# Table of contents



# Table of Figures

# Table of Tables

# Proposed Network

## Overview

This document will detail the steps required to reproduce the IPv4 / IPv6 dual stack network pictured in Figure 1. Where adjustments have to be made to suit individual organisations, they will be flagged up below. This document assumes that the network is already physically connected using Ethernet and Category 5 UTP or better.

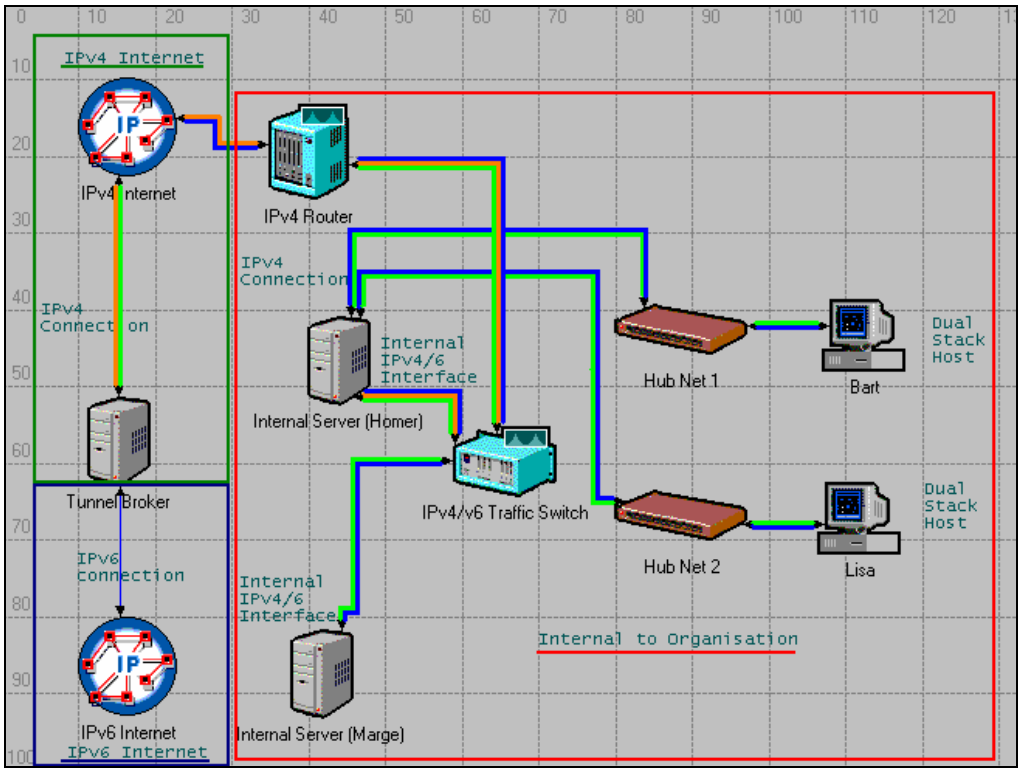


**Figure 1 - Topology**

Key:

Blue cable – IPv6

Green cable – IPv4

Orange cable – IPv6 over IPv4

Marge – Internal Server IPv4 / IPv6

Homer – Internal Server IPv4 / IPv6

Bart / Lisa – Dual Stack Hosts

Multi colour cable runs the chosen services over the same connection.

## Cost Estimation

The suggested distribution for this report is RedHat Linux version 8.0 (RedHat). The unit cost of RedHat is £34.07 incVAT. Cost involved in upgrading individual clients and servers is reflected in the choice of distribution. It includes IPv6 aware web browsers, IPv6 aware utilities and an Office suite, making it an inexpensive alternative to closed source operating systems. Cost savings in software and the additional resources available to the organisation when IPv6 connectivity is in place, should outweigh the loss of productivity whilst nodes are upgraded.

## Retraining

Staff may need retraining due to the move to a different operating system. The author does not see this as an issue that should concern the management of an organisation. RedHat uses the new "Blue Curve" desktop which will be very easy to navigate for anyone familiar with Windows or another Linux distribution. RedHat has a number of desktop interfaces from which the organisation can choose. All include easy-to-use applications needed for businesses use such as web browser(s), email clients and an office suite. Because of this the author considers "in house" training sufficient to aid the transfer of skills to RedHat.

# Installing the server

## Prerequisites

1 x Server Kickstart file located in appendix 6.6.2 which can be copied to floppy disk
1 x Set of RedHat 8.0 installation CD's
A computer with at least 2 Linux compatible Network Interface Cards (three cards used in this document)
A CD Rom drive
A floppy drive

# Server Installation:

## Pre-installation:

**Each configuration file is located in the appendices (section 6) of the main report for reference.**

Before the installation commences, the installer should make any changes to the boot floppy disk. Common changes might be to the server name or network device sections. The server name item is located in the "`--hostname` *`computer_name`*" parameter of the network line.
It is assumed that the installer will require the internal and external name of the server to be the same. If this is the case then both entries, for each network interface should be amended.
The server password for initial setup is "`simpsons`".

For the administrator to copy a file from the floppy disk, the floppy drive must be mounted. For speed and simplicity the administrator must be logged in as `root` to execute these commands.

1. Insert the floppy disk containing the source files into the floppy drive.

2. Click on RedHat (the RedHat symbol in the bottom left corner), System Tools, Terminal.

3. When the terminal has started ensure the `floppy` folder exists using `cd /mnt`, then `ls`. This command will display the folders currently within `/mnt`.

4. If the folder exists, use `mount /dev/fd0`.

5. The contents of the floppy disk are available in `/mnt/floppy`.

The author suggests copying the contents of the files in the appendices 6.2 to 6.8 to a floppy disk to be transferred to the server or clients. This will make installation and configuration much quicker.

