IBM

Token-Ring to Ethernet Migration

An insight into technical areas that need to be considered when migrating from a token-ring network to Ethernet

An understanding of a step-by-step process to follow in planning a network migration project

Practical considerations for migration based on real-life case studies

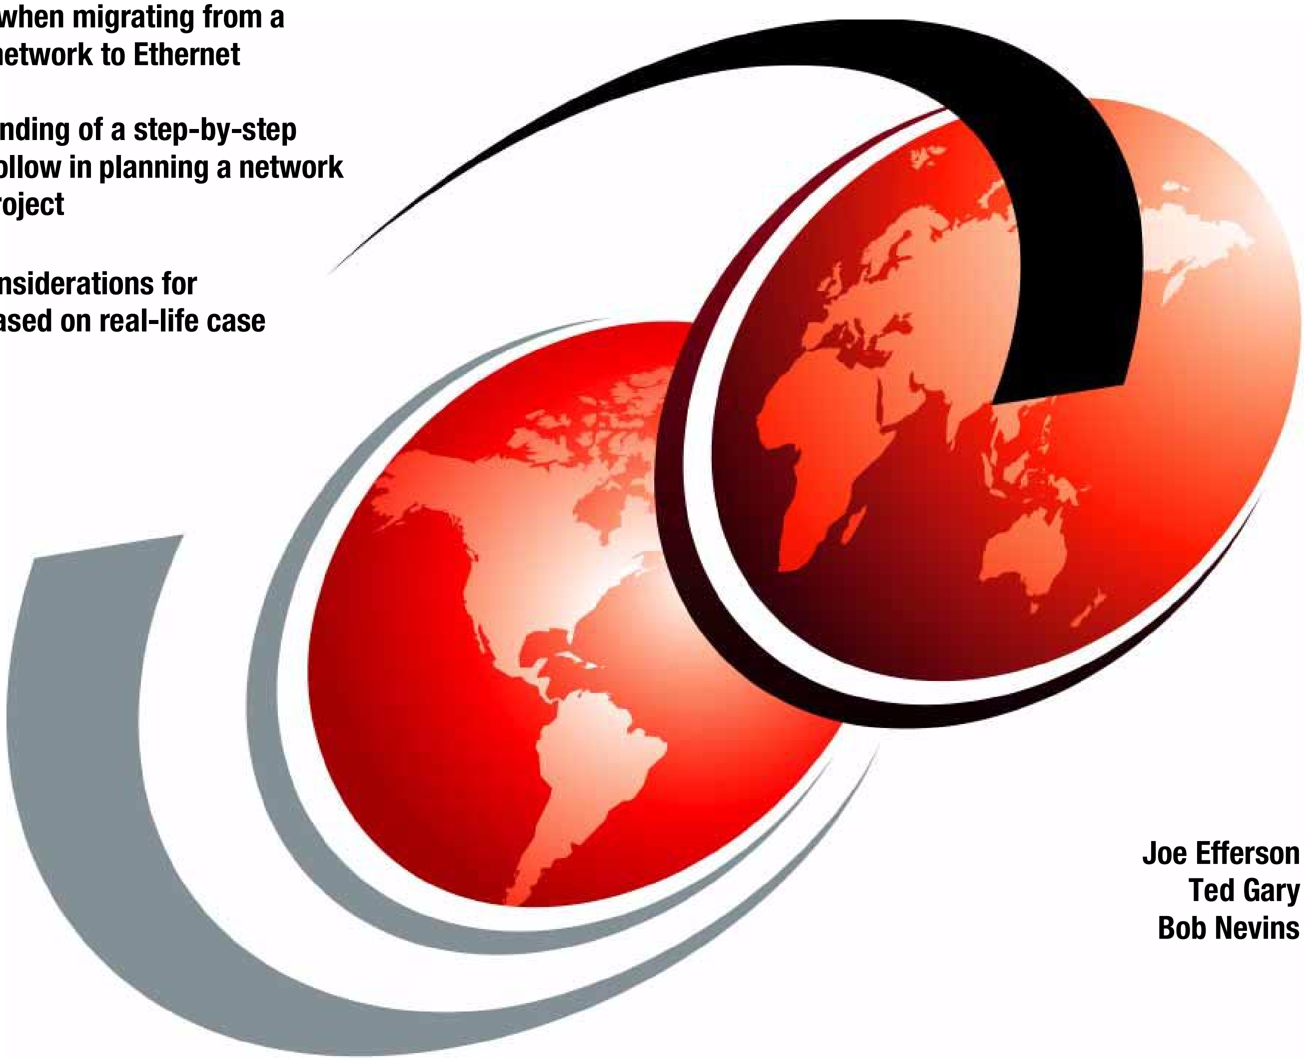

Joe Efferson
Ted Gary
Bob Nevins

Redpaper

ibm.com/redbooks

International Technical Support Organization

Token-Ring to Ethernet Migration

February 2002

Take Note! Before using this information and the product it supports, be sure to read the general information in “Special notices” on page v.

First Edition (February 2002)

This edition applies to real-life examples of Token Ring to Ethernet migration projects.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. HQ7 Building 662
P.O. Box 12195
Research Triangle Park, NC 27709-2195

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

Special notices . . . v

IBM trademarks . . . vi

Preface . . . vii
The team that wrote this Redpaper . . . vii
Notice . . . vii
Comments welcome . . . viii

Chapter 1. Executive summary . . . 1

Chapter 2. Economics of token-ring to Ethernet migration . . . 3

Chapter 3. Migration: Some things to consider . . . 7
3.1 Adapters and workstations . . . 8
3.2 Cabling . . . 8
3.3 Rings and stars . . . 10
3.4 Hubs and switches . . . 11
3.5 Layer 2/Layer 3 Network Design . . . 12
3.6 Servers . . . 13
3.7 Network management . . . 14

Chapter 4. A general approach to migration . . . 15

Chapter 5. Case study: American Stores . . . 19

Chapter 6. Case study: ManySite Corporation . . . 25

Special notices

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

IBM trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

AIX®	e (logo)®	Redbooks™
AS/400®	IBM®	S/390®
DB2®	Lotus®	SP™
e (logo)®	Redbooks (logo)™	Tivoli®

Other company trademarks

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc. in the United States and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States and/or other countries.

PC Direct is a trademark of Ziff Communications Company in the United States and/or other countries and is used by IBM Corporation under license.

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, and service names may be trademarks or service marks of others.

Preface

This Redpaper, based on experience gained from migration projects with many customers and clients, offers considerations, guidelines, and suggestions for approaching a Token Ring to Ethernet migration project. It highlights many of the technical considerations that must be addressed in a migration project, offering thoughts for how to stage such a migration, and provides a high-level, step-by-step outline for a sample migration. This Redpaper includes case studies that provide real-life examples of migration projects.

The team that wrote this Redpaper

This Redpaper was produced by a team of specialists from around the world working at the International Technical Support Organization, Raleigh Center.

Joe Efferson is a Consultant at the IGS Center for Infrastructure Solutions, in Raleigh, NC. He provides consulting services helping clients from around the world develop infrastructure strategies, architectures and migration plans. Five years prior to this assignment, Joe worked in La Gaude, France providing technical marketing support for the Networking Hardware Division.

Ted Gary is a Networking Consultant at the Center for Infrastructure Solutions, Research Triangle Park, NC. He writes extensively and has taught IBM classes in the USA on all areas of campus networking. Before joining the Center for Infrastructure Solutions five years ago, Ted worked in the Networking Hardware Division in New Jersey, USA as a marketing consultant.

Bob Nevins is a senior networking consultant with IBM Global Services, specializing in the application of networking and Internet technologies to solve business problems. He has 18 years of experience in consulting, marketing, and product development in the information technology industry.

Thanks to the following people for their contributions to this project:

Rufus Credle, ITSO Project Leader
International Technical Support Organization, Raleigh Center

Notice

This publication is intended to help those seeking considerations, guidelines, and suggestions when engaging in a Token Ring to Ethernet migration project. The information in this publication is not intended as the specification of any programming interfaces. See the PUBLICATIONS section of the IBM Programming Announcement for more information about what publications are considered to be product documentation.

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this Redpaper or other Redbooks in one of the following ways:

- Use the online **Contact us** review redbook form found at:

  `ibm.com/redbooks`

- Send your comments in an Internet note to:

  `redbook@us.ibm.com`

- Mail your comments to the address on page ii.

1

Executive summary

Token-ring technology entered the networking world in the mid-1980s. At the time, it offered performance, reliability, and scalability features far exceeding the capabilities of any other local area networking technology. Over the course of the past decade, however, advances in Ethernet and other technologies have surpassed the capabilities of token-ring. Today, although there are still many token-ring adapters shipped annually, the focus in networking has shifted almost exclusively to Ethernet.

Over the next several years, as vendors and suppliers continue to concentrate their efforts more and more on Ethernet, token-ring equipment and skills will become increasingly less available. Organizations that have not yet migrated to Ethernet would do well to begin making plans to do so.

