

My installation at Maple Park Development Corp.

U-verse ADSL (6 MB) →

(2) Grandstream HT702 VOIP adaptors

(4) Hardwired PCs (or devices)

(1) WRT54G (as a hub) →

3-6 wireless guests/other

@99.26.133.45 (Dynamic)

(All NATTED to 192.168.)

IPv6

In the Internet Protocol Version 6 (IPv6), the address block fe80::/10 has been reserved for link-local unicast addressing. The actual link local addresses are assigned with the prefix fe80::/64. They may be assigned by automatic (stateless) or stateful (e.g. manual) mechanisms.

Link-Local

Unlike IPv4, IPv6 requires a link-local address to be assigned to every network interface on which the IPv6 protocol is enabled, even when one or more routable addresses are also assigned. Consequently, IPv6 hosts usually have more than one IPv6 address assigned to each of their IPv6-enabled network interfaces.

IPv6

The link-local address is required for IPv6 sublayer operations of the Neighbor Discovery Protocol, as well as for some other IPv6-based protocols, like DHCPv6.

In IPv6, Stateless Address AutoConfiguration (SLAAC) is performed as a component of the Neighbor Discovery Protocol (NDP), as specified in RFC 4862. The address is formed from its routing prefix and the MAC address of the interface.

Network Size and Number of networks (The tasty version)



One IPv4 /24 -- 254 M&Ms

One IPv6 /64 -- Enough M&Ms to fill all 5 of the great lakes.



Full Address Space, One M&M per /24 covers 70% of a football field



Full Address Space, One M&M per /64 fills all 5 great lakes.



Comparison based on Almond M&Ms, not plain. Caution! Do not attempt to eat a /64 worth of any style of M&Ms.

Small size CDIR Prefix

- 2001:4978:000f:8640::/64
- |||| |||| |||| |||64 Single End-user LAN (default prefix size for SLAAC)
- |||| |||| |||| ||60 Some (very limited) 6rd deployments (ie, AT&T 6rd)
- |||| |||| |||| |56 Minimal end sites assignments (e.g. Home network)

Larger allocations

- |||| |||| |||48 Typical assignment for larger sites
- |||| |||| ||44
- |||| |||| |40
- |||| |||| 36 possible future Local Internet registry
extra-small allocations

Registries CDIR

- |||| |||32 Local Internet registry minimum allocations
- |||| ||28 Local Internet registry medium allocations
- |||| |24 Local Internet registry large allocations
- |||| 20 Local Internet registry extra large allocations
- ||12 Regional Internet Registry allocations from IANA[4]
- |8

IPv4 <--> IPv6

99.178.153.41

(99=6*16+3=63)

2001:4978:000f:8640:0000:0000:63b2:9929

2001:4978:f:8640::63b2:9929/64

IPv6 Routing

- [drf@maplepark ~]\$ ip -6 route
- 2001:4978:hhhh:hhhh::/64 dev sixxs proto kernel metric 256 mtu 1428 advmss 1368 hoplimit 4294967295
- 2001:4978:f:8640::/64 dev br0 metric 1024 mtu 1500 advmss 1440 hoplimit 4294967295
- default via 2001:4978:hhhh:hhhh::1 dev sixxs metric 1024 mtu 1428 advmss 1368 hoplimit 4294967295

Localhost

- [drf@maplepark ~]\$ ip addr
- 1: lo: <LOOPBACK,UP,LOWER_UP> mtu 16436 qdisc noqueue state UNKNOWN
- link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
- inet 127.0.0.1/8 scope host lo
- inet6 ::1/128 scope host
- valid_lft forever preferred_lft forever

LAN

- 2: eth1:
<BROADCAST,MULTICAST,UP,LOWER_UP>
mtu 1500 qdisc pfifo_fast state UNKNOWN
qlen 1000
- link/ether 00:19:66:ce:ee:1c brd ff:ff:ff:ff:ff:ff
- inet6 fe80::219:66ff:fece:ee1c/64 scope link
- valid_lft forever preferred_lft forever

Bridge

- 4: br0:
<BROADCAST,MULTICAST,UP,LOWER_UP>
mtu 1500 qdisc noqueue state UNKNOWN
link/ether 00:19:66:ce:ee:1c brd ff:ff:ff:ff:ff:ff
- inet 192.168.102.9/16 brd 192.168.255.255
scope global br0
- inet6 2001:4978:f:8640::63b2:9929/128
scope global
- valid_lft forever preferred_lft forever
- inet6 fe80::219:66ff:fece:ee1c/64 scope link
- valid_lft forever preferred_lft forever

SixXS Tunnel

- 6: sixxs: <POINTOPOINT,UP ...> mtu 1428
- inet6 2001:4978:hhhh:hhhh::2/64 scope global deprecated
- valid_lft forever preferred_lft forever
- inet6 fe80::4978:hhhh:hhhh::2/64 scope link
- valid_lft forever preferred_lft forever

Ayiya (UDP) (ME) 2001:4978:hhhh:hhhh::2 <--
> 2001:4978:hhhh:hhhh::1 (POP)

- aiccu 2127 root 6u IPv4 12451 0t0 UDP
99.178.153.41:54392->216.14.???.??:**ayiya**

Virtual Bridge

- 7: virbr0:
<BROADCAST,MULTICAST,UP,LOWER_UP>
mtu 1500 qdisc noqueue state UNKNOWN
- link/ether 52:54:00:d2:a7:29 brd ff:ff:ff:ff:ff:ff
- inet 192.168.122.1/24 brd 192.168.122.255
scope global virbr0
- 11: vnet0: mtu 1500 state UNKNOWN
- link/ether fe:54:00:ae:45:1d brd ff:ff:ff:ff:ff:ff
- inet6 fe80::fc54:ff:feae:451d/64 scope link

CentOS65

/etc/sysconfig/network-scripts/

- Ifcfg-eth0 ISP connection (2-wire 3800)
- Ifcfg-eth1 LAN (nVidia MCP61)
- route6-br0
2001:4978:f:8640::/64 dev br0
- Ifcfg-br0
 - TYPE=Bridge
 - V6INIT=yes
 - IPV6ADDR=2001:4978:f:8640::63b2:9929/128
 - IPV6FORWARDING=yes


```
# /usr/local/sbin/resetroutes
```

```
# D.R. Forrest 3/18/13
```

```
# resets IPv6 routes for maplepark.com:
```

```
# Necessary because we use fixed IPv6 for some  
hosts and rely on autoconf for others (visitors).
```

```
# CentOS6 does not yet easily establish the mixed  
use (both router w/IPv4 NAT and IPv6 workstation)  
of this machine. The radvd will not restart if  
forwarding is not enabled while V6 autoconf will be  
disabled if it is. Normal booting allows the proper  
sequencing but just doing the standard 6.4
```

```
service network restart makes radvd lose  
multicast communications and deletes our local  
static IPv6 route.
```

This routine is called in /bin/service network restart section by
/etc/rc.d/init.d/network after the rc=\$? (result code is set) like this:

```
# if [ -x /usr/local/sbin/resetroutes ]; then
```

```
    and restarts our router advertisement daemon, radvd, after  
# resetting IPv6 forwarding (on==1; does no harm if is not running)  
/bin/echo 1 > /proc/sys/net/ipv6/conf/all/forwarding      # on for forwarding as the router)  
/bin/echo 1 > /proc/sys/net/ipv4/ip_forward              # on for local NATing the W/S)
```

```
RADVD_ACTION="start"  
/usr/bin/pgrep radvd >/dev/null && RADVD_ACTION="restart"  
/sbin/service radvd $RADVD_ACTION
```

```
# if missing, set up a static route to dev br0 , so our local area network works.  
# A root crontab maintains it:
```

```
if [ /sbin/ip -6 route \|\bin/grep '2001:4978:f:8640::/64 ' &>/dev/null ]; then  
    /sbin/ip -6 route add 2001:4978:f:8640::/64 dev br0
```

DNS

(ISC BIND 9.10.0)

Our master external server (currently a view, "external") is dual-stacked, has A, AAAA, NS, MX, RRSIG, TXT (SPF), SPF, and LOC RRs from our firewall machine only, and gives out only authoritative information with no upward referrals. Zone transfer is allowed to our secondaries only. We do host a couple of virtual web sites which are also dual stacked. Currently, all public traffic enters through one IPv4 address, 99.178.153.41, (natted for locals), and IPv6 at 2001:4978:f:8640::/64.

What from Where?

-

<Queries via IPv6>

RR Type	Internal	External	Total	Internal	External
A	619567	95489	715056	619567	3491
AAAA	388164	39278	427442	388164	2202
DNSKEY	74	5216	5290	74	2044
MX	21514	8087	29601	21514	996
NS	19157	786	19943	19157	195
PTR	47067	422	47489	47067	79
TXT	32104	839	32943	32104	230
TOTALS	1127678	153243	1280921	1127678	9502
	88.0%	12.0%		100.0%	6.2%

-

Sendmail

We now receive SMTP via IPv6 from several domains without incident, including isc.org and gmail.com, but have had problems with some comcast.net mailers that required mailertable entries to force IPv4 SMTP sessions due to them rejecting my rDNS.

<http://owenhd.corp.he.net/ipv6/>

Owen DeLong

IPv6 Evangelist

Hurricane Electric

760 Mission Court

Fremont, CA 94539, USA

<http://he.net/>