3745 TOKEN-RING NETWORKING SESSION AVAILABILITY

February 22, 1991

Dore Rosenblum

Table of Contents

List of Illustrations

1.0 Introduction

An IBM Token-Ring Network has become a foundation for communications in many corporations. A Token-Ring network provides a high speed backbone through which users may access applications on other workstations, minicomputers and VTAM hosts on the network. Devices, functioning as SNA nodes on the Token-Ring Network, access VTAM applications through local (channel attach) or remote (link attach) gateways on the Token-Ring. The gateway may be a personal computer (e.g. PS/2), establishment controller (e.g. 3172, 3174), or enterprise controller (e.g. 3745). This document concentrates on an enterprise gateway, the 3745 Communication Controller, and more specifically, the 3745 Token-Ring network attachment subsystem (TRSS).

The purpose of this paper is to describe the availability characteristics of various 3745 configurations on a Token-Ring Network. The paper focuses on network availability and addresses how failure impact can be minimized for users of SNA Nodes on the Token-Ring Network. The network components studied can be grouped into two categories: the 3745 components and the Token-Ring Network components. The 3745 components are the hardware (e.g. 3745 Token-Ring Interface Coupler) and software (e.g. NCP) which support the Token-Ring connections. The Token-Ring Network components are the Bridges and Rings which make up the Token-Ring Network. Other components of availability, such as the host application and the end user workstation, are not discussed in this paper.

2.0 Terminology

This section describes the products and terms used throughout this paper. A basic understanding of Token-Ring Networking products and terms is assumed. Additional information may be found in the product announcements and other publications on Token-Ring Networking. The figure below shows some of the terms and products used.

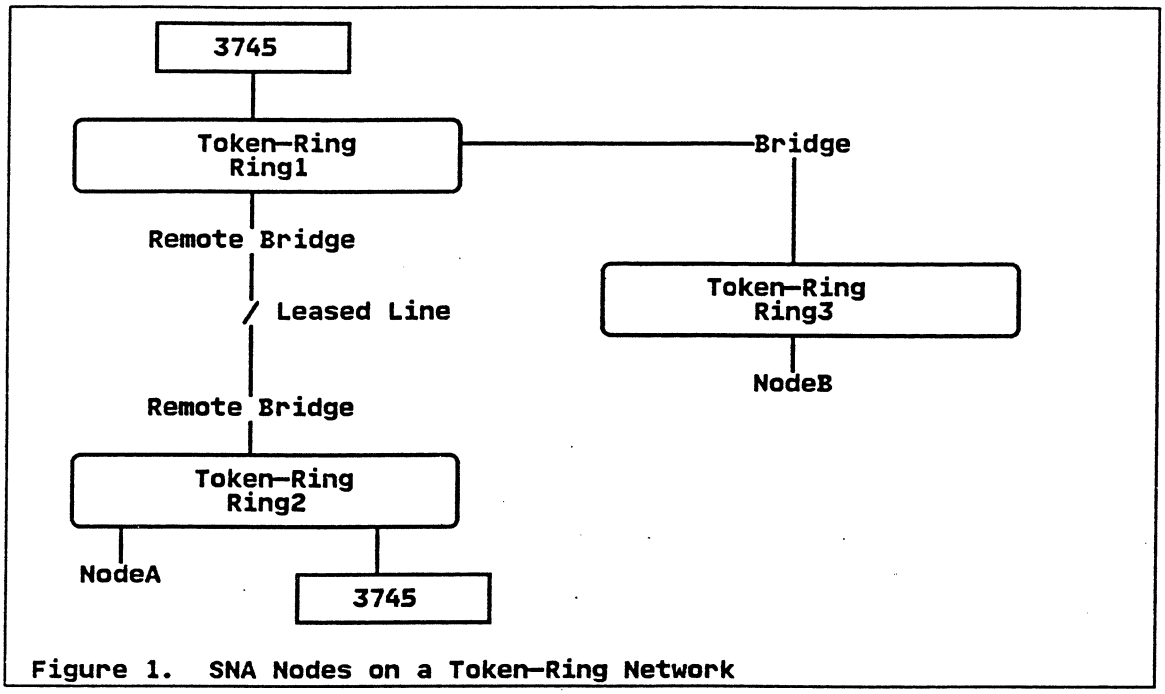

Figure 1. SNA Nodes on a Token-Ring Network

The figure above shows a **Token-Ring Network** consisting of three rings (Ring1, Ring2 and Ring3) connected together via bridges, which may be local or remote. A local bridge is a dedicated workstation or 8209 with two Token-Ring adapters - one adapter on each ring. A remote bridge consists of two dedicated workstations; each workstation has a Token-Ring adapter and a communication coprocessor which connects to a telecommunication line at speeds from 9.6 kbps to 1.344 Mbps. Rings connected using bridges form a single logical Token-Ring Network. Devices (e.g. PS/2s, 3174-x3Rs) on a Token-Ring Network run software to access applications on other devices (e.g. OS/2 LAN Servers, VTAM Hosts). A device on one ring may access applications on the other rings through bridges. Thus, NodeB on Ring3 may access the 3745 on Ring2 through the bridged network. In this paper, a **Bridge** refers to either a local or remote bridge connecting two rings.

Devices on the Token-Ring may run SNA software to access VTAM host applications through the 3745. In this paper, the Token-Ring devices running SNA software are called **SNA Nodes**. SNA Nodes communicate with the 3745 by sending frames to the 3745 Token-Ring address. Four types of SNA nodes connect to the Token-Ring Network today: Type 2.0, Type 2.1, Type 4 and Type 5. Type 2.0 nodes function like SNA 3174 controllers on the Token-Ring Network. Some examples of Type 2 nodes are a 3174-13R controller, OS/2 EE running 3270 emulation, or PC/3270 running 3270 emulation. Type 2.1 Nodes support additional functions such as peer communication across the SNA network. Some examples of Type 2.1 Nodes are a DOS PC running APPC/PC, an AS/400 running APPN, or a 3745 functioning as a Casual Connect Type 2.1 Node. In this paper, **Peripheral Nodes** are Type 2.0 and 2.1 SNA Nodes on the Token-Ring Network.

Type 4 Nodes are communication controllers (i.e. 3745, 3725, 3720) on the Token-Ring Network. Type 5 Nodes are VTAM Hosts in the SNA network. Today, most Type 5 Nodes connect to the Token-Ring through a communication controller. However, Type 5 Nodes may directly connect to the Token-Ring when supporting an integrated Token-Ring adapter (i.e. 9370) or when running VTAM Version 3.4 with a 3172 Interconnect Controller. A 3745 may establish Intermediate Network Node (INN) sessions with Type 4 and Type 5 Nodes across the Token-Ring Network. In this paper, **Subarea Nodes** are Type 4 or 5 SNA Nodes on the Token-Ring Network.

2.1 Abbreviations in this Document

Throughout this document, the following abbreviations are used:

- APPC: Advanced Program to Program Communication
- APPN: Advanced Peer to Peer Networking
- BNN: Boundary Network Node
- INN: Intermediate Network Node
- MLTG: Multiple Link Transmission Group
- NCP: ACF/Network Control Program
- NTRI: NCP Token-Ring Interface
- PTF: Program Temporary Fix
- SNA: Systems Network Architecture
- SLTG: Single Link Transmission Group
- TIC: 3745 Token-Ring Interface Coupler
- TRA: 3745 Token-Ring Adapter
- TRM: 3745 Token-Ring Multiplexor
- TRSS: 3745 Token-Ring Network Attachment Subsystem
- VTAM: ACF/Virtual Telecommunications Access Method
- XI: X.25 SNA Interconnection Program Product

3.0 3745 Token-Ring Gateway Description and Definitions

The 3745 Token-Ring Network Attachment Subsystem (TRSS) connects 3745 communication controllers to Token-Ring networks and supports VTAM host communication for SNA devices on the Token-Ring network. The 3745 TRSS consists of one or more Token-Ring Adapters (TRA's); the maximum number of TRA's varies according to 3745 model. Two types of TRA's exist: TRA Type 1 and TRA Type 2. TRA Type 1 is the original TRA and only connects to 4 Megabit Token-Ring networks. TRA Type 2 supports both 4 and 16 Megabit Token-Ring networks. Each TRA is made up of a Token-Ring Multiplexor (TRM) and two Token-Ring Interface Couplers (TIC's).

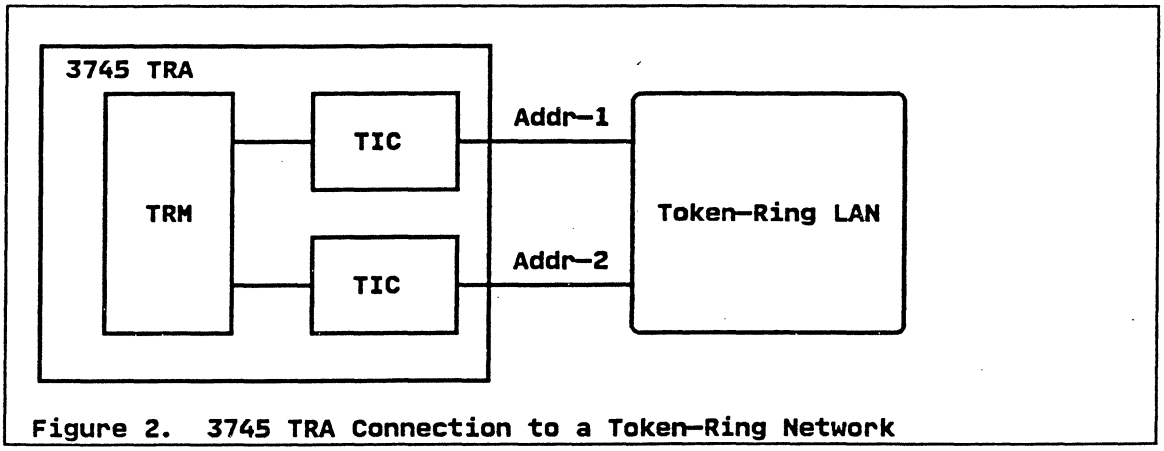

Figure 2. 3745 TRA Connection to a Token-Ring Network

The TRA Type 1 TICs are the same TICs as in the 3720 and 3725. Each TIC1 has 2392 bytes of storage for data received or sent to the Token-Ring. The TRA Type 2 TICs are only available in the 3745. The TIC2 has 62,880 bytes of storage for Token-Ring data. Thus, the TIC2 provides better performance across the Token-Ring Network by supporting larger frame sizes as well as higher speeds (16 Mbps).

3.1 Definition of 3745 Token-Ring Interface Coupler

Each 3745 TIC attaches to the Token-Ring network via an address specified in the NCP generation. As of NCP V5R2.1 with a PTF, the Token-Ring address may be the same or different from other TICs defined to an individual NCP. However, the 3745 must conform to the addressing conventions of the Token-Ring network (i.e. two TICs can not be active at the same address on a single Token-Ring).

The NCP Token-Ring Interface (NTRI) component of NCP supports the TRSS. In the NCP generation, the 3745 TIC is generated as a full duplex leased line. A sample definition is shown below of the TIC with Addr-1 in the above diagram.

```
GROUP ECLTYPE=(PHYSICAL,ANY) *Either Subarea or Peripheral Connections
LINE  ADDRESS=(ticpos,FULL),  *TIC Position in 3745
      ADAPTER=TIC2,           *TRA Type 2 (4/16 MBps)
      PORTADD=n,              *User assigned TIC number
      LOCADD=Addr-1           *Address on Token-Ring
PU
LU    ISTATUS=Inactive
```

3.2 Definitions of SNA Peripheral Nodes

As discussed in the overview, the 3745 communicates with Type 2.0, 2.1, 4 and 5 Nodes on the Token-Ring. These node types may be grouped into two categories: Peripheral Nodes (Type 2 Nodes such as 3174-13R or PC/3270, and Type 2.1 Nodes such as APPC/PC) and Subarea Nodes (Type 4 Nodes such as 37x5 or 3720, and Type 5 nodes such such as 9370).

Defining Peripheral Nodes on the Token-Ring Network

Individual definitions for each peripheral node are not defined in the NCP; a pool of definitions is generated for use by any Peripheral Node. This pool consists of logical point-to-point switched lines with PUs. Logical switched lines are associated with the TICs, and these definitions are dynamically used by the Peripheral Nodes which require host communications. An example of the logical definition for a Peripheral Node is shown below.

```
GROUP ECLTYPE=(LOGICAL,PERIPHERAL)
      PHYPORT=n               *Matches PortAdd of specific TIC
      CALL=INOUT              *VTAM or device can initiate call
      AUTOGEN=n               *Generate n Line/PU logical definitions
                              *for this group.
```

Peripheral Nodes are individually defined as VTAM switched major nodes (the same definitions used for devices that dial into the 3745). Each physical unit (PU) in the network requires a switched major node definition. Each Token-Ring device, functioning as a PU, has a switched major node definition which must be active before the device may access the host. An example of a switched major node definition for a PC with 3270 emulation is shown below:

```
VBUILD TYPE=SWNET,
       MAXNO=n                *Maximum number of dial numbers
       MAXGRP=n               *Number of TICs
PU     ADDR=04
       IDBLK=017              *PC 3270 emulation
       IDNUM=nnnnn            *Customized in emulator
       CPNAME=c               *Matches Type 2.1 Node Name
       MAXOUT=n               *Transmit Window Size
       ANS=CONT               *Don't drop sessions if owning host fails
       ....                   *Other Parameters
PATH   DIALNO=nnssaaaaaaaaaaaa,
                              *n=tic number,s=dest. SAP,a=Token-Ring addre ss
       GRPNM=g                *group name of logical NCP group
LU     LOCADDR=2              *Terminal
```

The IDNUM and IDBLK or CPNAME (optional for Type 2.1 Nodes) parameters uniquely identify SNA Nodes in the Token-Ring Network. The administrator customizes the peripheral node with parameters matching the corresponding VTAM definition. These parameters are passed to VTAM when the SNA Node requests communication with VTAM.

The ANS = CONT parameter is important in multiple-host environments. If ANS = CONT is not coded, a session with another host in the network is disrupted if the owning host (the host with the switched major node definition for the device) fails. With ANS = CONT coded, a session with another host in the network will not be disrupted if the owning host fails.

A PATH statement is only required when VTAM must initiate a connection to the Token-Ring SNA Node. The Path statement DIALNO must be a locally administered address on the Token-Ring; the burned-in addresses on the Token-Ring cards are not supported (VTAM does not support hexdecimal phone numbers).

Benefits of Defining Peripheral Nodes as Switched Major Nodes

Defining the Token-Ring Peripheral Nodes as switched major nodes provides several advantages to Token-Ring users. One advantage is that PUs may be **dynamically** added to VTAM; **no NCP changes are required to add a PU**. Another advantage is that individual SNA Nodes are not required to access the host through a specific TIC. A Token-Ring SNA node may be customized to access the host at any gateway TIC address having logical lines owned by a VTAM with a definition for that node. Thus, multiple gateways may be active to a single host (see "5.0 3745 High Availability Configurations with Peripheral Nodes" on page 18 for examples).

3.3 Definitions of SNA Subarea Nodes

Defining Subarea Nodes on the Token-Ring Network

Subarea Nodes on the Token-Ring Network are defined in the NCP generation. The NCP-NCP (INN) session is established across a specific TIC defined in the NCP. With the Token-Ring Adapter (TRA) Type 1, a TIC must be dedicated to Subarea traffic. The TRA Type 2 supports both Subarea and Peripheral traffic across the same TIC. Multiple Type 4 or 5 Node definitions are supported across a single TIC; therefore, through a single TIC, a NCP may have INN sessions with multiple Type 4 and 5 Nodes connected to the Token-Ring Network.