Although it is sometimes difficult to cost-justify a move to Ethernet because “one network is being replaced with another,” many businesses do so by pointing to improved network capabilities that can be deployed with the new network - higher bandwidth, multicast capability, or Quality of Service - or by combining the Ethernet deployment with another information technology (IT) project. Additionally, the move to a new network infrastructure often provides an opportunity to “clean up” the existing network - rationalizing the addressing structure and removing older protocols or applications -- which can dramatically improve the performance and manageability of the network.

This paper, based on experience gained from migration projects with many customers and clients, offers considerations, guidelines, and suggestions for approaching a token-ring to Ethernet migration project. It highlights many of the technical considerations that must be addressed in a migration project, offers thoughts for how to stage such a migration, and provides a high-level, step-by-step outline for a sample migration -- beginning with a thorough audit of existing infrastructure and proceeding through design of the new network, implementation, and coexistence considerations while the two networks are both in use. Case studies provide real-life examples of migration projects.

Of course, a general-purpose migration discussion can only go so far. Each organization will have its own unique challenges and requirements. The size of the organization's network, its existing infrastructure, and the time and resources available can make a migration either relatively straightforward or much more complicated. IBM's consulting and services teams have helped many clients successfully deploy and migrate to new networks. For more information about how IBM can help you, please contact your IBM representative.

2

Economics of token-ring to Ethernet migration

Three forces dictate the economics of all technologies: speed, cost, and market timing. Token-ring technology entered the technology marketplace with good benefits related to wiring requirements, shared LAN performance and network management. These exclusive benefits were chipped away by improvements in Ethernet technology. During an explosive growth phase in the networking equipment and connectivity marketplace, token-ring was hobbled by both a speed and cost disadvantage.

The marketplace for token-ring includes both manufacturers and purchasers that deploy this technology. Each has a set of decisions to make that are related to the task at hand and the money associated with doing the job. Manufacturers of token-ring equipment have been challenged with a declining population of purchasers, which in turn lowers the volume of products that must be built. The suppliers of chips and subassemblies are also seeking new sources of revenue. As the suppliers' attention has been on other opportunities, this has caused an even steeper decline in the availability of parts and materials from which token-ring products are made. These twin trends have conspired to turn a gently shrinking market downtrend into an accelerated decline.

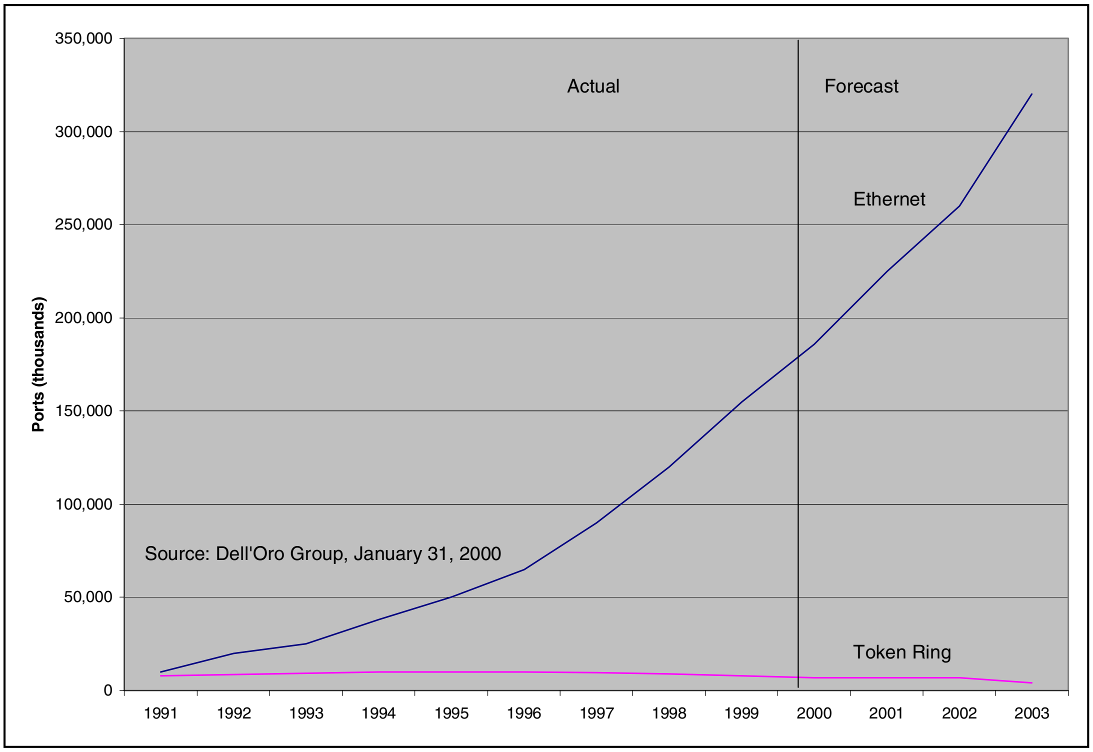

Figure 2-1 Ethernet versus token-ring shipments

The enterprises and small businesses that still have a token-ring install base must find both the will and the money to migrate to Ethernet technology. In many businesses, a financial analysis of the cost and benefits of a proposed investment is a routine part of the spending cycle. Conventional return-on-investment (ROI) analysis focuses on the efficiencies and cost-avoidance options that are derived from a technology change. Normally, this encompasses the “hard” savings related to a payback period that starts with the new technology showing its strengths. Hard savings can be found in smoother operations, lower people-related costs, and application enablement. On top of that there is usually a set of “soft” savings that purport to show an improvement in employee morale, customer goodwill, or opportunity for new business relationships.

Oddly, only one of the above ROI elements has any place in a token-ring to Ethernet migration. The return and benefits are all clustered around the area of application enablement. Of all the cost cases and business justifications of migration that we have reviewed, the businesses that migrated found that they had positioned themselves for unfettered network growth in the future, but had not saved any significant near-term money.

Applications that involve the use of newly available Internet appliances are the low-hanging fruit of an Ethernet migration. These devices almost exclusively have Ethernet interfaces and are packaged to be firewalls, perform sophisticated traffic management, supply QoS (Quality of Service) reports, and create Virtual Private Networks (VPNs). In addition, multicasting (whereby a single stream of IP traffic is multiplied at the routers or LAN switches to reach multiple user devices) is an integral part of several new applications. Token-ring LAN switches were in decline during the time that multicasting for LAN switches was being developed and so there are very few choices left in the token-ring marketplace if this is a network requirement.

As an organization prepares to undertake a migration from token-ring to Ethernet, the question of cost justification will usually be encountered. Because of the trends outlined above, this cost justification will rarely show significant monetary savings. Instead, it may focus on expanded capabilities of the new network infrastructure (higher bandwidth, multicasting, Quality of Service, etc.), or it may simply reflect the necessity of migration based on the increasingly higher cost of maintaining, supporting, and managing token-ring technology as both hardware and skills to support token-ring become harder to find.

Many organizations have successfully migrated from token-ring to Ethernet. The following sections of this paper will highlight issues that need to be considered in such a migration, based on experience gained from many clients. A sample migration plan will also be outlined, illustrating a high-level plan for undertaking an Ethernet migration project.

3

Migration: Some things to consider

Preparation for a migration from token-ring to Ethernet should include careful consideration of many elements of the network. In this section we will examine, in turn, each of the elements of the network shown in Figure 3-1: adapters and workstations, cabling, hubs and switches, Layer 2/Layer 3 network design, servers, and network management.

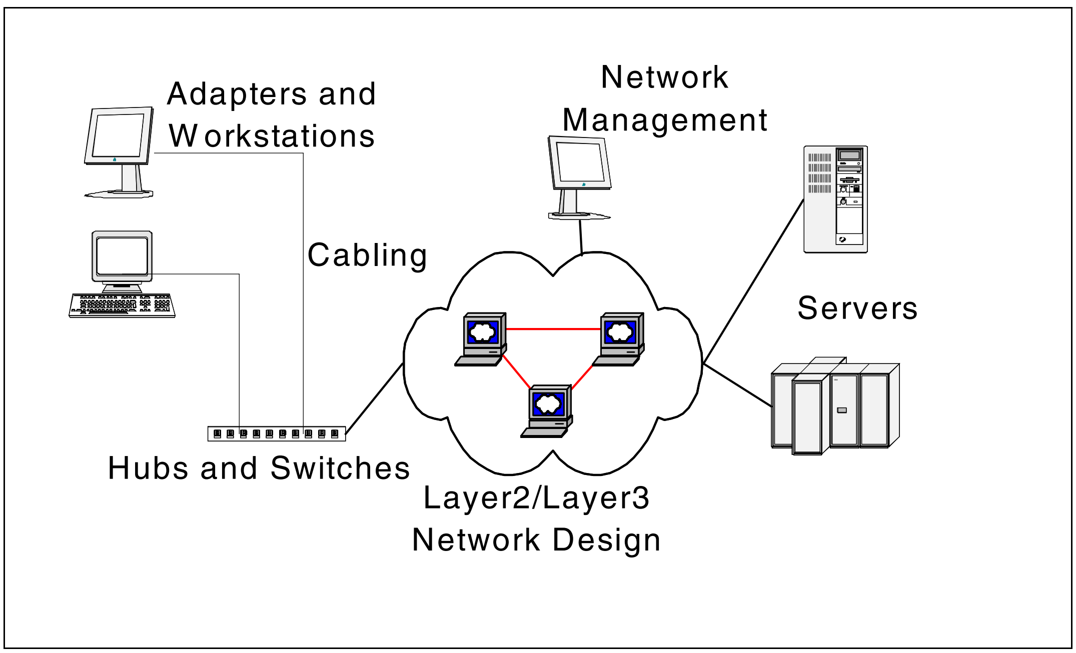

Figure 3-1 Areas to consider in migration

3.1 Adapters and workstations

Obviously, for a device to connect to an Ethernet network, it will require an Ethernet adapter. Many devices now come with an Ethernet adapter built in, and for those that do not, adapters are widely available in the marketplace. The relatively low cost of these adapters (for most devices) makes it fairly inexpensive, from a hardware point of view, to equip an end user's device for Ethernet.

