UNIX=y
   ```

```
CONFIG_INET=y
CONFIG_IPV6=m
CONFIG_IPV6_SUBTREES=y
CONFIG_IPV6_IPV6_TUNNEL=m
CONFIG_IPV6_MOBILITY=m
CONFIG_IPV6_MOBILITY_MN=m
CONFIG_IPV6_MOBILITY_HA=m
```

Since MIPL is still a work−in−progress you might want to enable:

```
CONFIG_IPV6_MOBILITY_DEBUG=y
```

With debug messages it is easier to figure out what happened when something goes wrong. Also, when reporting a bug, debug messages are very helpful.

To be sure you have all the correct options, you can run **`chkconf_kernel.sh`**, which is a small shell script included in the MIPL tarball.

5. Next you should compile and install your kernel.

   Hint: To easily distinguish this kernel from other kernels, you can change the "EXTRAVERSION" variable in the **`/usr/src/linux/Makefile`** to for example "−MIPv6−1".

Read the Linux Kernel HOWTO for detailed instruction on how to patch, compile and install your new kernel.

3.2. Userspace tools

The userspace tool **`mipdiag`**, config files and init scripts must be installed for the module to work correctly:

```
# cd /usr/local/src/mipv6−1.0−v2.4.22
# ./configure
# make && make install
```

3.3. MIPv6 device node

The MIPv6 module also needs a new device node entry. Issue the command:

```
# mknod /dev/mipv6_dev c 0xf9 0
```

3.4. Automatic startup

1. *Red Hat:*

   All init scripts are located in `/etc/init.d/`, which are sym−linked to the correct runlevel (`/etc/rcX.d/`). You can issue the command:

   ```
   # chkconfig −−add mobile−ip6
   ```

to enable MIPv6 at startup, or