INN sessions across Token-Rings are the same as INN sessions across SNA lines. For example, XI Nodes can communicate across a Token-Ring. When using Subarea networking across the Token-Ring Network, some limitations exist:

Duplicate TIC addressing (coding two TICs with the same address) is not supported for Subarea networking.

Multiple Link Transmission Groups are not supported for Subarea connections across the Token-Ring Network.

Frame Sizes up to 16K are supported across the Token-Ring. (The frame size may be limited by a remote bridge - IBM Bridge Program V2.2 supports a 4472 byte frame size).

PIU Blocking (feature of NCP V5.3) is not supported on Subarea connections across the Token-Ring Network.

However, subarea networking across a Token-Ring Network has advantages:

Single Network Connection for all Subarea Networking

All Connections Between Subareas on Token-Ring can be Adjacent Connections

Non-Disruptive Token-Ring Route Switching (see section 5.0)

First, a 3745 can communicate through a single TIC with all other 3745s and Type 5 Nodes on the Token-Ring Network.

Second, INN sessions across Token-Ring can be direct to the other subarea, instead of through an intermediate NCP. This simplifies NCP path definitions between subarea nodes on the Token-Ring network, because all paths can be directly to the destination subarea.

Finally, when alternate routes exist in the Token-Ring Network, non-disruptive route switching can bypass network failures (see "6.0 3745 Communication with Subarea Nodes" on page 27).

A sample Subarea NCP definition is shown below. A similar definition for this NCP must be defined in the remote NCP on the Token-Ring.

```
GROUP ECLTYPE=(LOGICAL,SUBAREA)
      SDLCST=(stprim,stsec),        *Primary/Secondary Select Table (defn. not shown)
      PHYPORT=n                     *Matches PortAdd of TIC used
LINE
PU    TGN=n,                        *Transmission Group Number
      MAXOUT=n,                     *Transmit Window Size
      Addr=ssaaaaaaaaaa             *s=Dest. SAP,a=Token-Ring address of remote NCP
```

3.4 Other NCP Token-Ring Parameters

Token-Ring Networking relies on Timers to ensure communications across the network. Three timers implemented by all SNA Nodes are:

"T1 Response Timer"

"TI Inactivity Timer"

"T2 Acknowledgement Timer"

First, a T1 Response Timer ensures that a message is delivered to the partner across the Token-Ring Network. When a data frame or command is sent to a partner, a response must be received within "T1 Response Timer" seconds. If the partner does not respond, then either the outstanding data frame or a poll is sent. If a response is still not received, then this process is repeated for (N2) number of retries. NCP implements two T1 Response Timers: local (LOCALTO) and remote (REMOTTO). The local timer is used when the partner SNA node is on the same ring as NCP, the remote timer when on a ring bridged (local or remote) to the NCP ring. The number of retries (N2) is specified on the NCP RETRIES parameter.

Second, "TI Inactivity Timer" ensures that both Token-Ring devices are still communicating. When no traffic is received or sent for "TI Inactivity Timer" seconds, a Token-Ring device sends a poll to the partner which must respond. The response shows that communications exists between the session partners. NCP implements a "TI Inactivity Timer" of 60 seconds.

Finally, a "T2 Acknowledgement Timer" enables more efficient communications across the Token-Ring Network. When a Token-Ring device receives a data frame, it can wait "T2 Acknowledgement Timer" seconds for more frames before sending a response. This timer works with the "N3" counter, which is the number of frames which can be received before sending a response. NCP Version 5.4 will support a "T2 Timer" and a "N3" Counter. NCP Version 5.3 (and previous releases) supports a "T2 Timer" of 0 and "N3" Counter of 1.

Other parameters must be coded in NCP to support Token-Ring SNA Nodes. This document is not intended to replace the official documentation available. Some references for coding NCP Token-Ring parameters are the NCP Resource Definition Reference (SC30-3448) and the Installation Guidelines for the IBM Token-Ring Network Products (GG24-3291). Tuning information is available in NCP Version 5 Network Performance and Tuning (GG24-3469).

3.4 Summary of Token-Ring Node Definitions

The Token-Ring Networks' workstations and SNA Nodes are defined in both VTAM and NCP. The following figure summarizes where the definitions are made.

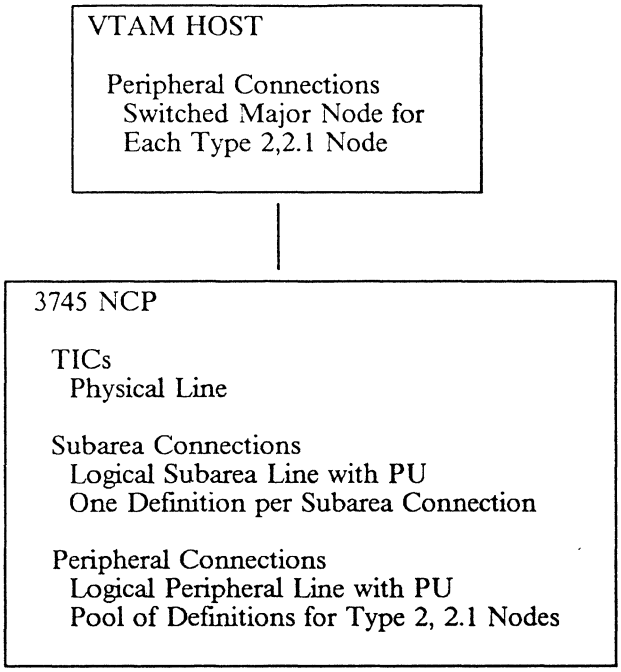

4.0 3745 Communication with Peripheral Nodes

Peripheral Nodes are Type 2.0 (e.g. 3174, PS/2 running 3270 emulation) and 2.1 Nodes (e.g.AS/400 APPN, PS/2 running APPC/PC) on the Token-Ring Network. The 3745 supports communications with the peripheral nodes across a Token-Ring Network. This section describes the flows to establish communications between the 3745 and peripheral nodes; it is divided into three parts:

Benefits of 3745 Peripheral Node Networking on Token-Ring

Peripheral Node Session Establishment - Single Route

Peripheral Node Session Establishment - Multiple Routes

4.1 Benefits of 3745 Peripheral Node Networking

The 3745 supports peripheral node access to SNA applications on hosts and peer (Type 2.1) nodes in the network. In traditional configurations, SNA Nodes were leased line attached to a 3745 or were channel attached to a host. Most applications ran on VTAM hosts in the network. In today's networks, Token-Ring Networks provide the transport to support SNA and non-SNA peer communications. A SNA Node can access VTAM host applications through a gateway, such as the 3745 controller. The figure below shows traditional and Token-Ring Network configurations.

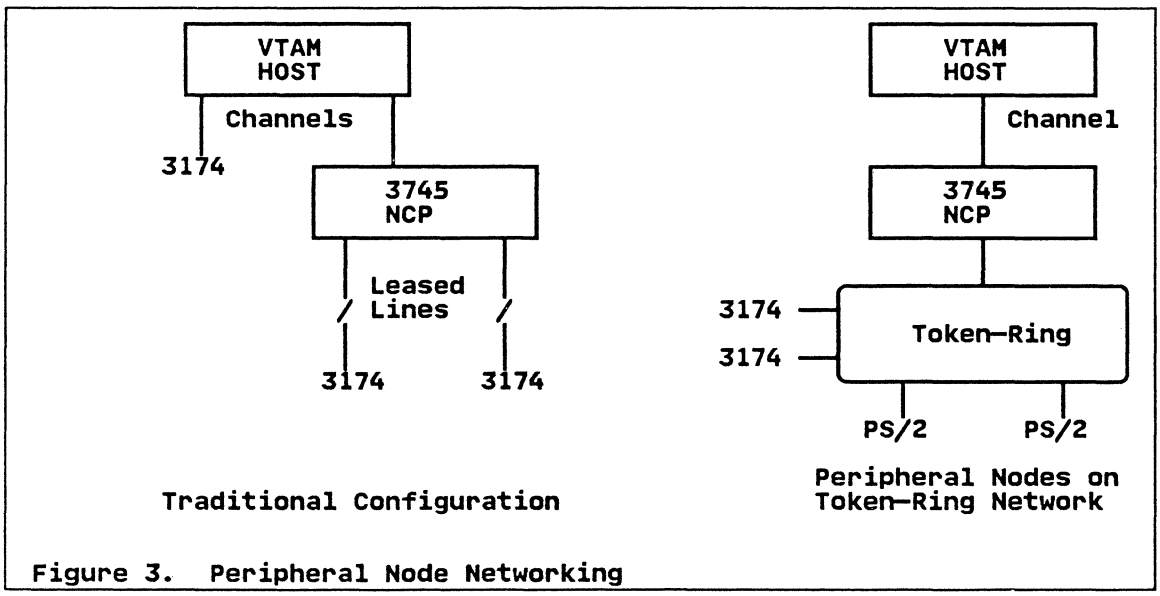

Figure 3. Peripheral Node Networking

The figure above shows two methods to provide establishment connectivity. The traditional configuration shows two types of control units: channel attached and leased line attached. The channel attached controllers can provide subsecond response to applications on the hosts. However, many users were not connected to channel attach controllers for various reasons - controller outside of channel cable lengths, higher cost of channel attach controller, etc.. When leased lines connected 3174s into the 3745, the user's network response time was based on the line speed and the number of controllers sharing the line. Lines were much slower than channels; therefore, response times are usually worse than on channel attach controllers.

In recent years establishment connectivity has shifted to high speed local area networks which provide connectivity between multiple application platforms. In the preceding figure, the 3174 controllers are moved to the Token-Ring Network to provide high speed connectivity to the host. The PS/2s are also connected to the Token-Ring for host access (i.e. 3270 emulation) and peer access (i.e. file server). Thus, the Token-Ring Network supports high speed communications between all devices in an establishment.

Some of the benefits of a Token-Ring Network with 3745s for establishment connectivity are:

Support for Simultaneous SNA and non-SNA Communications on a Token-Ring

Near Channel-Attach Network Response Time for Users on a Token-Ring

No 3745 Polling of Token-Ring SNA Nodes

Less 3745 Communication Ports Required

NCP Generations Not Required For Token-Ring Network Changes

NCP Routing to Multiple VTAM Hosts

A major benefit of the Token-Ring Network is the simultaneous support of multiple protocols across the network. Workstations may access peer applications (i.e. access a OS/2 File Server) at the same time a 3174 Workstation runs 3270 SNA protocol to access a host application. The Token-Ring Network is a data transport network that is protocol independent; it supports SNA and non-SNA communications concurrently.

Second, Token-Ring users receive near channel attach network response time when accessing host or peer applications. The Token-Ring Network runs at 4 or 16 million bits per second; workstations and 3174 controllers can attach directly to the Token-Ring Network. The network response time for users on the Token-Ring Network (excluding users accessing the network through remote bridges where the speed of the leased line between the bridges can affect response time) should be close to the network response time achieved through a channel-attach 3174 controller. Since the Token-Ring cabling supports distances which far exceed channel distance limitations, all controllers within an establishment support near channel attach speeds (not just the controllers within the data center).

Third, a 3745 supports many Peripheral Nodes through a single TIC, because the **3745 does not poll Peripheral Nodes on a Token-Ring Network like controllers on SDLC lines.** With traditional SDLC leased lines, Peripheral Nodes were continuously polled by the communication controller and could only send data when polled by NCP. Multiple Peripheral Nodes were supported on single line, but network response time degraded as more were added to a single line. On a Token-Ring Network, either NCP or Peripheral Nodes can initiate communications (NCP is not required to poll the Peripheral Node as it does on a SDLC line). All Token-Ring SNA Nodes (Type 2.0, 2.1, 4, and 5) have "TI Inactivity Timers" and "T1 Response Timers". NCP has a "TI Inactivity Timer" set to 60 seconds and a "T1 Response Timer" set by NCP Parameters - LOCALTO or REMOTETO. NCP maintains a separate set of timers for each session with a SNA Node. Peripheral Nodes (i.e. 3174 or PS/2) also have "T1 Response Timers" and "TI Inactivity Timers"; generally their "TI Inactivity Timers" are shorter than NCP's "TI Inactivity Timer" of 60 seconds. When either a NCP or a Peripheral Node timer expires, a poll is sent to the session partner to ensure it is still active and communicating. A Token-Ring trace will show polls, but these polls may be generated by either NCP or the peripheral node. These polls are much less frequent than polls on a SDLC line. Thus, many SNA Nodes can simultaneously communicate with the 3745 across a Token-Ring Network, because NCP has much less polling overhead and higher speed network connections (4 or 16 Mbps).

Fourth, less 3745 communications ports are required to communicate with peripheral nodes on the Token-Ring Network than on SDLC leased lines. With SDLC lines, communication ports connect the 3745 to a leased line to communicate with one or more link-attach SNA Nodes; many communication ports are usually necessary to communicate with the SNA Nodes. Adding link-attach SNA Nodes to the Network could require 3745 hardware upgrades to add ports, scanners, expansion frames, etc.. With the Token-Ring Network, a single TIC supports thousands of SNA nodes without any 3745 hardware changes to add new SNA Nodes.

Fifth, no NCP software changes are required to add/change/delete Token-Ring SNA Nodes. The NCP generation has a pool of definitions for all Token-Ring SNA Nodes. Individual SNA Node definitions are made in VTAM, as are all adds/changes/deletions. Thus, new SNA Nodes may be

dynamically added to VTAM, and the Token-Ring, without changing the NCP generation (assuming the pool of definitions in NCP is large enough to support the additional users).

Finally, NCP routes SNA traffic to multiple hosts in the network. Unlike the 3174 and 3172 local gateways, the 3745 is not required to send all SNA data to the the owning VTAM (VTAM in which the SNA Nodes are defined). After the owning VTAM establishes the SNA LU-LU session, NCP routes the SNA frames to the appropriate destination (e.g. channel attach VTAM host, another 3745, or Type 2.1 Node on the same 3745). Therefore, when multiple VTAM hosts exist in the network, 3745 SNA routing can offload processing in the owning VTAM.

4.2 Peripheral Node Session Establishment - Single Route

This section describes session establishment where only a single Token-Ring route exists between the peripheral node and the 3745. The figure below shows an example of peripheral nodes on a Token-Ring Network and is used to describe the establishment of communications between the 3745 gateway and the peripheral node.