The bigger cost in adapter/workstation migration is the labor involved. Each existing station that will be migrated to Ethernet must have an adapter installed and must be configured (or reconfigured) with the proper address, subnet mask, etc. to attach to the Ethernet network. Because of this labor-intensive process of “touching” each end-user device, some organizations attempt to “piggyback” the Ethernet migration process with some other project that involves touching the end-user devices. If this can be done even for a reasonable subset of the existing users, it can help alleviate the cost burden associated with the Ethernet migration project.

Another strategy for migration is to move end users to Ethernet on the “technology refresh” cycle. In other words, many organizations have a routine cycle on which users are given new workstations - for example, every three years or so. An organization might decide that at a certain point, all new workstations will be configured with Ethernet instead of token-ring. Thus, over a period of time, all (or at least most) end-user devices would be migrated through the normal upgrade process. Any remaining devices could then be migrated on an exception basis.

One consideration in adopting either of these approaches is that if users in a particular location (say, a floor of a given building) are not on the same technology refresh cycle, token-ring and Ethernet would both be required in that location for some period of time. Thus the wiring closets, cabling, etc. in that location would have to be equipped to support both Ethernet and token-ring during the migration period.

A different kind of issue arises if the existing token-ring network contains devices that for some reason cannot support an Ethernet adapter. Examples might include special-purpose devices (for example, a wireless base station with a built-in token-ring adapter) or older devices whose hardware or operating system cannot support Ethernet. In such a case, the only real alternatives are to replace the device or to maintain a small token-ring network segment for the remaining lifetime of the device. If there are multiple such devices, perhaps they could be combined or centralized on a single “legacy” token-ring segment.

3.2 Cabling

Originally, token-ring was designed to operate using the IBM Cabling System (ICS). The IBM Cabling System was a set of shielded twisted pair (STP) cable specifications (Type 1, Type 2, Type 6, etc.), connectors, and other hardware that IBM developed and tested to provide a complete building wiring system for token-ring.

The cable specifications for ICS token-ring cabling are different from those developed later for Ethernet using unshielded twisted pair (UTP Category 3, Category 5, etc.). There are several differences in these specifications:

- STP Cable has a characteristic impedance of 150 ohms, whereas UTP has a characteristic impedance of 100 ohms.
- The IBM Cabling System used a different connector form factor (an IEEE 802.5 data connector) from the RJ-45 connectors commonly used for Ethernet.

- Where RJ-45 connectors are used by various vendors for token-ring, the pin-out of the connectors is different from that used for Ethernet.
- The ICS data connector is self-shorting when unplugged from the end station, creating a loop in the cable from the hub/switch to the end station.

The good news is that all of these differences can be overcome, and that in most cases existing STP cable can be used quite successfully for Ethernet.

In 1999, IBM tested STP cable for use in Fast Ethernet (100BaseT) applications and announced that, in most cases, installations can use existing STP cable for Fast Ethernet networks. In fact, STP was shown to demonstrate superior performance for Fast Ethernet up to 170 meters, well beyond the standard 100-meter cable length for Ethernet. Additionally, the Fast Ethernet standard, 802.3u, supports both STP and UTP cable, so that a product that is 802.3u-compliant should support both STP and UTP. Of course, it is always worth verifying vendor product specifications to ensure STP support for a particular product, but in most cases existing STP cable will support Ethernet with excellent results.

The connector and pin-out differences between token-ring and Ethernet can be solved with the use of appropriate pin-swapping adapter cables, which are readily available in the marketplace.

Some guidelines should be followed in using STP for Ethernet networks:

- Long lengths of STP and UTP should not be mixed in a single cable run (UTP connected to STP). The impedance mismatch in the two types of cable causes signal reflections that can increase error rates. Generally, short UTP adapter cables connected to an STP cable lobe do not cause problems.
- Token-ring supports cable runs longer than 100 meters in some cases. Ethernet installations should adhere to the Ethernet-standard 100-meter maximum cable run. Installations with long token-ring cable lobes will need to test those cables to ensure adequate performance with Ethernet, or else redesign the wiring so as not to exceed 100 meters per cable run.
- Depending on the age of the wiring plant, it may be well worthwhile to conduct testing of a representative sample of the cable prior to migration, in order to ensure appropriate performance for Ethernet. Wiring deficiencies can cause problems that are notoriously difficult to troubleshoot, so if there is any doubt about the ability of the existing cabling to support Ethernet, it should be tested by a competent cabling consultant.

If ICS data connectors (not just STP cable) will be used in an Ethernet installation, the Ethernet switches should be checked to ensure there will be no problems with the self-shorting mechanism of these connectors. This mechanism shorts the transmit pair to the receive pair when the connector is not plugged into an end device, allowing the ring path to be maintained. With some Ethernet switch equipment, however, this mechanism can cause problems because the switch will receive its own frames. The problem can be alleviated by software in the switch, and most Ethernet vendors have in fact designed their Ethernet switches to handle this issue. However, as stated before, it is always worth verifying this capability with the Ethernet switch vendor.

Both token-ring and Ethernet support optical fiber cabling as well as copper. Gigabit Ethernet, however, which is frequently used as a backbone technology in Ethernet deployments, supports shorter maximum distances on multimode fiber than does token-ring. If long fiber optic cables are currently in use in a token-ring backbone network that will be migrated to Gigabit Ethernet, the wiring plan may have to be redesigned or single mode fiber may be required. Table 3-1 shows typical maximum distances supported by Gigabit Ethernet over various types of cable. These distances vary with specific types of cable and specific products, so verification of actual product and cable capabilities is essential.

Table 3-1 Typical supported distances for Gigabit Ethernet over various cable types

Cable Type	62.5/125 mm fiber	50/125 mm fiber	Singlemode fiber	Category 5 UTP
Gig Enet port				
1000BaseSX	220-275 meters	500-550 meters	N/A	N/A
1000BaseLX	550 meters	550 meters	5 km	N/A
1000BaseTX	N/A	N/A	N/A	100 meters

3.3 Rings and stars

Token-ring, as the name suggests, was originally conceived as a ring technology (see Figure 3-2). Ring topologies (seen in FDDI, token-ring, and some Metropolitan Area Network technologies) have inherent self-healing properties that provide survivability of certain network failure scenarios. Ethernet network design is based on a star topology, and redundancy is achieved using backup links and backup switches (see Figure 3-3).

