Fault Tolerant Network Configuration

This chapter outlines the configuration information needed to operate Cisco IOS for S/390 in a Fault Tolerant mode. It includes these sections:

- Fault Tolerant Overview

  Introduces fault tolerant networking
- Multiplexing

  Describes multiplexing in Cisco IOS for S/390
- Router Failures

  Describes how Cisco IOS for S/390 fault tolerant handles router failures
- Multihoming

  Describes the use of multihoming
- Virtual IP Addressing

  Describes Virtual IP Addressing
- GateD

  Introduces the GateD Daemon
- Improving Fault Tolerant Reliability

  Describes ways to improved reliability for fault tolerant networks
- Running the Routing Daemon (GateD)

  Describes how to set up and run GateD
- Example of a Fault Tolerant Configuration

  Describes samples of fault tolerant configurations

Fault Tolerant Overview

Cisco IOS for S/390 Fault Tolerant provides high levels of availability and reliability in network connections. Used with redundant network interface hardware, Fault Tolerant allows the user to maintain persistent sessions during a hardware failure or a routing outage or change. Redundant network interface hardware is required to use this feature.

The two main components of Fault Tolerant Cisco IOS for S/390 are:

- Multiplexing—the intelligent use of multiple controllers to handle hardware failures.
- Gateway Dæmon (GateD)—supporting Open Shortest Path First (OSPF) and Routing Information Protocol (RIP) to handle router failures or routing changes.

Fault Tolerant Cisco IOS for S/390 also provides a method to determine network outages by sampling network activity. In this way, if a network connection becomes unavailable due to a cable problem or wiring defect, Cisco IOS for S/390 Fault Tolerant addresses this and reroutes and/or redirects network traffic appropriately.

Multiplexing is two or more hardware interfaces bound to a single IP address. Multihoming is two or more hardware interfaces bound to multiple IP addresses and executing within the same Cisco IOS for S/390 address space.

Note In this case, multiplex or multihome network controllers are link level controllers.

The Fault Tolerant features of Cisco IOS for S/390 are completely automatic. Once configured and running, no operator intervention is required.

Fault Tolerant Limitations

Limitations of Cisco IOS for S/390 Fault Tolerant are listed below.

- Cisco IOS for S/390 GateD/OSPF does not support multicast.
- NSC HYPERchannel interfaces only recognize hardware outages. A network outage may go unreported due to the Internet Protocol (IP) router built in to these network controllers.
- Accurate network outage determination is possible with link level controllers supporting CETI and 3172 protocols.

Managing Controller Failures

Through the use of multiple controllers, Cisco IOS for S/390 Fault Tolerant software can increase availability and continue communication in cases where a Local Network Interface (LNI) fails and another LNI is available.

Managing Router Failures

Through the use of multiple network controllers and the gateway daemon (GateD), Cisco IOS for S/390 Fault Tolerant software can increase availability and communication in cases where network IP routers fail. GateD manages route changes and updates routes to your LAN and WAN within your Cisco IOS for S/390 address space. GateD also manages route changes in cases where channel attached routers are used. MEDIA statement parameters, including ARPIPTIMEOUT and IDLENET, control how long Cisco IOS for S/390 waits to determine outages or issue ARP attempts at the router layer.

Multiplexing

Cisco IOS for S/390 actively samples network activity to detect network outages. This can happen when a channel error occurs or a network defect is discovered (in other words, bad cable or defective hub). The Address Resolution Protocol (ARP) is used to dynamically map Internet (IP) and MAC (hardware) addresses. In addition, the SNMP agent of Cisco IOS for S/390 will send an "interface down" trap to a network management station when either of the above conditions are met. Read Chapter 15, SNMP Agent Configuration for information on configuring the SNMP agent.

When an interface fails in a multiplex configuration, Cisco IOS for S/390 notifies other stations on the LAN of the address of an active network interface that can be used. In this event, any existing sessions will be rerouted to use the active interface without session interruption. Hosts on the LAN should then update their ARP tables to point to the active interface (Figure 16-1). MEDIA statement parameters, including ARPTIMEOUT and IDLENET, can be configured to increase or decrease the length of time Cisco IOS for S/390 waits to determine network outages.

Figure 16-1 Multiplexing Environment

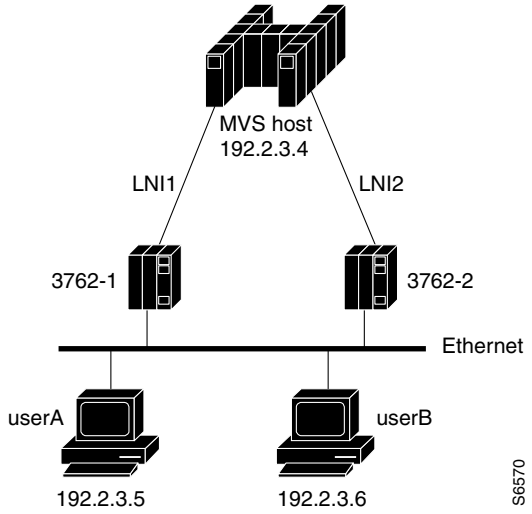

Configuring the TCPCFGxx Member for Multiplexing

For multiplexing, the TCPCFGxx member requires the following entries:

- 1 MEDIA statement
- 1 or more NETWORK statements
- 2 or more device statements
- 1 or more ROUTE statements

Note Some software packages for personal computers (PCs) never refresh their ARP tables. This can cause unexpected results if they are directly connected to a multiplexed device or host. In order to refresh these ARP tables you must reboot the PC.

Balancing I/O Traffic

In addition to improving availability, the fault tolerant features improve performance by balancing I/O traffic across multiplexed controllers.

For outbound sessions, such as large file transfers from the mainframe to network hosts, Cisco IOS for S/390 alternates file transfer output, using the multiplexed controllers in a round robin fashion. Using multiple controllers levels the data transfer load and improves throughput performance, especially for very large file transfers.

Router Failures

Through the use of multiple network controllers and the gateway dæmon (GateD), Cisco IOS for S/390 fault tolerant software can increase availability and communication in cases where network IP routers fail. GateD manages route changes and updates routes to your LAN and wide area network (WAN) within your Cisco IOS for S/390 address space. GateD also manages route changes in cases where channel attached routers are used. MEDIA statement parameters, including ARPIPTIMEOUT and IDLENET, control how long Cisco IOS for S/390 waits to determine outages or issue ARP attempts at the router layer.

Multihoming

GateD manages your routes in either a multiplexed environment or a multihomed environment. It can manage multiple routing protocols, including OSPF, and can instantaneously update routes and maintain sessions during router failures or channel attached router outages.

Note In order for a multihoming configuration to work in a fault tolerant environment, the host addresses must be on separate subnets. If they are not on separate subnets and you try to use Fault Tolerant, you will lose connections, since routes cannot be updated on the same local network. The network outage will be reported by Cisco IOS for S/390 in all cases.

Figure 16-2 illustrates how GateD appears in a mainframe multihoming configuration.

In this fault tolerant diagram, hostA and hostB (both are on the MVS host) are on the 191.2.0.0 class B network; hostA is on the 191.2.3.0 subnet, and hostB is on the 191.2.4.0 subnet:

Figure 16-2 GateD in a Mainframe Multihoming Configuration

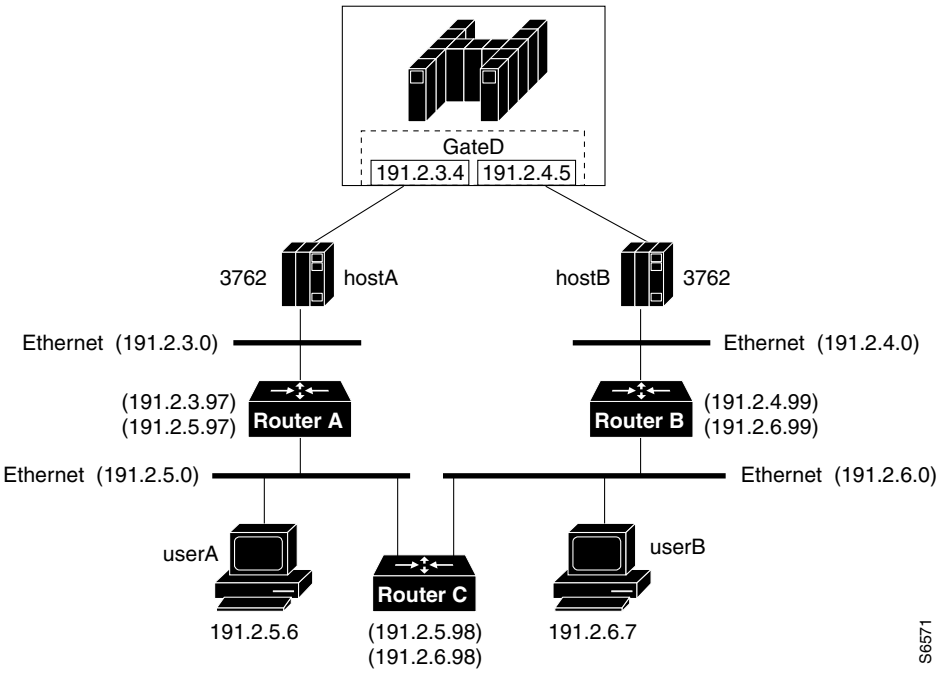

If the route to network 191.2.4.0 or ROUTERB fails, all traffic for userB will be routed to the mainframe host through network 191.2.3.0 via ROUTERC. GateD is responsible for broadcasting new routes to other network routers and hosts.

The GTDCFGxx member specifies the routing protocol to be used.

Note If RIP is used, propagation of new routes may take up to three minutes.

Example

The following configuration example shows some of the necessary parameters for the sample fault tolerant setup shown on the preceding page:

```
IP
            .
            .
   GATED(GTDCFG01)
   FORWARD
            .

MEDIA MSSOPT(NET) MSSDEF(1500) MTU(1500) NAME(ETHER1)
NETWORK IPADDRESS(191.2.3.4)
   SUBNET(255.255.255.0)
* LNI STATEMENTS FOR ETHER1
CETI
   .
   .
   CUTYPE(3762)
   MEDIANAME(ETHER#1)
   DEVADDR(86A)
   .
   .
* LNI STATEMENTS FOR ETHER2
MEDIA MSSOPT(NET) MSSDEF(1500) MTU(1500) NAME(ETHER2)
NETWORK IPADDRESS(191.2.4.5)
   SUBNET(255.255.255.0)
   NAME(ETHER2)
   .
   .
CETI
   .
   .
   CUTYPE(3762)
   MEDIANAME(ETHER#2)
   DEVADDR(8B0)
ROUTE DEST(0.0.0.0) ROUTE(191.2.3.97) MEDIANAME(ETHER1)
ROUTE DEST(0.0.0.0) ROUTE(191.2.4.99) MEDIANAME(ETHER2)
```

This is the GTDCFG00 member for accessing RIP:

```
traceoptions none;
/*=========================================================*
/*   GateD V3.0.3 default configuration file */
/*=========================================================*/
traceoptions parse;
rip yes      ( broadcast;
      sourcegateways 191.2.3.97 191.2.4.99;
      ) ;
traceoptions general kernel rip update;
```

Virtual IP Addressing

Virtual IP Addressing (VIPA) can be used with Cisco IOS for S/390 Fault Tolerant when Cisco IOS for S/390 is running in a multihomed environment. VIPA does not enhance the fault tolerant capabilities of Cisco IOS for S/390 running with a single interface or running multiplexed.

Benefits of Virtual IP Addressing

RunningCisco IOS for S/390 in a multihomed configuration with GateD had two deficiencies. The first is related to connectivity between hosts on a directly attached subnet. The second is related to routers which have a higher preference for attached subnets than host routes.

In the first case, most hosts do not use routing information to send packets when the destination host is on the same subnet. Instead, hosts often check to see if the destination IP address is on an attached subnet and if so, send the packet direct with no intervening routers. Because of this, when a Cisco IOS for S/390 interface goes down in a multihomed configuration, hosts on the same subnet cannot determine an alternate path to the IP address of the down interface and all connectivity to that IP address is lost.

In the second case, Cisco IOS for S/390 Fault Tolerant sends OSPF and RIP routing updates when an interface goes down. These routing updates tell the rest of the routers that Cisco IOS for S/390 can no longer reach the subnet with the down interface but can reach the host IP address of the down interface. When this routing information is processed by a router on the subnet with the down interface, the alternate route information for the down host is often ignored, because as long as the router's interface on the subnet is active, it believes that the best route to the down interface is to send the packet directly over the subnet.

VIPA solves these problems by creating a virtual subnet within the Cisco IOS for S/390 address space. A virtual subnet means that a NETWORK statement is defined in the Cisco IOS for S/390 configuration parameters (TCPCFGxx) but the subnet is not related to any real interface (LNI). The NETWORK statement specifies a host IP address on a subnet which does not exist anywhere else in the network.

When a remote host sends a packet to the virtual IP address, it is forced to use IP routing regardless of whether the host is on a direct attached subnet. Because the subnet of the destination is different than the remote hosts subnet, the remote host must use its routing tables to make a decision on the best path to the virtual IP address. For a host on a direct attached subnet, this may mean that the first packet in a connection is sent to a router, which in turn will either forward the packet or send the host an ICMP redirect if the Cisco IOS for S/390 interface on the direct attached subnet is still active. In either case, the host can establish a connection with Cisco IOS for S/390 regardless of whether the interface on the direct attached subnet is active.

Detecting Down Interfaces

When a router on a direct attached network detects that a Cisco IOS for S/390 interface is down (by a lack of routing messages), it recalculates the best path to the virtual IP address. This lets packets sent to this router be rerouted to the active Cisco IOS for S/390 interfaces.

Remote users connecting to Cisco IOS for S/390 should connect to the virtual IP address instead of the real IP address. The benefits of using VIPA are only realized when the Cisco IOS for S/390 IP address is the virtual IP address. The best way to implement this is to change the authoritative domain name server to reference the virtual IP address instead of a real IP address for the Cisco IOS for S/390 host name.

Cisco IOS for S/390 will always choose a virtual IP address when initiating a TCP connection. UDP and RAW datagrams will still be sent with the real IP address (a requirement for proper operation of GateD).

Using VIPA with GateD

The use of VIPA does not require GateD, but without GateD, IP routers must be configured with static paths to the virtual subnet, which might cause lost connectivity. Contact your routing vendor for advice on configuring any routers when GateD is not run.

Without GateD, VIPA does not guarantee session persistence for hosts on a direct attached subnet when the interface on that subnet fails. This is the result of few or no mechanisms in existence to notify the host that the Cisco IOS for S/390 interface is down, so the virtual IP address can only be reached via another route. This problem can be solved if the host can monitor routing messages on the subnet, allowing it to update its routing table as soon as the first router on the subnet detects the outage and advertises new routes.

GateD

Cisco IOS for S/390 uses the gateway dæmon (GateD) application to implement open systems interior and exterior routing protocols within the local network. With the use of GateD, Cisco IOS for S/390 functions as a router on the network and quickly detects changes in the routing environment and dynamically acts upon the information quickly enough to keep sessions from being interrupted or delayed.

This section describes what routing protocols are and provides information about the two routing protocols supported by GateD, OSPF and RIP.

Routing Protocols

The function of a routing protocol is to choose a path for sending data to a destination. Static routing uses a routing table to specify IP addresses for hosts and gateways. This type of routing is used for simple networks or networks whose configuration is rarely changed.

Dynamic routing protocols allow routing choices to be made according to current network conditions. GateD is best used on networks where node or gateway changes occur frequently, otherwise static routes are effective and use no CPU or bandwidth in route determination and propagation.

GateD is a dynamic routing application that collects information and makes routing choices based on that information. GateD uses an Interior Gateway Protocol (IGP) such as OSPF or RIP to accomplish this. An IGP is any protocol that interior gateways use to communicate with each other, exchanging routing and reachability information.

RIP

The Routing Information Protocol (RIP) is one of the most widely used IGPs. It was originally implemented in the routed dæmon program, a precursor to the GateD application. The routed dæmon was widely distributed with UNIX BSD systems and because of that, RIP has become a popular IGP. However, it does have some technical limitations.

RIP is a distance vector routing protocol. It sorts machines into two categories, active and passive, according to whether the machine advertises its routes (active) or simply listens and updates its routes based on those advertisements (passive). Generally, gateways are active and hosts are passive.

Active gateways advertise a routing table every 30 seconds. The routing table consists of pairs where a pair is an IP network address and the distance (metric) to that network. A hop count metric is used to measure the distance to a destination. The hop count would be the number of gateways that a datagram passes through along the path to its destination.

For RIP, a hop count of 16 means the destination is unreachable. After 3 minutes with no updates, a router is marked down and all routes through it are given metric 16. After another 2 minutes without updates, those routes are removed.

Note RIP does not pass subnet masks. If you have different subnet masks throughout your network, you should not use RIP.

RIP is documented in RFCs 1058 and 1723.

OSPF

Open Shortest Path First (OSPF) is a link-state routing protocol which uses the Shortest Path First (SPF) algorithm. During operation, OSPF tests the status of all neighbor gateways and periodically broadcasts this link status information to all other gateways. When link status information is received, each gateway updates its map of the network, marking links “up” or “down”. When a link status changes, the gateway automatically recomputes routes to determine the optimal paths to all destinations. The OSPF protocol is fully described in RFC 1583.

An overall picture of the networks is maintained in a topological database. It contains the link states (known as Link State Advertisements, or LSAs) from all the routers in the same area.

The actual link bandwidth consumed by OSPF is less than that consumed by RIP. OSPF consumes more memory than RIP, but the use of area partitioning can save memory. (Area partitioning breaks the network into a backbone and areas so active hosts and routers need to keep track of less of the network.) In general, OSPF consumes less CPU time than RIP, due to RIP’s frequent updates.

OSPF can operate within a hierarchy. The largest entity in the hierarchy is the Autonomous System (AS), which is a collection of networks under a common administration. An AS is divided into a number of areas. Each area is a group of contiguous networks and attached hosts. Sometimes the term “domain” is used to describe a portion of the network; this is often used interchangeably with AS.

Figure 16-3 is a simple, three-area OSPF network with a contiguous backbone:

Figure 16-3 Three-area OSPF Network

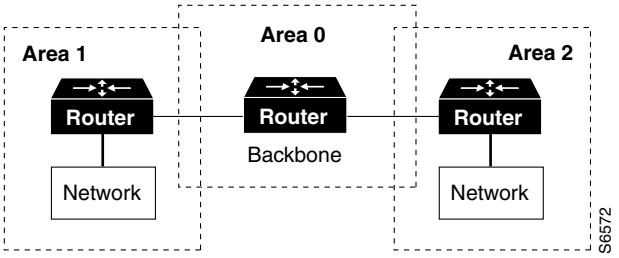

In order to exchange information between areas of a backbone, the areas must be contiguous. To be contiguous, all hosts and/or routers exchange information between autonomous systems using direct or virtual links (routers). If a direct link is unavailable, a virtual link can be used. A virtual link is an indirect link through backbone routers that share links between non-backbone areas.

The following diagram shows an OSPF network with a virtual link. To keep the backbone contiguous, a virtual link through Area 2 is defined such that if a link on the backbone were to crash, virtual links through other routers would keep traffic moving between Area 1, Area 3 and Area 4.

R1, R2, R4, and R5 are internal routers. These routers are directly connected to networks within their areas.

In Figure 16-4 R6 and R3, along with their links, comprise the backbone (Area 0).

Rx = routers
Nx = Networks (or subnetworks)
Hx = Hosts

The class B Network address for this autonomous system is 191.2.0.0

The Area 1 subnetwork address is 191.2.1.0
The Area 2 subnetwork address is 191.2.2.0 and 191.2.3.0
The Area 3 subnetwork address is 191.2.4.0
The Area 4 subnetwork address is 191.2.5.0

The R1 address is 191.2.1.1
The R2 addresses are 191.2.2.1 and 191.2.3.1
The R3 address is 191.2.6.1
The R4 address is 191.2.4.1
The R5 address is 191.2.5.1

H1 is FDDI at 191.2.1.10
H2 is token ring at 191.2.2.10
H3 is ethernet at 191.2.3.10

Figure 16-4 OSPF Network Configuration with Virtual Link

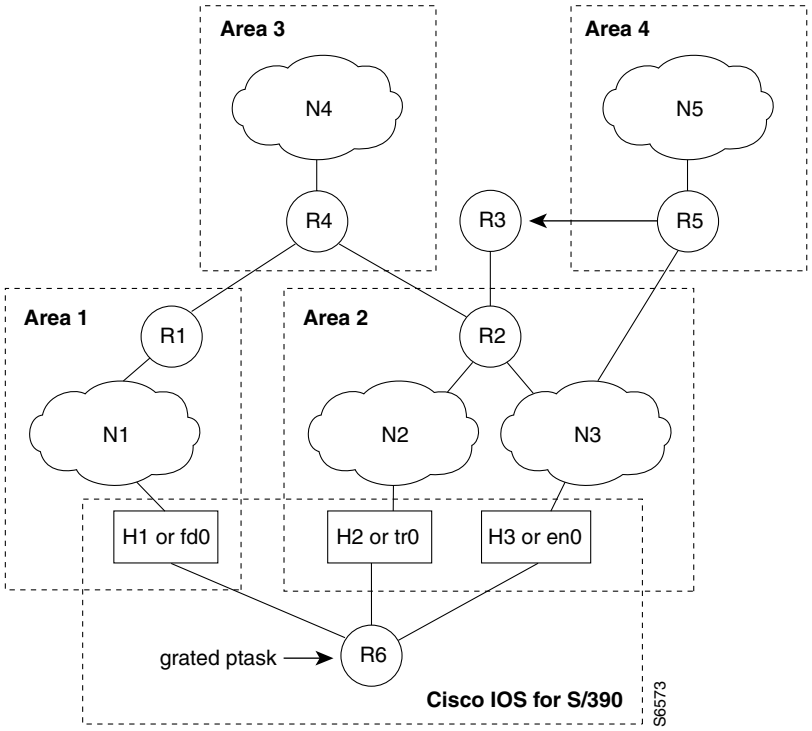

GateD Example

GateD is configured in member GTDCFG00 of the PARM data set. The sample shown here is member GTDCFG01, which was created and referenced in the TCPCFG*xx* configuration IP statement. For more information on configuring GateD, read Chapter 17, GateD Configuration—GTDCFGxx.

If Cisco IOS for S/390 is R6 in the diagram, here is a sample configuration:

```
traceoptions ospf general kernel mark ;
interfaces (
   options scaninterval 15;
);
rip off ;
ospf yes (
   area 0.0.0.1 (
      interface fd0 nonbroadcast (
         pollinterval 60;
         routers (
            191.2.1.1 eligible ;
         );
         retransmitinterval 5;  /* 5 sec. between LSA rexmits */
         hellointerval 6;       /* Issue hello every 6 sec. */
         routerdeadinterval 24; /* Dead if > 24 sec w/o hello */
      );
   );
   area 0.0.0.2 (
      interface tr0 nonbroadcast (
         pollinterval 60;
         routers (
            191.2.2.1 eligible ;
         );
         retransmitinterval 5;
         hellointerval 6;
         routerdeadinterval 24;
      );
      interface en0 nonbroadcast (
         pollinterval 60;
         routers (
            191.2.3.1 eligible ;
         );
         retransmitinterval 5;
         hellointerval 6;
         routerdeadinterval 24;
      );
   );
   backbone (
      virtuallink neighborid 191.2.6.1
         transitarea 0.0.0.2 ( /* virtual link through area 2 */
         retransmitinterval 5;
         hellointerval 6;
         routerdeadinterval 24;
      );
   );
);
```

Note fd0 reflects the interface for FDDI tr0 reflects the interface for token ring en0 reflects the interface for Ethernet

With the configuration shown on the previous page, if the route from R1 were to fail, GateD would update the routes to reflect the virtual link through R2 in Area 2 for all routes that were using R1. In this manner, connections that were made on the FDDI subnet (Area 1) could be routed across the Token Ring or Ethernet subnetwork (Area 2) via the R2 router to get to Areas 3 and 4.

Note area 0.0.0.1 specifies Area 1, not a host 0.0.0.1. It is not related to the actual network or subnet address. 0.0.0.1 is a valid OSPF area number.

Note The GateD configuration parameters retransmitinterval, hellointerval, and routerdeadinterval must all have the same value wherever OSPF is running in your network. If GateD fails to route properly, review the network router OSPF configurations for possible discrepancies.

NonBroadcast Multi-Access

Cisco IOS for S/390 does not presently support IP multicasting. Therefore, most interfaces supported by Cisco IOS for S/390 must be defined as nonbroadcast interfaces on a Non-Broadcast Multi-Access (NBMA) media. Since an OSPF broadcast media must support IP multicasting, any broadcast-capable media that does not support IP multicasting must be configured as a nonbroadcast interface.

A nonbroadcast interface supports any of the standard interface parameters, plus the two described below:

- pollinterval time—Before adjacency is established with a neighbor, OSPF packets are sent periodically at the specified pollinterval. This usually is a multiple of the hellointerval parameter value.
- routers—By definition, it is not possible to send broadcast packets to discover OSPF neighbors on a nonbroadcast interface. Therefore, all neighbors must be configured. The list includes one or more neighbors and an indication of their eligibility to become the designated router for the area.

Improving Fault Tolerant Reliability

There are two MEDIA statement parameters in the TCPCFGxx member of the PARM data set that can improve the reliability of your Fault Tolerant configuration.

TCP Parameters

These TCPCFGxx MEDIA statement parameters can be used to improve your Fault Tolerant implementation:

- ARPTIMEOUT specifies the time an ARP cache is considered valid. This is useful in a multiplex environment to change the length of time between ARP table refreshes.
- IDLENET changes the length of time the network will be probed or sampled for network activity. This timer can affect network outages that include bad wires or cables.

For more information on the MEDIA statement, read Chapter 4, MEDIA Statement.

Running the Routing Daemon (GateD)

GateD Routing Protocol

The GateD application implements multiple routing protocols allowing Cisco IOS for S/390 to function as a router in a multihomed environment. GateD also provides Cisco IOS for S/390 with improved knowledge of the state of routers and routes in the attached networks allowing for a faster response to routing outages, whether multihomed or not.

Note Some versions of GateD handle RIP, HELLO, BGP, EGP and OSPF (non-multicast mode) routing protocols; currently, Cisco IOS for S/390 officially supports RIP and OSPF, although HELLO, BGP, and EGP are in the product.

For complete information on configuring GateD, read Chapter 17, GateD Configuration—GTDCFGxx.

GTDCFGxx Member

The GateD configuration member (GTDCFGxx) is specified by the GATED parameter on the IP statement. You must modify it for your use. For more information, read GateD Configuration—GTDCFGxx. This is the default configuration for GateD:

```
/*============================================================*/
/* GateD V3.0.3 default configuration file                    */
/* Warning - Make sure no line numbers exist in               */
/* columns 73-80                                              */
/*============================================================*/
traceoptions general kernel;
rip yes;
```

This GateD configuration example demonstrates parameter settings:

```
OSPF and RIP example:
traceoptions ospf general mark protocol;
rip yes {
 broadcast;
};
ospf yes {
 backbone {
  interface fd0 nonbroadcast {
   routers {
    138.42.180.129 eligible;
   };
   pollinterval 10;
   routerdeadinterval 30;
  };
  interface tr0 nonbroadcast {
   routers {
    138.42.170.129 eligible;
   };
```

```
   pollinterval 10;
   routerdeadinterval 30;
  };
 };
};

static  {
 default gateway 138.42.236.233;
};
```

A complete description of the GateD configuration file, GTDCFGxx, is given in GateD Configuration—GTDCFGxx.

If only one interface is defined and GateD is configured to run RIP, GateD does not send data and becomes just a RIP listener; GateD does not broadcast.

Note GateD does not support SNMP or multicasting.

Configuring Virtual IP Addressing

Virtual IP Addressing (VIPA) support is configured in two simple steps:

Step 1 In TCPCFGxx, add a MEDIA statement that specifies VIRTUAL and a NETWORK statement that specifies an IP address which is on a non-existent subnet.

```
MEDIA NAME(VIRTUAL1) VIRTUAL
NETWORK    HOST(138.42.175.3)
    SUBNET(255.255.255.0)
```

Step 2 In GTDCFGxx, define the virtual interface as passive.

```
 interfaces {
  interface vr0 passive ;
 } ;
 rip yes {
    interface vr0 noripout ;
 } ;
```

Note that virtual interfaces are named with a prefix of vr, such that the first three virtual interfaces are named vr0, vr1 and vr2.

Update your name servers such that the IP address for the Cisco IOS for S/390 host name references the virtual IP address instead of the real IP address. Notify users of the change of IP address if they specify the IP address directly.

VIPA Configuration Examples

Below are excerpts of a TCPCFGxx member for VIPA:

```
*-------------------------------------------------------------------*
* Member: xxxx.xxxx.PARM(TCPCFG00)
* Description: TCP task group configuration
*-------------------------------------------------------------------*

* Optionally redefine LOOPBACK media parms.
```

```
MEDIA NAME(LOOPBACK)
      MTU(4352)
      MSSOPT(NET)
      MSSDEF(4352)

* DEFINE VIPA address

NETWORK HOST(138.42.183.9)
      SUBNET(255.255.255.0)
      NAME('VIRTUAL NETWORK')
      MEDIANAME(LOOPBACK)

* Define the real physical medium

MEDIA FDDI
      MTU(4352)
      MSSOPT(NET)
      MSSDEF(4352)
      NAME(FDDI)

* Define the host

NETWORK IPADRESS(138.42.180.9)
     SUBNET(255.255.255.0)

* Define the network interface

LCS DEV(2C00) NAME(LCS1) CUTYPE(3762)

LINK LCSNAME(LCS1)

* Define router

ROUTE DEST(0.0.0.0) ROUTE(138.42.180.129)


* Take the defaults for TCP, UDP, RAW
```

Below is the configuration member for GateD used in the above example:

```
/*=========================================================*/
 /* GateD V3.0.3 Virtual IP Support Configuration File
 /*                */
 /* Warning - make sure that no line numbers exist in
 /*  columns 73-80
 /*=========================================================*/
 interfaces {
     interface vr0 passive ;
 };
 routerid 138.42.183.9 ;
 rip yes {
     broadcast ;
     interface vr0 noripout ;
 };
 traceoptions mark kernel ;
```

Example of a Fault Tolerant Configuration

In the example in this section, the MVS system is running Cisco IOS for S/390 with GateD support and two Cisco 7000 channel attachments. The Cisco IOS for S/390 address space has three IP addresses associated with it: Interface A, Interface B, and VIPA. Each Channel Interface Processor

(CIP) has a network defined as well as one VIPA. Each host address (the VIPA and the two CIP addresses) is in a separate subnet on the 192.1.1.0 network. This class C network (192.1.1.0) is subdivided by various settings of the subnet masks.

The GateD configuration shows the routerid as 192.1.1.49 (VIPA) and the other networks that can be reached - 192.1.1.16, 192.1.1.32, 191.1.0.0 and 10.1.1.0.

In the following example, the Cisco routers are defined as eligible in the GateD configuration.

Figure 16-5 Eligible Cisco Routers in GateD Configuration

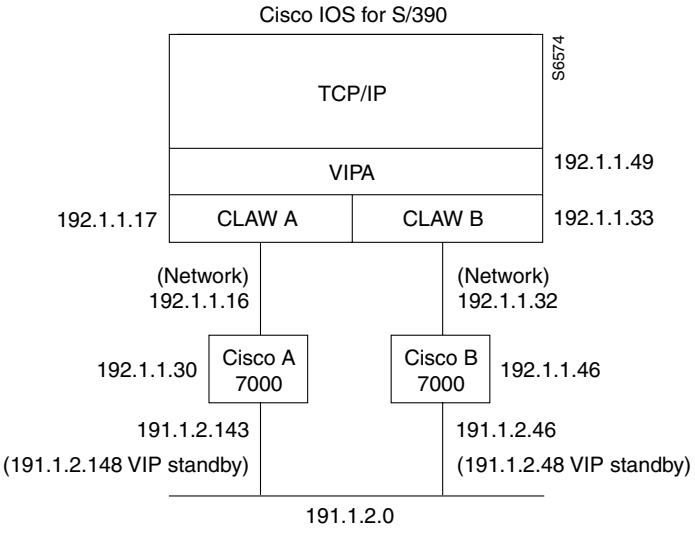

The Cisco router configurations show the real Token Ring interface IP addresses (191.1.1.143 (interface A) and 191.1.2.146 (interface B)), plus a standby address (191.1.2.148) on both A and B. Clients on the network would use this as a router address; it allows each of the Cisco routers to act as a "hot standby" for each other. In addition, the Cisco configuration is set up so that the routers filter IGRP (used on the network) to OSPF (used between the Cisco router and Cisco IOS for S/390).

Since Cisco IOS for S/390 does not support multicasting, these routes are explicitly defined in the routers to find the MVS VIPA network.

Host Addresses

These are the Cisco IOS for S/390 network addresses:

192.1.1.17—Interface A; specify this in the NETWORK and CLAW statements in the TCPCFGxx member.

192.1.1.33—Interface B; specify this in the NETWORK and CLAW statements in the TCPCFGxx member.

192.1.1.49—VIPA; specify this in the NETWORK statement in the TCPCFGxx member and in routerid in the GTDCFGxx member.

Subnet Addresses

These are the subnet addresses of the Cisco 7000 configuration:

192.1.1.16—Subnet address of interface A; specify this in the networks statement of the GTDCFGxx member.

192.1.1.32—Subnet address of interface B; specify this in the networks statement of the GTDCFGxx member.

192.1.1.48—Subnet address of VIPA; specify this in the ip route statement of the CIP configuration file.

Network and Gateway Addresses

These are the network and gateway addresses:

191.1.0.0—Token Ring network; specify this in the networks statement of the GTDCFGxx member.

192.1.1.30—Cisco CIPA interface; specify this in the router statement of the GTDCFGxx member and in the ROUTE statement for interface A in the TCPCFGxx member and in the Cisco CIPA interface statement configuration file.

192.1.1.46—Cisco CIPB interface; specify this in the router statement of the GTDCFGxx member and in the ROUTE statement for interface A in the TCPCFGxx member and in the Cisco CIPB interface statement configuration file.

Backbone Network Addresses

These are the backbone network addresses of the Cisco 7000 Configuration. Specify them in the CIP configuration file:

191.1.2.143—Cisco Token Ring interface on Cisco A router.

191.1.2.146—Cisco Token Ring interface on Cisco B router.

191.1.2.148—Cisco virtual standby address for token ring interfaces on Cisco A and Cisco B routers.

Cisco IOS for S/390 Configuration

This section provides an example of the TCPCFG*xx* configuration for use with the Cisco 7000.

```
/*   TCPCFGxx NETWORK, LNI and ROUTE statements        */

MEDIA VIRTUAL
      NAME(vipa)
      MTU(4096)
      MSSDEF(4096)
      MSSOPT(ALWAYS)
NETWORK IPADDRESS(192.1.1.49)
      SUBNET (255.255.255.240)

MEDIA CLAW
      NAME(TOKENA)
      MTU(4096)
```

```
      MSSDEF(4096)
      MSSOPT(ALWAYS)
NETWORK IPADDRESS(192.1.1.17)
      SUBNET (255.255.255.240)

CLAW DEVADDR(F90)
      HOSTNAME(MVSA)
      WSNAME(CLAWA)

MEDIA CLAW
      NAME(TOKENB)
      MTU(4096)
      MSSDEF(4096)
      MSSOPT(ALWAYS)
NETWORK IPADDRESS(192.1.1.33)
      SUBNET (255.255.255.240)

CLAW DEVADDR(F20)
      HOSTNAME(MVSB)
      WSNAME(CLAWB)

ROUTE     DEST(0.0.0.0)     ROUTE(192.1.1.30)     MEDIANAME(TOKENA)
ROUTE     DEST(0.0.0.0)     ROUTE(192.1.1.46)     MEDIANAME(TOKENB)
```

GateD Configuration

This section provides an example of GateD configuration for the Cisco 7000.

```
/* gated config for example of 2 cisco cips on class c network*/
traceoptions general mark protocol update ;
options noresolve ;
interfaces
{
   interface vr0 passive ;
} ;
routerid 192.1.1.49 ;
rip off ;
ospf yes
{
   backbone
   {
      networks
      {
         192.1.1.16      mask      255.255.255.240 ;
         192.1.1.32      mask      255.255.255.240 ;
         192.1.1.48      mask      255.255.255.240 ;
         191.1.0.0       mask      255.255.0.0 ;
         10.1.1.0        mask      255.255.255.0 ;
      } ;
      authtype 0 ;
      interface vr0 nonbroadcast cost 1
      {
         retransmitinterval 5 ;
         hellointerval 6 ;
         priority 1 ;
         pollinterval 6 ;
         routerdeadinterval 24 ;
      } ;
      interface c10 nonbroadcast cost 1
      {
         retransmitinterval 5 ;
         hellointerval 6 ;
```

```
         priority 1 ;
         pollinterval 6 ;
         routerdeadinterval 24 ;
         routers
         {
            192.1.1.30 eligible ;
         } ;
      } ;
      interface cl1 nonbroadcast cost 1
      {
         retransmitinterval 5 ;
         hellointerval 6 ;
         priority 1;
         pollinterval 6 ;
         routerdeadinterval 24 ;
         routers
         {
            192.1.1.46 eligible ;
         } ;
      } ;
   } ;
} ;
```

CIP Configuration Examples

This section provides examples of CIP configuration. Read the Cisco documentation for default values and specific CIP configuration information. The examples of CIP configurations shown here are for illustrative purposes only.

```
CIP-A

>CIP-A#wr t
>Building configuration...
>
>Current configuration:
>!
>version 11.2
>service udp-small-servers
>service tcp-small-servers
>!
>hostname CIP-A
>!
>enable password cisco
>!
>!
>interface TokenRing0/0
>ip address 10.1.1.168 255.255.255.0 secondary
>ip address 191.1.2.143 255.255.0.0
>no ip redirects
>ip route-cache
>ring-speed 16
>multiring ip
>standby 1 track Channel2/0
>standby 1 ip 191.1.2.148
>!
>interface TokenRing0/1
>no ip address
>shutdown
>!
>interface TokenRing0/2
>no ip address
>shutdown
```

```
>!
>interface TokenRing 0/3
>no ip address
>shutdown
>!
>interface Ethernet1/0
>no ip address
>shutdown
>!
>interface Ethernet1/1
>no ip address
>shutdown
>!
>interface Channel2/0ip address 192.1.1.30 255.255.255.240
>ip route-cache cbus
>ip ospf network non-broadcast
>ip ospf hello-interval 6
>no keepalive
>channel-protocol S4
>claw 0100 90 192.1.1.17 MVSA CLAWA TCPIP TCPIP broadcast
>!
>router ospf 1
>network 192.1.1.0 0.0.0.255 area 0
>network 191.1.0.0 0.0.255.255 area 0
>neighbor 192.1.1.17 priority 1
>default-information originate
>default-metric 50
>!
>router igrp 1
>redistribute ospf 1
>passive-interface Channel2/0
>network 10.0.0.0
>network 191.1.0.0
>default-metric 56 2000 255 255 1500
>!
>ip default-network 192.1.1.0
>ip route 192.1.1.48 255.255.255.240 192.1.1.17 200
>!
>!
>line con 0
>line aux 0
>transport input all
>line vty 0 4
>password cisco
>login
>!
>end

CIP-B
>Building configuration...
>!
```

```
>Current configuration:
>!
>version 11.2
>service udp-small-servers
>service tcp-small-servers
>!
>hostname CIP-B
>!
>enable password cisco
>!
>!
>interface TokenRing0/0
>ip address 10.1.1.169 255.255.255.0 secondary
>ip address 191.1.2.146 255.255.0.0
>no ip redirects
>ip route-cache
>ring-speed 16
>multiring ip
>standby 1 track Channel2/0
>standby 1 ip 191.1.2.148
>!
>interface TokenRing0/1
>no ip address
>shutdown
>!
>interface TokenRing0/2
>no ip address
>shutdown
>!
>interface TokenRing 0/3
>no ip address
>shutdown
>!
>interface Ethernet1/0
>no ip address
>shutdown
>!
>interface Ethernet1/1
>no ip address
>shutdown
>!
>interface Channel2/0
>ip address 192.1.1.46 255.255.255.240
>ip route-cache cbus
>ip ospf network non-broadcast
>ip ospf hello-interval 6
>no keepalive
>channel-protocol S4
>claw 0100 90 192.1.1.33 MVSB CLAWB TCPIP TCPIP broadcast
>!
>interface Channel2/1
>no ip address
>no keepalive
>shutdown
>!
>router ospf 1
>network 192.1.1.0 0.0.0.255 area 0
>network 191.1.0.0 0.0.255.255 area 0
>neighbor 192.1.1.33 priority 1
>default-information originate
>default-metric 50
>!
>router igrp 1
>redistribute ospf 1
>passive-interface Channel2/0
```

```
>network 191.1.0.0
>network 10.0.0.0
>default-metric 56 2000 255 255 1500
>!
>ip default-network 192.1.1.0
>ip route 192.1.1.48 255.255.255.240 192.1.1.33 200
>!
>!
>line con 0
>line aux 0
>transport input all
>line vty 0 4
>password cisco
>login
>!
>end
```

Note The hellointerval, pollinterval, and routerdeadinterval *must* have the same value in the GTDCFG*xx* member and in the CIP definitions for OSPF to operate properly.