## Installation

1. Ensure that the computer to be used is configured to boot from the CD Rom.

2. Copy the server version of `ks.cfg` as shown in 6.6.2 to a formatted floppy disk.

3. Boot the computer that will be the server with the Kickstart floppy in the floppy drive and CD 1 of RedHat in the CD Rom drive.

4. When installation screen appears (Figure 2) with the prompt "boot:", enter "`linux ks=floppy`" and press enter as show below in Figure 3.


**Figure 2 - Installation screen**

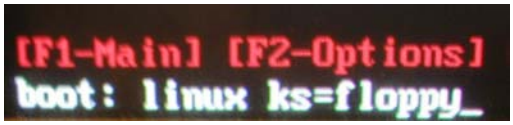


**Figure 3 – Kickstart installation command**

5. The automated installation file will be verified. Mistakes in configuration will cause the installation to error and stop. Any

mistakes must be resolved at this stage by reediting the file in a separate file editor and the process must be started again from step one.

6. The installation process will be completely automated apart from the hard disk partitioning section. This is due to the impact incorrect partitioning can have on the disk and performance.

7. If the installer is unsure how to partition the hard disk, it is suggested they consult the Official RedHat 8.0 Customisation Guide. This file is available from http://www.redhat.com/docs/manuals/linux/RHL-8.0-Manual/. Particular attention should be paid to section 1.5.

8. For this installation, it is assumed the server hard disk is a new, blank disk. In this situation, RedHat will automatically partition the disks as it sees fit. The installer is allowed to confirm the partitioning of the disk before the automated installation of the software begins. Should any alterations be needed, the following stages are where they should be made (see steps 9-15 below).

9. Select “Automatically Partition”, Next

10. When Figure 4 appears, select “Remove All Partitions” if any exist and click Next.

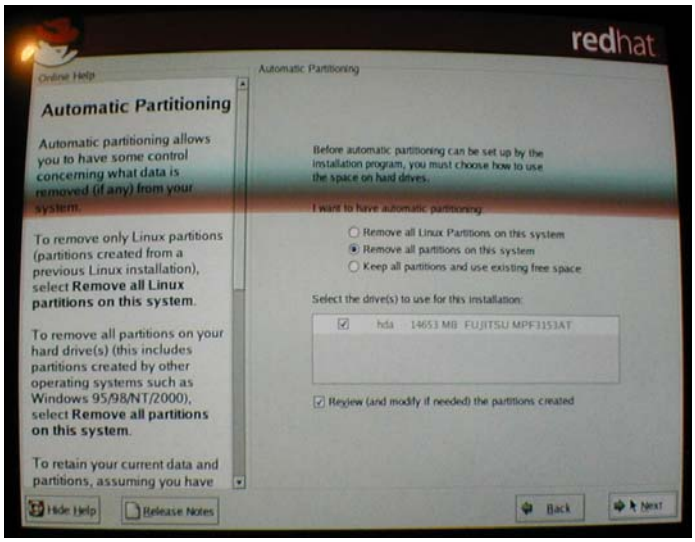


**Figure 4 - Hard disk partitioning**

11. Confirm "Yes" to removing all partitions from the hard disk(s). Click Next

12. Should any manual partitioning be required, it should be done on this page below (Figure 5). When partitioning has been confirmed, click Next.

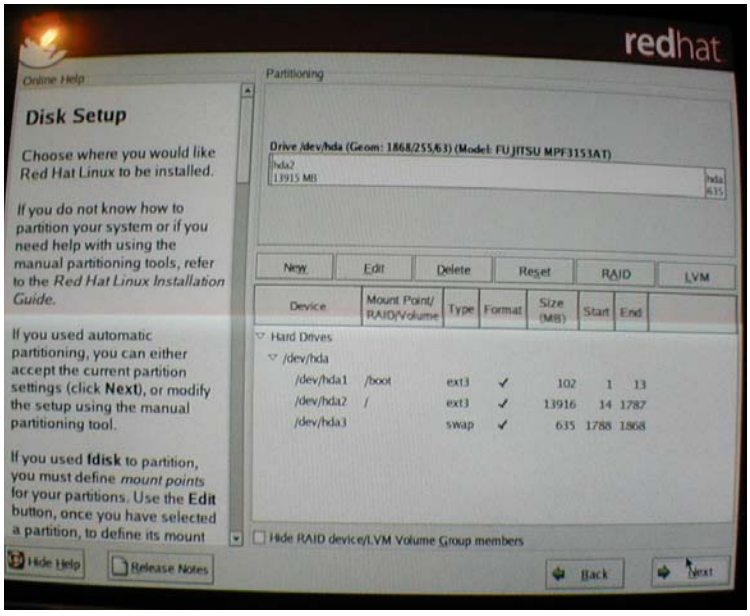


**Figure 5 - Hard disk partitioning confirmation**

13. The automated installation process will begin, installing the components needed for successful server deployment.

14. When required, insert the second and third installation CD's.

15. On completion of the installation (as shown below in Figure 6), click Exit. The server will reboot.

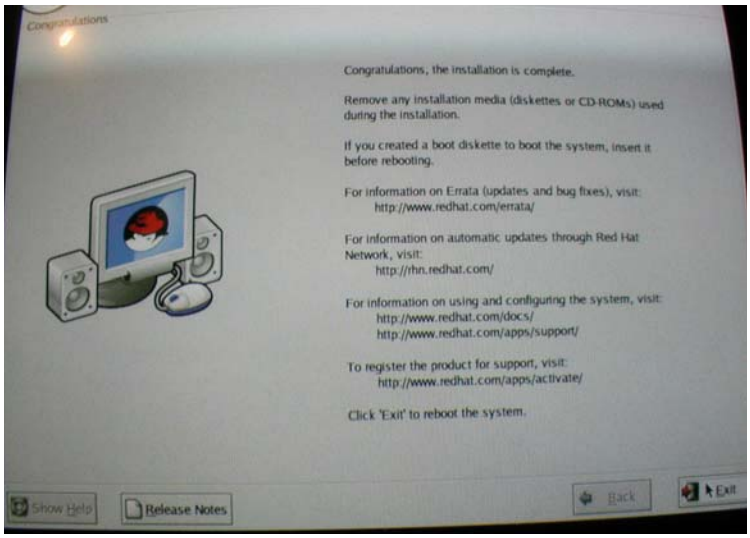


**Figure 6 - Completed Setup**

## Configuration

When the server login screen appears, login in as "`root`" with "`simpsons`" as the password.