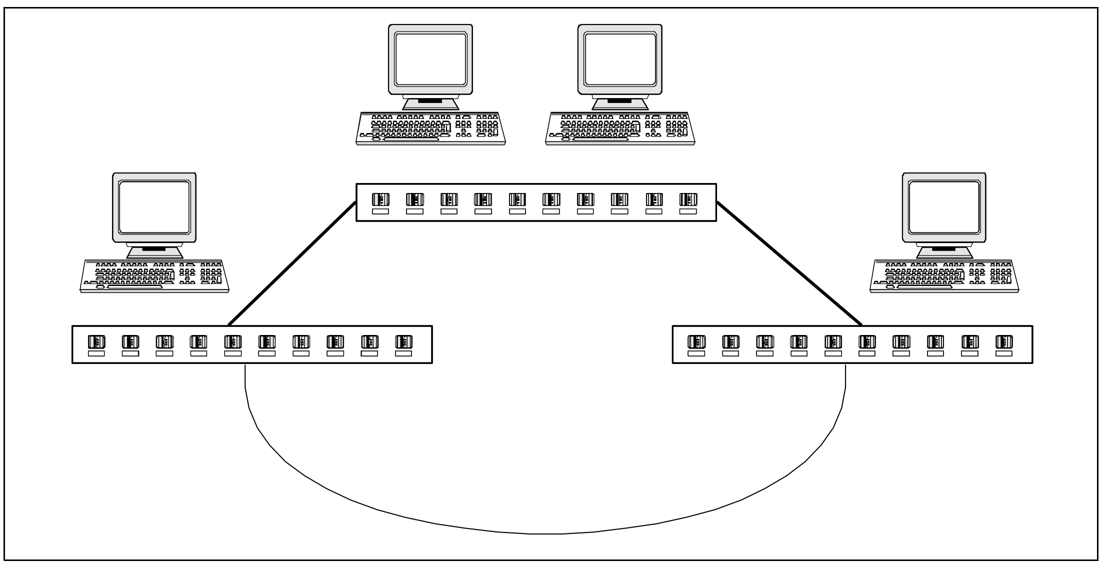

Figure 3-2 Ring topology in a token-ring network

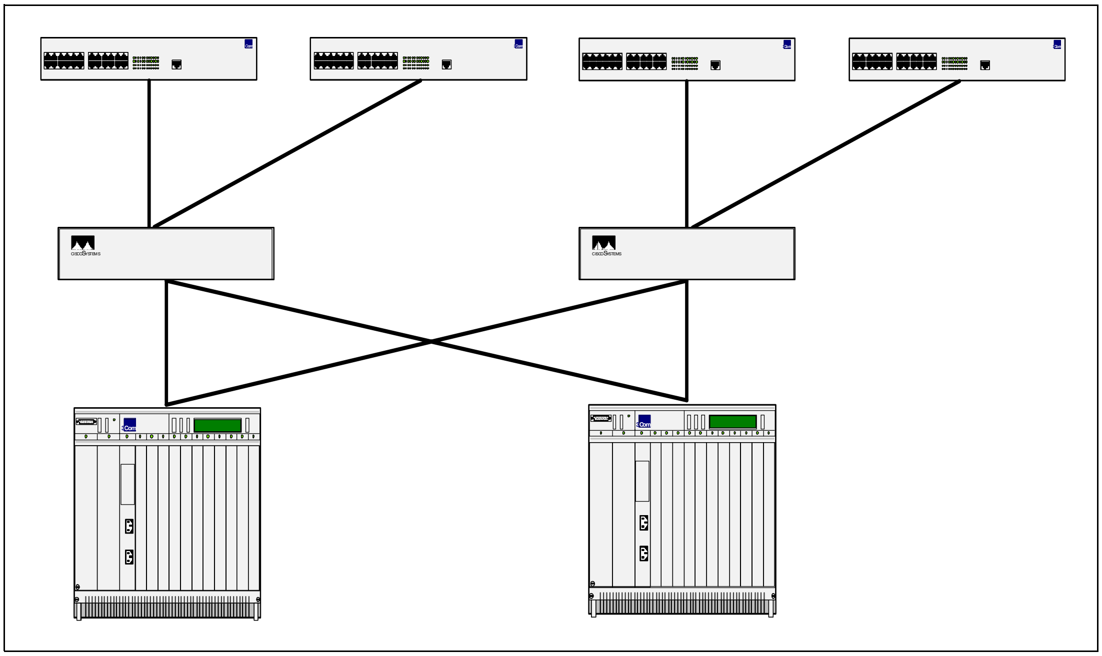

Figure 3-3 Star topology in an Ethernet network

Over the years, token-ring networks have been built in a variety of ways. In some cases, the actual physical wiring topology is in fact a ring. This topology can be used either within a building or across a campus (for example, a “distributed backbone” ring). In many other token-ring networks, however, the physical wiring topology is a star, and the logical ring path is achieved via dual cable runs out and back along the star-wired paths (for example, a “collapsed backbone”).

Ethernet networks should be designed using star topologies. One of the first tasks in planning for a migration to Ethernet is to analyze the existing token-ring cable topology. Depending on the existing physical cabling structure of the token-ring network (star, ring, or some combination thereof), achieving a star topology for the new Ethernet infrastructure may be straightforward or may require redesign and/or re-cabling. Appropriate backup links and redundant switches should be included in the design of the Ethernet infrastructure, as seen in Figure 3-3. The cabling design should dovetail with the Layer2/Layer3 design (IP subnet structure, placement of switches and routers, etc.) as will be discussed in a later section of this document.

3.4 Hubs and switches

Ethernet hubs and switches represent the wiring closet component of the Ethernet network. Given the state of the Ethernet hardware industry today and the price points of hubs and switches, switches are typically the preferred attachment technology for Ethernet end stations (there is little reason to use hubs, since basic switches normally don't cost much more).

In designing the wiring closet components of the Ethernet network, the biggest decision to be made is usually which vendor to select and which particular vendor products to use. Chassis-based switches, generally speaking, offer greater port density and the capability to add more function in the chassis (for example, Layer 3 technology, redundant power supplies, other specialized functions), but come at a higher price point than stackable switches.

Stackable switches, conversely, generally come at a lower price point and offer straightforward connectivity for end stations, the ability to uplink to core network components, and the ability to start small and add more ports over time by simply adding additional switches.

A thorough vendor selection process should be performed when designing and planning the Ethernet network. Selection criteria should include not only standard Ethernet product functions and pricing, but also other factors such as vendor support, compatibility with other network components, and any specialized functions required for the network (for example, Layer 3 technology, Quality of Service support, IP multicast support, etc.).

Migration planning should also include a thorough inventory of wiring closet environmental characteristics. First, is there enough space in the existing wiring closets to add Ethernet switches? Will new equipment racks be needed? New wiring closets? How long will the token-ring/Ethernet coexistence period be, during which time new Ethernet switches will have to be in the same wiring closets as existing token-ring equipment? Is there enough power in every wiring closet to support the equipment required? Is there enough air conditioning capacity? Many of these questions about the coexistence time frame are of course interrelated with the deployment philosophy that will be used for Ethernet, as discussed 3.1, “Adapters and workstations” on page 8.

3.5 Layer 2/Layer 3 Network Design

Layer 2/Layer 3 Network Design refers to all the issues surrounding the placement and interconnection of Layer 2 switches and Layer 3 switches or routers. This includes IP subnet structure, size of subnets, device addressing and subnet masking, choice of Layer 2 and Layer 3 protocols that will be supported in the Ethernet network, redundancy, and many other issues.

The trend in networking today is toward an IP-only network infrastructure. In the past, many different protocols were typically found in an average enterprise network, ranging from Layer 2 protocols such as SNA LLC and NetBIOS to Layer 3 protocols such as IP, IPX, AppleTalk, DECnet, and others.

Many enterprises use the migration to Ethernet as an opportunity to address other network issues, such as a move to an IP-only infrastructure. A frequent approach is to make IP the only supported protocol on the new Ethernet network, thereby enticing departments or user groups that currently use other protocols to upgrade in order to move to the new network. If enhanced services or higher bandwidth will be offered on the new Ethernet network, this will further encourage user groups to migrate.

Moving to an IP-only network involves determining the mix of protocols currently in use and addressing, one by one, the available alternatives for encapsulating those protocols in IP or changing to IP as the protocol for a particular application. Native NetBIOS can be moved to NetBIOS over IP; Native SNA can be moved to TN3270 or TN5250 or encapsulated using SNA Data Link Switching; NetWare can be upgraded to use IP rather than IPX. Similar approaches can be used for most other protocols found in networks today. If older platforms or applications exist in the network that cannot support IP, they can be addressed on an exception basis, left on token-ring, or upgraded to newer technology. The task of moving individual applications or protocols to IP can also be done ahead of the Ethernet migration, to simplify the changes involved during the actual migration.

For many organizations, the move to Ethernet also provides an opportunity to “clean up” the existing IP network structure. Often, the IP network has grown in a haphazard fashion and may have a mix of several different address structures, noncontiguous subnets that preclude effective address summarization, and even unregistered “illegal” IP addresses. The design of the new Ethernet offers a chance to lay out a new IP address structure, using either a registered IP address block or, more commonly, private addressing as defined in RFC1918. A “clean” IP network design often substantially improves network performance and manageability.

Token-ring networks were often built in a fairly “flat” manner, using more (source routing) bridges and fewer routers than Ethernet networks. With Ethernet networks, the typical design (with some exceptions) uses subnets no larger than a class C address space. This may be a substantial change from the existing token-ring network, involving smaller subnets and more routing (or Layer 3 switching). In planning for the move to Ethernet, network services such as Dynamic Host Configuration Protocol (DHCP) servers and Domain Name System (DNS) servers must be provisioned for the new network; the good news is that these tools (especially DHCP) can dramatically ease the migration tasks by automating the address administration process for workstations.

During the migration process, while both token-ring and Ethernet networks coexist, the two networks will normally need to be interconnected to allow enterprise-wise communication. The recommended approach for this interconnection is via routing or Layer 3 switching, rather than some form of translational bridging. Using a Layer 3 device to interconnect Ethernet and token-ring avoids many problems that arise from the differences between Ethernet and token-ring, such as bit-ordering, different bridging approaches, and frame-size issues.

The connection between the existing token-ring and the new Ethernet network is typically made between the two core networks (old and new). Other designs can be used, however, if the existing backbone network is overloaded. In any case, the connection point needs to be carefully sized (in terms of processing power and bandwidth) to prevent its becoming a bottleneck, since once the servers are moved to the new network, all user traffic from the old network to the new will go through that connection point.