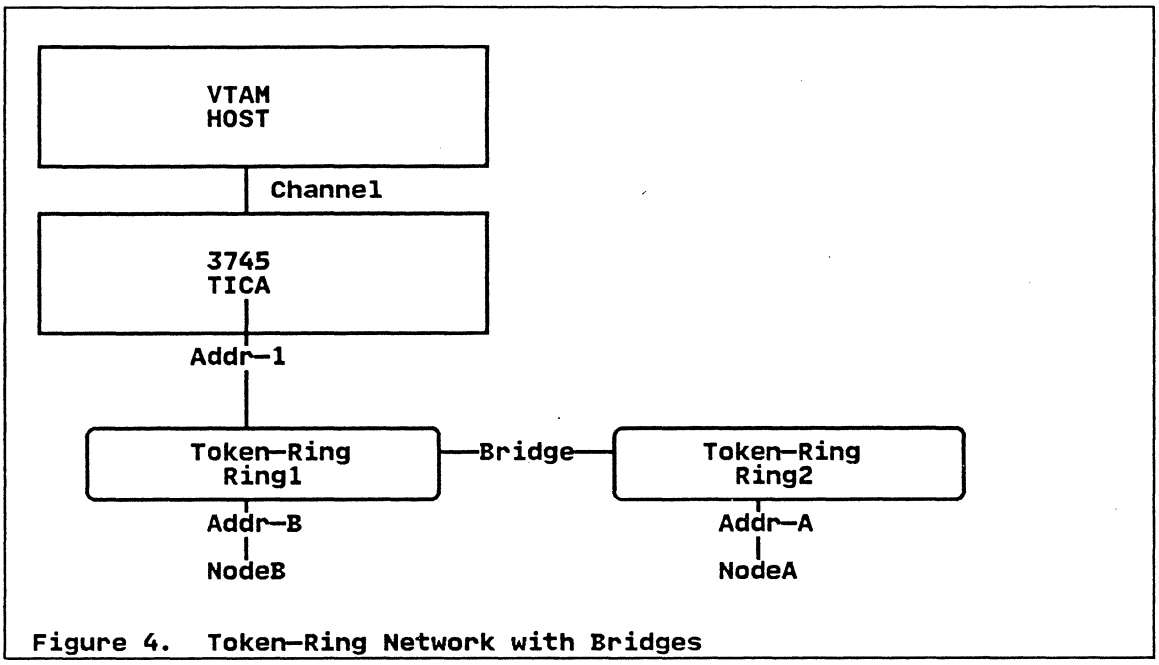

Figure 4. Token-Ring Network with Bridges

In the figure above, Ring1 and Ring2 are bridged together to create a logical Token-Ring Network. NodeA and NodeB function as Type 2.0 or 2.1 nodes on the Token-Ring Network. These devices have native SNA support (i.e. 3174-13R) or run SNA emulation programs (i.e. OS/2 EE Communication Manager emulating a 3174 SNA controller). NodeA attaches to the Token-Ring Network at address Addr-A; NodeB attaches at Addr-B. Both NodeA and NodeB have one route to the 3745 Token-Ring gateway through the Token-Ring Network (NodeB across Ring1, NodeA through the Bridge). The 3745 connects to the Token-Ring network at address Addr-1, which is specified in the NCP generation. Both NodeA and NodeB are customized to access the host through the gateway at Token-Ring address Addr-1 (the 3745 Token-Ring Interface Coupler address). Communications between the gateway and the node may be initiated by either the SNA Node or the VTAM host. In most cases, peripheral nodes will initiate communications with the 3745 gateway.

How do peripheral nodes initiate host communications?

Peripheral nodes on the Token-Ring Network use source routing to establish communications with the 3745 gateway. The IBM Multisegment LAN Design Guidelines (GG24-3398) provides an excellent description of source routing across bridged Token-Ring Networks. Most SNA nodes initiate communications with the 3745 gateway as described below. However, certain SNA nodes (e.g.3270 Emulation Program V3) send a single-route broadcast rather than an all routes broadcast. The manual, GG24-3398, describes the differences. Workstation Program V1.1, PC/3270 Communications V1.0, OS/2 EE Communications Manager, 3174-13R and other SNA Nodes use the basic communications flows shown below. The example below describes how source routing allows NodeA to initiate communications with VTAM. Please refer to Figure 4.

1. NodeA issues a token-ring frame (TEST RESOLVE or XID) which is sent to Addr-1 on the local ring (Ring2).

2. Since Addr-1 does not respond (it is not on Ring2), NodeA sends an all-routes broadcast (TEST RESOLVE or XID) which flows through all bridges, requesting Addr-1.

3. Each bridge on ring2 forwards the frame and adds the routing information to the broadcast frame. All bridges on ring2 perform this process which repeats itself until all bridged rings are searched or the maximum number (hop count) of bridges have been traversed.

4. TICA receives the broadcast frame from NodeA and responds to the peripheral node. The response retraces the route which was built by the bridge during the broadcast.

5. NodeA receives TICA's response, which contains the route to use for future communications between NodeA and TICA. NodeA then sends TICA a SNA XID in a Token-Ring frame, which notifies TICA of the route to use for the session. (Note: Type 2 Nodes send a SNA XID-0; Type 2.1 nodes send a SNA XID-3.) If multiple responses are received from TICA (i.e. if multiple routes exist), then the first response is usually selected as the route. All additional responses are discarded.

6. NCP receives the SNA XID and passes the information to VTAM.

7. VTAM receives the request from NCP and determines which device sent the frame from the IDNUM and IDBLK parameters. VTAM then establishes communications with the SNA Node.

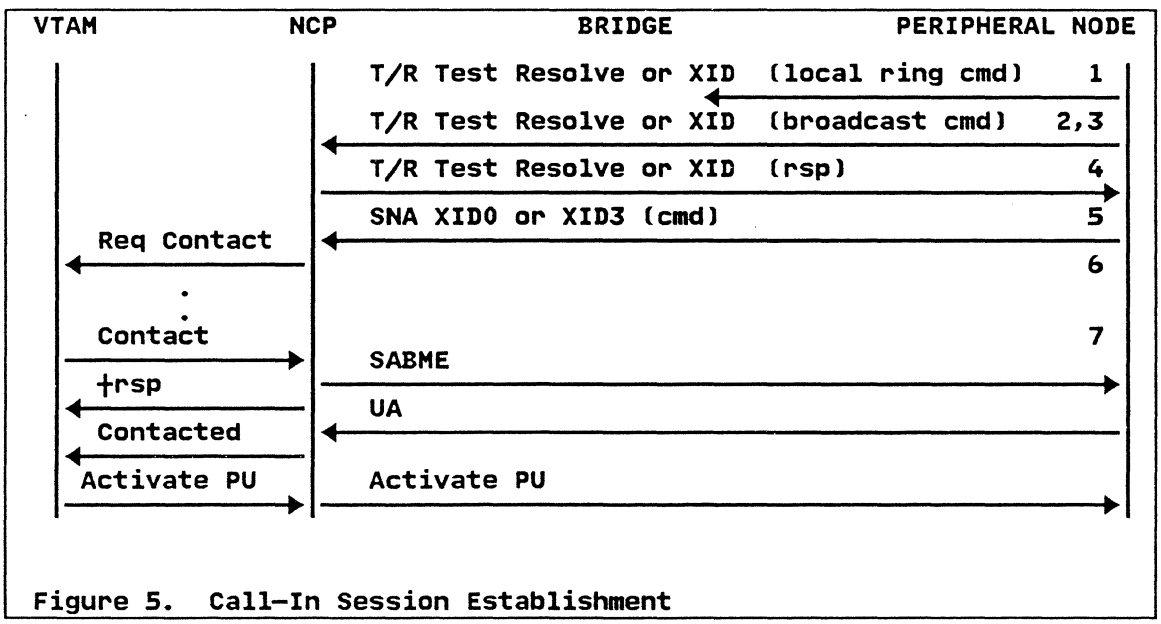

Figure 5. Call-In Session Establishment

(Note: The description above is not intended to describe the detailed flows between VTAM, NCP and the Node. It provides a basic understanding of the command flows which occur to make the Token-Ring node known to the VTAM host. The numbers under the peripheral node correspond to the text above the figure).

After source routing determines the Token-Ring route across which the 3745 TIC and NodeA communicate, all subsequent communications between the SNA Node and the TIC flow across the same route. Note that two different types of XID may flow: the LAN XID and the SNA XID. The LAN XID or TEST command is used to determine the route across the Token-Ring Network. The SNA XID is sent because NodeA appears to VTAM as a switched major node. Switched devices initiate communications with a SNA XID that tells VTAM which switched major node is requesting communications. If NodeA has been activated to VTAM, then VTAM sends the appropriate commands to initiate communications with the peripheral node.

The commands which flow after the SNA XID are used by VTAM to set up communication with the SNA Node. Not all of the VTAM to NCP commands are shown in the diagram. NCP sends a Request Contact which requests VTAM to activate the SNA Node on the Token-Ring. After several commands between NCP and VTAM (these are not shown), VTAM responds with a Contact command to NCP. NCP then sends a SABME command which sets up an Asynchronous Balanced Mode communication with the SNA Node on the Token-Ring. Asynchronous Balanced Mode communications allows either device to initiate communications - the 3745 does not poll the Token-Ring device to receive data. The UA command is the acknowlegement from the SNA device on the Token-Ring. NCP sends the Contacted response to VTAM which then activates the SNA Node with the ACTPU command.

NodeB Session Establishment

NodeB establishes communications with TICA using the same process. However, NodeB is not required to send a broadcast, because NodeB and TICA are on the same Token-Ring. Thus, steps 2 and 3 are not required for most SNA Nodes on the same ring as the TIC.

How does VTAM initiate communications with peripheral nodes?

VTAM may initiate communications with SNA Nodes on the Token-Ring. Normally, VTAM does not initiate communications with Type 2 SNA Nodes, such as a PC emulator. However, VTAM may initiate communications with Token-Ring nodes, such as a PC running APPC/PC, when the Path statement is coded in the switched major node. Host-initiated communications uses Token-Ring source routing to find the destination SNA Node. An example (see Figure 4 on page 11) of the flow that occurs when VTAM initiates communications with NodeA is shown below:

1. VTAM looks up NodeA's Token-Ring address in the switched major node definition (PATH statement).
2. VTAM requests NCP to locate NodeA at address Addr-A
3. NCP sends a Token-Ring frame (TEST RESOLVE) to address Addr-A on the local ring.
4. Since Addr-A does not respond (it is not on Ring1), NodeA sends an all-routes broadcast (TEST RESOLVE) which flows through all bridges, requesting Addr-A.
5. Each bridge on ring1 forwards the frame and adds the routing information to the broadcast frame. All bridges on ring1 perform this process which repeats itself until all bridged rings are searched or the maximum number (hop count) of bridges have been traversed.
6. NodeA receives the broadcast frame from TICA responds to TICA. The response retraces the route which was built by the bridge during the broadcast.
7. TICA receives the response from NodeA, and the response contains the route that will be used for future communications between NodeA and TICA. TICA then sends NodeA a SNA frame (null XID), which notifies TICA of the route to use for the session. Note: If multiple responses are received from TICA (i.e. if multiple routes exist), then the first response from the peripheral node is selected as the route for the session. All additional responses are discarded.
8. NodeA responds with a SNA XID, which requests communications with the host. (Type 2 nodes respond with a XID-0; Type 2.1 nodes respond with a XID-3.)
9. VTAM receives the SNA frame from NCP and determines which device made the request. VTAM then initiates communications with the Token-Ring Node.

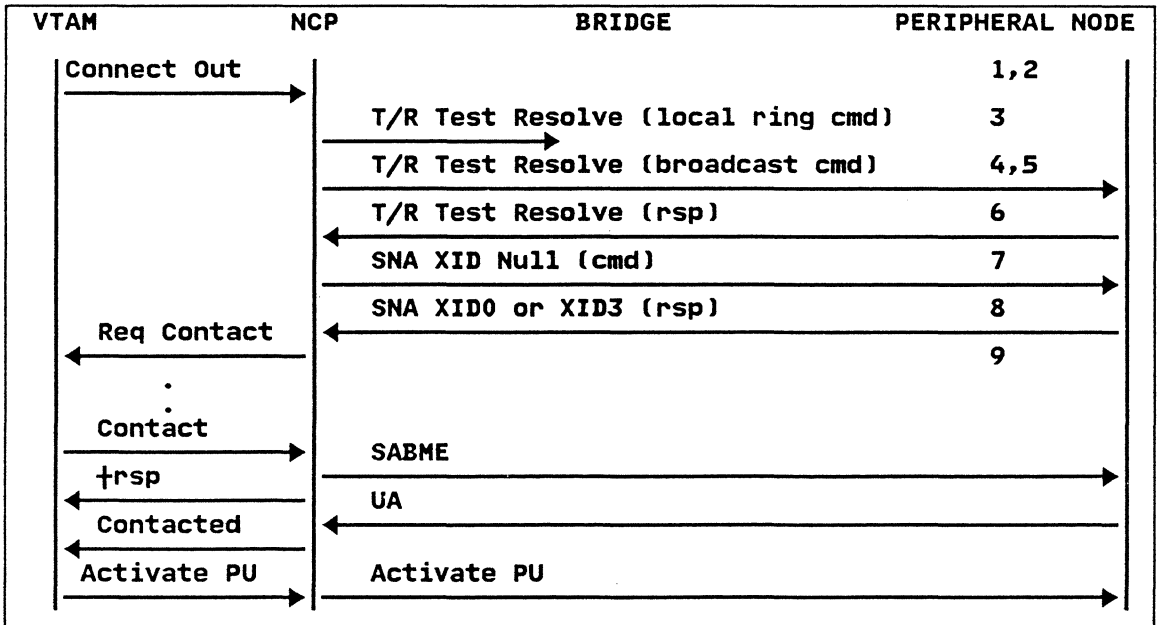

Figure 6. Call-Out Session Establishment

(Note: The description above is not intended to describe the detailed flows between VTAM, NCP, and the Node. It provides a basic understanding of the command flows which occur to make the Token-Ring node known to the VTAM host.)

4.3 Peripheral Node Session Establishment - Multiple Routes

The Token-Ring Network may be configured with parallel and alternate routes to provide high availability. The 3745 transparently supports these networks; no VTAM or NCP definitions are required to support parallel or alternate routes. Parallel routes are supported in the Token-Ring Network with two bridges which provide parallel routes between two rings. Bridges may also create alternate routes between two devices on a Token-Ring Network. The figure below shows a Token-Ring Network with parallel bridges and alternate bridge routes.

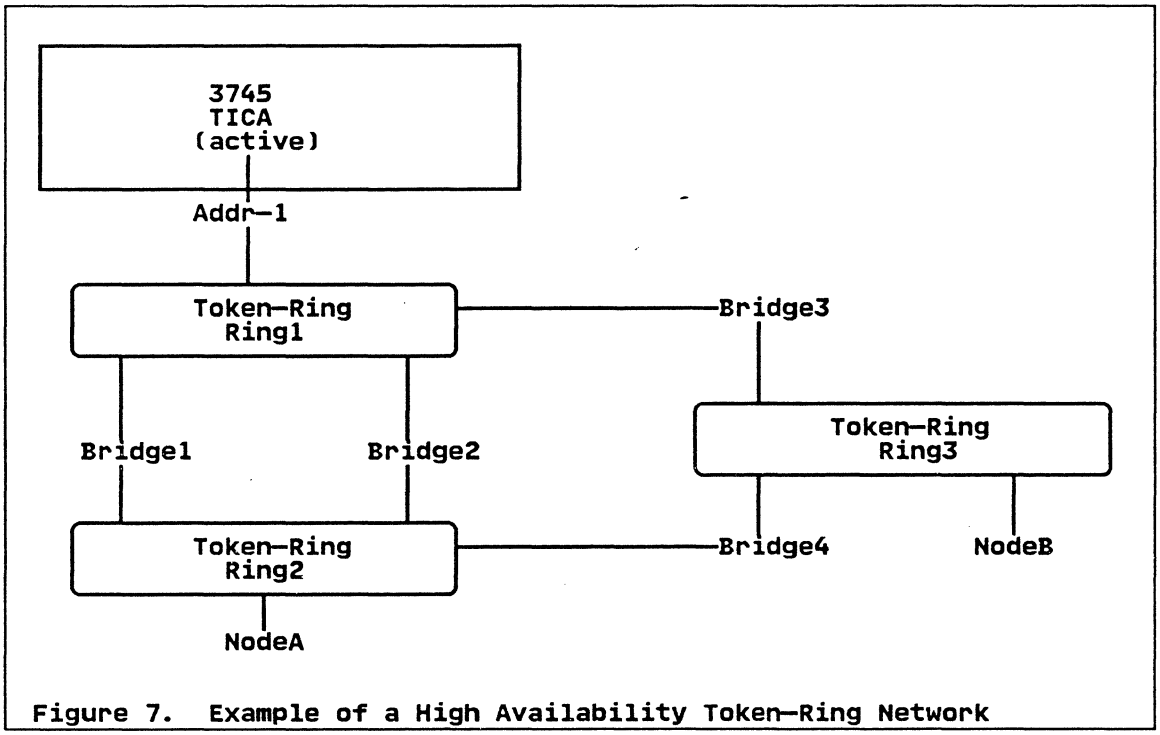

Figure 7. Example of a High Availability Token-Ring Network

Peripheral Nodes on any of the rings above may access the host through the 3745 gateway. Source routing is used by either NCP or the SNA Node to establish communications with the 3745 Gateway. Alternate and parallel routes are designed into the Token-Ring Network to provide higher availability. Bridge1 and Bridge2 are an example of parallel bridges which provide parallel routes between rings. Bridge3 and Bridge4 provide an alternate route from ring2 to ring1 through ring3. With source routing the fastest route, selected during the establishment of communications, is usually chosen for session communications.

An example of source routing with multiple routes is NodeA establishing communications with the 3745. This example assumes that the peripheral node is establishing the host connection and that no bridge filters or hop count limitations exist. If VTAM makes the connection, then the flows are similar with the peripheral node responding to multiple broadcast frames from the 3745.

1. NodeA issues a token-ring frame (TEST RESOLVE or XID) which is sent to Addr-1 on the local ring (Ring2).
2. Since Addr-1 does not respond (it is not on Ring2), then NodeA sends an all-routes broadcast (TEST RESOLVE or XID) which flows through all three bridges requesting Addr-1.
3. Each bridge on ring2 forwards the frame and adds the routing information to the broadcast frame. All bridges on ring2 perform this process which repeats itself (on ring1 and ring3) until all bridged rings are searched or the maximum number (hop count) of bridges have been traversed.
4. TICA receives three frames from NodeA - one frame for each route between the rings (route 1 - bridge1, route 2 - bridge2, route 3 - bridge4, bridge3). TICA responds to each of the frames, and the response retraces the route which was built by the bridge during the broadcast.
5. NodeA receives three responses from TICA, and each response contains a route that may be selected. The peripheral node usually selects the first response and uses the route in that response for future communications.
6. NodeA sends a SNA XID in a Token-Ring frame to TICA across the chosen route. (Type 2 Nodes send a SNA XID-0; Type 2.1 nodes send a SNA XID-3.)
7. NCP receives the SNA XID and passes the information to VTAM.
8. VTAM receives the request from NCP and determines which device sent the frame from the IDNUM and IDBLK parameters. VTAM then establishes communications with the SNA Node.

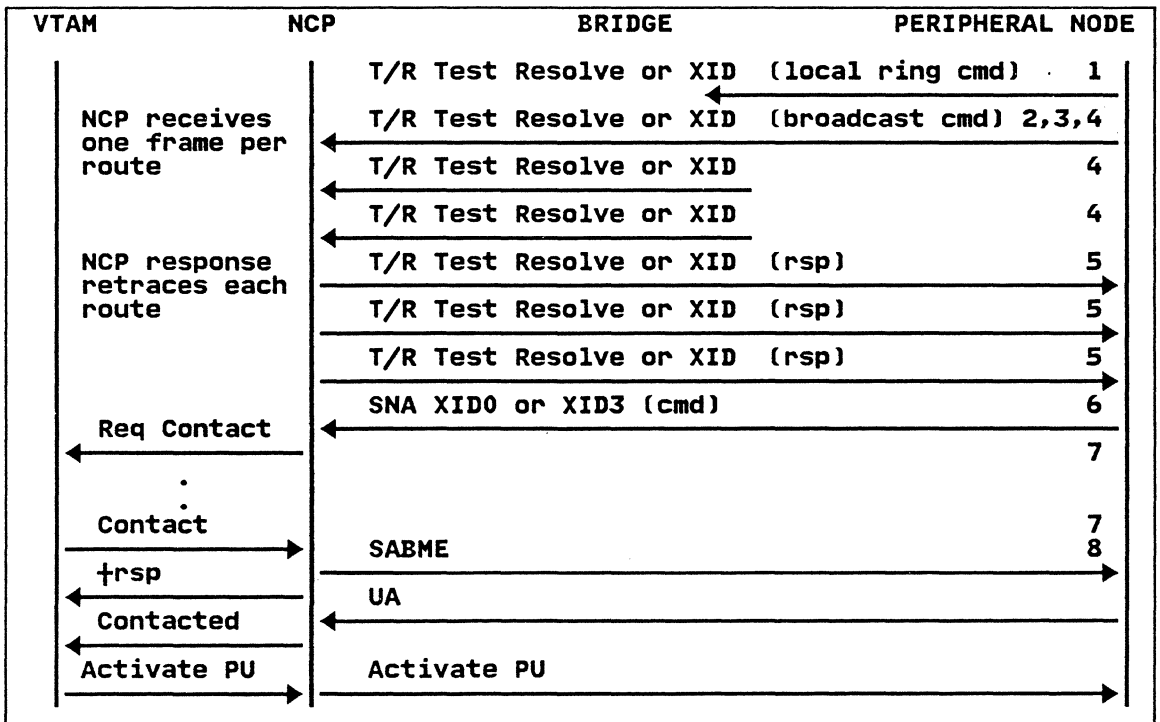

Figure 8. Call-In Session Establishment

The description above shows how multiple routes provide higher availability across the Token-Ring Network. When a SNA Node is activated, it usually chooses the fastest route available. In the

event of a failure, the SNA Node can select another route using the same process as the first selection. One issue with providing many routes is the number of broadcast messages created during activation of the SNA Node. However, SNA Nodes usually remain active until power off or termination of the emulation package. **SNA Nodes only send broadcast frames when activating the control unit to VTAM (SSCP-PU Session) not when a user logs on and off applications.**

Reducing Broadcast Frames

A Token-Ring Network with many routes to peripheral nodes creates multiple broadcast frames during SNA Node activation. The IBM Bridge Program can reduce the number of broadcast frames in three ways: ring reentry checks, hop count and filters. Ring reentry checks automatically occur in a bridge and ensure broadcast frames are only forwarded once across a ring (no loops). Bridges check each broadcast frame and will not forward the frame if it has already crossed the destination ring. Hop count is a bridge parameter which specifies the number of bridges that a broadcast frame may pass through before it is discarded. The maximum hop count is 7; lowering the hop count may reduce broadcast traffic without limiting communications. Filters are bridge programs which allow only certain frames to pass through to the adjacent ring. The filters can eliminate some broadcast traffic by not forwarding filtered frames.

5.0 3745 High Availability Configurations with Peripheral Nodes

The 3745 Token-Ring Subsystem supports multiple configurations which provide varying degrees of availability for the Token-Ring Network users. This section discusses a subset of the configurations, and the availability issues associated with each. Since each Token-Ring Adapter ships with two TICs, the simplest configuration is a single 3745 with two TICs. Obviously, these adapters could attach to two independent Token-Ring Networks and provide host access for the independent Networks. However, this paper assumes that the Network Designer desires some backup in the event of a failure. Therefore, this section only considers single and multiple 3745s with two TICs.

The last topic in this section discusses a local gateway configuration which could provide performance problems and should be avoided.

5.1 Single 3745 Ring, Duplicate TIC Address

A single 3745 provides high availability networking with a backup TIC. The first configuration consists of a Token-Ring Network with the 3745 TICs on a single ring. Both TICs are generated in the NCP at the same Token-Ring address. A single ring does not support duplicate addresses, therefore one of the TICs must be inactive. The figure below shows a possible configuration.

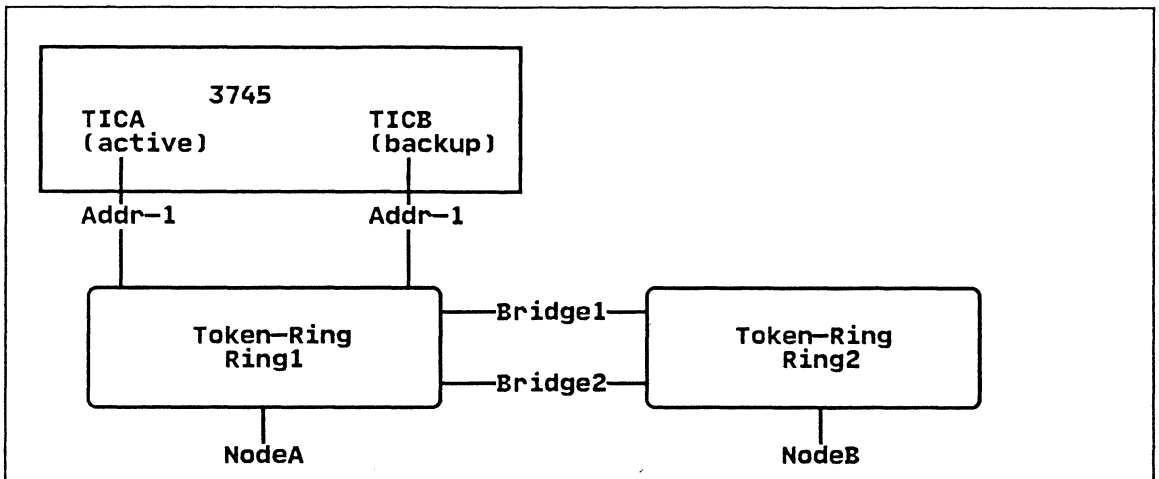

Figure 9. Example of a 3745 with Duplicate TIC Address

In the above figure, the 3745 is configured with two TICs - TICA and TICB. Both TICA and TICB are defined at the same address - Addr-1 (duplicate TIC addressing requires NCP V5.3, V5.2.1 w/PTF, or V4.3.1 w/PTF). TICA is the active TIC, and TICB is the inactive backup TIC. NodeA and NodeB are peripheral nodes which are customized to access the host gateway at address Addr-1.

Recovery from TIC Failures

A TIC failure disrupts Token-Ring users accessing the host through the failing TIC. If TICA (the active TIC) fails, then TICB must be activated by a VTAM command. After the backup TIC is activated, NodeA and NodeB may reestablish connections to the host. Token-Ring users may then log back onto their applications. Netview CLISTs could automate activation of the backup TIC.

Recovery From Token-Ring Route Failures

A Token-Ring route failure is the failure of a Token-Ring component in the route being used to access the host. For example, NodeB has two routes to the host - one route through bridge1 and the other route through bridge2. The route is chosen when the peripheral node is activated and remains the route until the device is deactivated. If NodeB is using the route through bridge1, then a failure of bridge1 would be a Token-Ring route failure.

Certain Token-Ring Route components are single points of failure. For example, the NCP TICs are attached to a single ring. If ring1 failed then no users could access the host. This paper assumes an alternate route is available for recovery (i.e. bypass bridge1 failure by crossing bridge2).

A Token-Ring route failure disrupts connections with peripheral nodes. All sessions across the failing route (with the possible exception of a subarea node - see "6.0 3745 Communication with Subarea Nodes" on page 27) are disrupted and the SNA Node times out. If alternate routes exist (i.e. a parallel bridge), then the peripheral node is able to re-access the host through the alternate route by re-establishing communications across the Token-Ring Network with source routing. Recovery time is dependent on how long it takes VTAM and the peripheral node to time-out and break the session.

Recovery from 3745 Hardware or NCP Failure

Both TICA and TICB utilize the same Token-Ring Adapter (TRA). Several 3745 components are common to both TICA and TICB, such as a Token-Ring multiplexor. Other 3745 components, such as the distributed power supply to the TRA and the NCP software are common points of failure for both TICs. If one of the common components failed, then the Token-Ring users would have no host access until the failing component was repaired.

5.2 Multiple 3745 Rings, Duplicate TIC Address

The previous configuration requires the activation of a backup TIC before Token-Ring devices may re-access the host. This configuration provides higher availability with an active alternate route. Another advantage of this configuration is that traffic is distributed across each TIC. Each TIC is active and may receive traffic, so the workload does not flow across one TIC as in the previous configuration (see discussion later in this section on why this occurs).

Token-Ring addresses must be unique on a given ring. However, different rings may have devices with the same address. These rings may be bridged together to create one logical Token-Ring Network, which has the same address on two different rings. The following configuration illustrates a way to bridge together Token-Rings when there are duplicate gateway addresses.

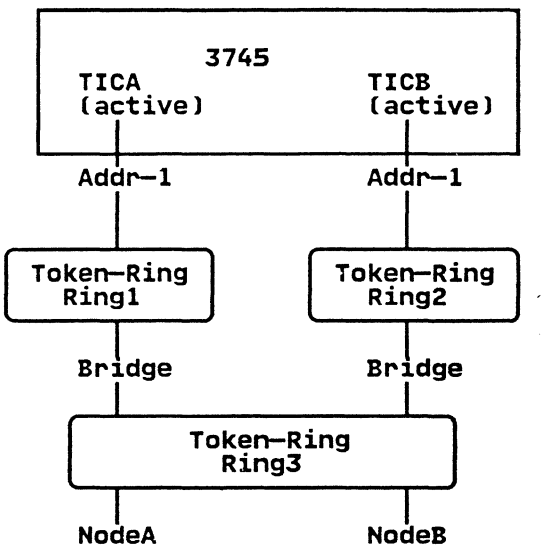

Figure 10. Example of Multiple TICs with the same Token-Ring Address

This configuration shows three rings bridged together in a single logical Token-Ring Network. TICA and TICB both have a Token-Ring address of Addr-1. These TICs are on separate rings, which satisfies the requirement of unique addresses on a given ring. Both peripheral nodes access the host through Token-Ring address Addr-1, which appears on two different rings. As discussed in Section 3, peripheral nodes access the 3745 with source routing. In the example above, NodeA uses the following flow to establish communications with the 3745 gateway:

1. NodeA sends a XID or TEST frame to Addr-1 on Ring3 (the local ring)
2. Addr-1 is not on the local ring, so NodeA sends a broadcast which flows through both bridges, requesting Addr-1. The bridges add routing information and pass on the request to the connected rings (Ring1, Ring2).
3. TICA and TICB both see the request. Both TICA and TICB independently respond to the request, and the response retraces the route which was built by the bridges.
4. NodeA receives the first response from one of the TICs and usually selects that route for host communications.
5. When NodeA receives the response from the other TIC, it discards the response, because a route has already been chosen.

In the figure above, two routes are available to TIC address Addr-1. Either route may be used for communications, although the fastest route at SNA Node activation is usually chosen. Each TIC responds to Node A's request for routing information. NodeA chooses the first response that is returned from a TIC. Assume the response from TICA is returned first and is the route used for NodeA's connection to the host. When NodeB establishes communication with the host, NodeB

broadcasts a request to Addr-1. Assume that TICB returns the routing information to NodeB before TICA, because TICA was busy with NodeA. NodeB selects the first response received (TICB's path) for the session route. Since the fastest path is chosen each time, the load should be naturally balanced across the two TICs.

This configuration also allows all SNA Nodes (PC gateways, 3174s) to access the host through the same gateway address. The traffic is split over multiple routes and TICs, so one TIC is not overloaded. (NOTE: 3174-x3R controllers sessions may not balance across multiple TICs. Since the 3174s establish host sessions when the network is first activated and little traffic exists on the network, the first 3745 TIC activated may be the route for all 3174 communications).

Recovery from TIC failures

This configuration backs up a TIC failure. All SNA sessions through the failing TIC are disrupted; all sessions through the operational TIC remain active. Users crossing the failing TIC may restart the SNA host connection (e.g. restart emulation package) through the alternate bridge and TIC. This capability is provided by source routing. The peripheral node issues the broadcast again, and the only available TIC responds to the broadcast. The peripheral node then restarts its session across the available path.

Recovery from Token-Ring Route Failures

A Token-Ring route failure disrupts connections with peripheral nodes. All sessions across the failing route (with the possible exception of a subarea node - see "6.0 3745 Communication with Subarea Nodes" on page 27) are disrupted and the PU times out.

This configuration provides additional Token-Ring routes to help availability. There is a backup NCP ring and backup bridge in case of failure. Therefore, NodeA and NodeB have alternate routes to the host in the event of either ring1, ring2 or bridge failure in the route to the host. Since two routes exist, NodeA may be disrupted while NodeB continues operation across the available route. NodeB could be restarted and access the host through the available Token-Ring route.

Recovery from 3745 Hardware or NCP Failure

Similar single points of failure exist in this gateway configuration as in the previous configuration. A 3745 failure, NCP failure, TRA power supply failure or TRA common component failure will require repair before a gateway path may be recovered.

5.3 Single 3745 Ring, Multiple 3745s

Multiple 3745s provide backup in the event of a 3745 failure. Each 3745 connects into the Token-Ring Network to provide alternate paths to the host. This document discusses the configuration options with two 3745 controllers; each 3745 is equipped with two Token-Ring Interface Couplers (TICs). Other configurations with more than two TICs in a 3745 and more than two 3745s provide additional backup possibilities. However, two 3745s show the basic configurations.

One configuration is to keep one 3745 active and the second 3745 as backup. A matrix switch is usually in place to switch telecommunication lines from the active 3745 to the backup 3745 in the event of a hardware failure. No matrix switch is necessary to backup the Token-Ring Connection. Both 3745s are connected into the Token-Ring Network, and only the active 3745 adapter receives Token-Ring traffic.

Another configuration (shown below) is two active 3745s with spare capacity for backup of critical resources. Peripheral nodes on the Token-Ring Network are assigned gateway addresses associated with one of the 3745s. Thus, the token-ring load is distributed across the 3745s. Each 3745 has an inactive TIC which backs up the Token-Ring Network connection.

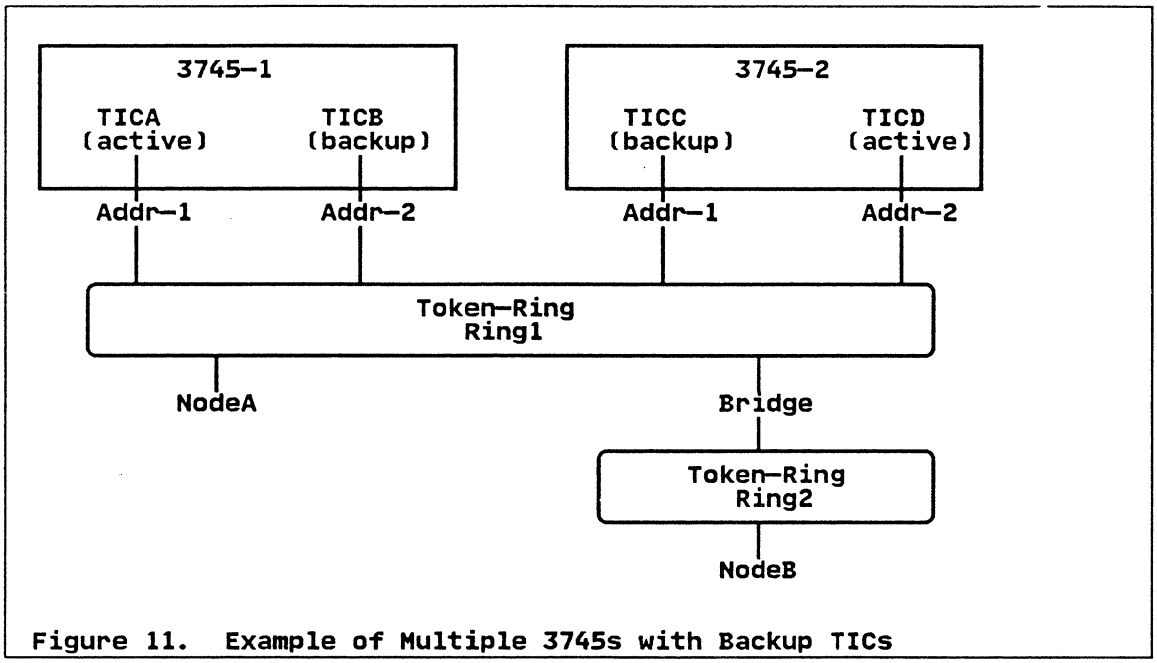

Figure 11. Example of Multiple 3745s with Backup TICs

In the figure above, each 3745 is active and has two TICs. One 3745 has TICA active at address, Addr-1 and TICB inactive at address, Addr-2. Conversely, the other 3745 has TICC inactive at address, Addr-1 and TICD active at address, Addr-2. In the example, NodeA and NodeB are customized to access the host through gateway address Addr-1. However, the Token-Ring Network administrator could split the network traffic across the two addresses.

Recovery from a TIC failure

A TIC failure disrupts the sessions of all SNA Nodes accessing the host through the failing TIC. The backup TIC must be activated on the other communication controller (VTAM command) to provide an alternate route to the host.

Recovery from Token-Ring Route failure

A Token-Ring route failure disrupts sessions with peripheral nodes going through the failing route (with the possible exception of Subarea connections). If alternate routes exist, then the peripheral node is able to re-access the host through the alternate route after session time-out.

Recovery from 3745 Hardware or NCP Failure

With two 3745s, a 3745 or NCP failure may be backed up by activating the TIC on the other 3745. After the backup TIC is activated, peripheral nodes may re-establish connections to the host. Thus, 3745 maintenance and NCP generations may be performed on one 3745 while the other runs the network. Since each 3745 is already connected into the Token-Ring, no manual switching must be done to activate the backup path. A VTAM command activates the backup path.

5.4 Multiple 3745 Rings, Multiple 3745s

Multiple 3745s may be connected into different rings to provide the highest availability. Backup routes are automatically in place to recover from most failures.

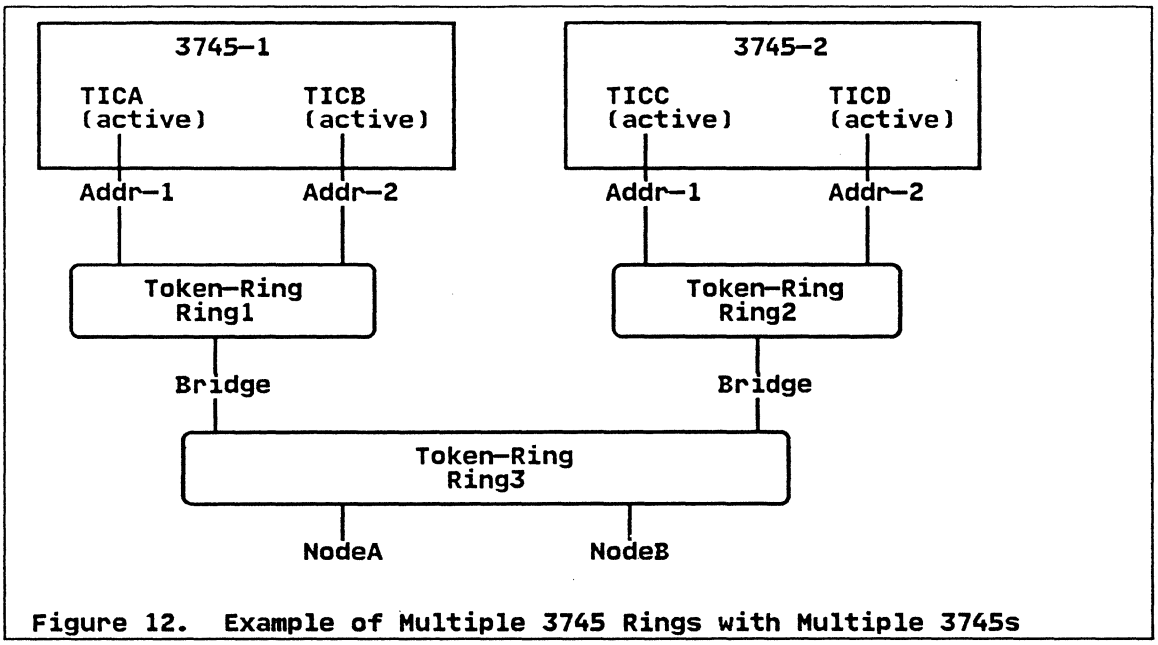

Figure 12. Example of Multiple 3745 Rings with Multiple 3745s

In the figure above, both 3745s have two concurrently active TICs. Both TICA and TICC are active at Token-Ring address Addr-1; TICB and TICD at Token-Ring address Addr-2. Therefore, TICA and TICC concurrently support traffic from SNA nodes accessing the host at gateway address Addr-1; TICB and TICD at address Addr-2.

Recovery from TIC failures

Any single TIC failure disrupts the users on peripheral nodes accessing the host through the failing TIC. However, the user can reestablish the host connection through the other communication controller at the same gateway address. Source routing provides this capability. No VTAM or NCP commands are required for the user to access the backup path.

Recovery from Token-Ring Route Failure

A Token-Ring Route failure is disruptive to peripheral nodes that are accessing the host through the failing route. However, alternate routes are available if ring1, ring2, or a bridge fails. After a failure, a peripheral node may reinitiate host communications across an alternate route (if an alternate route is available). For most failures, backup is automatically in place.

Recovery from 3745 Hardware or NCP Failures

If a 3745 fails, then all sessions going through that 3745 are disrupted. However, peripheral nodes may reestablish host communication through the other 3745 at the duplicate TIC address. Thus, a 3745 or NCP failure is backed up without operator intervention.

Another advantage of this configuration is that maintenance may be performed on one controller while the other controller runs the network. No external switching is required to provide backup during maintenance downtime.

5.5 *Multiple Ring Design Considerations*

Connecting the 3745 into the Token-Ring network requires planning to create a configuration which provides backup without negatively impacting Token-Ring performance. This document is not intended to cover design issues. However, one important point to consider in designing a network is the following:

When duplicate addresses exist in the network, SNA Nodes should not have a Token-Ring route to the host which crosses both rings with the duplicate address.

The key difference in the figure below as compared to Figures 7-10 is that the NCP rings are not bridged together. This diagram shows a configuration where performance problems could occur:

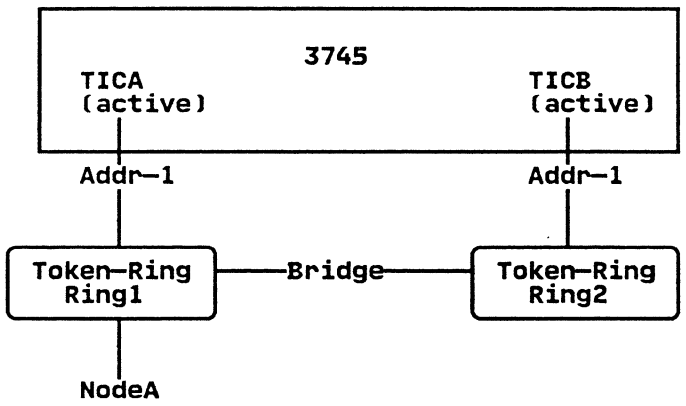

Figure 13. Example of Design with Potential Performance Problems

In the diagram above, NodeA normally accesses TICA (Addr-1) for host applications. If TICA were to fail (or be varied off-line), then NodeA could cross the bridge and access the host through TICB (Addr-1). The backup route may cause performance problems:

1. TICA is varied off-line
2. NodeA establishes communications with the host through TICB
3. TICA is varied back on-line

When NodeA crossed the bridge to access TICB, its route would cross Ring1 and Ring2 which both have the address Addr-1. Therefore, when TICA is brought back on-line, TICA would see its address (the same as TICB's) in the Token-Ring frames. TICA and TICB would both receive the frame, and NCP would process both frames of data. Thus, NCP unnecessarily processes the same data across two different TICs. The frames that are passed to NCP through TICA are discarded, but additional 3745 cycles are used to receive and to discard these frames.

6.0 3745 Communication with Subarea Nodes

Subarea Nodes are Type 4 Nodes (e.g. 3745) and Type 5 Nodes (9370 Integrated Token-Ring Adapter and VTAM Version 3.4 with 3172) on the Token-Ring Network. The 3745 communicates with local and remote subarea nodes on the Token-Ring Network. This section describes the flows and availability features of networking 3745s across a Token-Ring Network.

Previous sections of this paper showed figures with duplicate TIC addressing to provide high availability. **Duplicate TIC addressing is not supported with subarea networking, so the configurations in the previous section are not supported when communicating with subarea nodes.** A figure at the end of this section shows one way to provide high availability for peripheral nodes and subarea networking across a Token-Ring Network.

This section covers the following:

Benefits of 3745 Subarea Networking on a Token-Ring

Subarea Session Establishment

Subarea Non-Disruptive Route Switching

Subarea and Peripheral Node Networking

6.1 Benefits of 3745 Subarea Networking on Token-Ring

The 3745 on the Token-Ring Network provides benefits in both local (within a data center) and remote locations. In traditional configurations, the 3745 controllers are networked together with dedicated leased lines as the transport network. With the Token-Ring Network, the dedicated leased lines are replaced by the Token-Ring Network which supports high speed communications between all devices connected onto a ring. Thus, the Token-Ring Network provides a generic transport for all devices, not just the 3745 controllers.

Data Center Network

Within a single data center some customers have multiple 3745s which must be networked together to provide access to applications. Some reasons for connecting the 3745s together are:

Each 3745 Does Not Have Channel Connections to All Hosts

SNA Peer Networking without Passing through a Host

Many customers do not connect all of their application hosts to every 3745. Each 3745 may have several channel attachments, and other hosts are accessed through connections to other 3745s. Thus, a requirement exists for high speed networking between 3745s to access all hosts.

Another reason for connecting 3745s together is to provide peer access between SNA Type 2.1 Nodes. With Low Entry Networking a Type 2.1 Node may communicate with another Type 2.1 Node without sending the SNA traffic through a VTAM host. If the 3745s are connected together, then Type 2.1 Nodes may communicate across multiple 3745s without going through a VTAM host. The figure below shows a traditional configuration with leased lines and a configuration of 3745s connected with a Token-Ring Network.

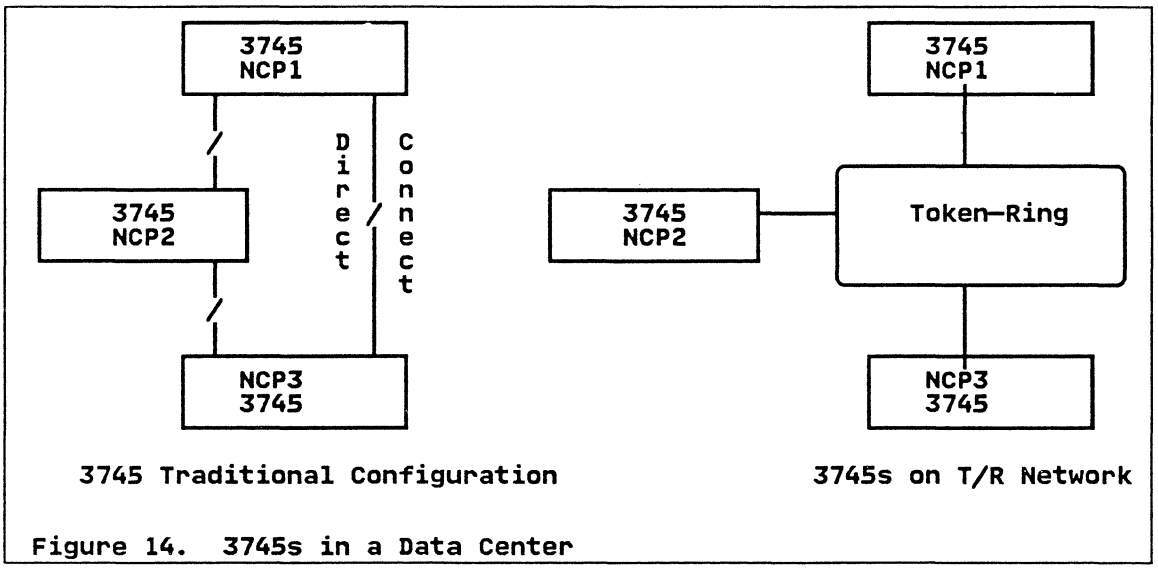

Figure 14. 3745s in a Data Center

An alternative to connecting the 3745s with direct connect lines is using Token-Ring for high speed connectivity. The 3745s are connected across a very high speed transport (4 or 16 Mbps) versus multiple high speed lines (up to 1.536 Mbps). In a traditional configuration with direct connections, each 3745 uses a dedicated communication port to communicate with another 3745. In the diagram above, each 3745 requires 2 communication ports. With a Token-Ring Network, each 3745 communicates with all other 3745s through a single TIC.

Some of the benefits of the 3745s on a local Token-Ring Network are:

Very High Speed Communication Between All 3745s

Single Network Connection vs. Multiple Lines

Less 3745 Hardware and Hardware Upgrades Required

A single 3745 TIC supports NCP-NCP (INN) sessions with all other 3745s connected into a 4 or 16 Mbps Token-Ring Network. Thus, a mesh network of 3745s is possible without running direct connections between each pair of 3745s. The Token-Ring connection reduces the need for separate communication ports on each 3745 controller.

Adding another communication controller to the Token-Ring Network requires no additional controller hardware on the existing controllers. An additional NCP definition is needed to add another 3745 onto the Token-Ring Network.

Remote 3745

In remote Token-Ring Network locations, the 3745 also provides benefits for the SNA users. Customers are installing the remote bridge to connect remote LANs into a backbone LAN at the data center. The key reason for bridging a remote site with a 3745 on a Token-Ring to a local site is:

Multiple Protocol Support

The remote bridge supports peer-peer networking between devices attached to a LAN at any location and is protocol independent for devices supporting source routing. Thus, the remote bridge supports SNA, TCP/IP, NetBIOS and other protocols simultaneously.

The Token-Ring Network provides flow control through dynamic windowing which is implemented in Token-Ring adapter cards (and by NCP). Dynamic windowing supports increasing the number of outstanding frames up to a maximum (MAXOUT parameter) during normal operation and reduces the number when a frame is rejected or a time out occurs. Congested conditions (e.g. bridge becomes congested) reduce the number of outstanding frames to 1 for each Token-Ring device.

Even when the remote bridge connects remote locations, a **remote 3745 still provides many benefits.** The figure below shows a configuration with a remote 3745 on a bridged Token-Ring Network.

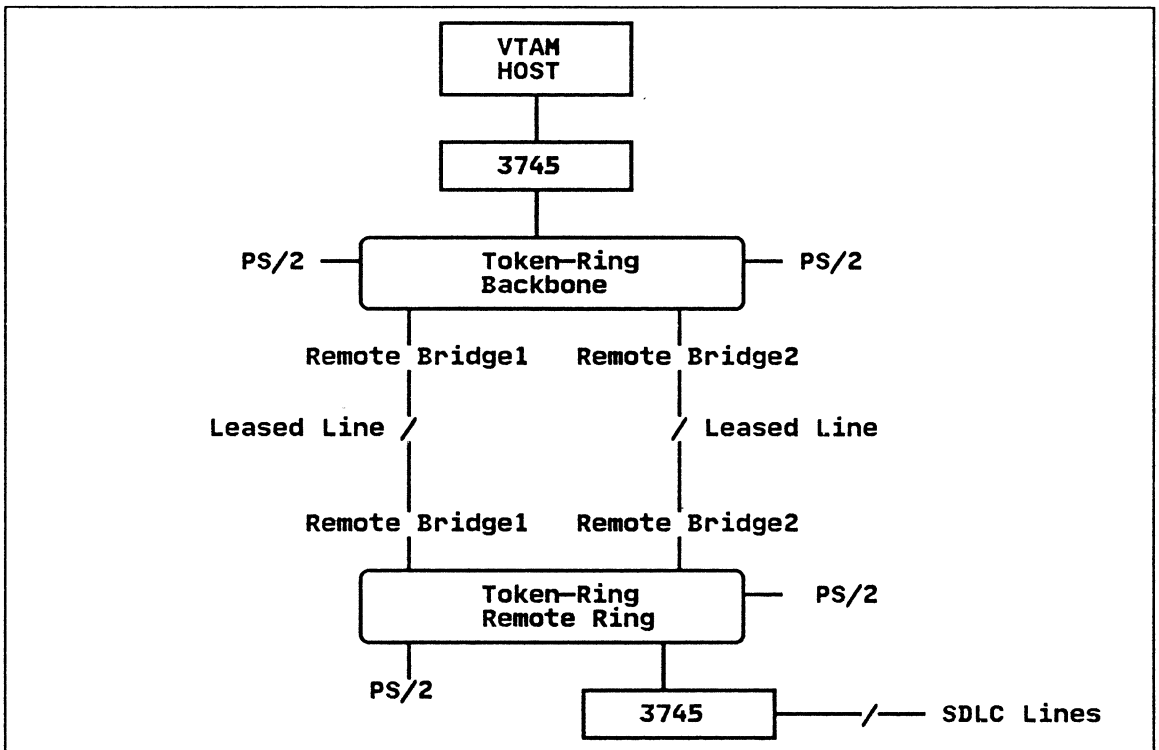

Figure 15. 3745 on a Remote Ring

In the figure above, a remote Token-Ring is bridged to the data center backbone Token-Ring. The Backbone and Remote Ring become a single logical Token-Ring Network. A device on either ring may access peer applications across the network. For example, a PS/2 on the remote ring can access files on an OS/2 LAN Server on the backbone ring or run TCP/IP to communicate with a PS/2 running TCP/IP on the backbone ring. SNA applications are supported through the remote 3745. A PS/2 on the remote ring can run 3270 emulation or APPC and access the host, or peer applications, through the 3745 on the remote ring.

Some of the benefits of the 3745 on the remote ring are:

Remote Concentration of Lines (SDLC,BSC,X.25)

Class of Service for SNA Traffic Across Bridged Network

Routing of SNA Traffic

Non-Disruptive Route Switching

Boundary Processing of SNA Nodes on Remote Rings

First, this configuration provides concentration of lower speed leased lines. Remote concentration may reduce telecommunication line costs and offloads the boundary processing cycles (polling, retries, etc.) from a central site 3745. SNA devices (i.e. 3274 controllers) which cannot natively attach to a Token-Ring may be leased line attached into the remote 3745. Other SDLC, BSC and X.25 (through X.25 SNA Interconnection Program Product) leased lines may also be connected into the remote 3745.

Second, NCP support of VTAM Class of Service prioritizes SNA traffic across the bridged network. Without a remote 3745, the bridge operates in FIFO mode (the first frame received is the first sent). If one workstation was accessing an interactive CICS application and a workstation printer was printing a CICS batch report, both sessions would have equal priority to send traffic across the bridge. Thus, the batch print traffic could slow down the response of the interactive application. With a remote 3745, all SNA traffic flows into the remote 3745 which sends the traffic across the Token-Ring to the peripheral node. Since the remote 3745 receives all SNA traffic, interactive traffic flows before batch (print, file transfer) traffic when using Class of Service (NOTE: SNA Nodes establish session priority during session establishment; a SNA network must be configured to support multiple classes of service for different types of traffic.). However, if non-SNA traffic exists, the 3745 SNA traffic has the same priority as the non-SNA traffic accessing the bridge.

Third, SNA frames are routed directly to host-attached 3745s across the Token-Ring Network. Without the remote 3745, all traffic from a SNA Node on a remote Token-Ring flows to a central site 3745, and the central 3745 must route the traffic to the application host or another 3745. Depending on the network configuration, traffic may be sent back and forth across the network before reaching the destination host. With a remote 3745, multiple INN sessions with other 3745s and hosts are concurrently supported from the remote 3745. SNA traffic is sent directly to a host-attached 3745 or application host on the Token-Ring Network.

Fourth, Token-Ring route switching is non-disruptive. Previously, the failure of a bridge link or bridge disrupted all SNA sessions across the failing route. With NCP V5.3 and later releases, NCP attempts to re-establish the session across an alternate route (i.e. bridge2 instead of bridge1). If an alternate route exists, then the 3745 switches the users' sessions across the alternate route without disrupting the session. (See description "6.3 Subarea Non-Disruptive Route Switching" on page 35) The 3745 tries to switch the route for all SNA Nodes (Type 2.0, 2.1, 4, 5) on a Token-Ring Network. However, only Type 4 nodes support non-disruptive route switching.

Finally, the remote 3745 offloads processing and traffic from the local 3745. When a SNA node on the Token-Ring Network communicates with a 3745, several timers are used to ensure that communications is maintained. The inactivity timer (TI timer) specifies the time that NCP (or the SNA node) waits before sending a poll. The response timer (T1 timer) specifies the time that NCP (or the SNA node) waits for a response after sending a data frame. Finally, the T2 acknowledgement timer (T2 timer) specifies the amount of time to wait before sending an acknowlegement; the T2 timer allows NCP (NCP V5.4 or later) to wait for multiple incoming frames before sending an acknowledgement. Supporting these timers requires 3745 cycles which can be offloaded to the remote 3745. Other processing offloaded from the central site 3745 is the overhead associated with supporting the virtual line definitions for the Token-Ring Nodes. This overhead is described and quantified in TR16PERF PACKAGE which is available on MKTTOOLS. Segmenting is also required when a Peripheral Node cannot support a single large frame of data. The remote 3745 performs segmenting of data to the Peripheral Node and allows the local 3745 to send larger frames of data to the remote ring. The larger frames may improve performance and line utilization across the remote bridge, in addition to offloading local 3745 processing.

Considerations for a remote 3745 on the Token-Ring are:

Limited NCP Load Module Support

Only Single Link TGs Supported

Initial loads of a remote 3745 are only supported from leased or switched SDLC lines or 3745 diskettes. The 3745 cannot be initially loaded across a Token-Ring network, because the 3745 microcode does not support Token-Ring without an active NCP. Loading a remote 3745 across the Token-Ring Network requires an active NCP in the remote 3745. The new load module must be loaded to the 3745 hard disk and then activated. The remote 3745 is designed to automatically load the NCP from hard disk during power up, so all loads after the initial load can be done across the Token-Ring Network.

With SDLC leased lines, parallel lines between two 3745s may be configured as a Multiple Link Transmission Group (MLTG). A MLTG automatically sends data across one or more of the parallel links utilizing the aggregate bandwidth. Failure of a single link is not disruptive (e.g. sessions are not broken) to the users. On Token-Ring Network connections, only Single Link Transmission Groups (SLTG) are supported between 3745s (multiple SLTGs are supported between 3745s, but each SLTG must be between a unique pair of TICs). The 3745 can bypass Token-Ring Network route failures with non-disruptive route switching but has no capability to utilize multiple Token-Ring routes as a MLTG.

Backup for Remote 3745 Controller Across Token-Ring Network

One issue with remote 3745s is backup in the event of a 3745 failure. When a Token-Ring backbone network is used to connect the 3745 controllers, the workstations which normally access the host through the remote controller may access the host through a local controller.

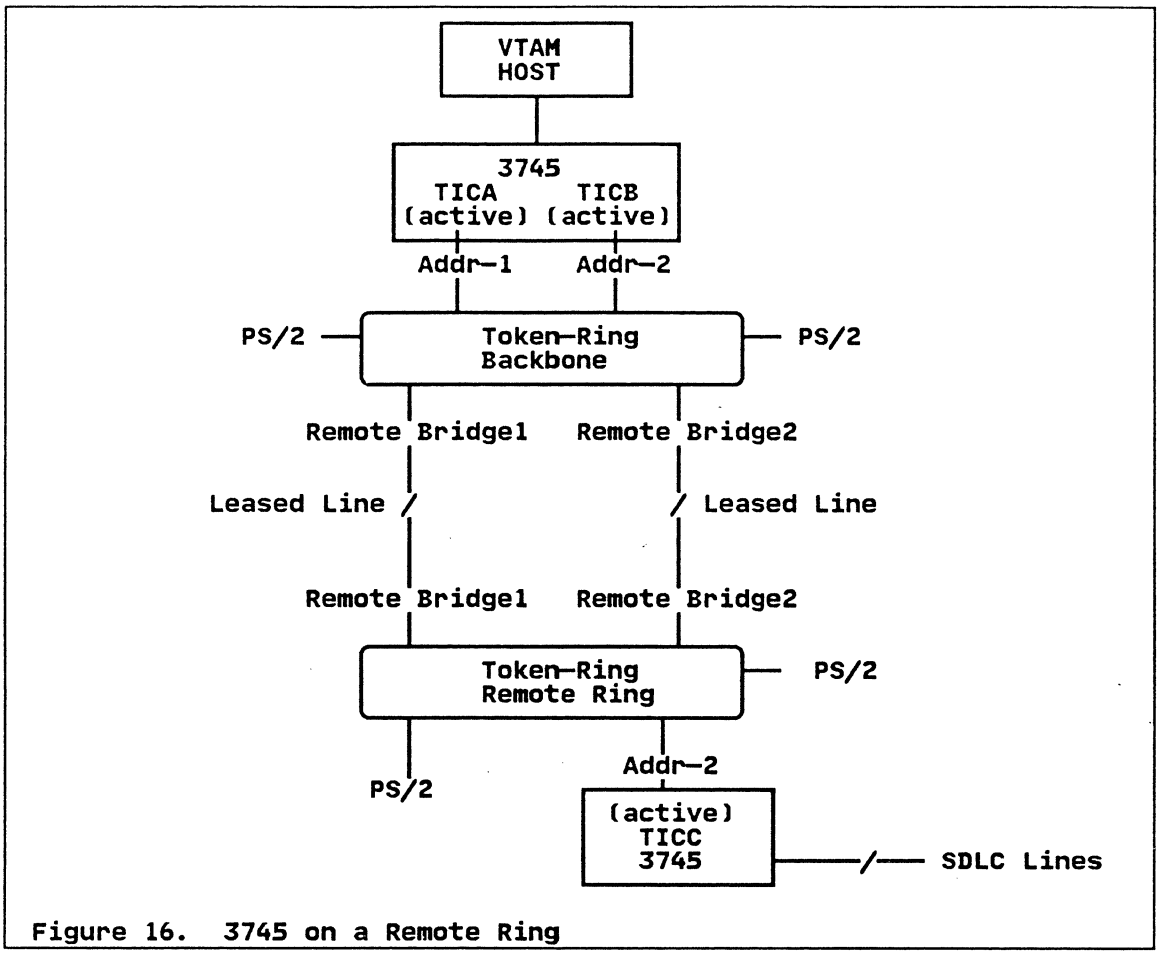

Figure 16. 3745 on a Remote Ring

In the figure above, the remote 3745 provides many of the benefits previously listed for the remote SNA devices. However, when all remote SNA Nodes access the host through the 3745, it becomes a single point of failure. A remote 3745 failure may be backed up by a 3745 at the central site (Only the SNA nodes on the Token-Ring have a backup path. The SDLC lines on the remote 3745 are not backed up). The central site 3745 supports a backup TIC which is defined at the same Token-Ring address as the remote 3745 TIC. Normally, traffic will flow to the 3745 on the remote ring (the SNA Node looks for address Addr-2 on the local ring first). If the remote 3745 fails, then the SNA Node will broadcast the request and find TICB at the central site (at address Addr-2). Thus, Token-Ring SNA Nodes on the remote 3745 are backed up by a central site 3745.

6.2 Subarea Session Establishment Across Token-Ring

Subarea Nodes also use source routing to establish communications across the Token-Ring Network. However, subarea sessions have two logical connections across the Token-Ring Network : a send connection and a receive connection. Each NCP independently establishes a send connection with the other NCP on the token-ring. The diagram below is used to illustrate Subarea networking across a Token-Ring Network.

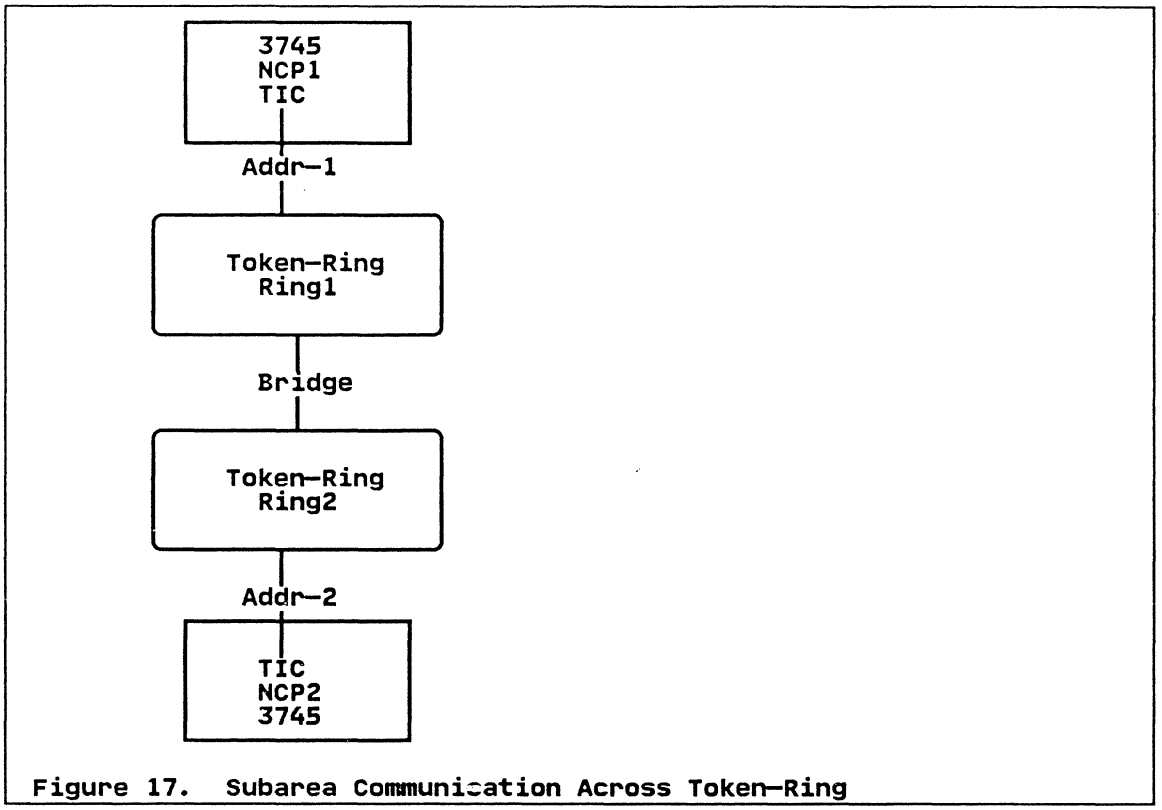

Figure 17. Subarea Communication Across Token-Ring

In the figure above, NCP1 connects into Ring1 at address Addr-1, and NCP2 connects into Ring2 at address Addr-2. Ring1 and Ring2 are bridged together and appear as one logical Token-Ring Network. (NOTE: The 3745s could be on the same ring, or the rings could be bridged with remote bridges).

The Subarea connection between the NCPs is established after the PU definition for the other NCP is activated. The basic flow that occurs when the definition for NCP2 is activated in NCP1 is described below:

1. Activate NCP1's logical line and PU for NCP2.
2. NCP1 issues a token-ring frame (TEST RESOLVE) which is addressed to NCP2's Token-Ring address on the local ring.
3. If not found, NCP1 sends an all-routes broadcast frame (TEST RESOLVE) which flows through all bridges, requesting NCP2's token-ring address (Addr-2).
4. Each bridge forwards the frame and adds the routing information to the broadcast frame. All bridges on ring1 perform this process which repeats itself until all bridged rings are searched or the maximum number (hop count) of bridges have been traversed.
5. NCP2 receives the Test Resolve Broadcast Frame and responds back to NCP1. The response retraces the route that was built by the bridge during the broadcast.
6. NCP1 receives the response and uses that route to send a SNA XID2 frame to NCP2.
7. Traffic from NCP1 to NCP2 flows across this route after the Subarea session is established.

The above flow establishes communication from NCP1 to NCP2. To complete the Subarea connection, NCP2 must establish communication to NCP1. Thus, NCP2 must initiate the same flow in the opposite direction to establish the Subarea connection. This connection is independent from the previous one and is initiated when the definition for NCP1 is activated in NCP2.

1. Activate NCP2's logical line and PU for NCP1.
2. NCP2 issues TEST RESOLVE on local ring looking for NCP1 address.
3. If not found, NCP2 broadcasts TEST RESOLVE frame.
4. Bridges add routing information.
5. NCP1 responds to NCP2 with the path information from the bridges.
6. NCP2 receives the response uses the route to communicate with NCP1. NCP2 sends a SNA XID2 frame to NCP1.
7. Traffic from NCP2 to NCP1 now flows across this path (after the Subarea session is established).

The diagram below shows the flow across the Token-Ring Network:

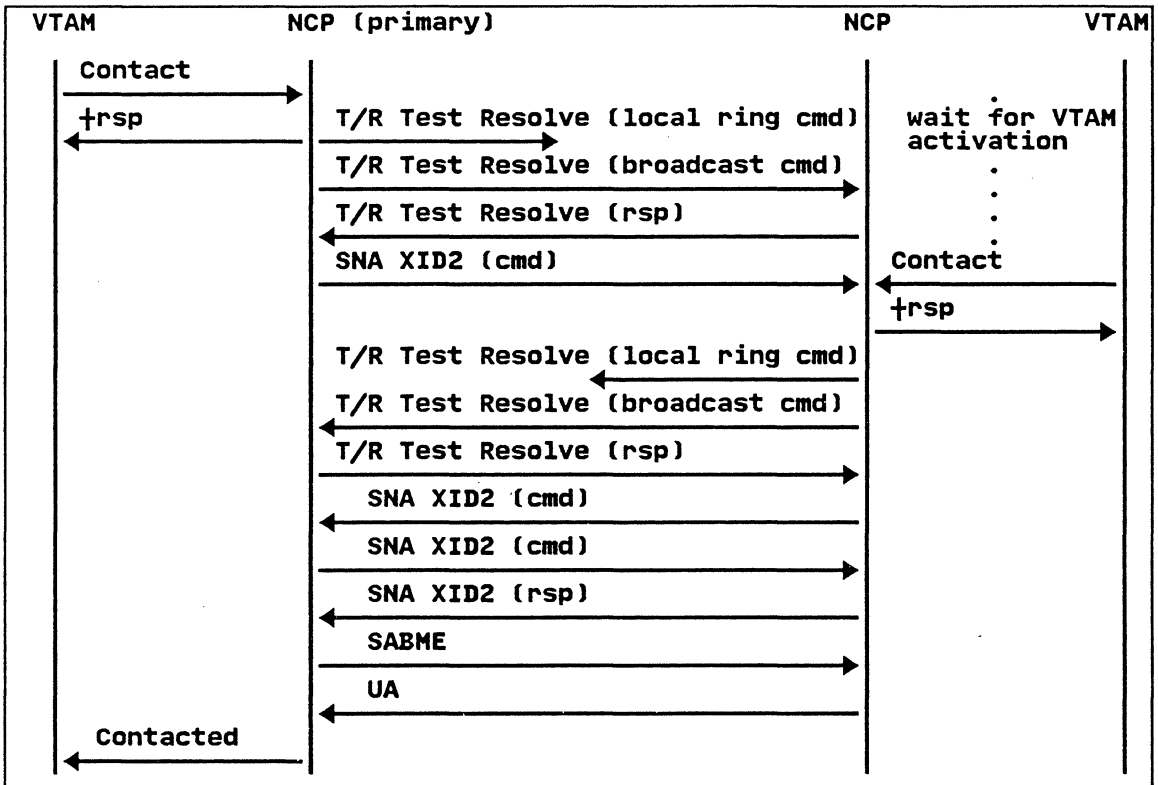

Figure 18. Subarea Session Establishment

After the subarea definition is activated in each NCP, half of the Subarea connection is established independently. Thus, NCP1 establishes a route to NCP2 after its definition is activated, and NCP2 establishes a route to NCP1 after its definition is activated. If multiple routes exist between NCP1 and NCP2 (i.e. parallel bridges), then NCP1 could send across one route and NCP2 could send across a different route (i.e. parallel bridge path). After both NCP definitions are activated, the Subarea connection is established between NCP1 and NCP2.

6.3 Subarea Non-Disruptive Route Switching

Section 4.0 discussed how Token-Ring route failures disrupt peripheral nodes with sessions across the Token-Ring Network. NCP Non-Disruptive Route Switching, a NCP Version 5.3 feature, is a method for dynamically switching NCP routes across the Token-Ring Network without disrupting sessions. The figure below shows a configuration where a bridge failure (either local bridge, remote bridge or line between remote bridges) is bypassed non-disruptively.

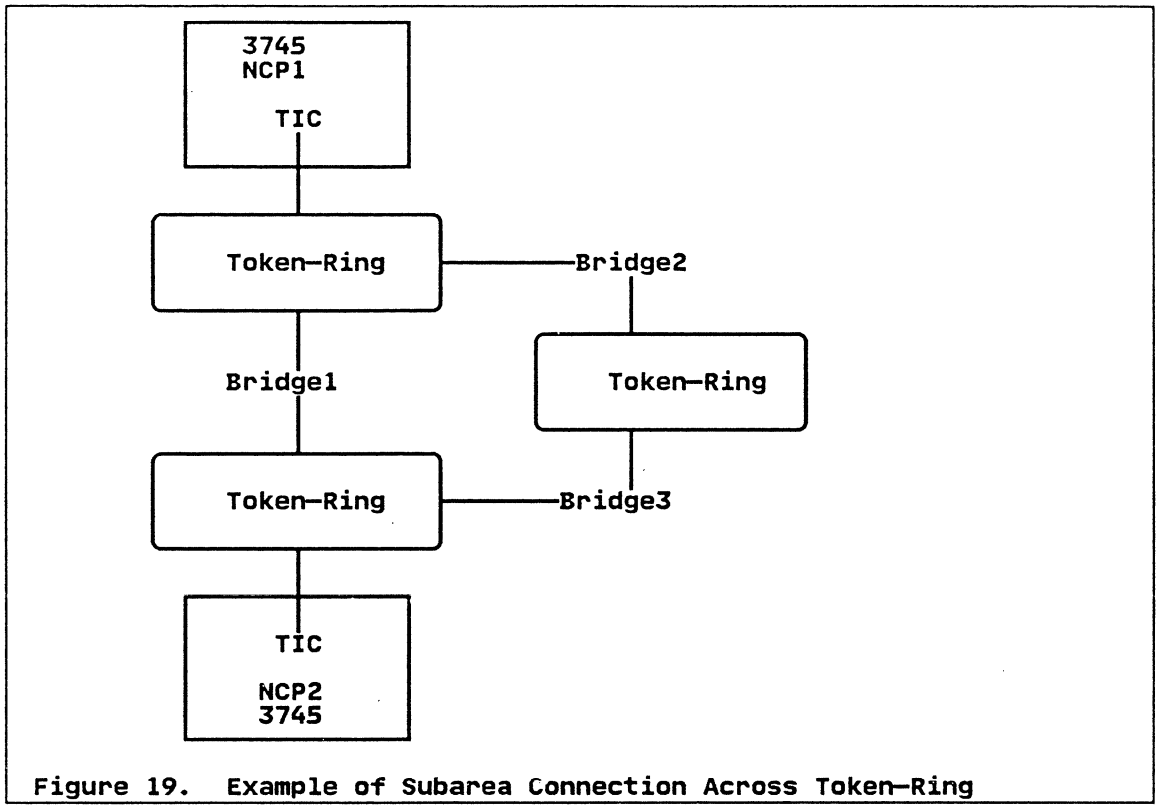

Figure 19. Example of Subarea Connection Across Token-Ring

In the diagram above, the 3745s have a Subarea connection across the Token-Ring Network. Two routes across the Token-Ring Network exist between the 3745s. The first route is across Bridge1; the second route crosses Bridge2 and Bridge3. The route used for the session is the fastest in each direction at the time of session establishment.

When NCP1 and NCP2 establish an INN session, the route for the session is independently established in each direction by NCP1 and NCP2. Assume the following:

1. NCP1 sends to NCP2 across Bridge1
2. NCP2 sends to NCP1 across Bridge3 and Bridge2

This is an arbitrary assumption. The fastest route at session establishment is chosen for the route, and this assumption shows that two different routes may be traversed for a single INN session.

Recovery from TIC Failure

Subarea definitions must be defined to a specific TIC. Duplicate TIC addressing is not supported for Subarea Networking. Therefore, the only way to recover the INN session is to use TIC swapping (see "7.0 TIC Swapping" on page 39). This recovery process is disruptive to the session. A different INN session could be active across another TIC in the same 3745; however, this is a separate single link Transmission Group.

NCP Non-Disruptive Route Switching

Pre-NCP Version 5.3

With NCP Releases prior to Version 5.3, a Token-Ring route failure caused a time-out on the Subarea session. NCP polled the subarea node for a user-specified number of retries. After the retries failed, NCP would break all of the INN sessions. The figure below shows how a pre-NCP V5.3 release supported a failure in the Token-Ring route (e.g. Bridge1 failure).

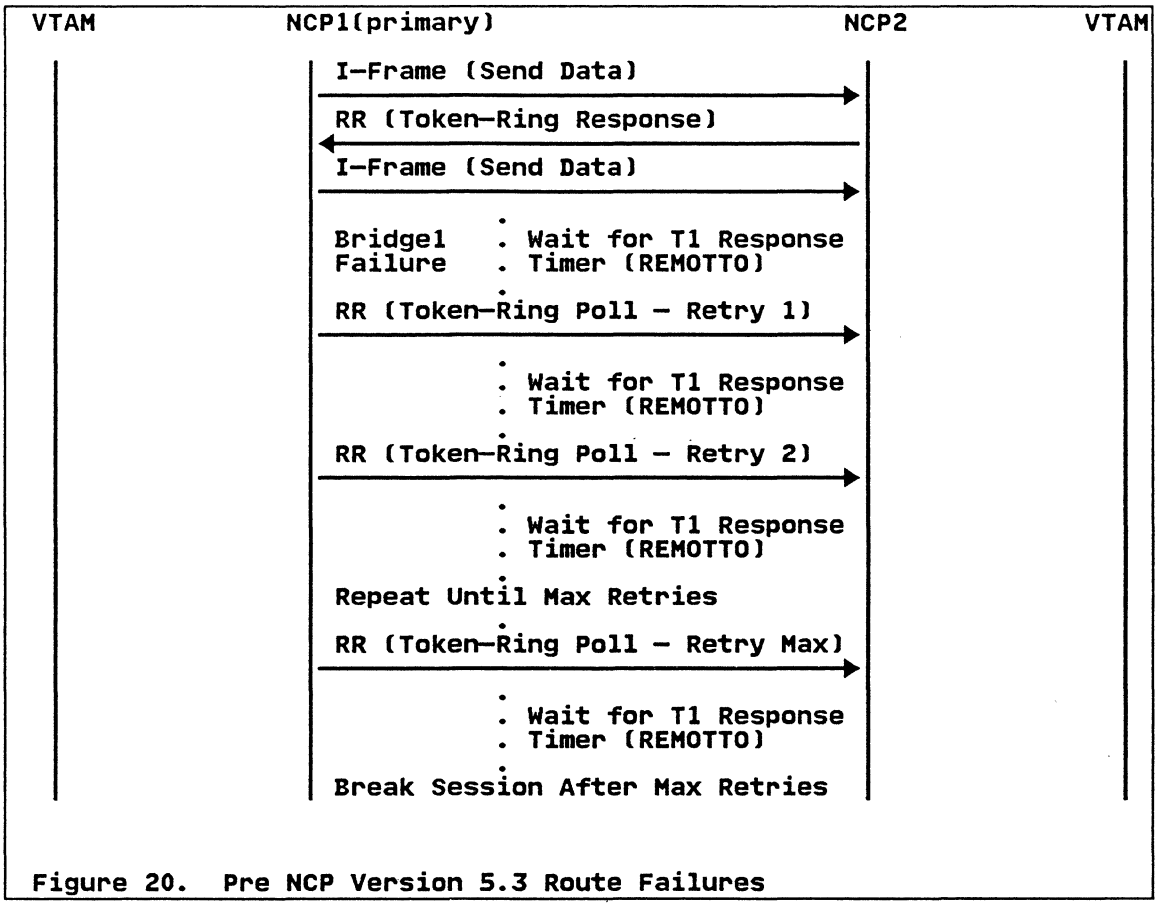

Figure 20. Pre NCP Version 5.3 Route Failures

Each frame across the Token-Ring is acknowledged within the timer specified in the NCP generation (REMOTTO or LOCALTO). If a frame is not acknowledged in time, then NCP polls the remote NCP. After the maximum number of retries is attempted, then NCP breaks the INN session.