Change the administrator password

1. Click RedHat (the RedHat symbol in the bottom left corner), System Tools, Terminal.

2. Enter the command `passwd`.

3. When prompted, enter a new password, and then confirm.

Kernel Support

The administrator should ensure that IPV6 kernel support is enabled. To do this, enter `test -f /proc/net/if_inet6 && echo "Kernel supports IPv6"` at a terminal prompt. If "Kernel supports IPV6" is returned configuration can continue. If the test was not successful, enter `modprobe ipv6` and execute the test again. The administrator should see an `inet6 addr` entry similar to the one shown below in Figure 7 when the administrator executes the `ifconfig` command. To allow IPv6 to be called upon when needed, edit `/etc/modules.conf` and add the line `alias net-pf-10 ipv6` if it does not already exist.

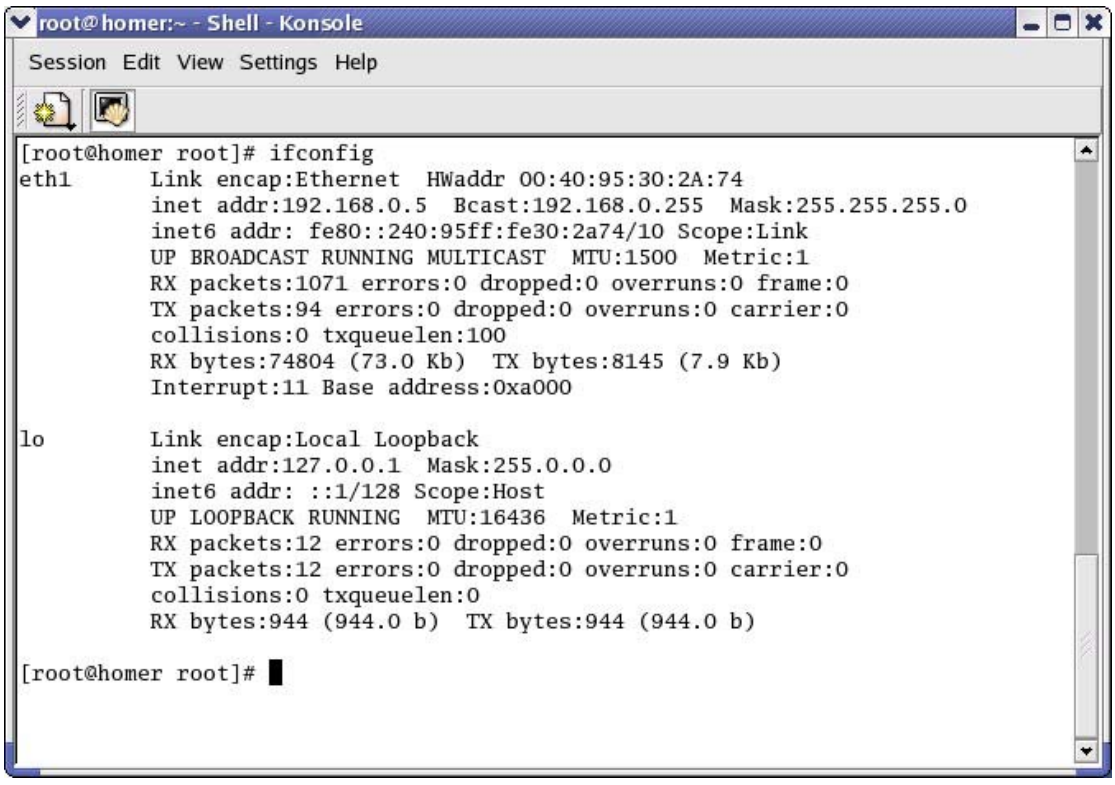


**Figure 7 - IPv6 link local address**

If `modprobe` is not successful, the administrator should investigate the resource at [Peter_Bieringer_IPv6] for instructions on how to recompile the kernel with IPv6 support. At this stage, the server is IPv6 ready, but without configuration information. The following section will detail the actions which must be taken to configure the server ready for commissioning.

Interfaces

1. The first action is to configure the interfaces with static IPv4 addresses through the network configuration tool. The network configuration tool is found by clicking on RedHat, System Settings, Network.

2. The installer must determine which interface will be used for internal communications and which will be used for external communications. Ideally, the quicker interface should be used internally because the WAN connection (across ADSL for example) is much slower. Therefore a slower 10Mbps connection would not have any detrimental effect on network performance. If a 10Mbps NIC is used on a 100Mbps internal network, network traffic to the server will be limited by the speed with which the server can receive files.

3. Once a choice has been made, each device in turn should be set with the correct IPv4 address, subnet mask and gateway address. Segmentation of the IPv4 networks is beyond the scope of this document though further information can be found at http://www.3com.com/other/pdfs/infra/corpinfo/en_US/501302.pdf. It is strongly recommended that each interface is assigned a network address from the private address range and NAT is used. In the example configuration below (Table 1) the following private addresses have been chosen

| Interface / Alias | IP/subnet mask | Gateway |
|---|---|---|
| Eth**0** | 192.168.**0**.5/24 | 192.168.**0**.5 |
| Eth**1** | 192.168.**1**.5/24 | 192.168.**1**.5 |
| Eth**2** | 192.168.**2**.5/24 | 192.168.**2**.5 |

**Table 1 - IPv4 Interface configuration**

4. The information given above is for use with the topology given in Figure 1. Each internal LAN connected to an interface is a Class C network with 254 available hosts per network.

5. IPv6 interfaces are assigned the following static addresses in `/etc/rc.d/rc.local` which allows them to be used with the IPv6 stateless client configuration section below (address on same line as command):

```
ifconfig eth1 inet6 add
fec0:0000:0000:0001:2e0:4cff:feec:d671/64
ifconfig eth2 inet6 add
fec0:0000:0000:0002:240:95ff:fe30:2a74/64
```

IPv4 / IPv6 Packet Forwarding

Before the server can act as a router and forward packets on to other networks or the internet it must be configured to do so. This is achieved through the addition of two commands to `/etc/rc.d/rc.local` (see appendix 6.7).

Adding `echo 1 > /proc/sys/net/ipv4/ip_forward` and `echo 1 > /proc/sys/net/ipv6/conf/all/forwarding` to this file will ensure the computer will act as a router after being rebooted.

Due to the sequence in which interfaces are initialised (eth0, eth1 then eth2)

default routes will be incorrectly configured. To resolve this issue, add `route add default gw 192.168.0.1 dev eth0` and `route del default gw 192.168.2.5 dev eth2` to `/etc/rc.d/rc.local`.

IPv4 DHCP

DHCP is used in this network to allow hosts to connect to existing IPv4 resources and obtain information needed to access IPv6 resources. Manual installation and configuration of DHCP can be found at http://www.isc.org/products/DHCP/dhcpv3-README.html.

1. If the server has been installed from the Kickstart floppy disk, the DHCP server software is already in place but not yet configured. To configure the software, copy the `dhcpd.conf` file from appendix 6.2.4 to `/etc/`. The file must then be edited to suit the organisation. A sample of the file has been reproduced below with a commentary to aid the configuration.

```
ddns-update-style interim;
subnet networkname netmask subnetmask{
      range low_IP high_IP;
      default-lease-time seconds;
      max-lease-time time>=default above(seconds);
      option subnet-mask subnetmask;
      option broadcast-address network_broadcast address;
      option routers default_router_IP;
      option domain-name-servers DNS_Server_IP;}
```

`ddns-update-style`, can be either ad-hoc or interim. Interim is the preferred update method, used in dynamic DNS updates.
`Subnet` defines the network name and corresponding subnet mask. This informs the DHCP application to which interface, information, enclosed within the braces, should be advertised to.

`Range` defines the range of IP addresses that can be assigned, between the lower and upper limit

`default-lease-time` gives the time taken before the lease expires (and should be renewed). An attempt is normally made to renew the lease before this time.

`max-lease-time` details the maximum time by which the lease must be renewed. If the lease is not renewed by this time, the IP address must be released

`subnet-mask` defines the subnet mask assigned to hosts obtaining an address from this server.

`broadcast-address` specifies the address used to communicate with all hosts on this sub network

`routers` details the address of the default router for this network

`domain-name-servers` specifies the address of the DNS server used for name / address resolution.

2. At a terminal prompt, enter `touch /var/lib/dhcp/dhcpd.leases`. This lease file is required by the DHCP server to record the leases currently assigned and must exist before DHCP can be started.

3. When the file has been configured with the information required, the DHCP service can be started by clicking RedHat, Server Settings, Services. Ensure run level 5 is to be edited. Tick dhcpd and click save.

4. In a terminal window, execute `/usr/sbin/dhcpd -d -f` and check for errors.

5. In addition, view the log file `/var/log/messages` for any error messages in the configuration.

6. Assuming the service starts successfully, the server is now able to assign addresses to clients as soon as the server is rebooted.

7. When a client is started, the administrator should check `/var/log/messages` to see a message such as the one below which confirms the server is assigning IPv4 addresses (see Figure 8).

```
Apr  3 21:14:47 homer dhcpd: DHCPREQUEST for 192.168.1.20 from 00:40:95:30:1d:09 via eth1
Apr  3 21:14:47 homer dhcpd: DHCPACK on 192.168.1.20 to 00:40:95:30:1d:09 via eth1
```

**Figure 8 - DHCP Example**

IPv6 stateless client auto configuration

Just as clients require configuration for IPv4, the same is true for IPv6. This can be accomplished most simply through the use of a router advertising application. Within the network above there are three networks which will need site local prefixes. A site local prefix is 64 bits in length with the first 48 bits set to fec0:0000:0000:$_{hex}$ and the following 16 bits used to define the site local network identifier. To ensure ease of administration, it is recommended that a relationship is made between the IPv4 sub networks and the IPv6 site local prefix. This is shown below in Table 2. Though different in size, this scheme should be helpful to administrators.

| IPv4 sub network | IPv6 site local prefix |
|---|---|
| 192.168.**0**.0/24 | FEC0:0000:0000:000**0**::/64 |
| 192.168.**1**.0/24 | FEC0:0000:0000:000**1**::/64 |
| 192.168.**2**.0/24 | FEC0:0000:0000:000**2**::/64 |

**Table 2 - IPv4 to IPv6 mapping**

In this test network, radvd has been pre-installed on the server as part of the Kickstart installation, though it still requires additional configuration. Manual installation and configuration of the software can be found at http://v6web.litech.org/radvd.

1. Copy the `radvd.conf` file from the floppy or appendix 6.2.1 to `/etc/`.

Below is a section of the file, detailing the configuration of eth1.

```
interface eth1
{
      AdvSendAdvert on;
      MinRtrAdvInterval 3;
      MaxRtrAdvInterval 10;
      AdvHomeAgentFlag off;
      prefix fec0:0000:0000:0001::/64
      {
            AdvOnLink on;
            AdvAutonomous on;
            AdvRouterAddr on;
      };
      prefix 2001:618:400:3afc::/64
      {
            AdvOnLink on;
            AdvAutonomous on;
            AdvRouterAddr on;
      };
};
```

The configuration defines all actions for eth1 within a set of braces. The first four lines configure the interface to send adverts and the time span between each transmission.

The items within each prefix line (within the braces) detail the prefix address to be advertised on this interface. The first prefix is used to assign a site-local prefix to all hosts accepting auto-configuration information on that network. The second prefix is used by the clients when accessing IPv6 services external to the organisation. The prefix used here is specific to each site and will be covered in the subsection on Tunnel configuration. Once the configuration is correct, ensure forwarding is enabled as discussed in IPv4 / IPv6 Packet Forwarding above.

2. Radvd can then be started by typing `radvd` at a terminal prompt. It is always good practice to ensure that `/var/log/messages` has not logged any errors.

3. The administrator can test the operation of radvd by entering `radvdump` at a terminal prompt. This executes an application that displays the adverts for each interface. The administrator can compare the output received with that provided in appendix 6.2.2.

4. Radvd can be started automatically by clicking RedHat, Server Settings, Services. Ensure run level 5 is to be edited. Tick radvd and click save.

5. In some situations, radvd will not run using the technique above, if this is the case, edit `/etc/rc.d/rc.local` and insert `radvd` as in `/etc/rc.d/rc.local` (see appendex 6.7).

6. To complete the setup, the server should be restarted.

NAT Configuration

Although after completing all of these steps IPv6 traffic can be routed to the internet, the same is not true for IPv4 traffic. Though the server is set up for routing, the default firewall rules do not allow the server to forward packets on all interfaces. To allow IPv4 traffic to be routed, and allow all hosts to reach the DNS / tunnel services on Homer, the following commands should be entered at a terminal prompt. NB: the last two entries are to confirm that IPv6 traffic should be allowed between Homer and asterix.ipv6.bt.com. The final command is used to save the firewall configuration.

```
Iptables -F
iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
iptables -N block
```

```
iptables -A INPUT -j block
iptables -A FORWARD -j block
iptables -A block -m state --state RELATED,ESTABLISHED -j
ACCEPT
iptables -A block -i ! eth0 -m state --state NEW -j ACCEPT
iptables –A INPUT –i eth0 –p ipv6 –s 193.113.58.80 –j ACCEPT
iptables –A OUTPUT –o eth0 –p ipv6 –d 193.113.58.80 –j ACCEPT

service iptables save
```

Tunnel configuration

To access the IPv6-enabled internet, the organisation will require an account with a tunnel broker. The administrator is advised to configure a tunnel to asterix.ipv6.bt.com.

This service is free of charge but activation of the tunnel end point may take a number of hours. The service, provided by BT, can be accessed via https://tb.ipv6.btexact.com.

1. The administrator should first set up an account with the tunnel broker. Once this is complete, tunnels can be defined.

2. When the user's account has been confirmed via email, log on to the tunnel broker service (using addresses above).

3. Select "create a new tunnel" and enter the organisation's current public IPv4 address (the address assigned to the device connecting the organisation to the IPv4 internet) and additional information specific to the organisation. From the page show in Figure 9, configure the tunnel as shown in Figure 10:

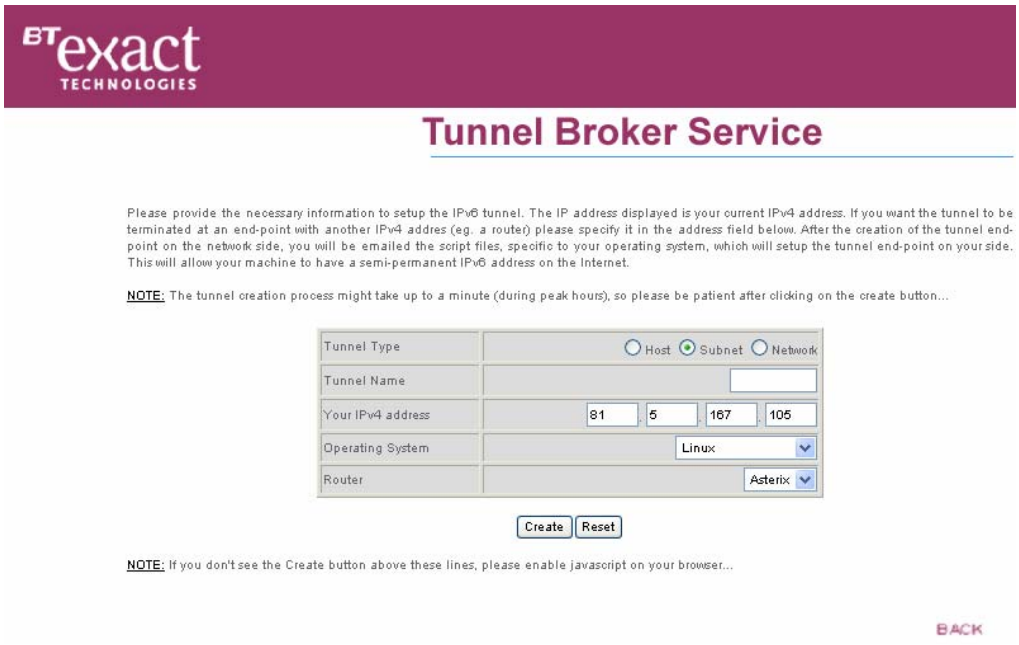


**Figure 9 - Tunnel Broker IPv6 tunnel creation**

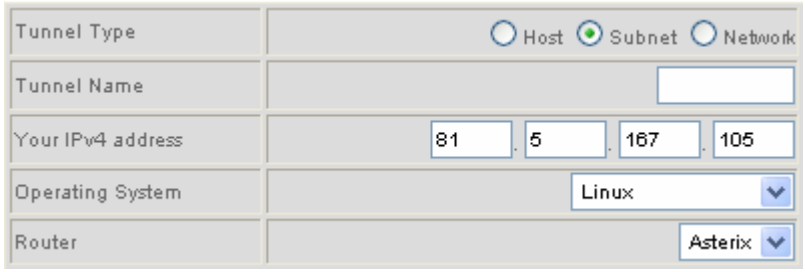


**Figure 10 - IPv6 tunnel settings**

4. Choosing “Subnet” allows the tunnel broker service to be applied to all hosts within the organisation. Clicking “Create” will create the tunnel ready for use.

5. When Figure 11 loads, click on the “send” button and the configuration script can be downloaded. This should allow the administrator to configure the server for tunnel services. The administrator should be aware the commands within the script often do not work. By viewing the contents of `/etc/rc.d/rc.local` the administrator can compare the downloaded file with the tunnel setup section in this file and make adjustments to the organisation’s tunnel entries in `/etc/rc.d/rc.local`. Below is a portion of this file with

commentary which details the role of each command to the administrator making configuration easier.

```
#TUNNEL SETUP
```

`ifconfig sit0 up` - Enable sit0 pseudo-interface

`ifconfig sit0 tunnel ::193.113.58.80` - Define asterix.ipv6.bt.com as the tunnel end point, using the public IPv4 address of asterix.

`ifconfig sit1 add 2001:618:400::515:a769/128` – Configure the local sit1 device with a static IPv6 global address.

`route -A inet6 add 2000::/3 gw fe80::c171:3a50 dev sit1` – Instruct sit1 to use the tunnel end point as the default route for all internet traffic (2000::/3) using the link local address of the end point. The default all-zeros address is not used due to issues in the kernel.

`ifconfig eth1 add 2001:618:400:3afc::/64` – Add the AGUA prefix to eth1.

`ifconfig eth2 add 2001:618:400:3afc::/64` – Add the AGUA prefix to eth2.

`radvd` – Start the radvd service to advertise routes to specific interfaces now

`bind` – Start BIND if it is not already running

`#END TUNNEL SETUP` - tunnel setup is complete

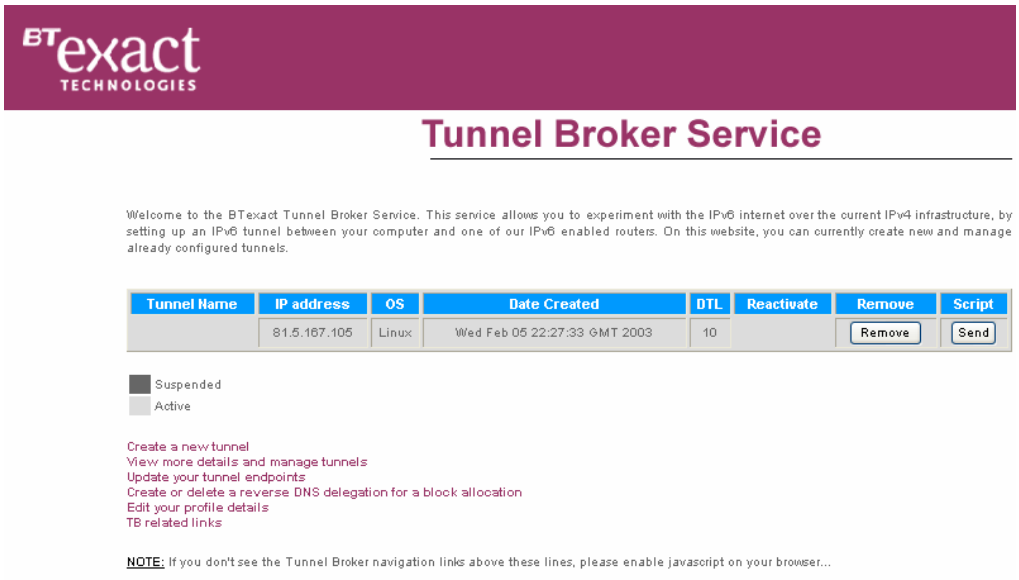


**Figure 11 - Tunnel Broker main menu**

6. On reboot of the server, the administrator should be able to access IPv6 resources from the server or internal networks which have been allowed external access. See Figure 12 below.

```
[root@homer root]# ping6 -I eth0 -c 3 www.kame.net
PING www.kame.net(apple.kame.net) from 2001:618:400::515:a769 eth0: 56 data bytes
64 bytes from apple.kame.net: icmp_seq=1 ttl=55 time=301 ms
64 bytes from apple.kame.net: icmp_seq=2 ttl=55 time=300 ms
64 bytes from apple.kame.net: icmp_seq=3 ttl=55 time=307 ms

--- www.kame.net ping statistics ---
3 packets transmitted, 3 received, 0% loss, time 2003ms
rtt min/avg/max/mdev = 300.810/303.383/307.616/3.082 ms
[root@homer root]#
```

**Figure 12 - Ping6 to www.kame.net**

7. The administrator can monitor usage of the tunnel service by logging onto the tunnel broker site, by choosing “View more details and manage tunnels”, and clicking on “Statistics” for the chosen tunnel. Below, in Figure 13, is the information for the tunnel between the authors existing network and asterix.ipv6.bt.com.

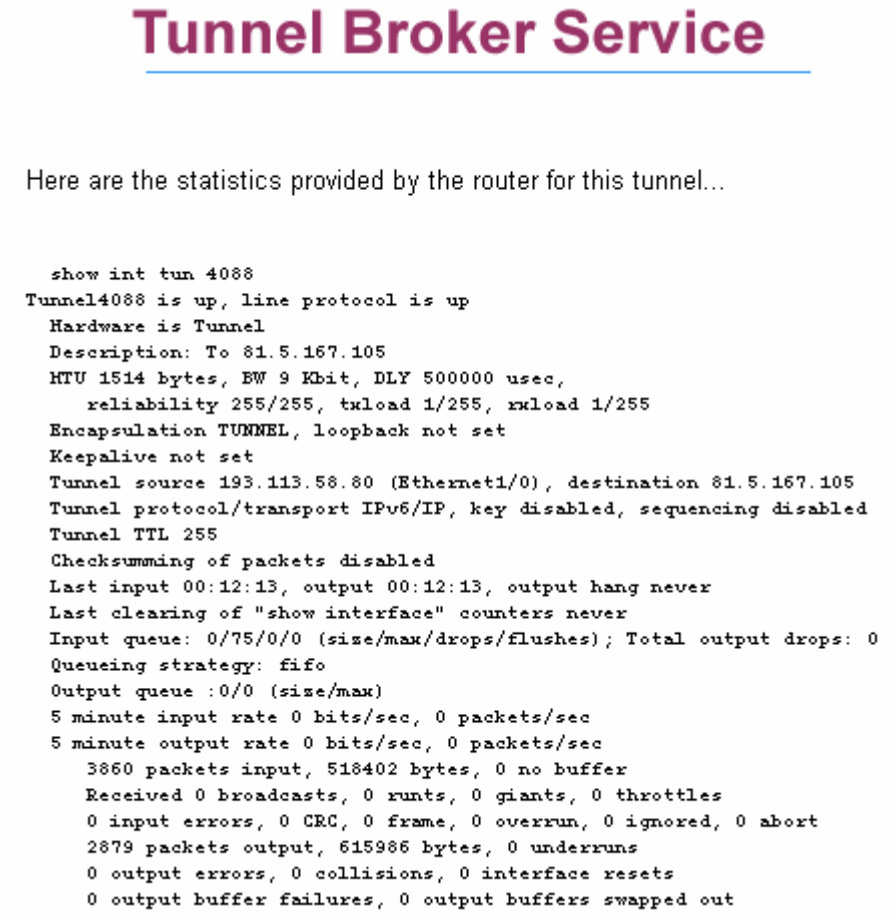


**Figure 13 - Tunnel statistics**

## BIND Configuration

To allow hosts to access IPv6 services by name either internally or externally BIND was installed as part of the automated installation. The provided `named.conf` in appendix 6.2.3 should be applied to the server and adjusted to suit. The configuration provided is for a caching primary name server. Manual installation and configuration of the software can be found at http://www.isc.org/products/BIND/bind9.html.

Configuration of the services is firstly through the `named.conf` file. The author will give an overview of each type of zone file which can then be adjusted to suit the needs of the administrator and the organisation. Due to the competing standards, the author has provided configuration information for

INT, ARPA, AAAA and A6 records. This should ensure which ever standard is adopted worldwide, the organisation will be ready.

**The "." Zone** – used to query worldwide root name servers for names / addresses of hosts the local server is currently unaware of. This entry should not be modified. IPv4 and IPv6 lookups are possible.

**The "localhost" Zone** – used to provide loop back services to the client which is in this case the server. Contains IPv4 and IPv6 forward lookup.

**The "simpsons.com" Zone** – used within the organisation to detail the name to address information for IPv4 and IPv6 hosts. The author is aware that it is not normal practice to enter DNS entries for hosts which are dynamically assigned addresses, though these are included for testing purposes.

**The "0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.c.e.f.ip6.int" Zone** – used for reverse lookup for hosts on the fec0:0000:0000:0000 site-local network using the INT format of request. Each site local network prefix will need an entry of this type in `named.conf`.

**The "\[xfec0000000000000/64].ip6.arpa" Zone** – used for reverse lookup for hosts on the fec0:0000:0000:0000 site-local network using the ARPA format of request. Each site local network prefix will need an entry of this type in `named.conf`.

The entry `listen-on-v6 { any; };` in the `named.conf` options allows BIND to listen for IPv6 DNS queries. Each zone has an associated file which contains information about that zone in `/var/named/` (as defined by the options header).

Using:

```
zone "simpsons.com" {
      type master;
```

```
      file "db.simpsons.com";
};
```

as an example, the entry "db.simpsons.com"; informs the software that when a query is made for the simpsons.com domain, the answer file "db.simpsons.com" should be checked for more information.

## Server conclusion

To conclude, the server should now be configured with IPv4 and IPv6 interfaces with static addresses, DNS, DHCP, Router advertising, NAT, forwarding and tunnelling. This should allow the hosts within the correct network to access IPv4 and IPv6 services. If configured correctly, the server should have a list of interface settings approximately the same as the one given below with any amendments made, as required, to suit the organisation.

```
eth0      Link encap:Ethernet  HWaddr 00:40:95:45:56:44
          inet addr:192.168.0.5  Bcast:192.168.0.255
Mask:255.255.255.0
          inet6 addr: fec0::240:95ff:fe45:5644/64 Scope:Site
          inet6 addr: fe80::240:95ff:fe45:5644/10 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:8552 errors:0 dropped:0 overruns:0 frame:0
          TX packets:7376 errors:0 dropped:0 overruns:0
carrier:0
          collisions:168 txqueuelen:100
          RX bytes:3337334 (3.1 Mb)  TX bytes:969352 (946.6 Kb)
          Interrupt:11 Base address:0xc000

eth1      Link encap:Ethernet  HWaddr 00:E0:4C:EC:D6:71
          inet addr:192.168.1.5  Bcast:192.168.1.255
Mask:255.255.255.0
          inet6 addr: 2001:618:400:3afc::/64 Scope:Global
          inet6 addr: fec0::1:2e0:4cff:feec:d671/64 Scope:Site
          inet6 addr: fe80::2e0:4cff:feec:d671/10 Scope:Link
```

```
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:4799 errors:0 dropped:0 overruns:0
carrier:0
          collisions:0 txqueuelen:100
          RX bytes:0 (0.0 b)  TX bytes:527474 (515.1 Kb)
          Interrupt:3 Base address:0x4000

eth2      Link encap:Ethernet  HWaddr 00:40:95:30:2A:74
          inet addr:192.168.2.5  Bcast:192.168.2.255
Mask:255.255.255.0
          inet6 addr: 2001:618:400:3afc::/64 Scope:Global
          inet6 addr: fec0::2:240:95ff:fe30:2a74/64 Scope:Site
          inet6 addr: fe80::240:95ff:fe30:2a74/10 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:6 errors:0 dropped:0 overruns:0 frame:0
          TX packets:4857 errors:0 dropped:0 overruns:0
carrier:0
          collisions:0 txqueuelen:100
          RX bytes:996 (996.0 b)  TX bytes:688496 (672.3 Kb)
          Interrupt:10 Base address:0x9000

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:18 errors:0 dropped:0 overruns:0 frame:0
          TX packets:18 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:1532 (1.4 Kb)  TX bytes:1532 (1.4 Kb)

sit0      Link encap:IPv6-in-IPv4
          inet6 addr: ::192.168.2.5/96 Scope:Compat
          inet6 addr: ::127.0.0.1/96 Scope:Unknown
          inet6 addr: ::192.168.1.5/96 Scope:Compat
          inet6 addr: ::192.168.0.5/96 Scope:Compat
          UP RUNNING NOARP  MTU:1480  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)
```

```
sit1      Link encap:IPv6-in-IPv4
          inet6 addr: 2001:618:400::515:a769/128 Scope:Global
          inet6 addr: fe80::c0a8:205/10 Scope:Link
          inet6 addr: fe80::c0a8:105/10 Scope:Link
          inet6 addr: fe80::c0a8:5/10 Scope:Link
          UP POINTOPOINT RUNNING NOARP  MTU:1480  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:1111 (1.0 Kb)  TX bytes:949 (949.0 b)
```

# Installing a client

## Prerequisites

1 x Client Kickstart file located in appendix 6.2.1 which can be copied to floppy disk
1 x Set of RedHat 8.0 installation CD's
A computer with 1 Linux compatible Network Interface Card
A CD Rom drive
A floppy drive

# Client Installation:

## Pre-installation:

The instructions from server pre-installation will apply to client installations where necessary.

## Installation

Installation of clients will follow the same Kickstart procedure as a server, though the Kickstart file for clients, as shown in appendix 6.6.1, will apply. On successful application of the Kickstart file, the client computer should be rebooted. The Administrator should log into the system and configure the users for this computer. This is achieved by clicking on RedHat, System Settings, Users and Groups, and using this tool to add additional users. The

administrator should also confirm the operation of the services as users are not permitted to carry out some of the necessary checks due to restrictions on their account.

## Client conclusion

If the checks are successful, the administrator should log off, and allow the user to log onto the workstation and test services and connectivity. The user / administrator can then install additional software and services, many of the IPv6 enabled services are listed in [IPv6_Enabled_Apps].

Administrators should be aware that depending on which record type is returned in a DNS query, different actions can take place. For example, when using Mozilla, if the address of site with both IPv4 / IPv6 addresses is requested, the IPv4 site address is often returned (as this is the first entry). If this is an issue, the user can manually ping6 the site, specifying the name, which will return the IPv6 address. This address can then be copied into the address bar and surrounded by braces as shown below.

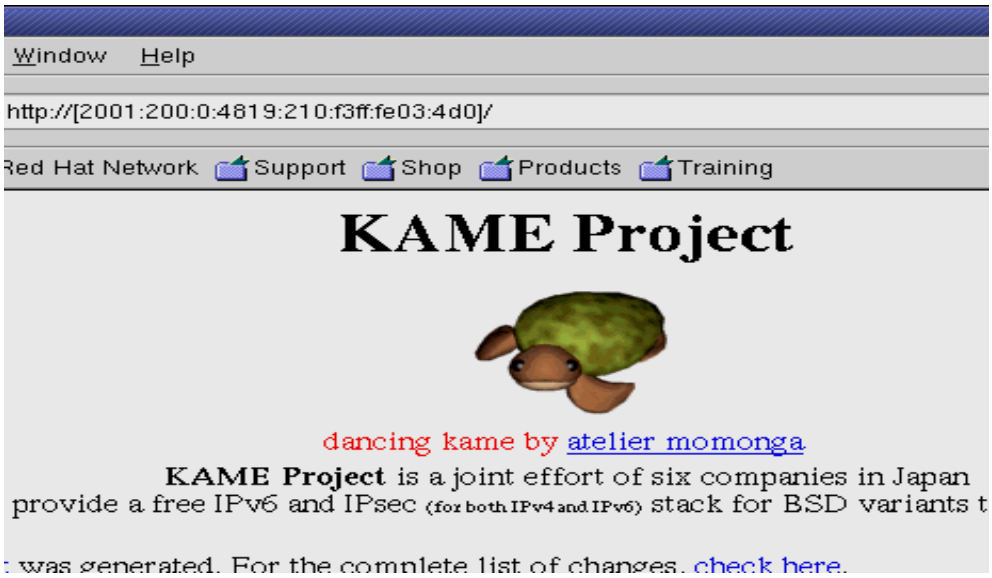


**Figure 14 - Example use of IPv6 address in Mozilla**