3.6 Servers

As with workstations, the first consideration for servers is provisioning Ethernet adapters. Unless the server is an older or specialized device, Ethernet adapters will normally be widely available. Each application and/or protocol that operates on a particular server should be inventoried, so that its suitability for operation over Ethernet can be determined. If the move to Ethernet is to be combined with a move to an IP-only network, each existing non-IP application or protocol will need to be migrated to use IP (as noted above, this can be done prior to the actual migration, rather than trying to do everything at once during the migration process).

Older servers or special-function token-ring devices that cannot be moved to Ethernet will of course require that token-ring connectivity be maintained through the lifetime of the device. This may involve special consideration in the design of the Ethernet network, such as defining a “legacy” token-ring subnet to which older devices can be attached.

During the migration planning (as discussed below), a server network design will be developed outlining the attachment points for each server, the addressing structure for the servers, the bandwidth required, and any redundancy considerations (for example, does a particular server require multiple attachments to one or more switches for high availability).

A common question in server migration is whether to attach the servers to both Ethernet and token-ring networks during the migration period or to move them wholesale to Ethernet and disconnect the token-ring attachment. Based on experience, the recommended approach is that whenever possible, the devices should be moved to Ethernet-only attachment. Using two adapters of different types in the same server can cause connectivity problems in certain cases, depending on the software in the server. For example, it is not always easy to control which adapter is used by clients. Traffic may come into the server over Ethernet but be sent back to the client over token-ring. In some cases this might work fine but if problems occur, diagnosis and correction can be very difficult. Moving the servers to Ethernet does require that traffic from token-ring clients flow across the token-ring to Ethernet interconnection device, but can save many headaches during the migration period.

A typical scenario in which it is desirable to maintain dual adapters in a server (token-ring and Ethernet) during the migration time frame is when native SNA devices exist in the token-ring network (which would typically connect to an S/390 or AS/400 host). Until the SNA devices are migrated to use IP, it is best to maintain token-ring connectivity so that the SNA data path can operate unchanged.

3.7 Network management

Network management should be a core element of the design of the new Ethernet network. Network management tools and platforms are widely available; a thorough discussion is beyond the scope of this document. Depending on the existing network management infrastructure (whether or not sophisticated network management tools are already in place in the token-ring network), a network management vendor evaluation project may be required as part of the Ethernet deployment.

In addition to network management tools, Ethernet network management skills will need to be developed. It is understandably preferable to develop these skills well before the migration project begins.

4

A general approach to migration

In this section we will discuss a generalized high-level project plan for an Ethernet migration project. Each organization, however, has particular environmental conditions that be must be taken into consideration. Probably the most important recommendation is that an organization carefully plan its migration to avoid potential problems during the process. The general migration plan as outlined here consists of eight steps.

Step 1 - Understand the existing environment

As in any migration plan, the first and most important step is to understand the existing environment. To understand the environment an inventory of workstations, cabling, closet hubs/switches, LAN backbone, server connectivity/protocols and overall design structure is required. For each of these areas we list some of the information that should be gathered.

From a workstation point of view, all end devices that have a token-ring connector need to be identified. As mentioned above, there may be some devices that do not offer any other type of networking adapter. This total number of devices will represent the minimum number of switched Ethernet ports deployed as part of the migration, not taking into consideration any growth.

The cabling within each building should be reviewed to document the type of connector, the cable type deployed, cable topology, the level of documentation identifying cable end points, and the length of cabling to the wiring closet. If the cable is older, it may be wise to test sample segments to verify the cable's ability to carry 10 Mbps or 100 Mbps Ethernet traffic.

In the wiring closets, identify the number and type of hubs or switches deployed. Also, verify the space and power availability in the closet to add additional equipment during the migration.

A backbone network for the existing token-ring campus would consist of the inter-building cabling and the switches or routers that provide the backbone traffic movement. In most cases the inter-building cabling will consist of fiber; an inventory of both used and unused fiber should be identified as well as the type of fiber available (single mode, multimode, etc.). The existing backbone network could be built on switches or routers using token-ring, FDDI, or ATM technology. Identify the type of equipment and number of available ports on the backbone equipment.

The servers could be connected with one of several types of adapters depending on the backbone technology used. Identify the adapter types, speed, and protocols supported by the server. Certain protocols, such as NetBIOS and SNA, will probably require additional changes prior to moving the server to Ethernet as discussed earlier.

Finally, the overall design structure of the existing network should be documented. This would consist of the topology used throughout the campus, the IP addressing scheme (the number and size of subnets and all address blocks in use), network management, and the use of DHCP for example.

Step 2 - Develop target network design for the Ethernet network

The second step is to develop the target network design for the Ethernet network. This design should take into consideration all the areas mentioned in step 1.

- For workstations and other end devices, suitable network adapters need to be selected. Devices that do not support Ethernet adapters will have to be replaced over time and will probably be the last devices migrated. To replace these devices, new machines need to be selected.
- In the cabling area the primary decision is whether the existing wiring can be used to support Ethernet. Any change to the cabling system is a major investment expense and will greatly lengthen the time required to complete the migration. Any changes required to connectors also need to be detailed.
- In most cases, Layer 2 Ethernet switches are used in the wiring closet. A particular vendor and model needs to be selected. Based on this product selection, space and power should be planned for each wiring closet.
- A backbone design should be established that takes into consideration the inter-building fiber connectivity, availability requirements, and site layout. Depending on the type and number of fibers installed, additional fiber or conduits may have to be planned. In most cases Layer 3 switches are used in the backbone, so a particular vendor and model needs to be selected. The method for interconnecting the token-ring and Ethernet networks should be defined based on the protocols supported.
- As with the edge devices, servers will require new Ethernet adapters, so specific adapter types need to be identified. Plans also need to be put in place for all protocols other than IP as mentioned in the section above.
- Overall design has at least one major requirement: the establishment of a new IP addressing plan for the new Ethernet network. If DHCP has not already been deployed, it may be best to migrate the workstations to DHCP as part of the Ethernet migration.

Step 3 - Decide on migration staging plan

The third step is to decide on how to stage the migration of token-ring edge devices to Ethernet. The servers will be discussed as a separate step below. As mentioned in the technology section above, there are several ways to migrate end users. The easiest plan would be to migrate as part of a technology refresh since it does not require any changes to the existing end device, but this requires deploying Ethernet switches to all wiring closets initially and putting in place an aggressive technology refresh program (usually completed in three years). A second approach is to deploy Ethernet as part of another project that requires changes to the end device. Another option is to target specific groups of end users -- for example, one building at a time -- so that deployment of equipment can be spread over a longer time frame.

There is also a possibility of migrating the workstations when users are moved from one office location to another (part of a move, add or change). In any case, the process for handling user moves, adds, and changes during the entire migration phase needs to be documented. Do token-ring users that move offices stay on token-ring after the move? Do they move to Ethernet automatically whenever they move? Is it different depending on the building?

One or more of the options can be combined depending on how fast a migration is desired. One thing to remember is that the longer the migration takes, the longer the mixed token-ring/Ethernet environment must be supported and maintained.

Step 4 - Prepare for the target network

The next step is to prepare for the target network. In this case the items listed in step 2 will be ordered or completed. Products are ordered (adapters and switches), racks and power outlets installed, and cable changes - within buildings and between buildings - are completed. Even if existing cabling is used, it is a good idea to test at least a portion of the cables and fibers to ensure they are within specs with proper connectors and patch panels in place.

Step 5 - Build the new network backbone

In step 5 it is time to start building the new backbone. This involves building a parallel network to the existing token-ring network with no users or servers attached. There are three main components: a server network, a core backbone network, and the building to core network that will all be set up with the new Ethernet IP addressing plan defined above. The logical and physical network design would specify the redundancy provided (backup links, switches, etc.) and the number of switches to install.

The server networks are built using Layer 2 switching technology, not hub technology. These switches are in place to provide direct 10 Mbps, 100 Mbps, and Gigabit Ethernet connections depending on the performance of the server itself.

Ethernet core networks are normally built on Layer 3 Gigabit Ethernet switches and placed in the data center location. Core switches would connect to the server switches and the building switches.

Depending on the migration plan chosen in step 3, building switches are deployed in all buildings or in particular ones, over time, according to the migration staging plan. These building switches connect to the core switches. Depending on the size of the building and the closet layout, there may be multiple switches within each building. When this parallel Ethernet network is up and running, test the new network by connecting a test server to the server network and a workstation to the building switch. If the network has redundant components in the system, test these components by forcing failures in the system. At this stage network management tools should be deployed that will be used to monitor and manage the new environment.

Step 6 - Tie the new Ethernet backbone to the token-ring backbone

Step 6 ties the new Ethernet backbone to the existing network as per the plan in step 2. Usually this is done by interconnecting the token-ring backbone to the Ethernet core backbone through a router or Layer 3 switch. The router would support the protocols required on the Ethernet-attached servers. When moving to Ethernet many environments will specify that only the TCP/IP protocol will be supported, and all other protocols must be encapsulated. This greatly simplifies the Ethernet network but will probably cause some changes to the workstation and server environments that are run today.

