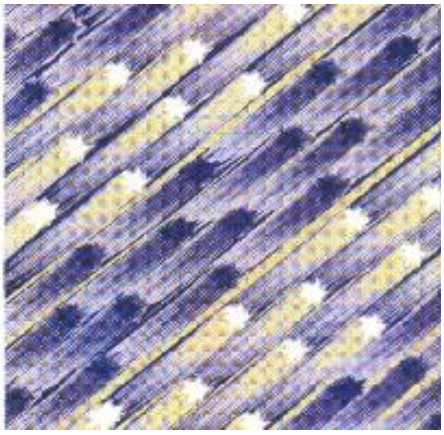

Implementing a Customer Migration from Token Ring to ATM and Ethernet - Phase 1

September 3, 1999

Jeff A. Ottman

IBM Networking Hardware Division
Product Installation Support
Cary, North Carolina

Abstract

Due to changes in the marketplace, many customers are migrating from legacy Token Ring networks to ATM and/or Ethernet networks. These requirements are based on the extensive use of Ethernet and cost considerations. The purpose of this paper is not to discuss the merits of one technology over other technologies, but to report on the implementation of a real customer situation in the process of a network migration. This is a phased migration and this paper documents the first phase from the existing token ring based network to an ATM backbone with Ethernet to the desktop. Existing network services remain on the token ring network and are reached via router access. Phase 2 will be documented in a future paper and show the expansion of the network and migration of network services to the Ethernet network.
(number of pages 26)

Table of Contents

Preface

The customer is in the process of converting a large multi-story building network from an existing token ring to Ethernet network. In the first phase, an ATM backbone will be built and connected to the existing token ring network. End users will be migrated from the existing token ring network to 100mb switched ethernet connections. During the first phase, network servers will remain on the existing token ring network. All migrated users will exist on a single ELAN and will access network servers through routers in the network during the transition period.

Keywords

MSS, ELAN, translational bridging, source-route bridging, DHCP, ATM, 802.1d, spanning tree, PNNI, Fast Ethernet, 10/100

Product List

8274,8265,MSS,8270

For questions related to this document, please send an e-mail to:

Jeffrey A. Ottman
jottman@us.ibm.com
919-301-4448

Chapter 1: Overview

Network Design

The customer's existing network is currently a source-route bridged token ring network. The desire of the customer is to move to all Ethernet to the desktop with an ATM backbone. This current design is a step in the migration of the network from token ring to Ethernet. The impetus for the change is the acquisition of PCs with Ethernet on the motherboard. As these PCs are brought in-house over several years, the end-users will be migrated from token ring to Ethernet. The end-users will be directly connected to 100mb switch ports on the edge switches. The diagram below shows the physical and logical aspects of the network in this migration phase. The focus is to prepare for a future to support high-bandwidth applications to the desktop.

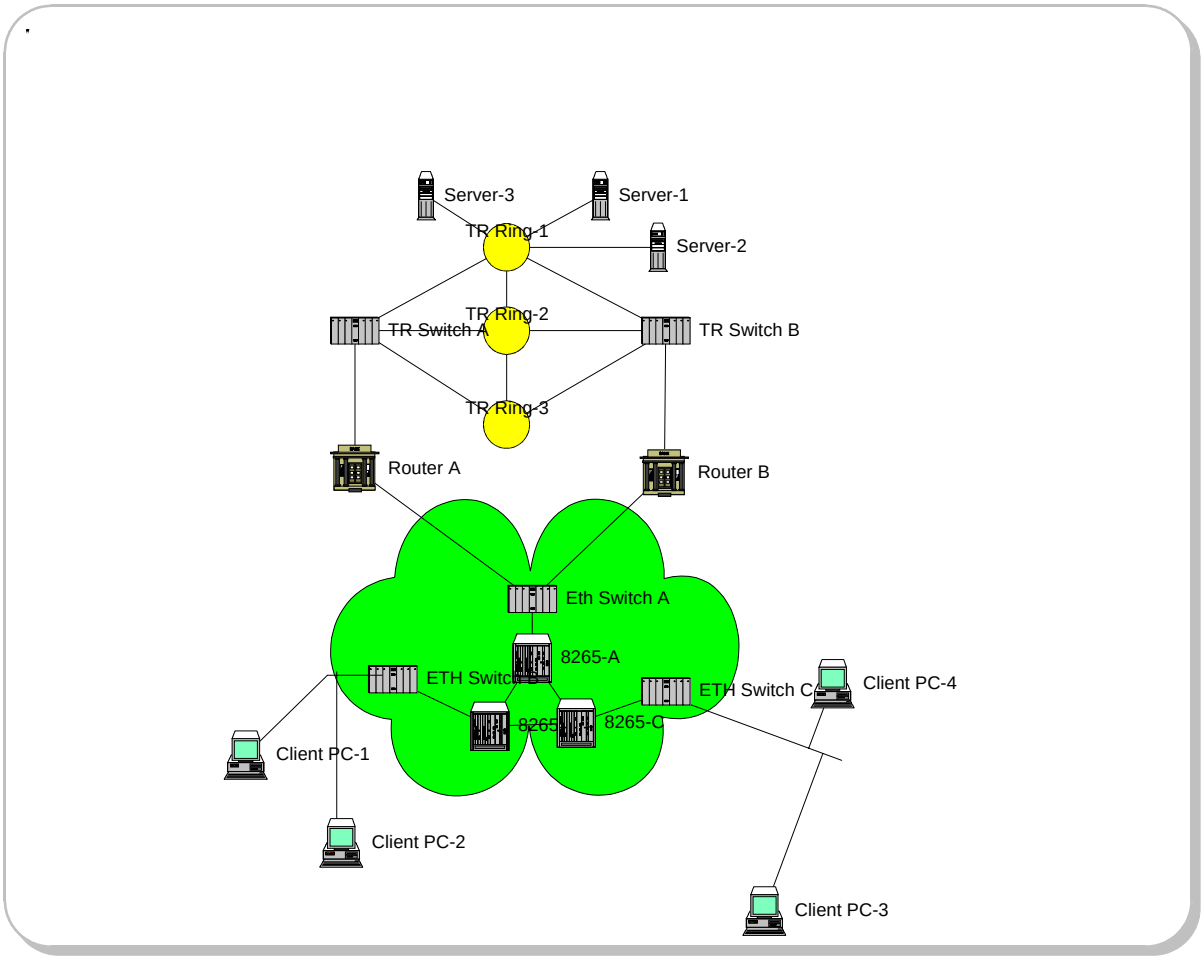

ATM Backbone

The network backbone consists of 3 8265s connected with 155mb PNNI links in a single hierarchy. Two MSSs are installed in the network for LAN Emulation services. One MSS is primarily used for the main network Elan while the other provides a backup LES/BUS and also an ELAN for management services.

Elans

In the current migratory phase of the network, end-users are expected not to exceed 500 before the second phase of the network in installed. Based on broadcast studies done on this network, the recommended number of devices per Elan is 500. This will also be true for the number of users on the same IP subnet. With 3000 devices expected for phase 2, the expectation is that 6 ELANs will be required for this network. Also, a management ELAN has been created on one of the MSSs to handle management traffic from the main ELAN to not expose the 8265 LAN Emulation Clients from processing unnecessary broadcasts. This should always be the case in networks larger than just a few clients.

Each of the above ELANs is backed up by another LES/BUS on a second MSS. Testing was performed to verify performance and network recovery characteristics. The main building ELAN contains multiple subnets on the same ELAN. In this phase of the design, routing is accomplished with the external routers.

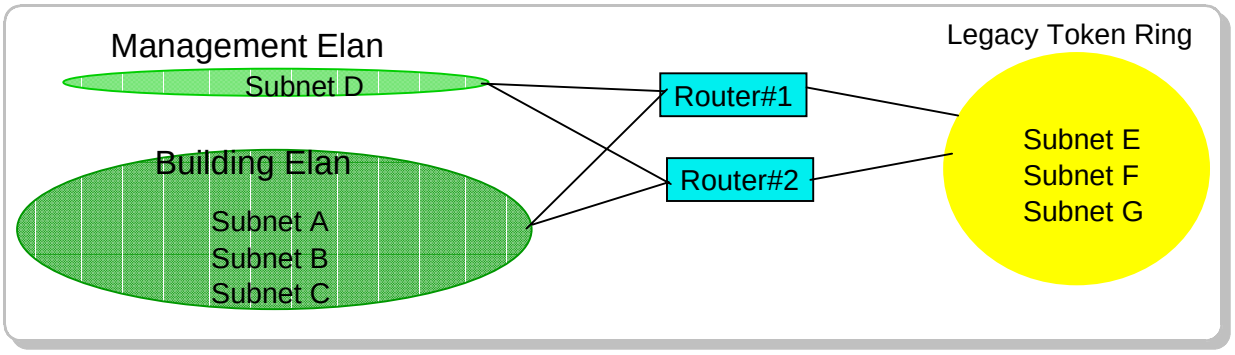

Edge Switches

The 8274 was selected to be the edge switch in the network based on the current technology at the time of this design. This was based on the current and future expected port density for 10/100 connections. Each 8274 is a 9-slot W93 chassis, with one MPM-1G processing module, one 2-port ASM 155mb ATM module, and between 2

and 7 12-port 10/100 ESM Ethernet modules. The intent in the future is to replace the ESM modules with the 32 port ESM Ethernet modules. Currently, both workstations and network servers will be directly attached to 12-port ESM modules. In the future, the desktop stations will be moved to the 32-port ESM modules, while servers will remain on the 12-port modules due to the higher throughput required for each attached station.

Each 8274 has 2 ATM uplinks installed and these are connected at 155mb to two different 8265s. An ATM service is configured on each uplink to join the primary backbone ELAN. Spanning tree will block one of these connections to this ELAN and all traffic will use one uplink. In the event of a failure of the link or the 8265, spanning tree will re-converge and the secondary uplink will handle the traffic.

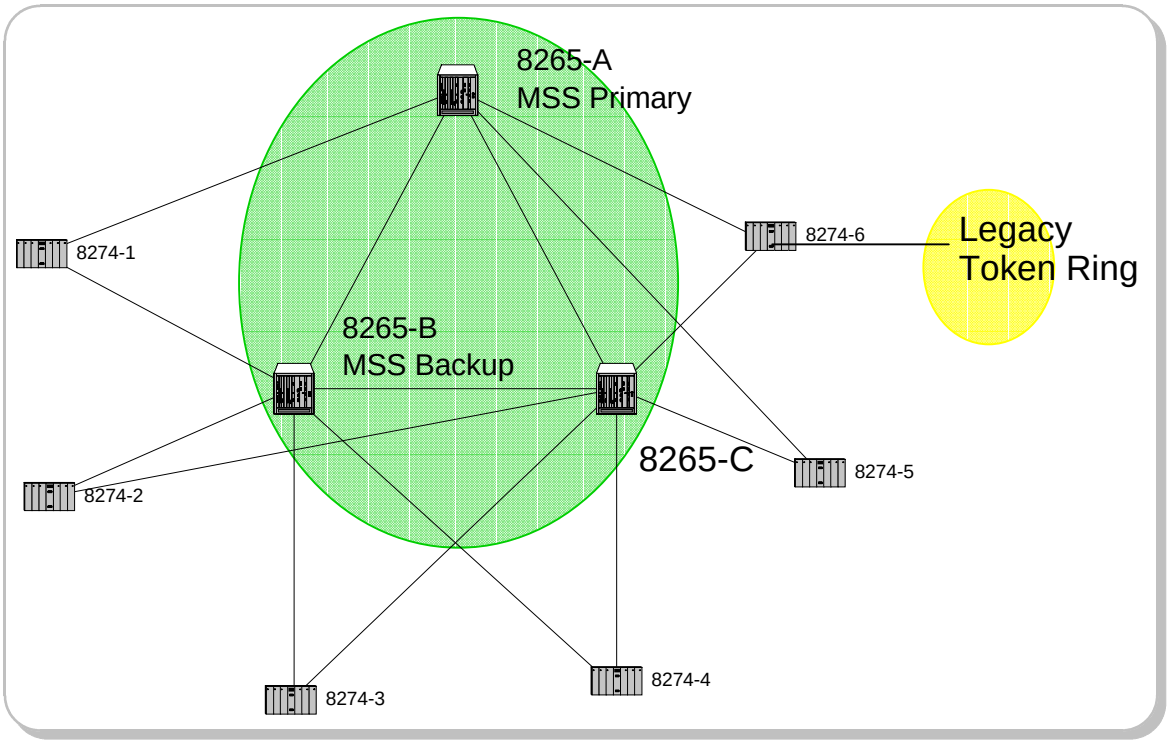

Servers

The currently installed network has over 160 network servers. These are mostly application servers running Lotus Notes and other customer applications. The servers are attached via Token Ring in the backbone network. The plan is to gradually migrate the servers from 16mb token ring to 100mb Ethernet as servers are rebuilt or replaced on the network. In this network scenario, most servers are still attached to the Token

Ring network. Access is accomplished through routers which are attached to both the token ring and ethernet networks. Based on the current number of users sufficient bandwidth is available, but as more users migrate, servers will need to be moved to the Ethernet network where they will be directly attached to 8274-GRSs in the network.

Network Management

The network is managed via an RS/6000 running NetView and Nways Campus Manager for ATM. This device is attached to both the ATM and Ethernet networks to manage the backbone ELAN and Ethernet segments. The ATM connection is configured to join the management ELAN (MGT) for management of the 8265s directly connected to this ELAN.

Chapter 2: Detailed Design and Implementation

Installation Experience

Equipment for this account was staged in over several months. The ATM backbone was deployed first and configurations completed on the 8265 and MSS. The network was installed in the building and the network was validated with a pilot consisting of 50 users.

Code Levels

Product	Function	Code Level
8265	Operational Boot	4.1.3 4.1.3
8274	All	3.4.4.6
MSS	Image Firmware	1.1.8 3.1
8270	Main Image Boot Code Microcode	4.03 4.03 4.03

8265 ATM Backbone Configuration

PNNI Configuration

The 8265s are installed on a single floor of the building in a triangle fashion.
Each 8265 has 2 155mb ports designated as PNNI to attach to the other hubs. In this and most networks, each PNNI port should be a different blade in the hub. Each hub is configured with an ATM address containing a unique hub number

Hub A - 39.99.99.99.99.99.99.00.00.99.99.01.01.40.82.65.01.01.01.00
Hub B - 39.99.99.99.99.99.99.00.00.99.99.01.02.40.82.65.02.02.02.00
Hub C - 39.99.99.99.99.99.99.00.00.99.99.01.03.40.82.65.03.03.03.00

Each hub contains the same Peer Group Identifier and this defines the switch assigned to the same peer group. After the hubs are connected, each will be a member of the same peer group and all nodes will be reachable within the network.

Peer Group ID : 39.99.99.99.99.99.99.00.00.99.99.01
Level Identifier : 96

These are default values. Further discussion of this is detailed in the *Redbook: IBM 8265 Nways ATM Campus Switch.*

Redundancy in the network is achieved by the dual paths available between the 8265s. Each 8265 can be configured with a path selection option. The option selected for this network is **shortest path.** This is used to limit the amount of connections on all switches though this does not do load balancing between the links. Since the switches are in a triangle setup this will eliminate traffic from making an extra hop in the network unless there is no available bandwidth.

In the event of one link or switch failure traffic will reroute through the other links. Disruption is dependent upon the type of traffic which is impacted by the failure. Virtual circuits connected to the LES/BUS may be disrupted requiring re-registration and some disruption to the traffic. This will cause impact to end devices in the network until the registration process is complete. For example, users currently in SNA sessions will be logged off and require them to login again to restart their session.

Port Configurations

The primary connections to the 8265 in this network are other 8265s, 8274s, and RISC/6000 workstations. A special note in this configuration is that with the 8274 code levels involved required that the 8265 ports connecting to the 8274 ASMs must be configured with UNI 3.1 and should not use the AUTO configuration option. Any future installations should inquire if this is still required with future code versions, but this is still the case as of the authored date of this document.

Set port *slot.port* enable uni signalling_verson:sign_3_1

LEC Configuration/Management ELAN

For management purposes each 8265 has an internal LEC joined to a unique management ELAN. In any network but the smallest, these 8265 LECs should not exist on the production ELAN. This would force the CPSW to process all broadcast packets existing on this LAN, but it is not necessary for the CPSW to see all broadcasts. By creating a separate management ELAN, the main broadcast traffic on the network does not need to be processed by the CPSW LECs. The LEC is configured with the following command:

Set device lan_emulation_client eth ip_address:*x.x.x.x*
subnet_mask:255.255.255.0 emulated_lan_name: MGT

MSS Configuration and LES/BUS Implementation

Two MSSs exist in the customer network. MSS#1 contains the primary LES/BUS (BUILDING) in the network and a backup LES/BUS for the management ELAN (MGT). MSS #2 contains the primary LES/BUS for the management ELAN (MGT) and the backup ELAN for the BUILDING.

The MSSs configurations are very basic for this network as there are no special characteristics. *Appendix A* has sample configuration screens from MSSs installed in the network.

8274 Configuration

8274s are installed in each floor of the building and are used to converge end-users to the network. Each 8274 contains a 2-port ASM module for 155mb ATM traffic and between 2 and 7 10/100 Ethernet modules depending on the floor location. Desktop PCs and other network devices are directly attached to the 10/100 switched ports.

ATM Configuration

Each 8274 has 2 ports connected to two different 8265s. Each port has one ATM service configured to join the primary (BUILDING) Ethernet ELAN. The ATM services along with all ports in the switch are configured to join the same group. Each ATM service is configured for 802.3 and to join the primary ELAN.

Spanning Tree

Since this network has redundant paths, it is important that 802.1d spanning tree is configured correctly to only allow traffic to pass through one path. One 8274 switch in the network is selected to be the root bridge. This is configured by assigning the lowest spanning tree priority in the network. In this network, one switch has been assigned a priority value of 100, a second switch is assigned a value of 200, and all others remain with the default of 32767. The *stc* command is issued to set the spanning tree characteristics and priorities. Once these are set, spanning tree will determine the best paths and handle re-converging the network in the event of a link failure on a switch or in the network backbone.

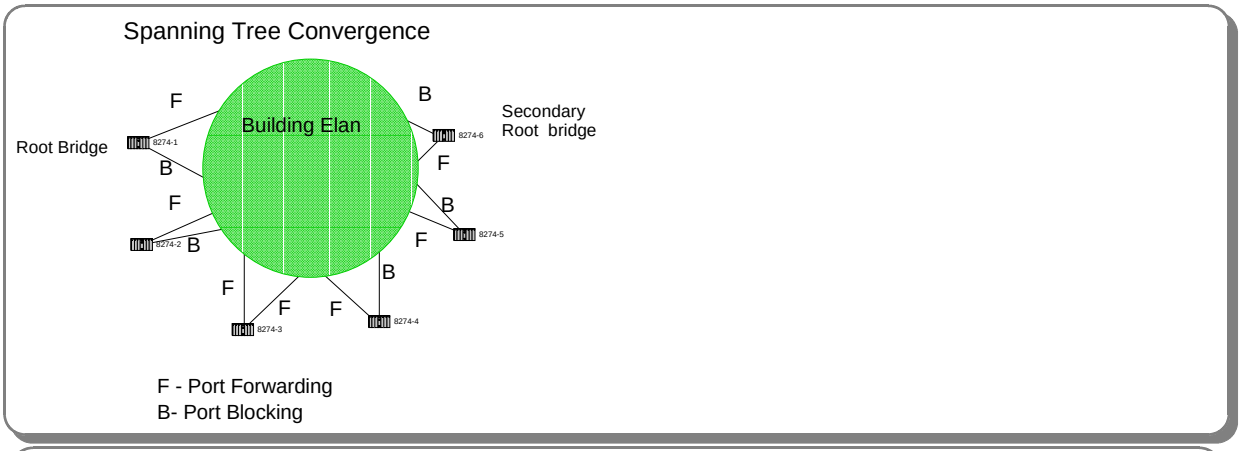


```
stc
Spanning Tree Parameters for Group 1 (Default GROUP (#1))

Spanning Tree is ON for this Group, set to OFF ?      (y/n) : y
New Priority (0..65535)    (current value is 32768[0x8000]) : 100
New Bridge Hello Time (1..10 secs)    (current value is  2) :  2
New Bridge Max Age (6..40 secs)       (current value is 20) :  20
New Bridge Forward Delay (4..30 secs) (current value is 15) :  15
Ageing Time (10..1000000 sec)    (current value is     300) :  300
Auto-Tracker VLAN Ageing Time (10..1000000 sec)(current value is 1200)
:1200

Save the new Spanning Tree Bridge parameters ? y/n : y
```

Legacy Connectivity

The connections between the existing Token Ring network and the new ATM backbone/Ethernet network is accomplished through routed traffic. Network routers are

connected to the 8274s on one floor of the building. These connections are assigned to the same group as the other floor ports. IP traffic not destined for other devices on the ethernet is routed to the legacy token ring network where all current network servers exist. The servers are connected on rings between the 8270s. In this network conversion, there were no special changes required on the 8270 or to the downstream token ring network.

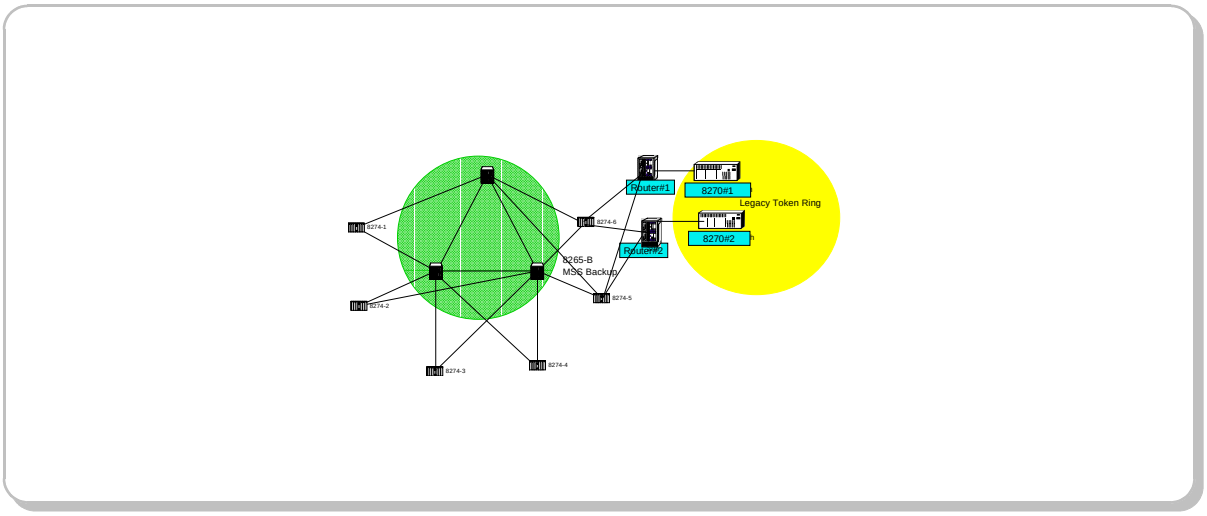

A goal of this network infastructure is to eliminate native SNA and Netbios traffic and limit the network to only IP traffic. Since this is a transistional network, support for these traffic types is included. The routers have also been configured to allow the native SNA and Netbios traffic to reach token ring attach servers in the legacy network. Source-Route Translational bridging (SR-TB) is configured on the routers to handle this traffic.

Special Characteristics and Considerations

This customer has requirements for a highly redundant network with backup capabilities in place for any switch or link failure. The goal is to have not more than a 1 to 2 minute disruption in the event of any failure on the backbone or of an edge switch. To accomplish this goal, this was integrated into the design and special testing was performed by IBM to verify recovery times in a controlled test network.

Testing was employed be creating traffic on the test network using various PC servers and workstations along with adding traffic using the Netcom Smartbits traffic generation tool. These tests and results are documented in the 'Test Procedures' section.

Test Procedures

The following test were performed on the test network by various methods. The first column describes the type of failure, the second is the time before full recovery of IP traffic was seen. Recovery times are measured by examining IP or SNA traffic recovery times on a client PC. These times reflect statistics from the test network and can vary based on the design of the network and applications. These are only presented to use as comparisons and guidelines in any future network design. Test were performed multiple times and all results are the average of the recovery times.

Test are done by either powering off a device or removing a link. The timings are based on measuring the ping response and recovery time from the end-station.

ATM Backbone

Test	IP Recovery
8265-A power-off	34 seconds
8265- A recovery	36 seconds
8265-B power-off	35 seconds
8265-B recovery	20 seconds
8265-C power-off	35 seconds
8265-C recovery	20 seconds
PNNI Link A drop	No traffic loss
PNNI LinkA recover	No effect
PNNI Link B drop	25 seconds
PNNI Link B recover	No effect
PNNI Link C drop	35 seconds
PNNI Link C recover	No effect
MSS Primary down	32 seconds
MSS Primary recover	34 seconds

Edge Devices

Test	IP Recovery
8274-1 Primary link drop	30 seconds
8274-1 Primary link recovery	22 seconds
8274-1 Secondary link down	No effect
8274-4 Primary link drop	20 seconds
8274-4 Primary link recovery	24 seconds
8274 Root Bridge down	18 seconds
8274 Root bridge recovery	20 seconds

Appendix A: Detailed Configuration Information

This section has actual screen captures taken from the test network. These have some information taken from representative devices in the network. This information is unique to this network, but some configuration information may be helpful to match with other installations.

8265 Screen Captures

Display current status of all ATM ports on an 8265

```
8265_Z> show port all

     Type Mode     Status
-------------------------------------------------------------------------------
 1.01:PNNI enabled  UP
 1.02: UNI enabled  UP
 1.03: UNI enabled  no activity
 1.04: UNI enabled  UP
     Type Mode     Status
-------------------------------------------------------------------------------
 3.01:PNNI enabled  UP
 3.02: UNI enabled  UP
 3.03: UNI enabled  no activity
 3.04: UNI enabled  no activity
    Type  Mode     Status                    Daughter Card Description
-------------------------------------------------------------------------------
 5.01:UNI  enabled  UP                         ATM MSS Server
     Type Mode     Status
-------------------------------------------------------------------------------
15.01: UNI enabled  UP
15.02: UNI enabled  no activity
15.03: UNI enabled  no activity
```

Display current 8265 configuration

8265_Z> **show device**
Boot EEPROM version: v.4.1.2
Flash EEPROM version: v.4.1.2 (PNNI)
Flash EEPROM backup version: w.3.3.4 (PNNI)
> Subnet ethernet: Up
IP address: 10.1.1.1. Subnet mask: FF.FF.FF.00
MAC Address: 0006291F8444 (BIA)
> Subnet atm:
IP address: 0.0.0.0. Subnet mask: 00.00.00.00
> Subnet lan emulation ethernet/DIX UP
Config ELAN Name :"MGT"
Actual ELAN Name :"MGT"
MAC Address: 408265010101
IP address : 148.161.32.67. Subnet mask: FF.FF.00.00
ATM address :39.99.99.99.99.99.99.00.00.99.99.01.01.40.82.65.01.01.01.00
Config LES addr:none
Actual LES addr:39.99.99.99.99.99.99.00.00.99.99.01.01.40.00.82.10.A1.FF.02
BUS ATM address:39.99.99.99.99.99.99.00.00.99.99.01.01.40.00.82.10.A1.FF.02
Config LECS add:none
Actual LECS add:47.00.79.00.00.00.00.00.00.00.00.00.00.A0.3E.00.00.01.00
LEC Identifier: 4. Maximum Transmission Unit: 1492
> Subnet lan emulation token ring
Not Started
Config ELAN Name :""
Actual ELAN Name :""
MAC Address: 000629770444
IP address : 0.0.0.0. Subnet mask: 00.00.00.00
ATM address :39.99.99.99.99.99.99.00.00.99.99.01.01.40.82.65.01.01.01.01
Config LES addr:none
Actual LES addr:00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00
BUS ATM address:00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00
Config LECS add:none
Actual LECS add:00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00
LEC Identifier: 0. Maximum Transmission Unit: 0
Default Gateway : OK
IP address: 148.161.32.1
ARP Server:
ATM address: 00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00.00
Device configured for PNNI port capability.
Device configured for Lan Emulation Servers.
Dynamic RAM size is 32 MB. Migration: off. Diagnostics: enabled.
Device defined as primary.
Memory profile: Mixed (32_P_M
)Duplicate ATM addresses are allowed
.Accounting is enabled.

8265 PNNI stastics

```
8265_Z> show pnni hierarchy
LEVEL  NODE NUM AT   PEER-GROUP-ID
ID     THIS SWITCH   THEN NODE-ID
96     0             60.399999999999990000999901
                  60.a0.399999999999999000099990101.408265010101.00
8265_Z> show pnni neighbor

NEIGHBORS OF NODE 0 CONTAINED IN PEER GROUP IDENTIFIED BY LEVEL ID 96:

 neighbor_1    node id:    60.a0.399999999999999000099990103.408265030303.00
 neighbor atm addr         39.999999999999000099990103.408265030303.00
              state:     Full
       slot.phys-port:     03.01

 neighbor_2    node id:    60.a0.399999999999999000099990102.408265020202.00
 neighbor atm addr         39.999999999999000099990102.408265020202.00
              state:     Full
       slot.phys-port:     01.01
8265_Z> show pnni crankback
Crankback:                    OFF
8265_Z> show pnni peer_group_members
VERTX  THIS   UP  CONN   IS  NODE_ID
NUMBR SWTCH LINK  CTED  PGL  THEN ATM_ADDR
   0   yes   no   yes   no  60.a0.399999999999999000099990101.408265010101.00
                        39.999999999999000099990101.408265010101.00
  11   no    no   yes   no  60.a0.399999999999999000099990103.408265030303.00
                        39.999999999999000099990103.408265030303.00
  10   no    no   yes   no  60.a0.399999999999999000099990102.408265020202.00
                        39.999999999999000099990102.408265020202.00
Peer group contains 3 member(s)
```

8274 Screen Captures

8274 ATM port configuration

```
Enabling service...
Switch1>mas 5/1 1
  Slot 5 Port 1 Service 1 Configuration
1)  Description (30 chars max)        : LAN Emulation Service 1
2)  LAN Emulated Group                : 1
   21) LAN type { 802.3 (1),
                  802.5 (2) }         : 802.3
   22) Change LANE Cfg { NO (1),
                         YES (2) }    : NO
3)  LECS Address (40-char-hex)        :
47007900000000000000000000000A03E00000100
4)  Admin Status { disable(1),
                   enable(2) }        : Enable
6)  Connection Type { PVC(1),
                      SVC(2) }        : SVC
   60) SEL for the ATM address        : 01
7)  BandWidth Group (1-8)             : 1
Enter (option=value/save/cancel) : 22=2
 Slot 5 Port 1 Service 1 LANE Configuration Parameters
1)  Proxy { NO (1), YES (2) }                 : YES
2)  Max Frame Size { 1516 (1), 4544 (2)
                     9234 (3), 18190 (4)    } : 1516
3)  Use translation options{NO (1), YES (2)   : Yes (use Swch menu to
set)
4)  Use Fwd Delay time { NO (1), YES (2) }    : NO
5)  Use LE Cfg Server (LECS){ NO (1), YES (2)}: YES
6)  Use Default LECS address { NO(1), YES (2)}: YES
7)  Control Time-out (in seconds)             : 10
8)  Max Unknown Frame Count                   : 10
9)  Max Unknown Frame Time (in seconds)       : 1
10) VCC Time-out Period (in minutes)          : 20
11) Max Retry Count                           : 2
12) Aging Time (in seconds)                   : 300
13) Expectd LE_ARP Resp Time (in seconds)     : 1
14) Flush Time-out (in seconds)               : 4
15) Path Switching Delay (in seconds)         : 6
16) ELAN name  (32 chars max)                 : BUILDING
```

8274 Spanning Tree Configuration Information for Root Bridge

```
Switch1>sts
  Spanning Tree Parameters for Group 1 (Default GROUP (#1))
Spanning Tree Status :                ON
Bridge Protocol Used :        IEEE 802.1D
Priority             :               100 (0x0064)
Bridge ID            : 0064-0020DA:A18110
Designated Root      : 0064-0020DA:A18110
Cost to Root Bridge  :                 0
Root Port            :              None
Hold Time            :                 1
Topology Changes     :                12
Last Topology Change : 3 minutes, 8 seconds ago
Bridge Aging Timer   :               300

                                  Parameters system uses when
  Current Parameters              attempting to become root
-------------------------------   -------------------------------
Max Age                  20 secs   System Max Age          20 secs
Forward Delay            15 secs   System Forward Delay    15 secs
Hello Time                2 secs   System Hello Time        2 secs

Switch1>stc
  Spanning Tree Parameters for Group 1 (Default GROUP (#1))

Spanning Tree is ON for this Group, set to OFF ?      (y/n) : n
IEEE spanning Tree for this Group, set to IBM ?       (y/n) : n
New Priority (0..65535)    (current value is   100[0x0064]) :
New Bridge Hello Time (1..10 secs)    (current value is  2) :
New Bridge Max Age (6..40 secs)       (current value is 20) :
New Bridge Forward Delay (4..30 secs) (current value is 15) :
Ageing Time (10..1000000 sec)    (current value is     300) :
Auto-Tracker VLAN Ageing Time (10..1000000 sec)(current value is
1200) :
```

MSS Screen Captures

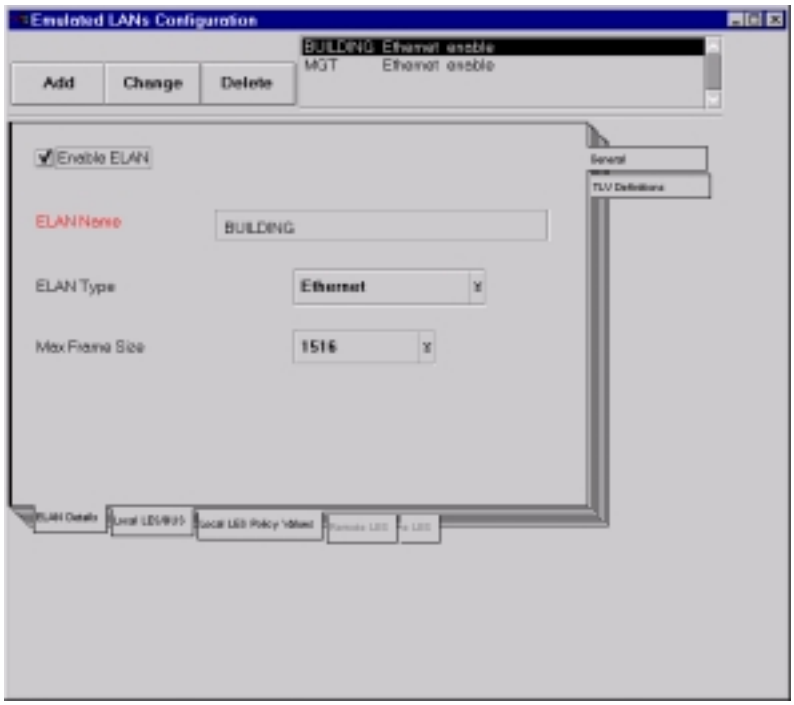

This screen defines the ELANs used in the network. The two listed ELANs are for the BUILDING elan and the MGT elan. These are both defined as Ethernet and are enabled on the MSS Sever.

Each ELAN is assigned a unique MAC address and selector byte.

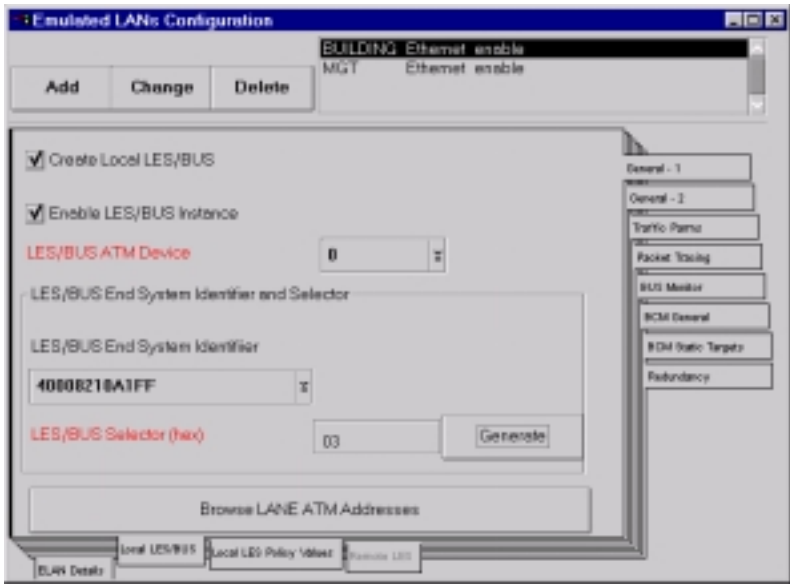

This screen indicates this is the definition for the ‘primary’ LES/BUS.
A backup is configured on the other MSS in the network.

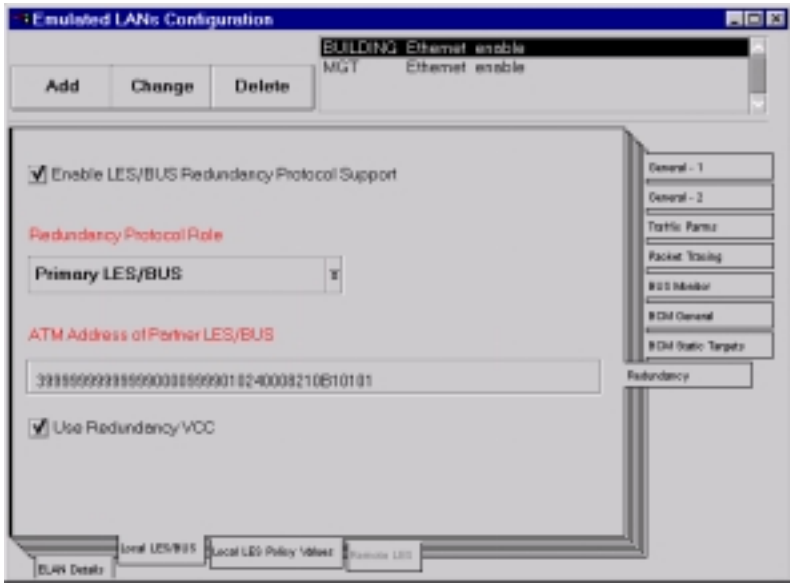

This is a capture from the second MSS showing the configuration of the backup LES/BUS for the BUILDING Elan.

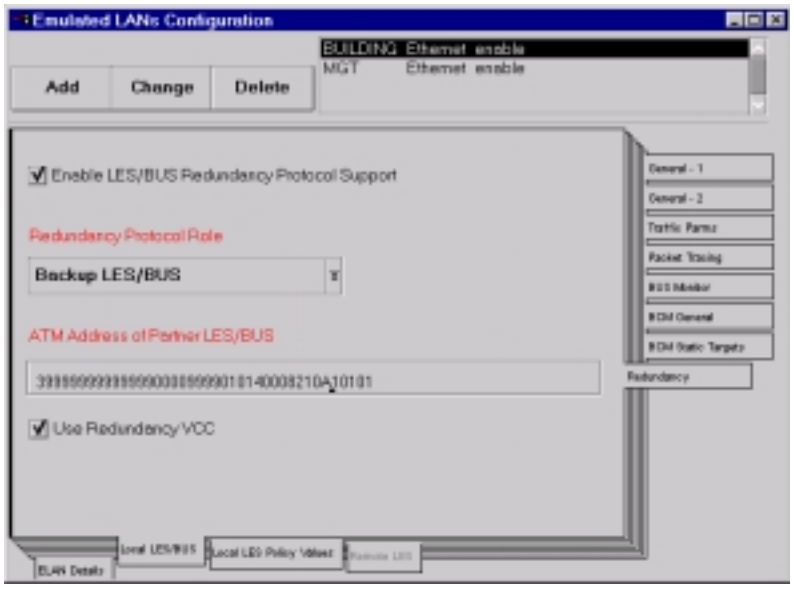

The ELANs in this network using the ‘by name’ method for LEC registration. The ELAN name is defined here.

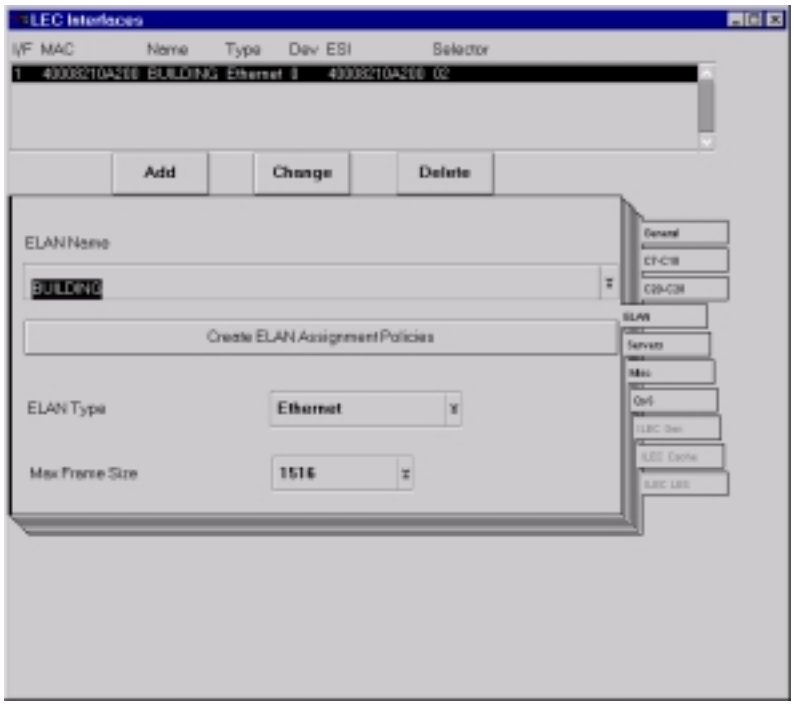

IP connectivity to the MSS is defined using Host services. TCP/IP host services are defined with the address and do not require internal MSS LECs to be defined with IP addresses. Bridging must be enable on the MSS for Host Services to work correctly.

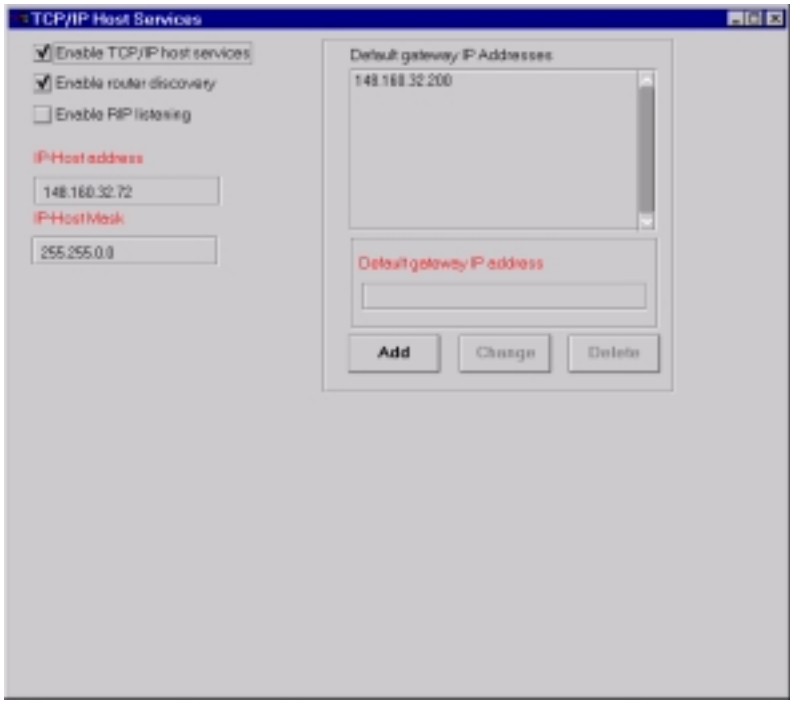

Bridging is enable on the MSS and also on the LEC for the internal MSS LEC. This is required for Host services.

The internal LEC for the BUILDING Elan has bridging enabled to allow TCP/IP host services to communicate with traffic on the Elan.

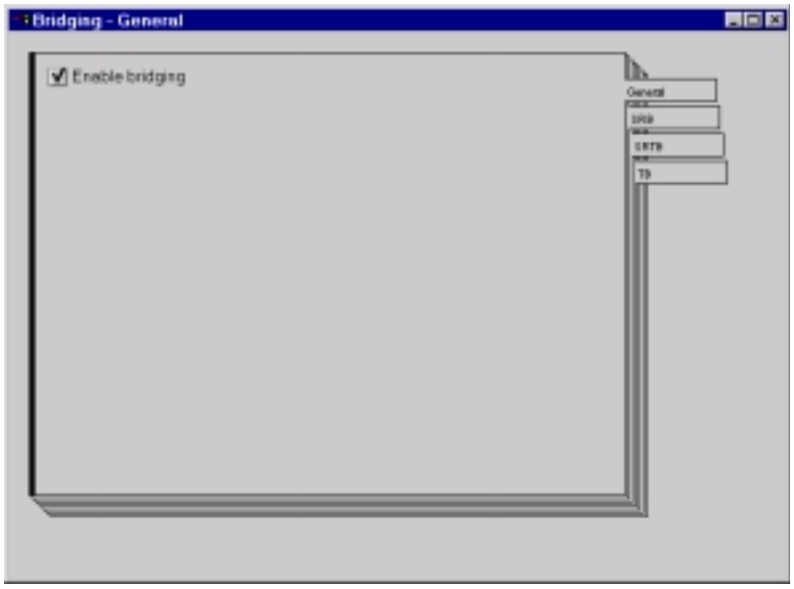

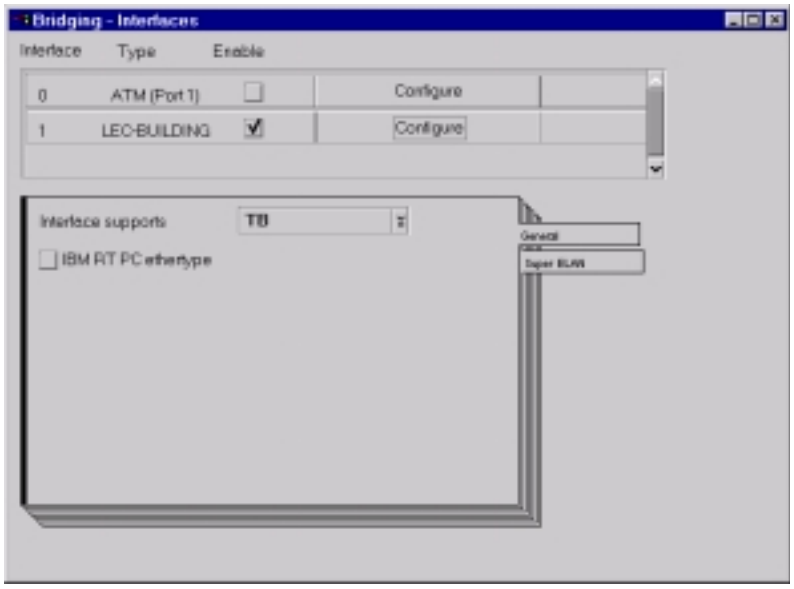

Appendix B: Glossary

BUS - **B**roadcast **U**nknown **S**erver to handle broadcast traffic on ATM emulated LAN

ELAN - **E**mulated **LAN** used to define a broadcast domain on ATM
LES - **L**AN **E**mulation **S**erver used to create ATM based emulated LAN
LEC - **L**an **E**mulation **C**lient
SRB - **S**ource-Route **B**ridging
SR-TB - **S**ource-**R**oute **T**ranslational **B**ridging
transparent bridging - method to connect Ethernet segments
VLAN - **V**irtual **LAN** defining a group devices in the same broadcast domain
802.1d spanning tree - algorithm used to create a loop-free path in a transparent bridged network.

Appendix C: References

IBM Redbook: IBM 8265 Nways ATM Campus Switch SG24-2004-00

IBM Redbook: Understanding and Using the IBM MSS Server SG24-4915-00

IBM Redbook: Understanding and Using MSS Release 1.1 and 2.0 SG24-2115-01

IBM Redbook: IBM Nways RouteSwitch Implementation Guide SG24-4881-00

IBM Redbook: Troubleshooting IBM LAN/ATM Campus Networks SG24-2105-00