

Adding Your Cray Research System to Your Network [11]

This chapter describes the minimal steps you must take to place your CRAY J90 system on an existing TCP/IP network. After your CRAY J90 system is a functional member of the network, you can do additional configuration. This chapter also includes a table that describes some of the most common TCP/IP configuration files.

After you have placed your CRAY J90 system on an existing TCP/IP network by using the information in this chapter, see the *UNICOS Networking Facilities Administrator's Guide*, Cray Research publication SG-2304, for additional networking configuration information; this publication also includes which menus to use for various networking tasks.

11.1 Related network information

The following publications contain additional information that will be of use to you:

- *UNICOS Networking Facilities Administrator's Guide*, Cray Research publication SG-2304
- *UNICOS Administrator Commands Reference Manual*, Cray Research publication SR-2022 (man pages):

arp(8)	inetd(8)	rexecd(8)
enstat(8)	initif(8)	rlogind(8)
fingerd(8)	mkbinhost(8)	route(8)
ftpd(8)	netstart(8)	rshd(8)
gated(8)	ntalkd(8)	sdaemon(8)
hyroute(8)	ping(8)	tcpstart(8)
ifconfig(8)		traceroute(8)

Note: To add your CRAY J90 system to an existing TCP/IP network, you must have several configuration files. The easiest way to create these files is to configure the system to talk to another host on the network, copy the necessary files from that machine to your system, and then change them. The first five steps of this procedure make the changes to allow you to talk to another host. You then copy files you need and change them for your system. This procedure assumes you are either an administrator of your network or that you will have a network administrator as a resource when you add your system to your existing TCP/IP network.

1. Verify that the network section of the `/sys/param` file contains the proper configuration for your system. See the Network section, Section 5.9.3.5.4, page 69, for details.
2. Create a minimal `/etc/hosts` file. (Do not overwrite the existing `/etc/hosts` file.)

The `/etc/hosts` file contains the database of all locally known hosts on the TCP/IP network. Create an `/etc/hosts` file that contains a local host entry, entries for the CRAY J90 system (one per interface), and an entry for at least one other host on the same network as the CRAY J90 system. The entry format is as follows:

IPaddress host_name annotations

Example:

```
# cat /etc/hosts
127.0.0.1      localhost loghost
(local host)
128.192.16.8   cray  cray-eth
(your CRAY J90 system)
128.192.16.125 cyclone cyclone-eth1
(other host)
```

Note: Contact your network administrator for the internet address of the Cray system.

3. Compile a binary hosts file.

Cray Research systems support a binary `/etc/hosts` file called `/etc/hosts.bin`. Create this file by using the `/etc/mkbinhost` command, as follows:

```
# /etc/mkbinhost
/etc/hosts.bin: 3 entries written
```

4. Update the `/etc/config/interfaces` file.

The `/etc/config/interfaces` file defines all network interfaces on the Cray Research system. Change the host name for each interface on your system to match those you chose in step 2; for additional information, see the `initif(8)` man page and the *UNICOS Networking Facilities Administrator's Guide*, Cray Research publication SG-2304. The entry format is as follows:

interface_name _family address ifconfig parameters

Example:

Note: All host names must be defined in your `/etc/hosts` file.

```
# cat /etc/config/interfaces

...some comment lines omitted...

lo0 - inet localhost -
en0 - inet cray.001 - netmask 0xffffffff00
fddi0 - inet cray.fddi - netmask 0xffffffff00
atm0 - inet cray.atm - netmask 0xffffffff00
```

5. Activate the changes by executing the `/etc/initif` script; an example follows:

```
# /etc/initif
Configuring all network interfaces: lo0 en0 fddi0 atm0
```

6. Create a default route.

This step creates a default route to let you communicate with hosts that are on different networks than the Cray Research system. To reach hosts that

are not on the same FDDI or Ethernet network as the CRAY J90 system, you must have a route. To create a route, execute the `/etc/route` command, as shown in the following example:

```
# /etc/route add defaultotherhost
add net default: gateway otherhost
```

The *otherhost* is the IP address (or *name* as shown in the `/etc/hosts` file) of a host that is on the same network as the CRAY J90 system and connects to one or more additional networks.

Place this command in the `/etc/tcpstart.mid` script so that it will be run automatically at system startup.

7. Test the network.

Test the network by using the `ping` command and view the configuration by using the `netstat` command. The `ping` command tests whether you can reach another host on the network. If `ping` succeeds, you can be confident that the hardware and routing works on all hosts and gateways between you and the system to which you are sending `ping`. The `netstat` command has many options. The `-i` option lets you view a table of cumulative statistics for transferred packets, errors, and collisions for each interface that was autoconfigured. The interfaces that are statically configured into a system but are not located at boot time are not shown; the `-r` option lets you view the routing table. The network address (currently Internet-specific) of the interface and the maximum transmission unit (mtu) in bytes also are displayed. You should become familiar with how these displays look on your system so that you will recognize changes and problems immediately.

Examples:

```
# /etc/pingotherhost
PING otherhost : 56 data bytes
64 bytes from 123.123.12.13: icmp_seq=0. time=10. ms
CONTROL-c
```

```
# netstat -i
Name      Mtu  Network      Address      Ipkts      Ierrs      Opkts      Oerrs
en0*      1496  cray-net     cray         0          0          2          0
fddi0     4352  crau-fddi-net  cray-fddi    249466    0          57636     0
fddi1*    4352  none        none         0          0          0          0
lo0       65535 loopback     localhost    264        0          264       0
```

Note: An * in the Name column of the `netstat -i` command output indicates that the interface is not configured up, so your CRAY J90 system cannot access that network.

8. Reboot the system.

Reboot the system to verify that all of your changes are handled automatically during system startup. (For booting procedures, see Chapter 3, page 19.) At some point, you should again test the network by using `ping` and view the configuration by using `netstat`.

9. Transfer full configuration files from another system, as shown in the following example.

Save copies of your original files and add the new entries for the CRAY J90 system to the files you transfer (for example, use `ftp` to transfer the `/etc/hosts` file from another system on your network).

Example:

```
# cd /etc
# cp hosts hosts.sav
# ftp cyclone
Connected to cyclone.cray.com.
220 fred FTP server (Version 5.2 Fri Feb 18 14:09:58 CDT 1994) ready.
Remote system type is UNIX.
Using binary mode to transfer files.
Name (fred:root): sam
331 Password required for sam
Password: <-----
Enter your password

230 User sam logged in.
ftp> get /etc/hosts hosts
200 PORT command successful.
150 Opening BINARY mode data connection for /etc/hosts (328758 bytes).
226 Transfer complete.
328758 bytes received in 0.6 seconds (5.3e+02 Kbyte/s)
ftp> quit
221 Goodbye.
# /etc/mkbinhost
/etc/hosts.bin: 2675 entries written
```

Note: You should check for entries for the J90 system in the `/etc/hosts` file. If they are not there, see Procedure 34, step 2, page 277.

Your CRAY J90 system should now be on the network.

11.2 Domain name service (DNS)

If you want to use domain name service (DNS) to perform host name lookup, you should configure your CRAY J90 system as a caching-only server. This should be done for the following reasons:

- A caching-only server is more efficient than a remote server (resolver only) because it maintains a cache of data and, therefore, requires less frequent network access.
- A caching-only server does not have authority over a particular zone, therefore, it does not have to answer queries from other authoritative servers.

- A caching-only server does not have to load configuration files from disk like a primary master server or across the network like a secondary master server, which gives it a faster start-up time.

To configure your Cray Research system as a caching-only server, you may configure both resolver and the local name server. You can perform most of this configuration by using the UNICOS Installation / Configuration Menu System.

Procedure 35: Configuring a caching-only server by using the menu system

Note: If you have not completed the previous procedure, "Adding a CRAY J90 system to an existing TCP/IP network," you should complete steps 3, 4, and 5 of that procedure before configuring a caching-only server.

To configure your CRAY J90 system as a caching-only server by using the UNICOS Installation / Configuration Menu System, complete the following steps:

1. Select YES for the "Use domain name service?" entry in the UNICOS Installation / Configuration Menu System. The menu system creates the /etc/hosts.usenamed file. The existence of this file indicates that the UNICOS system will use DNS, rather than the /etc/hosts file to look up host names. A sample TCP/IP Host/Address Lookup Configuration menu screen follows:

Configure System

```
->Network Configuration
    ->TCP/IP Configuration
        ->TCP/IP Host/Address Lookup Configuration
            ->TCP/IP Local Domain Name Server Config
```

TCP/IP Host/Address Lookup Configuration	
Use Domain Name (DN) service ?	YES
S-> DNS lookup (resolver) ==>	
Local DN server (named) ==>	

2. Configure the resolver, which consists of creating the /etc/resolv.conf file, which is created if you place information in the DNS lookup (resolver) menu. When you have a local name server (named process) running, you should have the local host address (127.0.0.1) as the first name

server; otherwise your local named will be bypassed, resulting in decreased performance and increased network traffic.

The following is an example of the menu screen:

```

TCP/IP Domain Name Service Lookup (resolver) Configuration
Local domain name                cray.com
Address for Domain Name server #1: 127.0.0.1
S-> Address for Domain Name server #2: 128.162.19.7
Address for Domain Name server #3: 128.162.19.13
Address for Domain Name server #4:
Address for Domain Name server #5:
Address for Domain Name server #6:
Address for Domain Name server #7:
Address for Domain Name server #8:
Address for Domain Name server #9:
    
```

3. Configuring a caching-only local name server consists of creating the `named.boot`, `root.cache`, and `localhost.rev` files. You can do this by using the Local DN server (named) menu. Perform the following steps to create these files:

- a. The `named.boot` file is read when `named` starts up. It tells the server what kind of server it is, over which zones it has authority, and where to get its initial data. The following is an example of the menu screens:

```

Configure System
->Network Configuration
  ->TCP/IP Configuration
    ->TCP/IP Host/Address Lookup Configuration
      ->TCP/IP Local Domain Name Server Config
    
```

```

TCP/IP Host/Address Lookup Configuration

Use Domain Name (DN) service ?          YES
DNS lookup (resolver) ==>
S-> Local DN server (named) ==>
    
```

```

TCP/IP Local Domain Name Server Configuration

S-> Directory for name server files           /etc/named.d
    Address of forwarding name server #1     128.162.19.7
    Address of forwarding name server #2     128.162.19.13
    Address of forwarding name server #3     128.162.1.1
    Slave server?                            NO
    Root name server cache file              root.cache
    Root name server cache ==>
    Primary zones ==>
    Secondary zones ==>
    
```

- b. The local name server also needs configuration information for the zones for which it is the primary server (the zones for which it has authority). On a caching-only server, the only zone for which the local named has authority is the 0.0.127.IN-ADDR.ARPA zone. This information is stored in the localhost.rev file; you can configure its name, but not its contents, by using the menu system:

```

Configure System
->Network Configuration
    ->TCP/IP Configuration
        ->TCP/IP Host/Address Lookup Configuration
            ->TCP/IP Local Domain Name Server Config
                ->TCP/IP Root Nameserver Cache Config
    
```

```

TCP/IP Root Nameserver Cache Configuration

Server name      Server address  Time To Live
-----
E-> earth.cray.com 128.162.3.55   1000000
    
```

```

TCP/IP Root Nameserver Cache Configuration

S-> Server name           earth.cray.com
    Server address        128.162.3.55
    Time to live          1000000
    
```

- c. The local name server must know the name of the server that is the authoritative name server for the domain. The root.cache file is used

to "prime the cache" with this information, and you can configure it by using the menu system:

```
Configure System
->Network Configuration
  ->TCP/IP Configuration
    ->TCP/IP Host/Address Lookup Configuration
      ->TCP/IP Local Domain Name Server Config
```

```
TCP/IP Domain Name Service Primary Zones

Name          File          Account  Serial #
-----
E-> 0.0.127.IN-ADDR.ARPA localhost.rev
```

```
TCP/IP Domain Name Service Primary Zones

Zone name          0.0.127.IN-ADDR.ARPA
File to contain zone information localhost.rev
S-> Account name of responsible party
Serial number for zone
```

4. Start the named daemon by executing the following command:

```
# /etc/sdaemon -s named
```

Procedure 36: Configuring a caching-only server without using the menu system

Note: If you have not completed the previous procedure, "Adding a CRAY J90 system to an existing TCP/IP network," you should complete steps 3, 4, and 5 of that procedure before configuring a caching-only server.

To configure your CRAY J90 system as a caching-only server without using the UNICOS Installation / Configuration Menu System, complete the following steps:

1. Enable the domain name system by creating the `/etc/hosts.usenamed` file. The existence of this file indicates that the UNICOS system will use DNS, rather than the `/etc/hosts` file to look up host names.

2. Configure the resolver by creating the `/etc/resolv.conf` file.

When you have a local name server (named process) running, you should have the local host address (127.0.0.1) as the first name server; otherwise your local named will be bypassed, resulting in decreased performance and increased network traffic. The following is an example

`/etc/resolv.conf` file:

```
## Domain name resolver configuration file
#
domain cray.com
#
nameserver 127.0.0.1
nameserver 128.162.19.7
nameserver 128.162.1.1
```

3. To configure a caching-only local name server, you must complete the following steps:

- a. Create the `/etc/name.boot` file. This file is read when named starts up. It tells the server what kind of server it is, over which zones it has authority, and where to get its initial data.

The following example shows a sample `named.boot` file. The `directory` line tells the server that all file names referenced are relative to the `/etc/named.d` directory. The `forwarders` line tells the server to forward requests that it cannot resolve to the server at 128.162.19.7. The `cache` line tells the server to load the `root.cache` file (in the `/etc/named.d` directory) as its initial cache entries. The `primary` line tells the server that it has primary authority for the `0.0.127.IN-ADDR.ARPA` domain. The `domain` line tells the server that its default domain is `cray.com`.

```
directory      /etc/named.d
forwarders     128.162.19.7
cache          .                root.cache
primary        0.0.127.IN-ADDR.ARPA  localhost.rev
domain         cray.com
```

- b. Create the `localhost.rev` file. The local name server also needs configuration information for the zones for which it is the primary server (the zones for which it has authority). On a caching-only server, the only zone for which the local named has authority is the `0.0.127.IN-ADDR.ARPA` zone. This information is stored in the

localhost.rev file. The following is an example of a localhost.rev file:

```
$ORIGIN 127.IN-ADDR.ARPA.
@           IN      SOA    localhost.cray.com. tas.cray.com. (
                        86400
                        3600
                        36000000
                        86400
                        )
1.0.0      IN      NS     localhost.cray.com.
1.0.0      IN      PTR    localhost.cray.com.
```

- c. Create the root.cache file. The local name server must know the name of the server that is the authoritative name server for the domain. The root.cache file is used to "prime the cache" with this information. The following is an example root.cache file:

```
$ORIGIN .
                        1000000 IN      NS     earth.cray.com.
earth.cray.com. 1000000 IN      A      128.162.3.55
```

4. Start the named daemon by executing the following command:

```
# /etc/sdaemon -s named
```

11.3 Common TCP/IP configuration files

After you have the Cray Research system on the network, you should do a few additional things to make sure it is a fully functional member of your network, including configuring `inetd`, adding additional routes, and updating other configuration files. If you are using the menu system, you should update these files by using the menu system options. Table 3 describes some of the most common TCP/IP configuration files. The *UNICOS Networking Facilities Administrator's Guide*, Cray Research publication SG-2304, describes all of these files.

Table 3. TCP/IP configuration files

File (relative to /etc)	Description	Change
config/daemons	Lists system and network daemons to start at system boot.	Probably
gated.conf or tcpstart.mid	Contains routes to be installed at system boot.	Yes
config/hostname.txt	Contains text host name for TCP/IP.	Probably not
hosts	Maps Internet addresses to host names.	Yes
hosts.equiv	Lists trusted hosts for rlogin, rsh, and so on.	Optional
hycf.xxx	Maps physical addresses to Internet addresses for nonbroadcast media (for example, HYPERchannel and HIPPI).	Yes, if you have HYPERchannel or HIPPI
inetd.conf	Lists network services to be handled by inetd.	Probably not
config/interfaces	Lists network interfaces and their characteristics.	Yes
networks	Maps network names to network Internet addresses.	Optional
protocols	Maps protocol names to protocol numbers.	No
\$HOME/.rhosts	Lists trusted users for rlogin, rsh, and so on.	Optional
services	Maps protocol and port numbers to service names.	Probably not
shells	Lists shells allowed for ftpd.	Probably not