Step 7 - Migrate the servers to Ethernet

In step 7, Ethernet adapters would be installed in the existing servers and the servers would be connected to the new Ethernet network. Depending on the server, either it will be connected with only an Ethernet adapter or it will have both an Ethernet adapter and a token-ring adapter, as discussed above. Depending on which protocols are to be supported on the server, there may be changes required to the server configuration (TN3270, TN5250, NetBIOS over IP, NetWare 5.0, etc.).

Servers need to be moved and tested for connectivity by the user groups still on token-ring. At the completion of this step, servers should now have an Ethernet connection to the new Ethernet network.

Step 8 - Migrate user groups

Finally, in step 8, users can be migrated to the new Ethernet network based on the plan already defined above. Depending on the method used in the migration plan, new adapters will be installed, some software changes may be required for protocol changes, and the switch ports in the closets need to be available. As the deployment continues it will be important to monitor the use of switch ports in the closets and the utilization of links to the backbone network to ensure that network capacity is appropriate.

5

Case study: American Stores

American Stores is a food retailer that has a token-ring-based infrastructure both in its corporate offices and at each of the over 120 store sites. The migration to Ethernet was initiated by a plan to improve customer service by installing newer models of store computers (controllers) that perform price lookup and product tracking for warehousing and reordering. The new store controllers run IP-based applications, and therefore a decision was made to ensure the maximum simplicity of the new environment by installing all new Ethernet devices as IP-only.

Understanding the American Stores environment

The corporate offices supported approximately 700 users spread among three buildings. The three buildings (all within 700 meters of one another) were connected in a “V” configuration via fiber optic cables on IBM token-ring hubs. Three large routers on the token-ring hubs performed the duties of WAN connectivity and campus subnetting for the token-ring LAN segments.

Ten token-ring segments were in the building at the apex of the “V”, and the two outlying buildings each supported one ring segment. All of the endpoint user devices were token-ring connected using Category 5 cabling. This simple fact reduced the migration time because all eight of the cable conductors had been properly connected and terminated in the main wiring closet. All of the cable runs were less than 100 meters and therefore presented no impediment to a normal star-wired Ethernet topology. Secondarily, there was sufficient room in the wiring closet to accommodate the new Ethernet equipment racks without disturbing the existing token-ring equipment.

The apex building housed a token-ring-based server farm of approximately 20 Windows NT servers and 5 AIX servers. There was a mixture of native NetBIOS and IP traffic on the Windows NT servers, but a small amount of protocol stack work made it possible to use NetBIOS-over-IP, maintaining American Stores' vision of an IP-only deployment for Ethernet.

Target network design for American Stores

While it is obvious that the store controllers will need to support Ethernet, it was less obvious that a new set of IP addresses must be reserved for new device types that represent the future of retailing. Store infrastructures in the future will include IP-addressable refrigeration units, meat scales, wireless shelf pricing displays and remote barcode readers for customer convenience. Each of these new devices must be merged into a disciplined network management scheme that would readily identify the device and its operational impact should there be a problem with it. Together, the device types and network management/monitoring requirements suggested the need for a “clean sheet” IP addressing scheme.

Two weeks of work were allocated strictly to the task of forming the correct address mask boundaries and number ranges to allow the network management console to identify the store that is having a problem simply from three of the digits within the IP address, and a simple table lookup thereafter could specify the device type. At the corporate offices, the user machines would get their new IP addresses via DHCP as the machines were installed on the Ethernet segments for their department.

This type of address planning will have a long-term payoff in smoother operations and faster problem identification as the network grows.

After IP addressing had been firmly established, American Stores turned their attention to availability of the new Ethernet-connected servers. A redundant-path server network isolated from the user subnets by an L3 switch formed the basis of the new server farm. VRRP (Virtual Router Redundancy Protocol) was an integral part of the new server subnet design. VRRP allows the IP address of a failed interface to appear on another physically reachable interface so that data from the servers can “fail-over” and always find a path to the client without any reconfiguration.

A wide variety of equipment choices were investigated as options for connecting the old and new networks. For example, Cisco's proprietary ISL (token-ring frames on Fast Ethernet Inter-Switch Links) was considered as a possible migration strategy. Since these products were fairly old, we decided not to invest in them even if it were an inexpensive way to do a migration. Other options included allowing the WAN router to do double duty as campus routers. Although this would have worked, there were network management and service demarcation problems that would arise depending on which interface was thought to be the source of the problem. Despite the RBOC's willingness to support the dual-purpose router (the wide area network was to be outsourced to the RBOC), a design that cleanly separated WAN routing from campus routing was chosen as much more acceptable.

The chosen campus network topology for American Stores' headquarters buildings is shown in Figure 5-1. The new Ethernet infrastructure is at the top of the diagram, and is tied to the existing token-ring infrastructure through the routers in the middle.

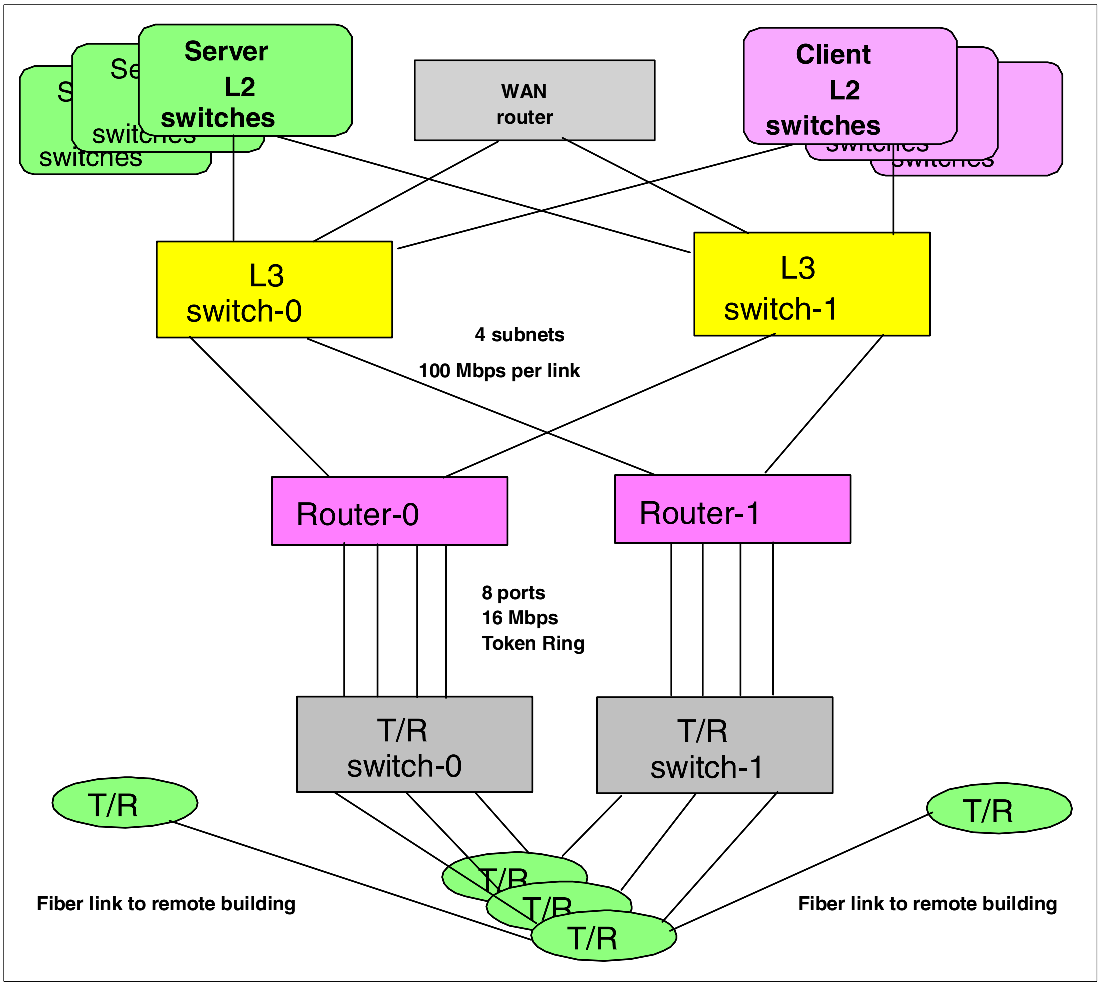

Figure 5-1 Campus network topology

Migration staging for American Stores

New Ethernet LAN switches with both Gigabit and Fast Ethernet interfaces were purchased. Conventional routers were also purchased that had token-ring and Ethernet interfaces. These routers provided the temporary campus connectivity between token-ring and Ethernet during the time when both technologies are present. Once the migration was complete, the routers were sent to some of the stores for wide area connections to the corporate office.

As the new Ethernet LAN switches with L3 function were unpacked, the IP address values for the management interfaces were configured, and a small “sandbox” for connectivity testing was built to ensure that all the equipment was ready for production work. At this time, microcode levels and hardware revision levels were checked and compared to the order forms. Last-minute configuration changes are readily fixed during this phase, and the American Stores migration team was able to fix a potentially devastating interoperability problem on the Fast Ethernet link between the LAN switches and the routers.