NCP Version 5.3 and Later Releases

NCP Version 5.3 performs an additional function before breaking the sessions. NCP Version 5.3 re-issues the TEST RESOLVE token-ring frame which is broadcast through the Token-Ring Network. If an alternate route exists (see Figure 19 on page 35) then the remote NCP receives the TEST RESOLVE and responds back across the route that the TEST RESOLVE was received. If NCP receives the TEST RESOLVE response before its timer expires (REMOTTO parameter), then it updates its tables and begins communicating across the new route.

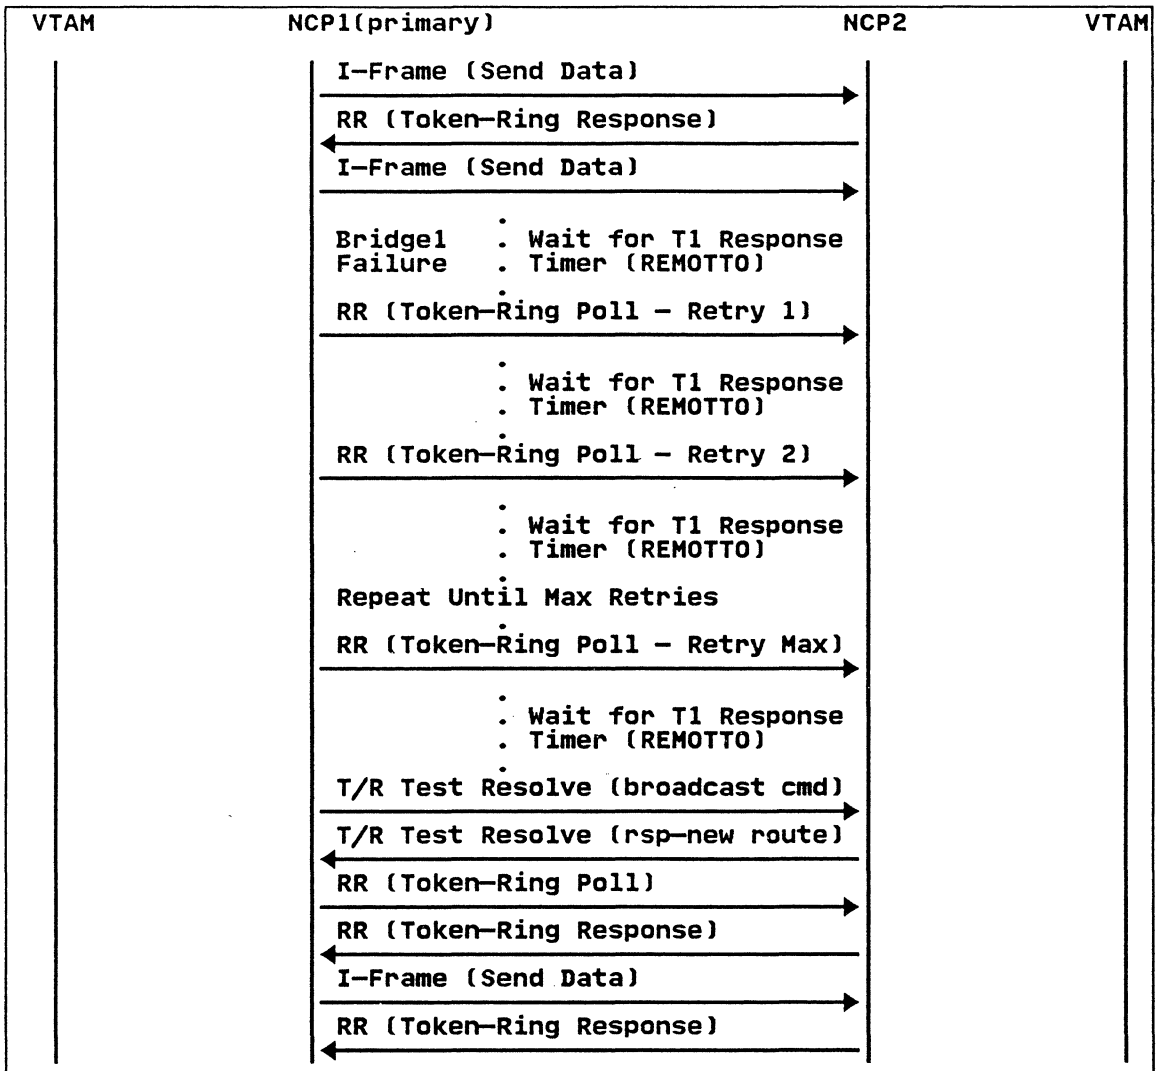

Figure 21. NCP Version 5.3 Route Switching

In the diagram above, NCP1 holds up its sessions while it tries to find a new Token-Ring route by resending the Test Resolve Token-Ring Frame. When it gets a response to the Test-Resolve, NCP1 updates its tables and polls the remote NCP. After receiving a positive response, NCP1 continues sending data.

For the example above, the assumption was made that NCP1 sends across one route and NCP2 sends across a different route. If both NCPs transmitted across the same route, then both NCPs would time-out and issue a Test Resolve frame to discover a new route.

6.4 Subarea and Peripheral Networking

An example network which provides high availability to peripheral nodes and networking between subareas is shown below:

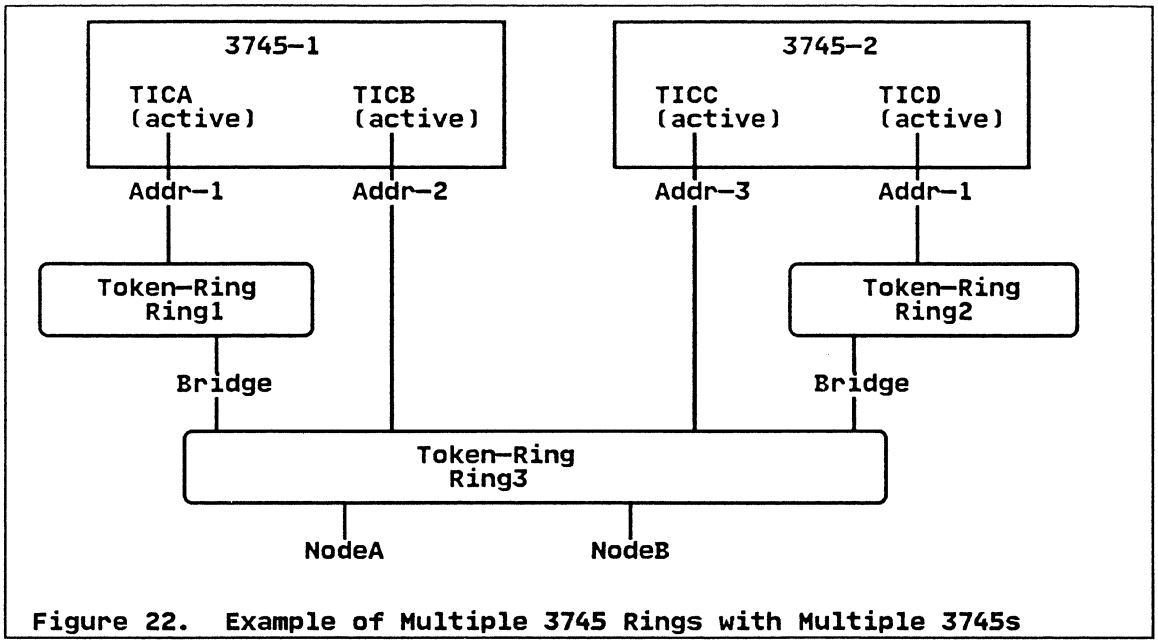

Figure 22. Example of Multiple 3745 Rings with Multiple 3745s

Peripheral Networking

In the figure above, each 3745 has two active TICs: one TIC for Peripheral Node networking and the other for Subarea networking. Peripheral Node traffic is supported through TICA and TICD which are active at Token-Ring address Addr-1. NodeA and NodeB are Peripheral Nodes which are customized to access the host at address Addr-1. The Token-Ring Network provides high availability for the Peripheral Nodes. TICA and TICD concurrently support traffic from the Peripheral Nodes, and each TIC can backup the other in the event of a failure. Therefore, the only single point of failure in the network is the Peripheral Node ring (ring3).

Subarea Networking

TICB and TICC have unique addresses and support the NCP-NCP session between the 3745s. Duplicate TIC addressing is not an issue, because each TIC has a unique address on the Token-Ring Network. TICB and TICC are connected into a single ring (ring3) to reduce traffic through the bridges.

7.0 TIC Swapping

In the configurations discussed in sections 4 and 5, each active TIC was backed up by a separately generated backup TIC which was either active or inactive. These configuration worked well when additional TICs were available and duplicate TIC addressing was supported. However, Subarea connections do not support duplicate TIC addressing. TIC swapping can back up multiple TICs in a 3745 with a single TIC; no NCP changes are required to support TIC swapping.

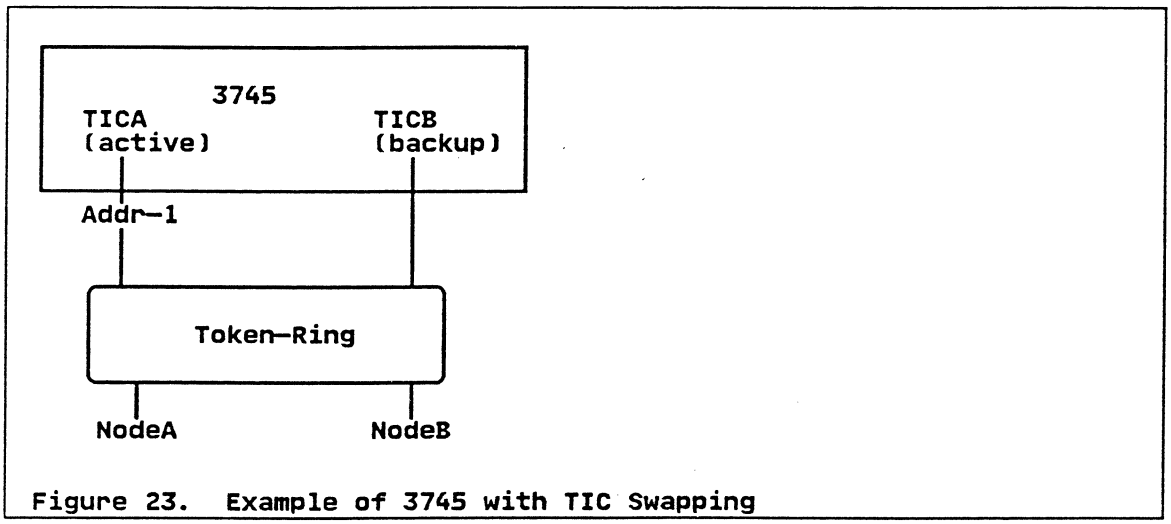

Figure 23. Example of 3745 with TIC Swapping

In the above figure, the 3745 is configured with two TICS - TICA and TICB. TICA is defined at address - Addr-1. TICB is not defined to the NCP. NodeA and NodeB are both customized to access the host gateway at Addr-1. Since TICB is an inactive backup, an operator may use TIC swapping to backup any of the TICs in the same NCP. Thus, TIC swapping provides 1 to many backup for TICs.

Failure Scenarios

Recovery from TIC Failure

A TICA failure will disrupt all Token-Ring users accessing the host through the failing TIC. In the figure above, if NodeA is accessing the host through TICA, then the failure of TICA will cause NodeA to break its connection with the host. All LU-LU sessions from NodeA will be disrupted.

After the failure of TICA, TIC swapping allows TICB to take over for TICA. TIC swapping is a manual process performed at the MOSS console. After TICB is TIC swapped for TICA, NodeA and NodeB can reestablish connections to the host at Addr-1, and Token-Ring users may log back onto their applications.

8.0 Availability Summary and Conclusion

Several high availability configurations are described in this document. The figure below summarizes how a user can recover from specific failures.

	TIC Failure	Token-Ring Route Failure	3745 Failure
Single 3745 Ring, Single 3745, Dup TIC Figure 9 on page 19	Act Backup TIC Restart Node	Restart Node	Wait for Repair
Single 3745 Ring, Multiple 3745s Figure 10 on page 21	Act Backup TIC Restart Node	Restart Node	Act Backup TIC Restart Node
Multiple 3745 Rings, Single 3745, Dup TIC Figure 11 on page 23	Restart Node	Restart Node	Wait for Repair
Multiple 3745 Rings, Multiple 3745s Figure 12 on page 25	Restart Node	Restart Node	Restart Node
Subarea Node Session Figure 19 on page 35	Swap in TIC Reactivate TIC Restart Session	NON-DISRUPTIVE	Wait for Repair

NOTES:
Dup TIC: Duplicate TIC Addressing
Act Backup TIC: VTAM Command to Activate Backup TIC
Restart Node: Restart SNA node (e.g. restart 3270 emulator)
Wait for Repair: No backup available; must wait for hardware repair
Token-Ring Route Failure: Assumes a backup route exists.

Figure 24. Availability Summary

In conclusion, the 3745 provides considerable flexibility on the Token-Ring Network. Network Designers can use the 3745 as part of their high availability Token-Ring Network designs. With proper planning, many 3745 and Network failures can be bypassed - some failures non-disruptively.

9.0 Bibliography

The following IBM publications were referenced:

GG24-3398 IBM Multisegment LAN Design Guidelines

SC30-3448 NCP Resource Definition Reference

GG24-3291 Installation Guidelines for the IBM Token-Ring Network Products

GG24-3469 NCP Version 5 Network Performance and Tuning

MKTTOOLS TR16PERF PACKAGE: Host Attach Token-Ring Performance Information

Heading ID's

id	File	Page	Heading References
highav	3745TR1	18	5.0 3745 High Availability Configurations with Peripheral Nodes 6
subarea	3745TR1	27	6.0 3745 Communication with Subarea Nodes 6, 19, 22
switch	3745TR1	35	6.3 Subarea Non-Disruptive Route Switching 30
tswap	3745TR1	39	7.0 TIC Swapping 35

Figure ID's

id	File	Page	Figure References
fig2	3745TR1	11	4: 14
fig2x	3745TR1	12	5:
fig2y	3745TR1	14	6:
figx	3745TR1	15	7:
fig20x	3745TR1	16	8:
fig9	3745TR1	19	9: 40
fig10	3745TR1	21	10: 40
fig11	3745TR1	23	11: 40
fig12	3745TR1	25	12: 40
fig29y	3745TR1	34	18:
fig19	3745TR1	35	19: 37, 40
fig87	3745TR1	36	20:
fig88	3745TR1	37	21: