

Using Virtual IP Addressing
with Bus-Tech TCP/IP Controllers

August 2001

The
Solution
is simple. **BUS-TECH**
INC.

Contents

Introduction.....	3
VIPA Requirements	3
TCP/IP Passthrough.....	4
LPAR Support.....	5
Multicasting	5
Using RIP-2	5
Using OSPF	8
Link Definition.....	9
TCP/IP Direct Connect	10
LPAR Support.....	11
Using RIP-2	11
Using OSPF	15
Link Definition.....	18

Introduction

When a client accesses an IBM mainframe using TCP/IP it uses an IP address associated with a network controller attached to the mainframe. Unfortunately when the controller fails the client loses connection to the mainframe even though the client, network, and mainframe may all still be operational.

To provide support for redundancy in such an environment IBM has added Virtual IP Address (VIPA) support to their OS/390 and VM TCP/IP stacks. VIPA allows a customer to define a virtual device to the IP stack and associate an IP address to that device. The virtual device is always active. OS/390 and VM-based routers are used to define routes to the virtual IP address using real devices attached to the mainframe. Failure of a controller causes the router to route traffic through an alternate controller; allowing clients using the virtual IP address continued access to the mainframe.

The purpose of this white paper is to discuss use of Bus-Tech TCP/IP products with IBM's Virtual IP Address support.

Bus-Tech offers two distinct TCP/IP access products. TCP/IP Passthrough, used with Mainframe Appliance for Networking and as the default access method in NetShuttle controllers, provides a LAN Control Station (LCS) connection to the Mainframe's TCP/IP stack. TCP/IP Direct Connect, included with Bus-Tech ESCON adapters and as option for NetShuttles, provides a Point-to-Point TCP/IP connection between the Mainframe's TCP/IP stack and a Windows 2000 or Windows NT TCP/IP stack running on an ESCON-attached server. Both TCP/IP Passthrough and TCP/IP Direct Connect may be used in conjunction with VIPA.

VIPA Requirements

Virtual IP Addressing (VIPA) uses standard TCP/IP routing in order to map the virtual IP address to a real IP address and controller. And although the use of static IP routes will allow access to the mainframe using VIPA, dynamic re-routing of TCP/IP traffic following a mainframe controller outage requires a router daemon be active on the mainframe.

In OS/390 IBM provides two separate router daemons. OROUTED is the older of these daemons and provides RIP-1 and RIP-2 routing protocol support. OMROUTE provides support for RIP-1, RIP-2, and OSPF routing protocols. IBM documentation indicates that over time OROUTED will be completely replaced with OMROUTE.

In VM the router daemon is known as ROUTED. ROUTED supports RIP-1 and RIP-2 protocols.

VIPA works with either RIP-2 or OSPF and can operate with any one of the above router daemons.

Since VIPA uses routing to re-direct traffic to the real devices in the network, virtual IP addresses need to be in their own subnet. Routers generally do not perform bridging. If the VIPA is in the same subnet as the real IP address the router will not build the necessary route in the route tables.

TCP/IP Passthrough

TCP/IP Passthrough is Bus-Tech's oldest, simplest, and most-successful TCP/IP product. TCP/IP Passthrough, as its name implies, takes TCP/IP traffic destined for the mainframe off the network and passes it, unprocessed, to the channel connection. The channel connection (ESCON or Bus and Tag) passes the traffic directly to the Mainframe's TCP/IP stack.




Figure 1. VIPA - TCP/IP Passthrough Example

Because it is a pass through product rather than a router, TCP/IP Passthrough needs only a single IP address to provide a connection to the mainframe. Notice in Figure 1 that each connection has only a single IP address (10.10.11.1 and 10.10.11.2). These addresses are defined to the mainframe's TCP/IP stack. The controllers themselves have no IP address and do not use a TCP/IP protocol stack in Passthrough connections.

Implementing VIPA with Passthrough-based controllers generally requires three separate IP subnets. The virtual IP address is in one subnet so that the mainframe router daemon can route traffic to one of the controller's addresses. The Passthrough controllers (10.10.11.1 and 10.10.11.2) and one or more neighboring routers are in the second subnet. The mainframe router daemon will send route information for the virtual IP address to neighboring routers. Third and subsequent subnets support the clients.

Clients are not normally placed in subnet 1 (above). Since Passthrough controllers are transparent to the network, clients in subnet 1 would require their own route tables to get to the mainframe. Those route tables would not automatically be updated when one of the controllers became unavailable. Clients in subnet 1 would lose access to the mainframe.

LPAR Support

Bus-Tech TCP/IP Passthrough controllers (Mainframe Appliance for Networking and NetShuttle) provide up to 32 separate LCS connections per channel interface. Typically customers configure multiple LCS connections in order to communicate with multiple Logical Partitions (LPARs). When multiple connections are defined, Passthrough requires IP Packet Filtering. IP Packet Filtering uses the mainframe's IP address to filter out traffic not destined for an LPARs TCP/IP stack. When VIPA is in use both the virtual IP address and the real IP address assigned to the connection must be included in the passthrough filtering list.

Multicasting

RIP-2 and OSPF protocols use multicasting (rather than broadcast) to send route updates between routers. When an OS/390 router daemon (OMPROUTE or OROUTED) is started it will attempt to join the necessary multicast groups to support routing on each of its interfaces. For RIP-2 the daemon must join multicast groups 224.0.0.1 and 224.0.0.4. For OSPF the daemon must join multicast groups 224.0.0.1, 224.0.0.5, and 224.0.0.6. Bus-Tech's TCP/IP Passthrough controllers recognize and respond to requests from the mainframe to join multicast groups.

Using RIP-2

Following are sample TCP/IP profile and OMPROUTE configuration statements needed to set up a VIPA environment using RIP-2 protocols.

In this example Controller 1 is installed using devices 300 and 301. Controller 2 is installed using devices 220 and 221. VDEV1 is defined as the virtual device. Controller 1 has an IP address of 10.10.11.1. Controller 2 has an IP address of 10.10.11.2. The virtual IP address is 10.10.10.1 in a different IP subnet than the real IP addresses assigned to the controllers.

The SOURCEVIPA option tells TCP/IP that mainframe-based IP applications are to use the virtual IP address when establishing connection to other systems in the network. If you do not specify this option outbound sessions may not be correctly recovered if a controller outage occurs.

DATAGRAMFWD is required in the TCP/IP Profile to allow packet forwarding. Make sure you do not have the NOFWD option specified.

Notice that 10.10.11.2 ETH2 is defined first in the HOME section following the virtual IP address. This makes Controller 2 the primary link associated with the virtual address. Controller 2 will be used for all traffic going from the mainframe to the network.

```

;
; TCPIP PROFILE
;
IPCONFIG
    SOURCEVIPA
    DATAGRAMFWD
;
; Hardware definitions:
;
    DEVICE VDEV1 VIRTUAL 0
    LINK VLNK1 VIRTUAL 0 VDEV1
;
    DEVICE NTS1 LCS 300 AUTORESTART
    LINK ETH1 ETHERNET 0 NTS1
;
    DEVICE NTS2 LCS 220 AUTORESTART
    LINK ETH2 ETHERNET 0 NTS2
;
; OMPROUTE Ports
PORT
    520 UDP OMVS ; OMPROUTE
    520 UDP OMPROUTE NOAUTOLOG ; OMPROUTE Server
;
; IP Addresses
HOME
    10.10.10.1 VLNK1
    10.10.11.2 ETH2
    10.10.11.1 ETH1
;
; Start the defined devices.
;
START NTS1
START NTS2

```

10.10.11.2 (Controller 2) has been assigned an Out_Metric of 2 on the RIP_Interface definition in the OMPROUTE configuration file. This causes OMPROUTE to advertise 10.10.11.2 to other routers with a cost of 2. As a result, inbound traffic will arrive through Controller 1 (10.10.11.1) because it has a lower cost (0).

Together this TCP/IP Profile and OMPROUTE configuration file result in traffic being split between the two controllers. Traffic will come in through Controller 1 and go out through Controller 2. But if either controller becomes unavailable, OMPROUTE will make the necessary route updates to restore connectivity through the remaining unit.

```
; OMPROUTE Config File -- for RIP-2
;
RIP_Interface
  IP_Address=10.10.11.1
  Name=ETH1
  Subnet_Mask=255.255.255.0
  Out_Metric=0
  MTU=1492
  RIPV2=Yes;
;
RIP_Interface
  IP_Address=10.10.11.2
  Name=ETH2
  Subnet_Mask=255.255.255.0
  Out_Metric=2
  MTU=1492
  RIPV2=Yes;
;
RIP_Interface
  IP_Address=10.10.10.1
  Name=VLNK1
  Subnet_Mask=255.255.255.254
  MTU=1492
  RIPV2=Yes
;
```

Using OSPF

The same TCP/IP profile used for RIP-2 can also be used for OSPF. Following is a sample OMPROUTE configuration file to support VIPA with OSPF.

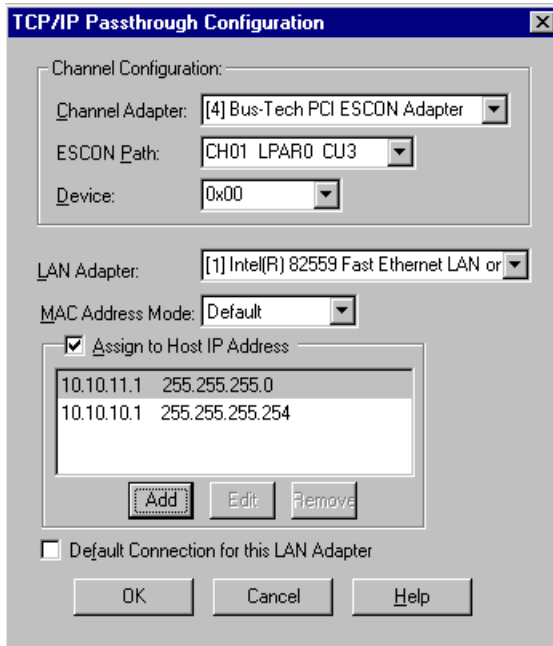
```
; * OSPF configuration
;
RouterID=10.10.10.1;
;
Area
    Area_Number=0.0.0.0
    Stub_Area=NO
    Authentication_Type=None;
;
OSPF_Interface
    IP_Address=10.10.11.2
    Name=ETH2
    Subnet_Mask=255.255.255.0
    Attaches_To_Area=0.0.0.0
    MTU=1492
    Cost0=2
    Parallel_OSPF=Primary;
;
OSPF_Interface
    IP_Address=10.10.11.1
    Name=ETH1
    Subnet_Mask=255.255.255.0
    Attaches_To_Area=0.0.0.0
    MTU=1492
    Cost0=2
    Parallel_OSPF=Backup;
;
OSPF_Interface
    IP_Address=10.10.10.1
    Name=VLNK1
    MTU=1492
    Subnet_Mask=255.255.255.0
    Destination_Addr=192.188.10.1
    Attaches_To_Area=0.0.0.0;
```

OMPROUTE advertises the last interface defined in the configuration file for a subnet to external routers. In this sample configuration 10.10.11.1 will be advertised to neighboring routers as the route for 10.10.10.1. Since both interfaces have equal cost, OMPROUTE will use the primary interface defined in the HOME list in the TCP/IP profile for outbound traffic. Using the HOME list shown on page 6, traffic will be sent out through 10.10.11.2. The Parallel_OSPF statements define which interface will be used to send OSPF route updates out over when more than one interface has equal cost. In this example 10.10.11.2 will be used.

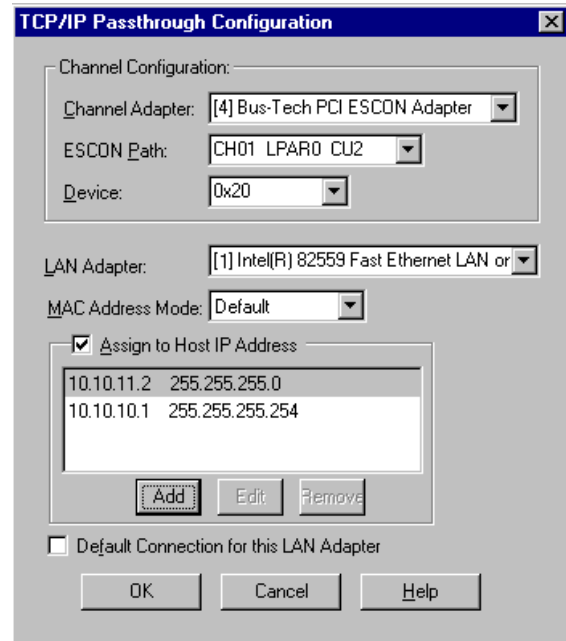
This OMPROUTE configuration file and TCP/IP profile (page 6) cause traffic to be split between controllers. Traffic will come in to the mainframe using Controller 1 (10.10.11.1). Traffic will go out from the Mainframe using Controller 2 (10.10.11.2). OSPF route updates will be multicast through Controller 2 because it is defined as the primary parallel OSPF interface. If either Controller becomes unavailable, OMPROUTE will alter route tables in order to restore connectivity through the remaining unit.

Link Definition

Following are the link definitions for Bus-Tech's TCP/IP Passthrough for the examples discussed in the previous 2 sections.



Controller 1



Controller 2

Note that IP filtering has been configured. If each controller only has one LCS connection defined there is no need to configure IP filtering. You may just leave the **Assign to Host IP Address** box unchecked and VIPA will work without entering IP addresses into the configuration. For VIPA to work properly when IP filtering is turned on you must include both the controller's real IP address and the virtual IP address in the filter list as shown above.

TCP/IP Direct Connect

TCP/IP Direct Connect provides a Network Driver Interface Specification (NDIS) compliant driver for Bus-Tech's PCI ESCON adapter. This driver allows a Point-to-Point connection between OS/390's TCP/IP stack and Microsoft's Windows 2000/NT TCP/IP stack (see Figure 2).

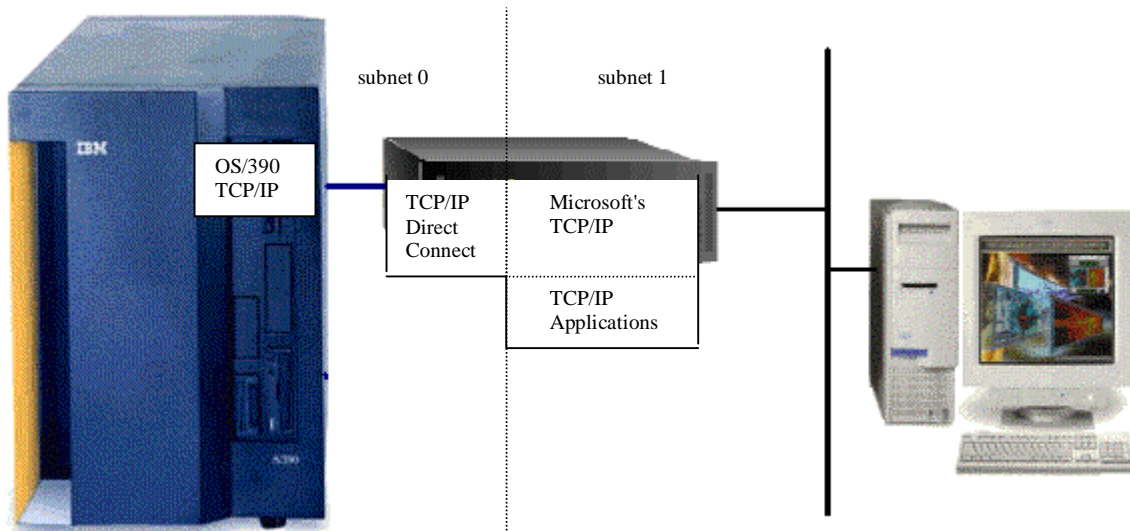


Figure 2. TCP/IP Direct Connect Overview

TCP/IP Direct Connect allows applications, such as IBM DB/2 Connect Enterprise Edition, to run directly on an ESCON-attached Windows 2000/NT system and gain access to the Mainframe over the channel. When TCP/IP Direct Connect is used to create an attachment between a network and the mainframe it takes on the appearance of an ESCON-attached router connecting the mainframe subnet to the LAN subnet.

If you intend to use VIPA with multiple ESCON-attached Windows systems running TCP/IP Direct Connect you will need to use Microsoft's Routing and Remote Access Services (RRAS) software on these servers. RRAS is a standard component shipped with Windows 2000 Server. RRAS is available as a free download from Microsoft for Windows NT Server.

RRAS provides a router supporting RIP-1, RIP-2, and OSPF routing protocols. VIPA requires either RIP-2 or OSPF in order to operate properly. Without RRAS installed the network's other routers will never see the route advertisements from the Mainframe.

Once routing services have been installed both your LAN adapter and TCP/IP Direct Connect interfaces must be added to the list of interfaces supported by the routing protocol (RIP or OSPF).

LPAR Support

TCP/IP Direct Connect is installed in Windows 2000/NT as a standard network adapter. A single install of TCP/IP Direct Connect supports only one connection and therefore one mainframe LPAR. However, TCP/IP Direct Connect can be installed multiple times on the same server with each install providing a single connection to the mainframe. Any TCP/IP Direct Connect adapter properly configured to the routing service on the Windows server will support VIPA traffic.

Using RIP-2

Figure 3 shows the Windows NT Routing Service properly configured for RIP-2. Notice that both the LAN Adapter and TCP/IP Direct Connect have been defined as RIP interfaces.

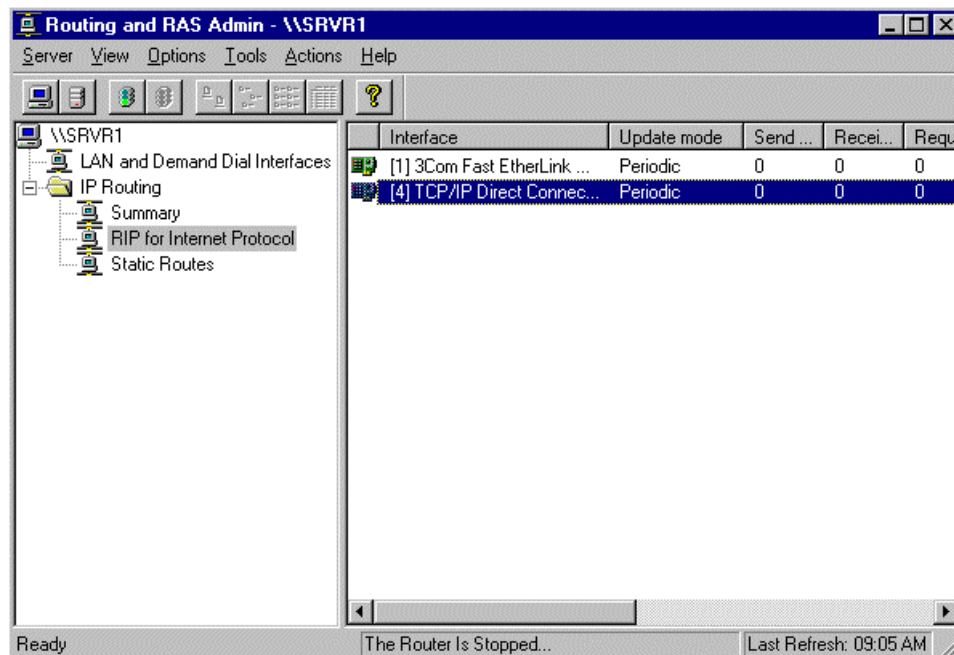


Figure 3. Configuring RIP in RAS

Make sure the TCP/IP Direct Connect Interface is configured for Periodic route updates and RIP-2 Multicast (see Figure 4).

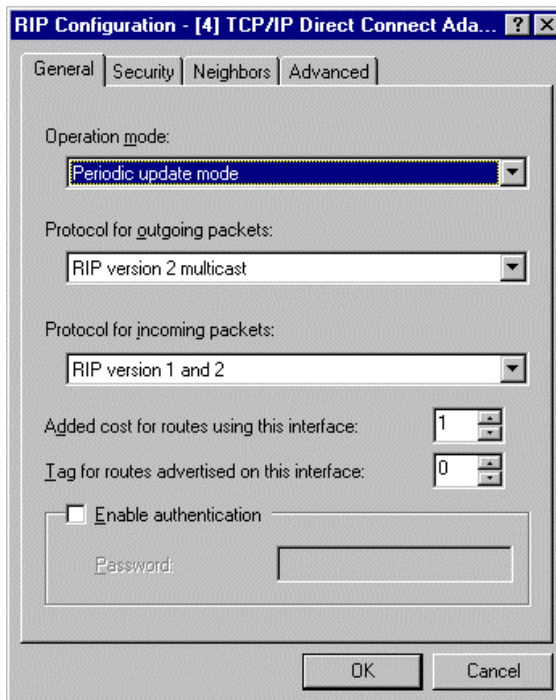


Figure 4. RRAS Interface Configuration

Once RRAS is configured on the NT servers, configure the Mainframe router to multicast route information to both NT servers.




Figure 5. Sample TCP/IP Direct Connect Configuration

Figure 5 shows a sample configuration for TCP/IP Direct Connect. Each connection to the mainframe uses two devices which must be consecutive. For this example device pairs 300, 301 and 220,221 will be used. Following are the control statements required in the Mainframe's TCP/IP Profile and OMROUTE Configuration file necessary for VIPA to work in this configuration using OMROUTE as a RIP-2 router.

```

;
; TCPIP PROFILE
;
IPCONFIG
    SOURCEVIP
    DATAGRAMFWD
;
; Hardware definitions:
;
    DEVICE VDEV1 VIRTUAL 0
    LINK VLNK1 VIRTUAL 0 VDEV1
;
    DEVICE DIR1 CTC 300 IOBUFFERSIZE 65535
    LINK ETH1 CTC 0 DIR1
;
    DEVICE DIR2 CTC 220 IOBUFFERSIZE 65535
    LINK ETH2 CTC 0 DIR2
;
; HOME Internet (IP) addresses of each link in the host.
;
HOME
    10.10.10.1 VLNK1
    10.10.11.1 ETH1
    10.10.11.3 ETH2
;
GATEWAY
    10.10.11.2 = ETH1 1492 HOST
    10.10.11.4 = ETH2 1492 HOST
;
; Start all the defined devices.
;
START DIR1
START DIR2

```

Note that in the GATEWAY section two static HOST routes have been defined for the IP addresses used for TCP/IP Direct Connect in the NT Servers. These static routes are necessary to complete the point-to-point connections. 10.10.11.2 and 10.10.11.4 are the IP addresses assigned to the TCP/IP Direct Connect Adapter in the Windows IP stacks.

```

;
; OMPROUTE Configuration
;
RIP_Interface
    IP_Address=10.10.11.1
    Name=PTHRU2L
    Subnet_Mask=255.255.255.254
    MTU=65525
    In_Metric=1
    Out_Metric=0
    Destination_Addr=10.10.11.2
    RIPV2=Yes;
;
RIP_Interface
    IP_Address=10.10.11.3
    Name=PTHRU1L
    Subnet_Mask=255.255.255.254
    MTU=65525
    In_Metric=1
    Out_Metric=0
    Destination_Addr=10.10.11.4
    RIPV2=Yes;
;
RIP_Interface
    IP_Address=10.10.10.1
    Name=VIPA1L
    Subnet_Mask=255.255.255.254
    MTU=65525
    RIPV2=Yes;

```

Notice that the `Destination_Addr` parameter is coded on the interface statements in the OMPROUTE configuration file. The IP addresses are the addresses assigned to the TCP/IP Direct Connect Adapters in Windows. `Destination_Addr` tells OMPROUTE that there is a Point-to-Point connection to the external router.

Using this OMPROUTE configuration file RIP-2 route updates will be multicast over both TCP/IP Direct Connect links. If one of the NT Servers fails, the remaining server will process the route updates necessary to re-route traffic away from the failure.

Using OSPF

In order to use OSPF and TCP/IP Direct Connect to provide VIPA support, Point-to-Point connections must be configured between OS/390 and the Windows system.

The following TCP/IP Profile and OMPROUTE Configuration file provide the necessary connections under OS/390.

```
;
; TCPIP PROFILE
; =====

IPCONFIG
    SOURCEVIPA
    DATAGRAMFWD

AUTOLOG 5
    OMPROUTE ;OSPF Router
ENDAUTOLOG

PORT
    520 UDP OMVS ; OMPROUTE
    520 UDP OMPROUTE NOAUTOLOG ; OMPROUTE Server

; VIPA device definition
DEVICE VIPA1 VIRTUAL 0
LINK VIPA1L VIRTUAL 0 VIPA1

; Server 1 TCP/IP Direct Connect
DEVICE DCON1 CTC 500 AUTORESTART IOBUFFERSIZE 65535
LINK DCON1L CTC 0 DCON1

; Server 2 TCP/IP Direct Connect
DEVICE DCON2 CTC 400 AUTORESTART IOBUFFERSIZE 65535
LINK DCON2L CTC 0 DCON2

HOME
    10.10.10.1 VIPA1L
    10.10.11.1 DCON1L
    10.10.11.3 DCON2L

; Empty Gateway Section
GATEWAY

; Start the Devices
START DCON1
START DCON2
```

```

; OSPF configuration

RouterID=10.10.10.1;

Area
  Area_Number=0.0.0.0
  Stub_Area=No
  Authentication_Type=None;

OSPF_Interface
  IP_Address=10.10.11.1
  Name=DCON1L
  Subnet_Mask=255.255.255.0
  Attaches_To_Area=0.0.0.0
  MTU=65525
  Cost0=2
  Destination_Addr=10.10.11.2
  Router_Priority=1;

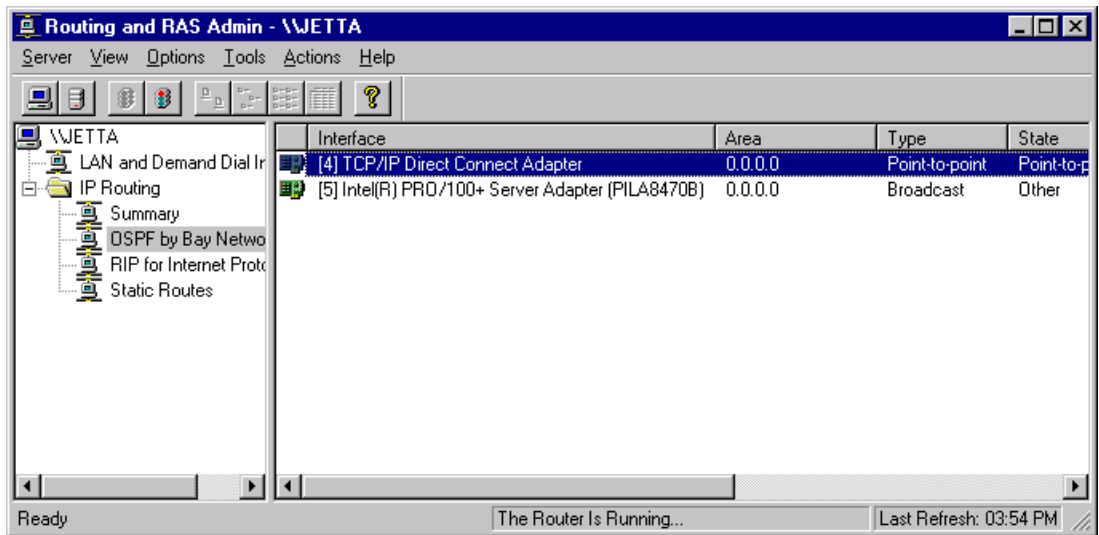
OSPF_Interface
  IP_Address=10.10.11.3
  Name=DCON2L
  Subnet_Mask=255.255.255.0
  Attaches_To_Area=0.0.0.0
  MTU=65525
  Cost0=2
  Destination_Addr=10.10.11.4
  Router_Priority=1;

OSPF_Interface
  IP_Address=10.10.10.1
  Name=VIPA1L
  MTU=65525
  Subnet_Mask=255.255.255.0
  Attaches_To_Area=0.0.0.0;


```

Notice that the `Destination_Addr` parameter is coded. Coding this parameter tells OMPROUTE that point-to-point connections exist to the external routers. The IP addresses are the addresses assigned to the TCP/IP Direct Connect Adapter in Windows.

On the Windows systems TCP/IP Direct Connect is configured as an OSPF Interface.



Be sure to enable the TCP/IP Direct Connect interface for OSPF and configure it as Point-to-Point connection as shown below.

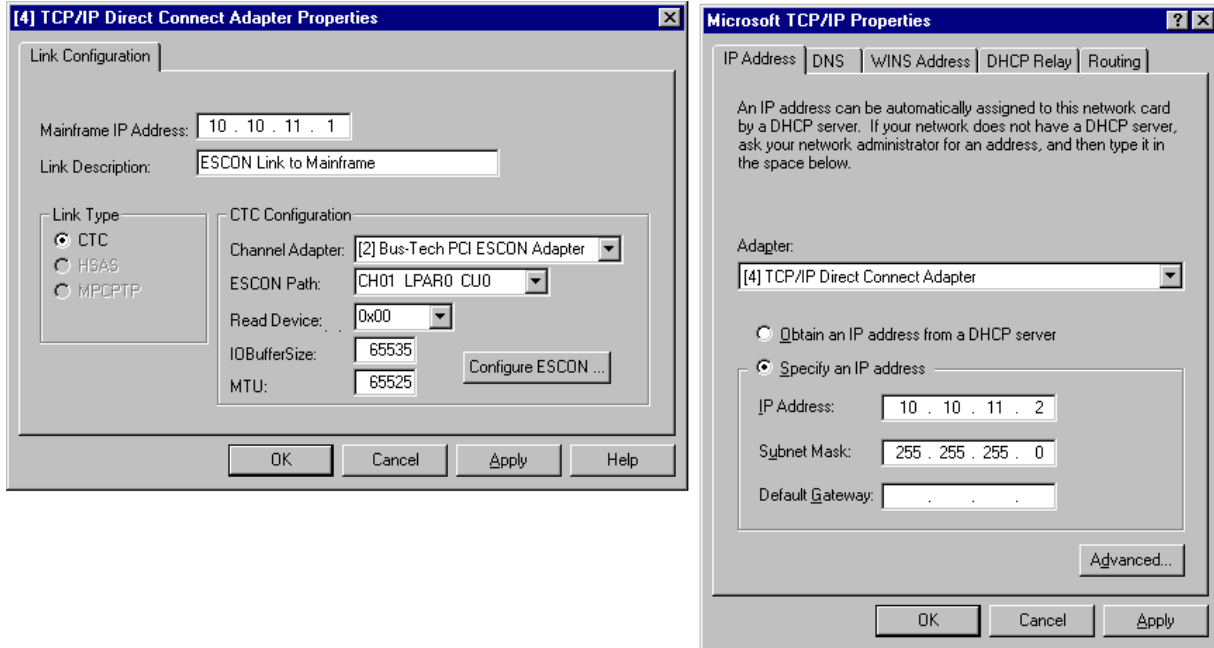


When OMPROUTE and the ESCON-attached routers are all started they find each other in the network and exchange route information.

Link Definition

TCP/IP Direct Connect link definitions supporting both the RIP-2 and OSPF examples above are as follows:

Windows NT Server 1



Windows NT Server 2