The DNS servers were reviewed to verify that they had correct entries for the IP addresses that would be hard-coded into the servers when moved from token-ring to Ethernet. Test client machines belonging to the network migration team were put on small L2 Ethernet switches as the first phase of simulated user migration. The help desk was also alerted to the possible problem of users being unable to reach the servers if they were not using the DNS but were instead using hard-coded addresses or host files.

Preparing American Stores' target network

At this point it was time to make the leap from “sandbox” to production readiness. Extra cable runs for each of the new Ethernet adapters on the servers were prepared and tested. The new Ethernet adapters were checked for device driver compatibility with the Windows NT servers, and the AIX servers were investigated the same way.

The HVAC system in the server room was being upgraded during this time, and the physical rearrangement of the room necessitated a new set of patch cables between the servers and the wiring closet. Somehow during this process, a crossover patch cable was substituted for a standard straight-thru cable and cost the team two hours of diagnostic time to correct the problem when the new adapter had been connected to the network.

Building the American Stores network backbone

During the production week, a test server was connected to the Ethernet L3 switch and its reachability tested from one of the outlying buildings at the end of a fiber optic cable. Also at this time, the cables that formed the redundant paths between the LAN switches were sequentially unplugged to test the VRRP recoverability, and to see if the network management tools properly showed an alert from the anticipated place. This is when the team discovered that there was some significant work required to refine the views that the network management station had of the new network.

Merging the American Stores Ethernet backbone to the token-ring backbone

The routers that connect the token-ring segments to the Ethernet LAN switches were allowed to learn the routes and distribute DHCP addresses to the test client machine.

Migrating the American Stores servers to Ethernet

Earlier in the process of planning the migration a decision was made to avoid the complexity of running two adapters (one token-ring and one Ethernet) per server. Although this would have lessened the critical nature of the migration of the servers from one network to the other, it would have raised the complexity of running the servers and correcting routing problems that sometimes occur on a campus.

During a scheduled outage, each server had its new IP address hard-coded onto its new Ethernet adapter interface and placed on a port in the Ethernet L3 switch. After the test clients were proven to reliably connect to the servers, the help desk personnel were moved onto Ethernet.

Help desk personnel were thought to be the best candidates for this early phase of user migration since they ought to be the ones most tolerant of changes and subsequent troubles. What had not been anticipated was the rapid backlash within a help desk department that can occur when the help desk people are given time to rest while their machines are migrated or replaced. Apparently the call and problem record workloads were so high within the department that even a small rise in call volume can overwhelm those who are still on duty and caused workload problems. After the initial experience with these problems, we decided that it was better to migrate the help desk during the low-volume off hours.

Migrating American Stores users to Ethernet

The backbone and server connection to Ethernet represented the least visible portion of the migration, but it was also the most demanding. Now that the heart of the network was reliably working on Ethernet, the next phase was to move the rest of the large corporate site and all the store sites to Ethernet.

At the corporate site, the unit of change was an L2 switch within the wiring closet. Units of 24 users were notified that they were to be migrated sometime during the next two working days. By notifying two different groups of 24 for each pair of days, it was easier for the change team to keep up a good pace, and it gave the users time to make adjustments (by replying to the notice) if one day was more convenient than another. After some initial missteps related to desks that had been rendered immovable by added bookcases, and user protests over touching rather elaborate cubicle decorations, the corporate site change for nearly 700 users was completed in two months.

The store migration was dramatically more difficult and time-consuming. The store sites had wireless token-ring access points that had to be replaced with Ethernet versions. Moreover, the store sites had a much wider range of variability in the closet layouts, physical accessibility, and user temperament.

Each store was physically surveyed for site readiness, and a “punch list” of change items were generated that would bring the store into our standardized “profile” for a store site. The profile dictates logical attributes such as IP address range based on the corporate store number and the IP devices installed, as well as physical attributes for the position of the LAN switches, routers and store controllers in the racks. Coordination with the local RBOC for the frame relay WAN connections became the biggest project management challenge.

As the RBOC listed its intended “turn-up” date for the frame relay service, the roving team would plan its site visit for six business days ahead of that date to bring the store into “profile”. A stubborn phone line or a missed service call by the RBOC rendered useless the roving team's visit to the store. As time went on, it was obvious that the rollout schedule had to be made less aggressive in order to avoid wasted trips.

Each store controller had to have its database migrated from the old hardware to the new hardware, so a special transfer utility was coded for the controllers. A migration router (that had both Ethernet and token-ring interfaces) was issued as part of the standard kit for the roving teams. The migration router, configured for translational bridging, was only used for the database movement. After the data was moved, the Ethernet-only WAN router was permanently installed by the RBOC.

All 120 stores were completed in 4.5 months, more slowly than the two months of the original plan, but victoriously less than the six months predicted by the naysayers.

6

Case study: ManySite Corporation

ManySite Corporation is a large worldwide manufacturing and consulting company that has a token-ring-based infrastructure in all parts of its business operations. All of the corporate offices, all of the 50 manufacturing sites, and all but a few of the over 1000 sales and marketing branches use token-ring technology along with the shielded twisted pair wiring that was part of the IBM Cabling System (ICS). The migration initiative started after a grudging acknowledgment that the future of the company depended on a truly future-ready data communications network, and a real-world demonstration to its consulting clients that the ManySite Corporation was ready to lead them into the next millennium of sophisticated digital solutions. Since ManySite deals almost exclusively with corporate clients who never see the speed of the applications (but merely receive the output of the billing system), the typical justifications for the migration expense had even less applicability than usual for ManySite.

Understanding the ManySite environment

ManySite is run as a consortium of semi-independent business units and subsidiaries that are free to choose their core technologies and suppliers for practically every function. The local country managers have a technical team that takes directional suggestions from a US-centric planning group. The implementation is left to local interpretation because of the great disparity between the wide area network cost structures for data carrier services and the extreme range of the quantity of users that must be supported in a given site's network. Despite the freedom to interpret network design parameters, the 100,000 users that are supported on the worldwide network can interact with a rich variety of applications regardless of whether their initial connection is via token-ring, Ethernet or dial-up networking.

Mirroring the experience of other corporations, the success of the installed network froze all earlier attempts to change more of the network to Ethernet. Nobody could show a compelling reason to change something that worked so well, and nobody could show a payoff for the expense of the change. Nevertheless, as the techno-political environment changed with the acquisition of more and more Internet-savvy customers, the danger of appearing to be married to a sunset technology became in itself a good reason to change.

ManySite had developed three different network topology standards. Like the Goldilocks fairy tale, these standards were simply known as “Small-Site”, “Medium-Site”, and “Large-Site and Manufacturing.” The site design guidelines for each size had been disseminated to the various country teams, and no specific feedback for monitoring the country's decision was required. ManySite's tradition of hiring high-quality personnel was the glue that kept these loose policies from trapping it in failed implementations. Every site had focused on availability, migration, cost management, and performance well enough that it was quite rare to have complaints from any sector of the user population.

Manufacturing sites had the highest availability requirements and therefore spent significantly more on redundant paths and configuring lightly loaded token-ring segments.

Unlike the corporation profiled in the first case study, almost all of the endpoint devices were token-ring connected using the ICS. ICS offers a structured cabling plan with tremendous scalability. Part of the ICS specification is 150-ohm shielded twisted pair (STP) wire. A major obstacle to the earlier migration initiatives had been the huge per-drop (wall jack to switch/hub port) expense of changing the STP wire and ICS connectors to the industry-standard unshielded twisted pair (UTP) with RJ-45 jacks. While the ever-smaller cost of Ethernet LAN switch ports each year had reduced the potential expense of migration, the costs associated with changing the cabling were stagnant and more than an order of magnitude higher than the port costs of the closet hardware.

Fortunately, ambitious testing by IBM firmly established that the STP could be re-used for Fast Ethernet with the added bonus of increased link distances (170 meters versus 100 meters) with no increase in bit error rates. Suddenly the cost of migration, while onerous, could be absorbed without invoking multi-year write-offs.

The dearth of international standards within the ManySite Corporation made some network analysts unsure that all of the drop distances could be guaranteed to fall within the generous 170-meter distance. This led some of them to proffer a standard without any Fast Ethernet connections for users. Despite the lingering doubts, the decision was made to forge ahead with a Fast Ethernet standard and to manage distance or speed exceptions as they were encountered.

ManySite immediately became a victim of its own efficiency as nearly 50% of the wiring closets were found to be too full of token-ring equipment to concurrently accommodate any other devices required for Ethernet. This forced the migration team to create temporary wiring closets to “seed” the Ethernet ports within a site, while the token-ring ports ceded their domination of the network.