```
# chkconfig --del mobile-ip6
```

to remove MIPv6 from startup.

2. *Debian:*

   If you are so lucky to be running Debian, you can issue the command:

   ```
   # update-rc.d -n mobile-ip6 start 75 3 4 5 . stop 05 1 2 6 .
   ```

   to set up all the necessary links.

3. *Slackware:*

   Slackware users have all their startup/runlevel scripts in `/etc/rc.d`. Since 'configure' doesn't check for "/etc/rc.d", you can add *INIT_SLACK="/etc/rc.d"*, and then INIT_SLACK to INITDIRS in 'configure' (search for INITDIR in configure). Since you are running Slackware, you probably know this already. The following command should then do the trick:

   ```
   # echo '/etc/rc.d/mobile-ip6 start' >> /etc/rc.d/rc.local
   ```

   If you don't hack the Makefile, the `mobile-ip6` script is installed at '/' (you may then move it to /etc/rc.d/).

---

4. Test bed

Now you should have a working MIPL patched kernel, installed userlevel tools and enabled automatic startup at boot. If anything went wrong, go through the above sections carefully.

4.1. Testcase

The addresses we are using in our test−bed are site−local. You may as well use global addresses, but do *note that link local addresses won't work!* Our test−bed consist of four nodes; see figure "Mobile IPv6 testbed".

1. *HA − Home Agent:* The HA is located at the home network with address **`fec0:106:2700::2`**, with one wireless interface.
2. *MN − Mobile Node:* When MN is on the "home network", it has address **`fec0:106:2700::4`**. When MN travels to another network, it generates a new "care−of" address.
3. *R − Router:* This is the router from the home network to the internet. It has one wireless interface with address **`fec0:106:2700::1`** and a wired interface with address **`fec0:106:2300::2`**.
4. *AR − Access Router:* The link between AR and R is our "internet" − but in this testcase only a cross−cable (can be any network). The AR has two interfaces; the wired interface to R has address **`fec0:106:2300::1`**, the wireless has address **`fec0:106:1100::1`**.

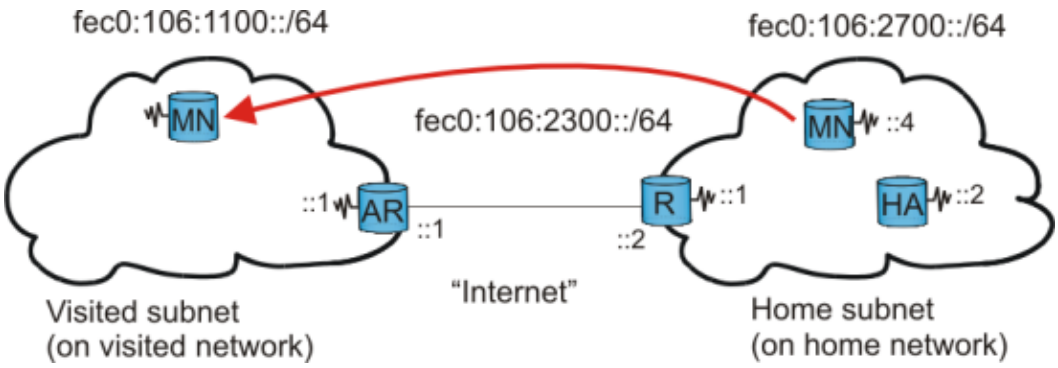

Mobile IPv6 testbed

4.2. Step−by−step configuration

4.2.1. Setting up a fully functional IPv6 network

Before we can start testing mobile IP, we need a fully functional IPv6 network. All the nodes should be able to ping each other. *This is a crucial part.* If, for example, AR is not able to ping HA, then there will be no binding update.

I will give a brief instruction to get our network up and running using IPv6. For more info on setting up an IPv6 network, you can read Peter Bieringer's excellent Linux IPv6 HOWTO.

I've turned off encryption for simplicity − *NOTE that you should ALWAYS use encryption when dealing with wireless networks!*

Also note that the different wireless networks have different ESSIDs!

1. *MN:* The Mobile Node has one wireless interface. Forwarding should be turned off, but should accept autoconf and ra's:
```
# iwconfig eth0 mode ad-hoc essid homenet enc off
# ifconfig eth0 inet6 add fec0:106:2700::4/64
# echo "0" > /proc/sys/net/ipv6/conf/eth0/forwarding
# echo "1" > /proc/sys/net/ipv6/conf/eth0/autoconf
# echo "1" > /proc/sys/net/ipv6/conf/eth0/accept_ra
# echo "1" > /proc/sys/net/ipv6/conf/eth0/accept_redirects
# /etc/init.d/mobile-ip6 start
```
2. *HA:* The Home Agent has one wireless interface. It should have forwarding turned on because it uses normal routing to deliver packets captured from a physical interface to the virtual tunnel interface. *Note: You must add a default route or else HA will have problem contacting the MN on visited LAN's. One possible solution is to use HA as the default router of the home network.*
```
# iwconfig eth0 mode ad-hoc essid homenet enc off
# ifconfig eth0 inet6 add fec0:106:2700::2/64
# echo "1" > /proc/sys/net/ipv6/conf/eth0/forwarding
# echo "0" > /proc/sys/net/ipv6/conf/eth0/autoconf
# echo "0" > /proc/sys/net/ipv6/conf/eth0/accept_ra
# echo "0" > /proc/sys/net/ipv6/conf/eth0/accept_redirects
# ip route add ::/0 via fec0:106:2700::1
# /etc/init.d/mobile-ip6 start
```
3. *R:* The (home) Router has two interfaces; one wireless and one line. The Router must have forwarding turned on.
```
# ifconfig eth0 inet6 add fec0:106:2300::2/64
# iwconfig eth1 mode ad-hoc essid homenet enc off
# ifconfig eth1 inet6 add fec0:106:2700::1/64
# echo "1" > /proc/sys/net/ipv6/conf/all/forwarding
# echo "0" > /proc/sys/net/ipv6/conf/all/autoconf
# echo "0" > /proc/sys/net/ipv6/conf/all/accept_ra
# echo "0" > /proc/sys/net/ipv6/conf/all/accept_redirects
# ip route add fec0:106:1100::/64 via fec0:106:2300::1
```
4. *AR:* The Access Router (on a foreign network) also has two interfaces; one wireless and one line. Forwarding must be turned on.
```
# ifconfig eth0 inet6 add fec0:106:2300::1/64
# iwconfig eth1 mode ad-hoc essid visitnet enc off
# ifconfig eth1 inet6 add fec0:106:1100::1/64
# echo "1" > /proc/sys/net/ipv6/conf/all/forwarding
# echo "0" > /proc/sys/net/ipv6/conf/all/autoconf
# echo "0" > /proc/sys/net/ipv6/conf/all/accept_ra
# echo "0" > /proc/sys/net/ipv6/conf/all/accept_redirects
# ip route add fec0:106:2700::/64 via fec0:106:2300::2
```

Instead of modifying proc variables, you can use *sysctl*.

Note: We are setting static routes on our test-bed. You should now be able to ping all the hosts from every host.

4.2.2. Configuring Mobile IPv6

The last configuration is MIPv6 settings in `network-mip6.conf`. In Debian/Slackware the file is found under `/etc/`. (RedHat the file is found under `/etc/sysconfig/`.) The file should be pretty self-explanatory.

1. *HA:* The HA config file should contain these settings:
```
# cat /etc/network-mip6.conf

# Home Agent configuration file
FUNCTIONALITY=ha
DEBUGLEVEL=1
MIN_TUNNEL_NR=1
MAX_TUNNEL_NR=5
TUNNEL_SITELOCAL=yes
```
2. *MN:* The MN config file should look like this:
```
 # cat /etc/network-mip6.conf

 # Mobile Node configuration file
 FUNCTIONALITY=mn
 DEBUGLEVEL=1
 TUNNEL_SITELOCAL=yes
 MIN_TUNNEL_NR=1
 MAX_TUNNEL_NR=3
 HOMEDEV=mip6mnha1
 HOMEADDRESS=fec0:106:2700::4/64  # MN's home adress
 HOMEAGENT=fec0:106:2700::2/64    # HA's address
```
3. Next, start mobile-IP:
```
# /etc/init.d/mobile-ip6 start
Starting Mobile IPv6: OK
```

You can verify that it started by doing a **`ifconfig`** on HA. If the tunnel(s) comes up, `ip6tnl1`, mobile-ip6 is started:

```
# ifconfig
eth1      Link encap:Ethernet  HWaddr 00:02:2D:2D:DE:79
          inet6 addr: fec0:106:2700::2/64 Scope:Site
          inet6 addr: fe80::202:2dff:fe2d:de79/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:618 errors:6 dropped:6 overruns:0 frame:6
          TX packets:1485 errors:22 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:100
          RX bytes:87914 (85.8 KiB)  TX bytes:252596 (246.6 KiB)
          Interrupt:3 Base address:0x100

ip6tnl1   Link encap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00 ❶
          UP POINTOPOINT RUNNING NOARP  MTU:1460  Metric:1
          RX packets:6 errors:0 dropped:0 overruns:0 frame:0
          TX packets:6 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:576 (576.0 b)  TX bytes:624 (624.0 b)

ip6tnl2   Link encap:UNSPEC  HWaddr 00-00-00-00-00-00-00-00-00-00-00-00-00-00-00-00 ❷
          UP RUNNING NOARP  MTU:1460  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
```

```
        RX packets:8 errors:0 dropped:0 overruns:0 frame:0
        TX packets:8 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:560 (560.0 b)  TX bytes:560 (560.0 b)
```

❶ The tunnel is up and ready for connections.

❷ Another tunnel ready.

You will also see the mipv6 kernel modules are loaded (MN):

```
# lsmod
Module                  Size  Used by    Not tainted
mip6_mn                59888   0  (unused)
ipv6_tunnel            11448   1  [mip6_mn]
mip6_base              40728   0  [mip6_mn]
ipv6                  179764  -1  [mip6_mn ipv6_tunnel mip6_base]
...
```

4.2.3. Configuring radvd on AR

When MN comes to a new network, it does a link−local address configuration, going to the next phase if that succeeds. I'll let [RFC2462] (IPv6 Stateless Address Autoconfiguration) describe the next phase:

"The next phase of autoconfiguration involves obtaining a Router Advertisement or determining that no routers are present. If routers are present, they will send Router Advertisements that specify what sort of autoconfiguration a host should do. If no routers are present, stateful autoconfiguration should be invoked."

"Routers send Router Advertisements periodically, but the delay between successive advertisements will generally be longer than a host performing autoconfiguration will want to wait. To obtain an advertisement quickly, a host sends one or more Router Solicitations to the all−routers multicast group." −−− page 8

This is where we use RADVD.

Read [RFC2462] more more details concerning IPv6 Stateless Address Autoconfiguration.

We'll configure RADVD on AR's wireless interface. The `radvd.conf` file should contain this:

```
# cat /etc/radvd.conf
interface eth1
{
      AdvSendAdvert on;
      AdvIntervalOpt on;

      MinRtrAdvInterval 3;
      MaxRtrAdvInterval 10;
      AdvHomeAgentFlag off;

      prefix fec0:106:1100::/64
      {
          AdvOnLink on;
          AdvAutonomous on;
          AdvRouterAddr on;
```

```
        };
};
```

We then start it:

```
# /etc/init.d/radvd start
```

You should now be able to use **`radvdump`** to see that the radvd messages really are being sent periodically:

```
# radvdump
Router advertisement from fe80::202:2dff:fe54:d1b2 (hoplimit 255)
Received by interface eth1
      # Note: {Min,Max}RtrAdvInterval cannot be obtained with radvdump
      AdvCurHopLimit: 64
      AdvManagedFlag: off
      AdvOtherConfigFlag: off
      AdvHomeAgentFlag: off
      AdvReachableTime: 0
      AdvRetransTimer: 0
      Prefix fec0:106:1100::/64
              AdvValidLifetime: 2592000
              AdvPreferredLifetime: 604800
              AdvOnLink: on
              AdvAutonomous: on
              AdvRouterAddr: off
      AdvSourceLLAddress: 00 02 2D 54 D1 B2
```

Note! When using radvd on HA and enabling "autoconf" (in proc), you will also get an autogenerated IPv6 address on MN (which is superfluous) in addition to your static address:

4.2.4. Configuring radvd on HA

To enable the MN to know when it's home, HA should also be sending out RAs. We should therefore enable RADVD on the HA as well. The `/etc/radvd.conf` file should contain:

```
 # cat /etc/radvd.conf
 interface eth0
 {
    AdvSendAdvert on;
    MaxRtrAdvInterval 3;
    MinRtrAdvInterval 1;
    AdvIntervalOpt off;
    AdvHomeAgentFlag on;
    HomeAgentLifetime 10000;
    HomeAgentPreference 20;
    AdvHomeAgentInfo on;
    prefix fec0:106:2700::2/64
    {
      AdvRouterAddr on;
      AdvOnLink on;
      AdvAutonomous on;
      AdvPreferredLifetime 10000;
      AdvValidLifetime 12000;
    };
```

```
    };
```

Also do a **radvdump** on HA to check whether radvd messages are beeing sent:

```
# radvdump
Router advertisement from fe80::202:2dff:fe54:d11e (hoplimit 255)
Received by interface eth0
    # Note: {Min,Max}RtrAdvInterval cannot be obtained with radvdump
    AdvCurHopLimit: 64
    AdvManagedFlag: off
    AdvOtherConfigFlag: off
    AdvHomeAgentFlag: on
    AdvReachableTime: 0
    AdvRetransTimer: 0
    Prefix fec0:106:2700::2/64
            AdvValidLifetime: 12000
            AdvPreferredLifetime: 10000
            AdvOnLink: on
            AdvAutonomous: on
            AdvRouterAddr: on
    AdvSourceLLAddress: 00 02 2D 54 D1 1E
    AdvHomeAgentInfo:
            HomeAgentPreference: 20
            HomeAgentLifetime: 1000
```

```
# ifconfig eth0
eth0  Link encap:Ethernet  HWaddr 00:90:7D:F3:03:1A
      inet6 addr: fec0:106:2700:0:290:7dff:fef3:31a/64 Scope:Site ❶
      inet6 addr: fec0:106:2700::4/64 Scope:Site                  ❷
      inet6 addr: fe80::290:7dff:fef3:31a/64 Scope:Link           ❸
      UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
      RX packets:513 errors:89 dropped:89 overruns:0 frame:85
      TX packets:140 errors:41 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:100
      RX bytes:56084 (54.7 Kb)  TX bytes:19212 (18.7 Kb)
      Interrupt:3 Base address:0x100
```

❶ A new (superfluous) autogenerated address. Since we are setting **autoconf** in **/proc/sys/net/ipv6/conf/eth0/autoconf** to **1**, MN will generate a new adress combined with HA's prefix and it's own MAC address. I do not think is it possible to avoid having this address generated.

❷ Our original static IPv6 address

❸ The link−local address generated at boot.

5. Doing some tests

5.1. Pre−test

Do every configuration as shown above; it's especially important to have a different ESSID on the home net and visited network.

When you start mobile−IPv6 on MN, you will see multicasting router solicitations messages:

```
# tcpdump −i eth0 −vv ip6 or proto ipv6

...
13:32:54.681763 fe80::202:a5ff:fe6f:a08a > ff02::2: icmp6: router solicitation \
(src lladdr: 0:2:a5:6f:a0:8a) (len 16, hlim 255)

13:32:55.681763 fe80::202:a5ff:fe6f:a08a > ff02::2: icmp6: router solicitation \
(src lladdr: 0:2:a5:6f:a0:8a) (len 16, hlim 255)

13:32:57.681765 fe80::202:a5ff:fe6f:a08a > ff02::2: icmp6: router solicitation \
(src lladdr: 0:2:a5:6f:a0:8a) (len 16, hlim 255)
...
```

5.2. Movement detection

Generic movement detection uses Neighbor Unreachability Detection to detect when the default router is no longer bi−directionally reachable, in which case the mobile node must discover a new default router (usually on a new link).

To easily see whats going on, you should have one xterm window for each of these commands:

```
# watch ifconfig eth0
# watch route −A inet6
# tcpdump −i eth0 −vv ip6 or proto ipv6
```

To "travel" to another net, you can issue the command on MN:

```
# iwconfig eth1 essid visitnet
```

The MN is then on the other wireless network, and since it is sending out "router solicitation" (multicast), our AR will respond with it's prefix. MN will then configure itself with at new IPv6 address with the received prefix and it's own MAC address. If you type **ifconfig eth0** you will see the new IPv6 address:

```
# ifconfig eth0
eth0  Link encap:Ethernet  HWaddr 00:90:7D:F3:03:1A
      inet6 addr: fec0:106:1100:0:290:7dff:fef3:31a/64 Scope:Site ❶
      inet6 addr: fec0:106:2700:0:290:7dff:fef3:31a/64 Scope:Site ❷
      inet6 addr: fec0:106:2700::4/64 Scope:Site                 ❸
      inet6 addr: fe80::290:7dff:fef3:31a/64 Scope:Link          ❹
      UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
```

```
          RX packets:854 errors:154 dropped:154 overruns:0 frame:148
          TX packets:293 errors:58 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:100
          RX bytes:96536 (94.2 Kb)  TX bytes:44664 (43.6 Kb)
          Interrupt:3 Base address:0x100
```

❶ The new "foreign" address, generated by combining AR's prefix and MAC−address

❷ The superfluous home network address (because of HA radvd messages and MN autoconf set to "true").

❸ The "original" (home) address

❹ The link−local address generated at boot

Almost at the same time, the MN will perform a binding update to HA. In your tcpdump window, you will see several packets destined to HA. To verify that the binding update has been sent and acknowledged from MN:

```
      # mipdiag −s
      Mobile IPv6 Statistics
      NEncapsulations           : 0
      NDecapsulations           : 0
      NBindUpdatesRcvd          : 0
      NBindAcksRcvd             : 1    ❶
      NBindNAcksRcvd            : 0
      NBindRqsRcvd              : 0
      NBindUpdatesSent          : 1    ❷
      NBindAcksSent             : 0
      NBindNAcksSent            : 0
      NBindRqsSent              : 0
      NBindUpdatesDropAuth      : 0
      NBindUpdatesDropInvalid   : 0
      NBindUpdatesDropMisc      : 0
      NBindAcksDropAuth         : 0
      NBindAcksDropInvalid      : 0
      NBindAcksDropMisc         : 0
      NBindRqsDropAuth          : 0
      NBindRqsDropInvalid       : 0
      NBindRqsDropMisc          : 0
```

❶ One binding ACK received.

❷ One binding UPDATE sent.

You can also verify the binding with the following command (on MN):

```
     # mipdiag −l
     Mobile IPv6 Binding update list
     Recipient CN: fec0:106:2700::2
     BINDING home address: fec0:106:2700::4 care−of address: fec0:106:1100:0:290:7dff:fef3:31a
             expires: 936 sequence: 0 state: 1
             delay: 3 max delay 32 callback time: 736
```

You can also verify it on HA with the statistics option (−s) and with the "binding cache" (−c) option:

```
# mipdiag −c
Mobile IPv6 Binding cache
Home Address      Care−of Address                    Lifetime  Type
fec0:106:2700::4  fec0:106:1100:0:290:7dff:fef3:31a  971       2
```

5.3. ping6

From the MN, you can try to ping AR's eth1 (fec0:106:1100::1):

```
# ping6 fec0:106:1100::1
PING fec0:106:1100::1(fec0:106:1100::1) from fec0:106:2700::4 : 56 data bytes
64 bytes from fec0:106:1100::1: icmp_seq=1 ttl=62 time=8.01 ms
64 bytes from fec0:106:1100::1: icmp_seq=2 ttl=62 time=8.02 ms
...
```

By using tcpdump, you can see how the packets travel:

```
12:13:51.789688 fec0:106:1100:0:202:a5ff:fe6f:a08a > fec0:106:2700::2: \  ❶
fec0:106:2700::4 > fec0:106:1100::1: icmp6: echo request \ ❷
(len 64, hlim 64) (len 104, hlim 255)

12:13:51.797675 fec0:106:2700::2 > fec0:106:1100:0:202:a5ff:fe6f:a08a: \ ❸
fec0:106:1100::1 > fec0:106:2700::4: icmp6: echo reply \
(len 64, hlim 62) (len 104, hlim 253)
```

❶ The packet first goes from MN to the HA using MN new IPv6 address.

❷ Then from HA to AR.

❸ The AR then responds to HA and tunnels the packets to MN.

You can now see the statistics have been updated (on MN):

```
# mipdiag −s
Mobile IPv6 Statistics
NEncapsulations           : 56
NDecapsulations           : 25
...
```

5.4. Kernel IP routing table

One interesting thing MIPv6 does is change the default route to a tunnel. The new default route becomes:

```
# route −A inet6
Kernel IPv6 routing table
Destination      Next Hop       Flags Metric Ref    Use Iface
::/0             ::             UD    64     0        0 ip6tnl1
....
```

If it doesn't add a default route, you may add it manually:

```
# ip route ::/0 via dev ip6tnl
```

5.5. Travelling through several foregin LAN's

To travel to several visited networks, is no different than travel to *one* network. The only thing you must have in mind is that you will generate a new address for each visited network.

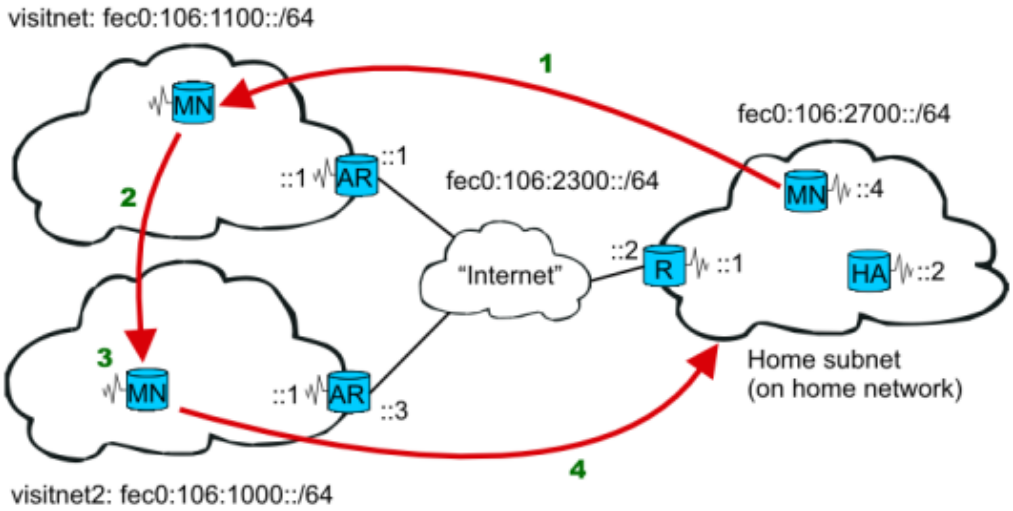

MN travelling through severeal different LANs.

1. MN first visits 'visitnet', as we have been through above.
2. MN is then travelling from 'visitnet' to 'visitnet2'.
3. When at 'visitnet2', MN generates a new IPv6 address and do a new binding update to HA.
4. MN then travels back home. (Se next section.)

The AR at "visitnet2", is configured exactly as the other AR (at "visitnet"), except using address **`fec0:106:1000::/64`** instead of **`fec0:106:1100::/64`**.

To make the mobile node travel from 'visitnet' to 'visitnet2', issue the command (on MN):

```
 # iwconfig eth0 essid visitnet2
```

You will then see the MN configures itself to the new network:

```
# ifconfig eth0
eth1  Link encap:Ethernet  HWaddr 00:90:7D:F3:03:1A
      inet6 addr: fec0:106:1000:0:290:7dff:fef3:31a/64 Scope:Site ❶
      inet6 addr: fec0:106:1100:0:290:7dff:fef3:31a/64 Scope:Site
      inet6 addr: fec0:106:2700:0:290:7dff:fef3:31a/64 Scope:Site
```

```
          inet6 addr: fec0:106:2700::4/64 Scope:Site
          inet6 addr: fe80::290:7dff:fef3:31a/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:1073 errors:212 dropped:212 overruns:0 frame:204
          TX packets:371 errors:72 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:100
          RX bytes:120340 (117.5 Kb)  TX bytes:56912 (55.5 Kb)
          Interrupt:3 Base address:0x100
```

❶ The new autoconfigured address at 'visitnet2'.

Note! You may have to restart mobile−ipv6 on MN when coming to a new network!

```
# /etc/init.d/mobile−ip6 restart
Stopping Mobile IPv6: OK
Starting Mobile IPv6: OK
```

The MN will then perform a new binding update to HA. Notice the new "care−of address":

```
# mipdiag −l
Mobile IPv6 Binding update list
Recipient CN: fec0:106:2700::2
BINDING home address: fec0:106:2700::4 care−of address: fec0:106:1000:0:290:7dff:fef3:31a
    expires: 973 sequence: 14 state: 1
    delay: 3 max delay 32 callback time: 773
```

You can also see the "binding cache" on HA has been updated:

```
# mipdiag −c
Mobile IPv6 Binding cache
Home Address      Care−of Address                    Lifetime  Type
fec0:106:2700::4  fec0:106:1000:0:290:7dff:fef3:31a  943       2
```

5.6. Returning home

To make the MN return home, you can just issue the command:

```
# iwconfig eth0 essid homenet
```

The MN will know it is back home, since HA is sending out radvd messages with the HA−bit set (AdvHomeAgentFlag), see Section 4.2.4

You can see the MN "is back home", since the binding cache information at HA is flushed (empty):

```
Mobile IPv6 Binding cache
Home Address      Care−of Address                     Lifetime  Type
```

5.7. Real life testing − smooth handover

To really get the feel on how mobile IP works, fire up GnomeMeeting (See the figure GnomeMeeting and start a netmeeting. Note! You must use the latest GnomeMeeting to get support for IPv6! Then do a "travel" and you can see an almost smooth handover.

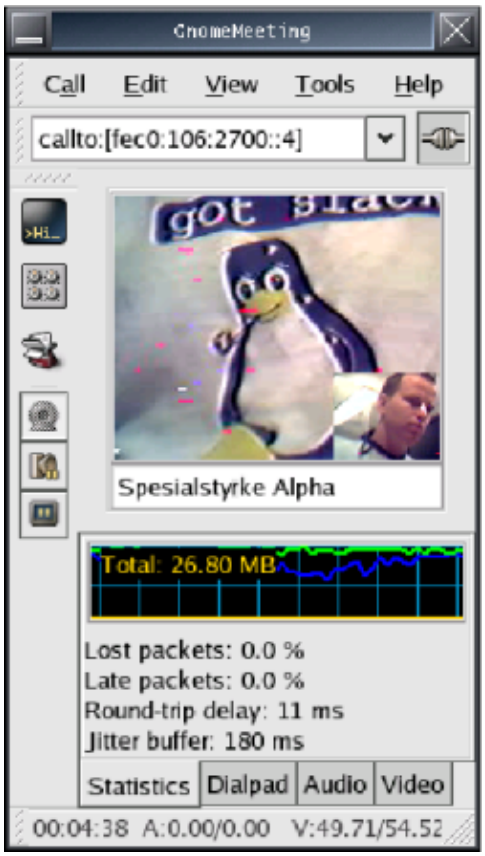

Using GnomeMeeting with IPv6 to test roaming between two wireless networks

6. FAQ

1. *Q: Why do we have to create the `/dev/mipv6_dev` entry?*

   A: The dev file is mainly so that the userspace tool, mipdiag, can make modifications to the kernel parameters using ioctl calls through the device file. **mknod** creates the special device file with paramters recognizable by the mobile−ip6 module.
2. *Q: Is there any support for kernel 2.6.x?*

   A: Here is the answer from Henrik Petander on the MIPL mailinglist:

   "Here is a short overview of the status of MIPL for 2.6 kernel series:"

   "We have finished the kernel infrastructure for Mobile IPv6 in cooperation with the USAGI project. The infrastructure does route optimization, tunneling and policy routing."

   "We are now working on the userspace daemon which handles the MIPv6 signaling and controls the operation of the kernel part. The userspace part is also progressing nicely. However, the protocol logic is still missing, so there isn't really anything for users to test yet. We should have a well working and tested prototype ready and by the end of March."
3. *Q: Does MIPL support IPSec?*

   A: There is no support IPSec on 2.4.x. MIPL for 2.6 series will have IPSec support from the start. You may use a third−party IPSec implementation.
4. *Q: How can I control the type of routing used for communication between the MN and a CN (through HA tunnel or by direct communication using binding update/acks)?*

   A: You can control this through:

   `/proc/sys/conf/net/ipv6/mobility/accept_return_routability`

   If you do not want to use return routability and route optimization, set it to 0 with:

   `# echo 0 > /proc/sys/..../accept_return_routability`

   Then MN will communicate with CNs only through the home tunnel.
5. *Q: Can different wireless networks have different ESSIDs/WEP keys?*

   A: Yes, but you must change this upon arrival to the new network. MIPv6 from MIPL can't do this automatically.
6. *Q: If MN has travelled through several visited LAN, and then returning home; the interface still has all the autogenerated IPv6 addresses from all the visited networks! Is there any way to "flush/delete" these addresses?*

   A: No, I do not know of any automatic way these adresses can be removed, but you can delete them manually:

   `# ifconfig eth0 inet6 del <ipv6−address>`
7. *Q: Host B has two interfaces with two different subnets assigned. When I ping B from host A, it does not answer! Why not? Host A knows where host B (subnets) are!*

A: The host B doesn't know where host A is (B doesn't know where A's net is), so you must add a route entry:

```
# ip route add fec0:106:2700::/64 via fec0:106:2300::1
```

or

```
# route -A inet6 add fec0:106:2700::/64 gw fec0:106:2300::1 dev eth0
```

8. *Q: How do I set a default gateway in IPv6?*

   A: You do that using the traditional "route":

   ```
   # route -A inet6 add default gw <ipv6-host>
   ```

   or the newer "ip" command:

   ```
   # ip route ::/0 via <ipv6-host>
   ```

9. *Q: Why does the host send a multicast address rather than an anycast address, requesting for router solicitation?*

   A: Because the host wants an answer from every router, not from just any router. The idea is to be able to get all parameters and to choose the "best" default router.

10. *Q: Why doesn't MN notice that it has moved?*

    A: It thinks that it's previous router is still reachable. This may result from very large lifetimes in router advertisements. Check the configuration of the program sending router advertisements in the router. If the program supports router advertisement intervals, you can use this to help MN in movement detection by setting the use of interval to `on`. See **man radvd.conf** for details.

---

7. Useful Resources

1. Mobile IPv6 for Linux http://www.mipl.mediapoli.com/
2. Mobile IP Working Group (IETF) http://www.ietf.org/html.charters/mobileip−charter.html
3. Mobile IPv6 draft http://www.ietf.org/internet−drafts/draft−ietf−mobileip−ipv6−24.txt
4. IPv6 Working Group (IETF) http://www.ietf.org/html.charters/ipv6−charter.html
5. RFC2460 Internet Protocol, Version 6 (IPv6) Specification http://www.ietf.org/rfc/rfc2460.txt
6. RFC2461 Neighbor Discovery for IP Version 6 (IPv6) http://www.ietf.org/rfc/rfc2461.txt
7. RFC2462 IPv6 Stateless Address Autoconfiguration http://www.ietf.org/rfc/rfc2462.txt
8. Peter Bieringer's Linux IPv6 HOWTO (en) http://ldp.linux.no/HOWTO/Linux+IPv6−HOWTO/
9. Current Status of IPv6 Support for Networking Applications http://www.deepspace6.net/docs/ipv6_status_page_apps.html
10. Linux Kernel HOWTO http://www.ibiblio.org/pub/Linux/docs/HOWTO/other−formats/html_single/Kernel−HOWTO.html

8. Copyright, acknowledgments and miscellaneous

8.1. Copyright and License

8.2. How this document was produced

This document was originally written in LaTeX using Emacs. HTML version created with latex2html. Later it was converted to DocBook XML.

An up−to−date version of this document can be found at:

HTML: http://www.tldp.org/HOWTO/Mobile−IPv6−HOWTO/

8.3. Feedback

Suggestions, corrections, additions wanted. Contributors wanted and acknowledged. Flames not wanted.

I can always be reached at <lars at unik no>

Homepage: http://www.gnist.org/~lars

8.4. Acknowledgments

This document was produced as a part of Interoperable Networks for Secure Communications (INSC task 6)

Thanks to Andreas Hafslund (andreha [at] unik.no) for initial support. Also thanks to UniK (University Graduate Center) http://www.unik.no and FFI (Norwegian Defence Research Establishment) http://www.ffi.mil.no for hardware support.

Thanks also to the other HOWTO authors whose works I have referenced:

Linux IPv6 HOWTO (en) by Peter Bieringer

A. GNU Free Documentation License

Version 1.2, November 2002

A.1. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondarily, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

A.2. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world−wide, royalty−free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front−matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front−Cover Texts or Back−Cover Texts, in the notice that says that the Document is released under this License. A Front−Cover Text may be at most 5 words, and a Back−Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine−readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard−conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine−generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

A.3. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

A.4. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front−Cover Texts on the front cover, and Back−Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine−readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer−network location from which the general network−using public has access to download using public−standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

A.5. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
D. Preserve all the copyright notices of the Document.
E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
H. Include an unaltered copy of this License.
I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front−matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties−−for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front−Cover Text, and a passage of up to 25 words as a Back−Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front−Cover Text and one of Back−Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

A.6. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

A.7. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

A.8. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

A.9. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

A.10. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

A.11. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See http://www.gnu.org/copyleft/.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

A.12. ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

> Copyright (c) YEAR YOUR NAME. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front−Cover Texts, and no Back−Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".

If you have Invariant Sections, Front−Cover Texts and Back−Cover Texts, replace the "with...Texts." line with this:

> with the Invariant Sections being LIST THEIR TITLES, with the Front−Cover Texts being LIST, and with the Back−Cover Texts being LIST.

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.