Protocol issues were equally problematic despite the aggressive elimination of IPX that had occurred in recent years. There were still critical applications that depended on SNA and NetBIOS, while the routing of IP traffic within ManySite was managed with a mix of OSPF, RIP and BGP. Organizational issues and IP re-addressing nightmares sapped the momentum of the migration team every time a new business unit chimed in about how their business plan could not be inconvenienced by any changes to its far-flung server farms.

Months of meetings and several white papers finally ironed out the difficulties among the business units. As a result, network protocol rules were established for all of the ManySite Corporation. A sampling of the rules reveals severe restrictions such as all LU6.2 traffic must be bridged within a single campus, never crossing to another campus, and all LU2 traffic must be completely converted to TN3270. The new standards also demand that all native NetBIOS be either changed to NetBIOS-over-IP or confined to individual departmental LANs.

Target network design for ManySite Corporation

Server farm migration had begun years before the Ethernet connectivity for end users was ever considered. This is because token-ring bandwidth at the servers was quickly swamped by each new Web-based application. ManySite had deep experience with ATM connectivity and token-ring LAN Emulation, ATM Classical IP, FDDI, and a smattering of Fast Ethernet interfaces on servers. Each of those non-token-ring technologies had been applied to overcome the 16 Mbps ceiling on token-ring performance. As one subsidiary of ManySite had built new server farms as part of its foray into Web hosting, all of those hosting servers had Fast Ethernet connections on a Gigabit Ethernet campus backbone.

This experience at the hosting subsidiary had just the right combination of client visibility and corporate approval to become the linchpin of the Ethernet migration. Lessons learned by the hosting teams were quickly disseminated to the technical teams of other business units via a combination of white papers and informal internships. The new expertise brought home by the technical teams allowed these high performance individuals to bring a shared vision and experience to widely disparate situations. The overwhelming majority of migration work would be centered on end-user connectivity with very few servers involved in the change.

Each technical team used its experience at the hosting sites and its knowledge of the local topology to build and submit a customized migration plan for their site. The migration plans were then sorted and funded by a rank based on windows of opportunity, required completion dates, cost, available personnel, and of course, executive demand. The ManySite network management center was expanded to become the project management center for worldwide migration. Local technical teams were required to submit a detailed project plan that denoted port counts, IP addressing schemes, sequence of changes, and any WAN-related topology changes.

With so much attention on the core and edge of the network, user issues had taken a back seat. ManySite, through attrition, had replaced conventional desktop machines with portable laptop computers. Users had been issued token-ring PCMCIA adapters and were quite comfortable with removing them while traveling between sites. Now the help desk must be geared with expertise to handle the inevitable problems that come with four different supported operating systems, and two different network cards that can be swapped at will. Windows 98, Windows NT, Linux and Windows 2000 each presented a variety of symptoms depending on the state of the machine when the swap was made, and each symptom required a unique recovery procedure.

In contrast, equipment choices were remarkably easy despite the apparently complex relationships between ManySite's consulting divisions and several popular vendors of networking equipment. A single-vendor solution for ManySite was understood by everyone involved (except the vendors) to be completely untenable for political reasons. Despite the bleating from ManySite's equipment partners who each wanted 100% of the equipment contract, each vendor got to put the strongest part of its product line in the ManySite network.

This decision to mix brands of equipment greatly lengthened the “sandbox” certification time for each device. Interoperability pitfalls and command-line differences made many of the implementers wish for the politically impossible single-vendor design.

Migration staging for ManySite

New Ethernet LAN switches with Gigabit uplinks were purchased for the new user migration closets. A migration closet was created in three steps. First, a building was surveyed to find two directly adjacent qualifying token-ring wiring closets. Usually these were closets that were stacked directly atop one another in a multi-story building. Alternately it could be two closets that could be linked to many (more than 10) Ring-Out and Ring-In links between token-ring hubs. In either case, they'd qualify if one of the two closets was less than 50% full. Second, the adjacent closets were physically merged as much as possible to create an “empty” closet. Thirdly, the empty closet had its Ethernet equipment installed in it for the end-user stations. Server connections or router connections were also installed at this point, and basic connectivity tests were performed.

Although Fast Ethernet would eventually be the standard user interface, each user switch port would be initially configured to support the lower 10 Mbps Ethernet speed. The chosen vendor had to certify that their LAN switches could handle the MAC addressing table problem caused by the automatic pin loopback when an ICS connector is removed from its socket.

All of the new Ethernet LAN switches were L2 switches. Only the IP address values for the management interfaces had to be explicitly configured, because the DHCP servers at the server farm were already reachable via either token-ring or Ethernet devices. Once again, the multi-vendor environment became an extra load for the network operations staff as new spreadsheets of supported and interoperable microcode levels, and hardware revision levels were checked and tested.

Testing scenarios had to be created, reviewed by experienced team members to affirm their viability and usefulness, and finally run and documented. Since skill levels of people varied depending on the equipment involved, this consumed much more time than anyone in the migration chain wanted to tolerate. Test management was seen as a low value-add and a hindrance to a job that needed every hint of swiftness it could get. Although cooler heads prevailed, and the testing was completed, several people's efforts were not appropriately appreciated because it was “just testing”.

Test client machines belonging to the network migration team were attached to the migration closets as the first phase of simulated user migration. At this time, each of the recovery scenarios for the three Windows operating systems and Linux were documented. The help desk received a copy of these recovery scenarios, and even had a chance to practice them as their laptop computers were moved to the Ethernet migration closets.

Migrating ManySite users to Ethernet

As the migration closets were filled, it was now time to shift the users to Ethernet one closet at a time. The patchwork of departments and interdependent functions that shared a single wiring closet threatened to kill any hope of ever migrating anybody to Ethernet. Downtime was tantamount to destruction, and no manager wanted to authorize the devastation of their department, even if it was for their ultimate good. When interviewing managers about the downtime they could tolerate as their conversion cables were installed and workstations rebooted, so many efficiency problems surfaced that the migration teams made a decision to establish a weekend-only policy for migration (except for manufacturing sites).

As mentioned before, the server farms had separately handled their migration work well in advance of the end-user migration. Although this simplified the overall migration, it created a new network performance problem: Routers that had previously been handling shared segments of 60 token-ring users now were getting switched Fast Ethernet connections that had switched 10 Mbps Ethernet users. The size of the traffic bursts on the Fast Ethernet connections caused intermittent packet loss problems even though there was no change in the number of users. The peak-to-average traffic ratios that the routers encountered had changed as the increased bandwidth exposed quite a large amount of pent-up demand at the user workstations. The higher peaks required more routing capacity all along the return path from the servers to the users in order to prevent future packet loss problems.

All of the data collection and analysis that confirmed the peak-to-average hypothesis was done with a new set of network management tools that had been acquired as part of the migration. The network migration teams were pleased to see that the work they had put into interoperability testing for network management had an immediate payoff by allowing them to quickly assess and respond to a routing capacity problem. This quick and decisive response to a problem also saved the migration from getting a “bad name” within the corporation. Momentum for the migration was maintained, and this problem boosted the importance of testing and network management.

Manufacturing sites had a unique set of problems. While the unit of change at research, marketing, and consulting sites was the wiring closet, the high availability requirements at the manufacturing sites did not allow for the weekends-only strategy to be used. Instead of that, each manufacturing site was saddled with the agonizing choices of which component or device lines could be stopped and changed.

For that reason, all of the manufacturing facilities scheduled their network changes around the major changes in the products being built. These change windows were months apart, and in other cases were as long as a year apart. Clearly, something had to be done to make the migration possible within a calendar year, but an occasion for change had to be artificially created.

The key event that allowed the manufacturing lines to be migrated actually was an exogenous event. The product businesses had come to a state of having too much capacity because of changes in the demand patterns across the world. This change in demand meant that manufacturing would shift load from one country to another temporarily so that equipment and connectivity changes could be made. Managers didn't have to make the awful downtime decision, but instead were able to gracefully schedule their production movements in a way that silenced the lines in a disciplined and predictable fashion.

As the manufacturing changes proceeded, problems with the staging of the disposal of decommissioned token-ring equipment became an issue. While the small branch offices and outlying offices were able to simply use the local trash disposal for their old boxes, manufacturing sites and larger development sites had rules and a procedure for disposal. These disposal procedures became a large expense that required management attention and long hours of tracking to make sure that neither the accountants nor the EPA (Environmental Protection Agency) would have any protests that would arise in years to come.

Conclusion

As the migration continues smoothly across the world at ManySite, the expectation of the users is that it'll be no more intrusive than a light bulb being changed in an office. That is the mark of a process that is properly tuned and managed.

Token-Ring to Ethernet Migration

Redpaper

An insight into technical areas that need to be considered when migrating from a token-ring network to Ethernet

An understanding of a step-by-step process to follow in planning a network migration project

Practical considerations for migration based on real-life case studies

This Redpaper, based on experience gained from migration projects with many customers and clients, offers considerations, guidelines, and suggestions for approaching a token-ring to Ethernet migration project. It highlights many of the technical considerations that must be addressed in a migration project, offering thoughts for how to stage such a migration, and provides a high-level, step-by-step outline for a sample migration. This Redpaper includes case studies that provide real-life examples of migration projects.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks