8371 Networking Multilayer Ethernet Switch and the 8371 (MLS) Modules for the IBM 8265 Nways ATM Switch

Lunch & Learn

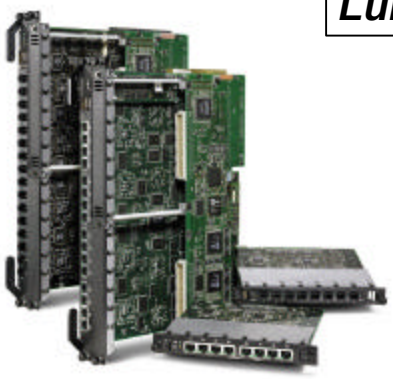

Hugh S. Fish

Networking education

8371 Multilayer Ethernet Switch

Notes:

The correct name for the products are:

The IBM Networking Multilayer Ethernet Switch - 8371 Model A16
The 8371 Multilayer Ethernet Switch Modules for IBM 8265 Nways ATM Switch

The author of this document is Hugh S. Fish of IBM NHD Education

I wish to thank James Forbes and Peter Russell for their assistance in creating this document. Much of this material was obtained from documents created by Cedell Alexander of the 8371 Development Team.

8371 Networking Multilayer Ethernet Switch and the 8371 (MLS) Modules for the IBM 8265 Nways ATM Switch

- **Topics covered:**
  - 8371 MLS Introduction
  - 8371 MLS Switch Hardware Components
    - 8371 MLS Switch
    - 8371 MLS Blades
    - 8371 MLS UFCs
  - 8371 MLS Switch Hardware Architecture
  - 8371 MLS Switch Software Components
    - VLANs
    - MPOA Client
    - Self Learning IP
    - VLANs, LECs and Uplinks
    - Network Management
    - Configuration & Console Interfaces

8371 Multilayer Ethernet Switch

Notes:

8371 Introduction

- **A Family of Ethernet switches**
  - The IBM Networking Multilayer Ethernet Switch - 8371-A16
  - 8371 Multilayer Ethernet Switch Modules for the IBM 8265 ATM Switch
- **Common attributes:**
  - 10/100 Base-TX and 100 Base-FX plug-in feature cards
  - ATM uplinks - 2 - ATM OC-3 in standalone switch and OC-12 backplane connection in 8265
  - 5 Gbps (10 Gbps full duplex) Layer-2 and Layer-3 hardware switching engine
  - Built-in software that supports LANE V2 and MPOA Clients with the ATM uplinks
  - Provides local wire-speed IP routing between ports on the same switch module and MPOA one-hop routing over an ATM backbone - in conjunction with the MSS Server
  - 8371 switches all traffic (Transparent Bridges) - no routing in this Release
  - Self Learning IP for integration with current routers - standalone 8371 only
  - MSS Dynamic Protocol VLANs
  - IP Multicast IGMP Snooping

Networking education

8371 Multilayer Ethernet Switch

Notes:

The hardware and software are similar in architecture and functions, the biggest differences are in packaging (i.e., no ATM UFC in the Blade, no base 100Base-FX in the standalone, etc.). The 8371 Switch has 16 switched ports in the chassis with it's own power supply. The UFCs are optional in both the blade and the box.

The 8265 has had the capability of having either the 8271 or 8272 blades installed which has made it a "Hybrid Switch" but the 8271 and 8272 blades were 8260 blades with very limited port density due to the four blade restriction of 8260 blades in 8265s. So the 8371 blades are the first true non-ATM media that can be installed in the 8265 that give it meaningful port density and native 8265 support, ie., 622 Mbps to the backplane.

MPOA and LANE V2 (LANE V2 is a pre-req for MPOA) are included in the 8371 MLS software. MPOA is automatically pre-configured when there is an ATM uplink in the 8371 Module and Blade. IP Autolearn is automatically pre-configured when there is no ATM uplink.

The 8371 switch is capable of 5 Gbps half-duplex or 10 Gbps Full-duplex which enables non-blocking operation in all configurations.

8371 Multilayer Ethernet Switch Family

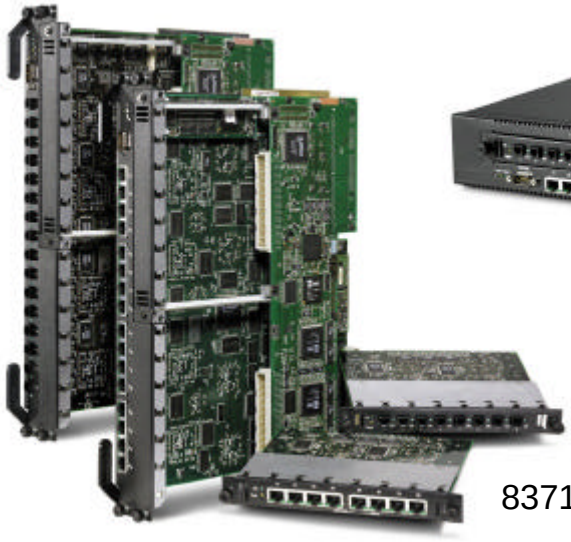

8371 MLS Blades for 8265

8371-A16 Ethernet Switch

8371 Multilayer Ethernet Switch

Notes:

Universal Feature Cards (UFC) are shown next to the 8265 Blades.

Pictured are the 16-port 10/100 Base-TX Module, the 16-port 100 Base-FX Module, the 8-port 10/100 Base-TX UFC, the 8-port 100 Base-FX UFC and the standalone 8371 Model A16 Switch. The UFC installed in the right-hand slot of the 8371 shown here is the two-port ATM OC-3 UFC.

8371 MLS Switch
Hardware Components

Networking education

8371 Multilayer Ethernet Switch

Notes:

8371 MLS Switch - Model A16

- A 16-port 10/100 Base-TX switch expandable to 24 or 32 ports
- All ports can be half or full duplex, auto-negotiated or configured
- All 10/100 Base-TX ports can operate at 10 Mbps or 100 Mbps - auto-negotiated or configured
- A 5 Gbps (10 Gbps full duplex), non-blocking switch fabric
- Two optional feature slots for additional Fast Ethernet and/or ATM OC-3 ports
- A built-in auto-ranging power supply and three cooling fans
- Standard RS-232 Serial Port for configuration

Networking education

8371 Multilayer Ethernet Switch

Notes:

With the 16 base ports, the switch can be configured for the following:

- 16-10/100 Base TX ports
- 24-10/100 Base TX ports
- 32-10/100 Base TX ports
- 16-10/100 Base TX and 8-100 Base FX ports
- 24-10/100 Base TX and 8-100 Base FX ports
- 16 10/100 Base TX and 16-100 Base FX ports
- 16-10/100 Base TX, 8-100 Base FX and 2 ATM ports
- 24-10/100 Base TX and 2 ATM ports

Only one ATM Feature card is configurable

Throughput for the 8371 Switch and 8371 MLS Blades for the 8265 is 4.7 Mpps (Million packets per second)

8371 MLS 8265 Blades

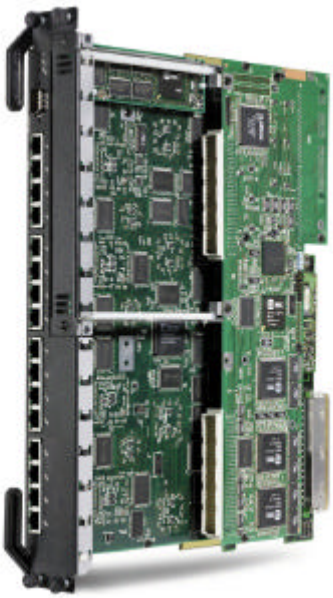

16-port 10/100 TX Blade

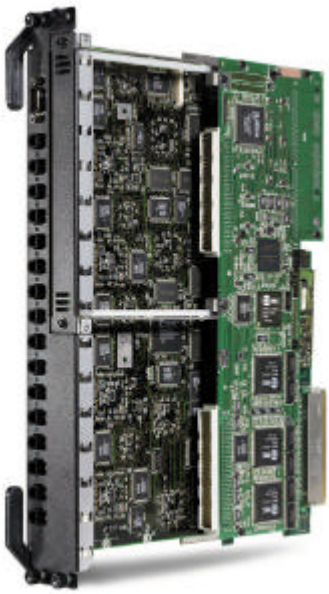

16 port 100 FX Blade

Networking education
Version 1.0 © Copyright IBM Corp. 1998

8371 Multilayer Ethernet Switch

Notes:

Notice that the blades have only one front panel UFC but two connectors for UFCs. The 622 Mbps module is installed on the blade version but it is backplane connected so it does not appear on the front of the blade.

The copper or fiber versions of the UFCs can be installed in the one UFC slot on either the copper or fiber blades. Mix and match in any combination.

8371 MLS Blades - Hardware Overview

- **Two slot wide blades - can be installed in any available slots in 8265**
- **Two versions**
  - 16-port 8265 Blade with 1 Universal Feature Card (UFC) slot
    - 10/100 Mbps TX Ethernet ports (RJ-45)
  - 16-port 8265 Blade with 1 Universal Feature Card (UFC) slot
    - 100 Mbps FX Ethernet ports (Mini MT-RJ Fiber connectors)
- **All TX (copper) ports can operate at 10 Mbps or 100 Mbps - Auto-negotiated or configurable**
- **All ports can be full or half duplex - Auto-negotiated or configurable**
- **Built-in OC-12 (622 Mbps) backplane connection**
- **Standard RS-232 Serial Port for configuration**
- **Same hardware and software as 8371 standalone switch**
- **Hot Swappable - if swapped card is of same type**

8371 Multilayer Ethernet Switch

Notes:

16-port 10/100 Base-TX Module	**FC 6616**	**PN 90G0500**
16-port 100Base-FX Module	**FC 6617**	**PN 44L2461**
8-port 10/100 Base-TX UFC	**FC 6626**	**PN 90G0530**
8-port 100Base-FX UFC	**FC 6627**	**PN 35H9203**

8371 MLS Modules in an 8265

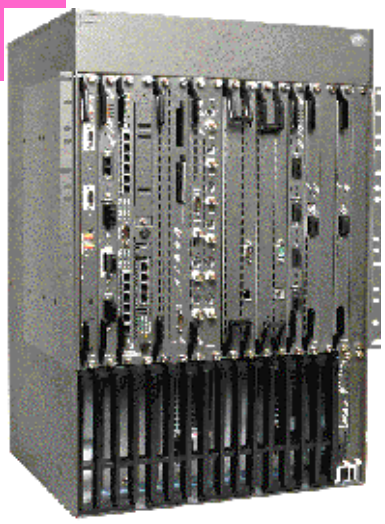

- Each 8371 blade is equivalent to a standalone 8371
- Add additional blades for more ports/switches
- Mix and match with other ATM 8265 modules
- Will need ATM ports or MSS blade for MPOA solution

8371 Multilayer Ethernet Switch

Notes:

Theoretically, up to 168 ports can be installed in the 8265. In this case, blades would be installed in positions 1, 3, 5, 7, 12, 14, and 16. There is no restriction on which slots the blades can plug into except that slots 9 and 10 are reserved for the CPSW and slot 11 has no ATM backplane connections for a non-CPSW blade.

However, for communication to another 8265, 8260 or 8285, ATM ports would be required which would decrease the number of 8371 blades that can be installed. An MSS module or external 8210 would also require an available slot where an MSS blade or ATM blade could be installed. An MSS blade or external connection to an MSS Server is necessary to provide the MPOA Server capability built-in and automatically pre-configured by the 8371 MLS blades.

8371 blades in the 8265 require that the 8265 be at operational code level 4.1.x or later.

8371 MLS UFCs

2-port ATM OC-3 UFC

8-port 10/100 TX UFC

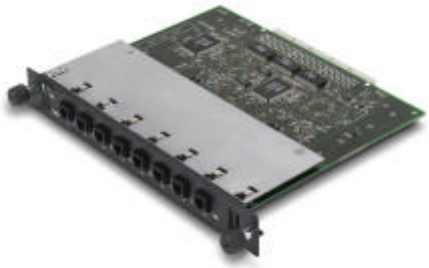

8-port 100 FX UFC

Networking education

8371 Multilayer Ethernet Switch

Notes:

The 2-port ATM OC-3 UFC is available for the 8371 standalone box but it is not available for the 8371 blades since there is an embedded OC-12 uplink card built into the blade version.

8371 MLS UFCs - Hardware Overview

- **Three Universal Feature Cards (UFC)**
  - 2-port ATM OC-3 - 8371 Standalone only - SC connectors
  - 8-port 10/100 Mbps TX Ethernet ports (RJ-45)
  - 8-port 100 Mbps FX Ethernet ports (Mini MT-RJ Fiber connectors)
- **Ethernet UFC's can be installed in either blade**
- **Same capability as fixed ports**
- **32 Fast Ethernet ports maximum on 8371**
- **24 Fast Ethernet ports maximum per 8265 blade**
- **Blades have built-in ATM OC-12 adapter, therefore, no ATM UFC is supported**

8371 Multilayer Ethernet Switch

Notes:

The MT-RJ connector is new. It is a dual strand fiber cable with a single cladding encasing the two strands. The connector is very small which gives the vendor the capability of high port density which has been a problem with SC, ST and MIC connectors in the past.

By installing a Copper RJ-45 MLS Blade and a Fiber MT-RJ UFC, this combination would be a good solution for attaching the copper ports to the user workstations and the fiber ports to servers in a closet or on a raised floor. All ports are capable of the same function, regardless of media type, and can be used with both varieties of VLAN capability.

Cabling Requirements

- **RJ-45 10Base-T connections**
  - UTP/STP category 3, 4, and 5
  - 100 Meters
- **RJ-45 100Base-TX connections**
  - UTP category 5 only
  - 100 Meters
- **MT-RJ 100Base-FX**
  - Multimode Fiber
  - 412 Meters - Half-duplex
  - 2000 Meters - Full-duplex
- **SC connectors for ATM155 Feature**
  - Multimode Fiber - 62.5 micron
  - 2000 Meters
- **Management Port**
  - Null-modem adapter to standard serial cable

8371 Multilayer Ethernet Switch

Notes:

Shipping Package contains:

- CD-ROM containing documentation
- An 8371-A16 Networking Ethernet Switch or 8265-L3S
- A rack mounting kit (not included with 8265 blade)
- Operations Reference Card
- A card tray to hold Reference Card (not included with 8265 blade)
- Safety Manual

8371 MLS Modules
Hardware Architecture

Networking education

8371 Multilayer Ethernet Switch

Notes:

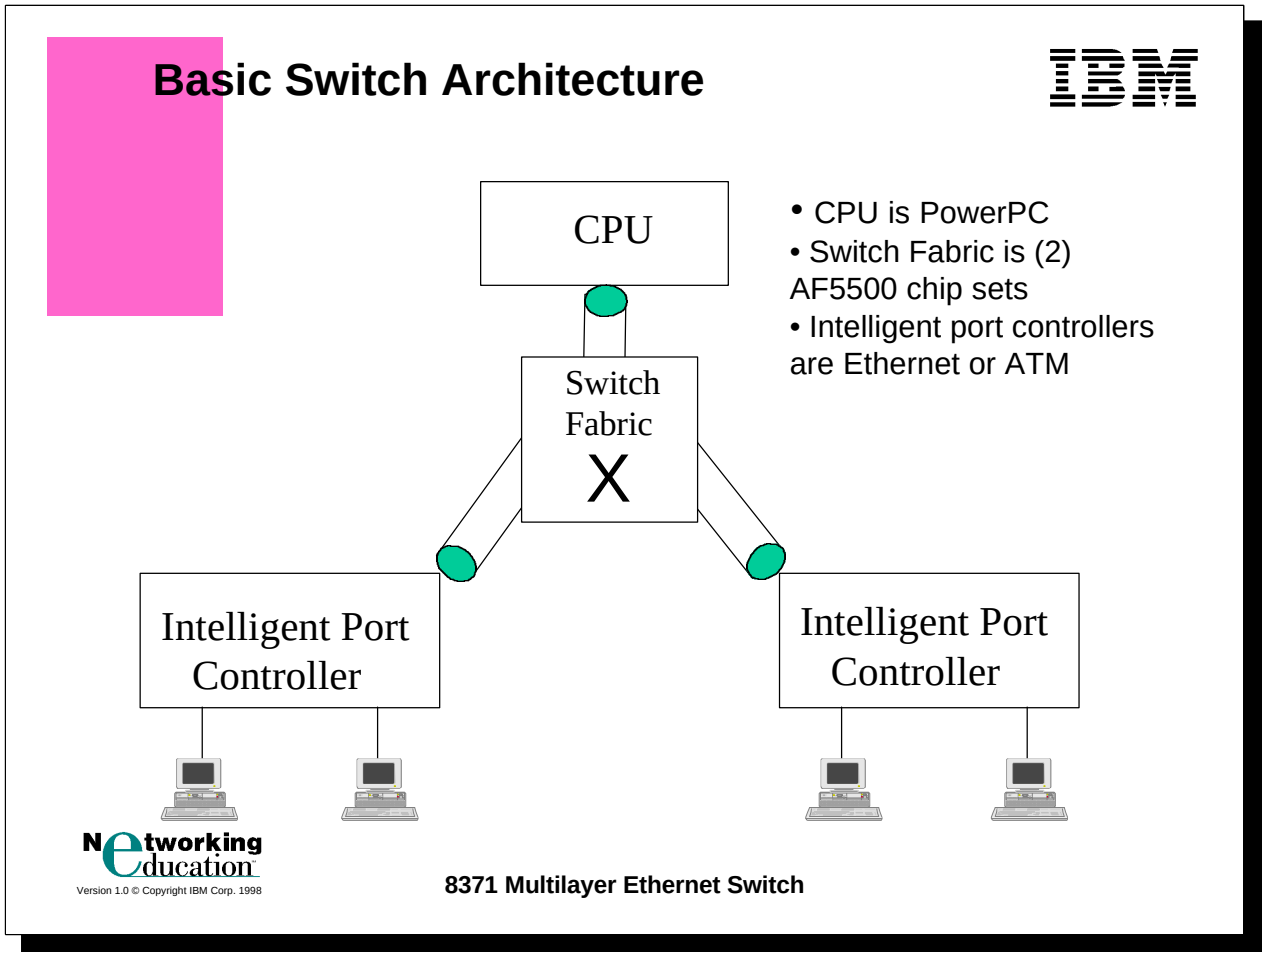

Notes:

CPU - 50 MHz Motorola 860T PowerPC

Switch Fabric - consists of two nP5500 Network Processors that manage a shared pool of packet/cell memory.

Intelligent Port Controllers - Implemented with EPIFs (Ethernet Port Interface Bitstream Processors) for LAN ports and PIF2 (Port Interface Interface Processors) for ATM connections. Each PIF2 can support one OC-12 channel or up to 4 OC-3 channels. The EPIF is similar in that it provides connectivity between the switch core and Ethernet physical layers. However, the EPIF differs from the PIF2 by providing additional functionality in the form of a microcontroller and the hardware search engine. Each EPIF can support four 10/100 Ethernet channels.

Basic Switch Architecture (continued)

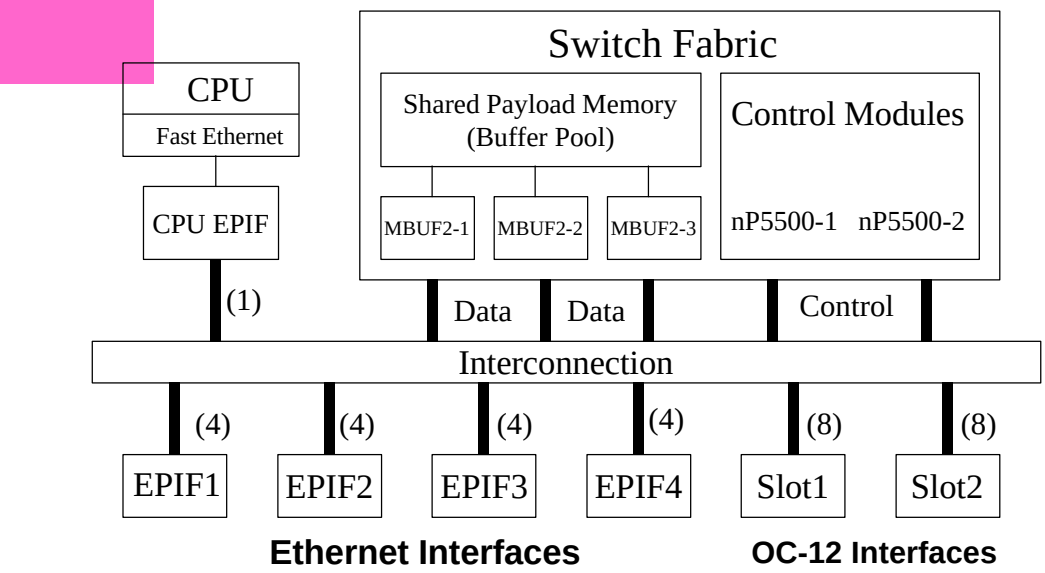

8371 Multilayer Ethernet Switch

Notes:

Memory Access Buffer 2 (MBUF2) chips provide the interface between the shared memory and the bitstream processors. Three MBUF2s are employed.

A Fast Ethernet Controller is attached to the CPU subsystem. The Fast Ethernet controller is used to provide a 100 Mbps interface between the CPU and the switch fabric via an EPIF.

The nP5500 is fundamentally an ATM switch that can operate in both cell and packet modes. Conceptually, each nP5500 supports 16 OC-3 switch legs to bitstream processors. Additionally, the two nP5500s also support one CPU port. The Parenthesized digits above indicate the number of 155 Mbps switch legs devoted to the associated EPIF/slot

Basic Switch Architecture (continued)

- **Basic Switch CPU**
  - 32 Mbytes of CPU DRAM for program and data storage
  - 12 or 16 Mbytes of Flash Memory
    - Two code images
    - Two Configuration files
  - 3 Mbytes of switch payload memory
  - Sufficient nP5500 control memory to support 4000 ATM VCCs
  - 16,000 MAC addresses at box level
- **LAN EFIPs - each**
  - 512 Kbytes of memory
    - Enough for 8000 database entries - Layer-2 and Layer-3 entries
  - 2 Mbytes backbone access memory
  - Four 155 Mbps switch legs dedicated to each EPIF
  - Four slicers which divide frames into AAL5 cells for transmission into the switch fabric
  - One search engine
- **Feature card slots**
  - Eight 155 Mbps switch legs dedicated to each slot

8371 Multilayer Ethernet Switch

Notes:

The search engine on the EPIF implements Patricia Tree databases in hardware. The database contains Layer-2 and Layer-3 addresses. Patricia trees are ideal for efficiently performing both exact and longest-prefix matches. The search machine executes commands issued by the microcontroller, such as Insert, Delete, and Search. Searches are performed on keys that are up to 48 bits in length.

On the 8371, each LAN EPIF has 512 Kbytes of external memory, which is enough for approximately 8000 database entries. Note that those 8000 entries are used for both Layer-2 and Layer-3 functions, and that the entries are shared among the four Ethernet ports.

The EPIF instructions reside in a 4 Kbyte on-chip program store, which is enough for 1000 instructions. The EPIF cannot execute out of external memory; however, an applet capability is provided that can ease space restraints. Applets are segments of code that can be downloaded (by the CPU) and executed by the EPIF. After an applet completes execution, the program space that it occupied is freed. Thus, the on-chip program store can be partitioned into two areas: one for resident functions and one that is reserved for applet execution.

Although most of the preceding hardware information is equally applicable to both the standalone and switch versions of the 8371, there are a few differences. The major difference is that the blade version only has one slot for user-installed feature cards. Slot 2 of the blade is occupied by an ATM card. This card supplies 622 Mbps connection to the backplane of the 8265 ATM Switch.

8371 MLS Modules
Software Components

Networking education

8371 Multilayer Ethernet Switch

Notes:

8371 MLS Modules Software Overview

- **Based on Nways Common Code**
  - Already shipping in MSS Server, MSS Client, 22xx Family
  - New functions for 8371 MLS
    - IP Autolearn
    - MPOA Client support for multiple ATM Adapters (8371 Standalone)
    - RMON-Lite
    - Various Device Drivers
- **Re-used and adapted functions**
  - ATM Signaling, Ethernet LANE and ILMI
  - Transparent Bridging and Spanning Tree
  - Protocol Virtual LANs for IP, IPX, NetBIOS
  - IP Multicast VLANs
  - MPOA Client for IP and IPX
  - IP Host functions - TELNET and PING
  - Command Line and Web Browser Interfaces
  - SNMP support and MIBs
  - Box Services such as code image and configuration management

8371 Multilayer Ethernet Switch

Notes:

Nways Common Code is the IBM code that was developed for a family of products that have common functions. Those products are the MSS Server (both 826x blade and 8210 versions), the 22xx Router products, the 8270 MSS Client, the 3746 MAE (Multi-Access Enclosure) and now the 8371. Common Code functions are enabled or disabled depending on which hardware platform it resides in.

VLANs

- **Two types supported:**
  - **Protocol VLAN (PVLAN)**
    - Similar to MSS Server and MSS Client (8270)
    - Uses Dynamic Protocol Filtering (DPF)
    - Supports IP, IPX, and NetBIOS
    - A single NetBIOS PVLAN is supported
    - Runs on central CPU rather than EPIFs so frames are sent to CPU for processing
    - If PVLAN is split across two or more switches or blades, static configuration of ports may be necessary
  - **IP Multicast VLAN**
    - Identical to MSS Server function
    - Packets only forwarded to ports that are involved with the Multicast group
- **Both VLANs types Plug-and-Play**
  - Add static definitions only when necessary
    - Such as adding ports for "quiet" devices, adding ports for PVLANs, changed timer defaults, etc.

8371 Multilayer Ethernet Switch

Notes:

A VLAN is a logical group of hosts that form a broadcast domain. The logical grouping is independent of the physical network topology with VLAN membership governed by a set of rules or policies. VLANs have several advantageous characteristics, three of which are listed below.

1. Broadcast domains are controlled. The scope of a broadcast frame originating within a VLAN is limited to LAN (and ELAN) segments containing hosts that are members of the VLAN. Most protocols use broadcasts for address resolution; therefore, if the scope of the broadcast is limited, the scope of host-to-host communication is also limited.

2. Intra-VLAN communication is typically very efficient, with VLAN members communicating directly over a switched hardware path.

3. Host moves, adds and changes are simplified since VLAN membership is independent of physical topology. No reconfiguration is required to retain VLAN memberships when stations move to new physical locations. Similarly, no wiring modifications are needed to move stations from one VLAN to another.

PVLAN Example

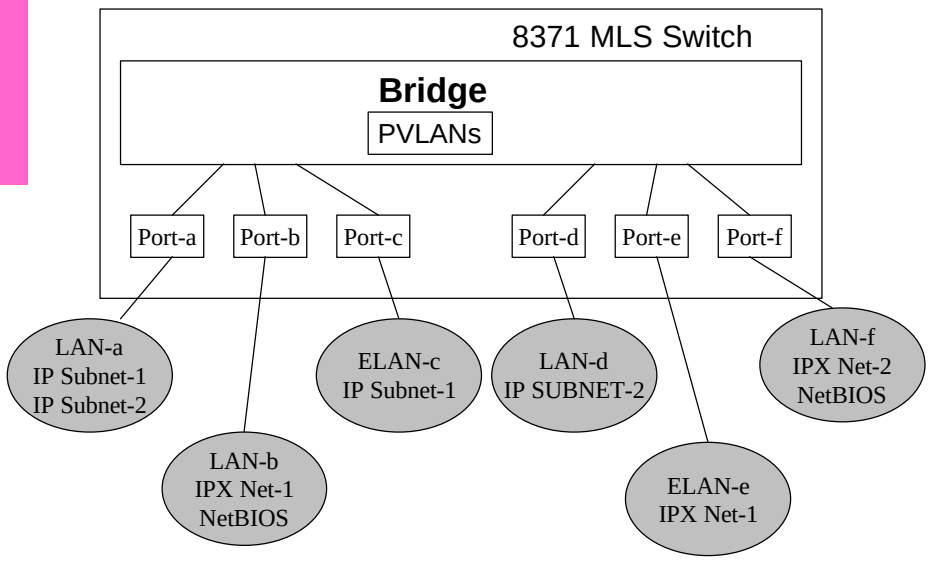

- IP Subnet-1 traffic is restricted to Ports a & c
- IPX Net-1 traffic is restricted to Ports b & e
- etc.

8371 Multilayer Ethernet Switch

Notes:

Membership in a PVLAN is determined by network protocol or the combination of protocol and network address (e.g., a particular IP subnet or IPX network). The 8371 creates broadcast domains by dynamically learning the set of PVLANs active on each LAN or ELAN segment. Broadcasts are then limited to the segments containing stations that are members of the PVLAN.

The 8371 utilizes a technique called *Dynamic Protocol Filtering* (DPF) to dynamically partition the bridged network into several PVLANs. DPF then uses the PVLANs to limit the scope of broadcast/multicast frames that are normally forwarded over all active bridge ports.

DPF monitors the broadcast/multicast traffic received over each bridge port to learn the protocols and subnets being used on that segment. In Release 1.0 of the 8371, the user controls the set of PVLANs to be managed by DPF via configuration. DPF manages the forwarding domain for each configured PVLAN, where the forwarding domain is a subset of bridge ports on which traffic for that PVLAN is being received. Transmission of broadcast/multicast frames for a particular PVLAN is limited to the forwarding domain for that PVLAN. Thus, the scope of broadcast/multicast packets for a particular protocol/subnet is reduced to those segments that are actually utilizing the protocol/subnet. More specifically, IP ARP traffic only travels over LANs/ELANs associated with that IP subnet, IPX broadcast/multicast traffic only travels over LANs/ELANs associated with that IPX network, and NetBIOS broadcast/multicast traffic only travels over LANs/ELANs using NetBIOS.

In the absence of traffic, a port's membership in a PVLAN forwarding domain will age out after a configurable length of time. The default aging times are 10,000 minutes for IP PVLANs, 10 minutes are IPX PVLANs, and 5000 minutes for NetBIOS. Alternatively, aging times can be set to infinity if the administrator chooses, so that a port's membership in a PVLAN is only learned once.

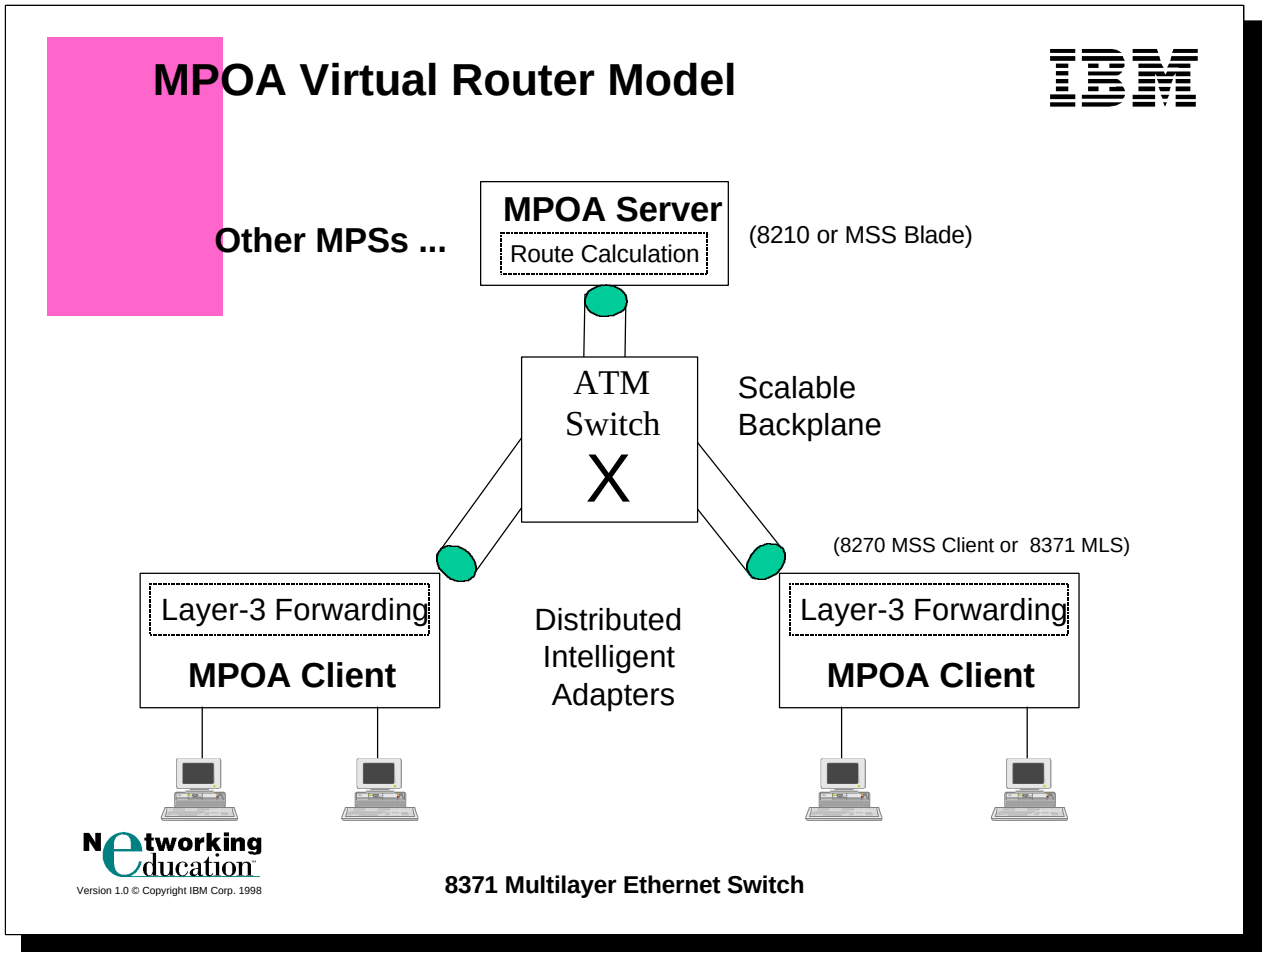

Notes:

The *Multiprotocol Over ATM* (MPOA) specification standardizes the virtual router model depicted in the above chart. The virtual router is implemented with client-server protocols, where the MPOA clients issue requests to MPOA Servers (MPSs). Internetworking layer route calculation and forwarding functions are partitioned in the model, with MPSs performing the route calculations, MPOA Clients (MPCs acting as distributed intelligent adapters devoted to high-speed forwarding, and the ATM network providing a scaleable backplane for the virtual router.

The virtual router model provides both scalability and manageability benefits, with improved manageability being the bottom-line value. The management benefits are based on the scalability attributes of the model, which allow performance to be increased in an incremental, vendor-independent manner. For example, cumulative forwarding capability can be increased by simply adding more MPCs. Similarly, backplane throughput can be increased with the addition of ATM network capacity.

The complexities of administering multiple distributed edge routers is significantly reduced with MPOA which includes auto-configuration and device discovery protocols that minimize device-specific configuration. Furthermore, the number of devices participating in routing topology protocols is reduced in the virtual router model, which is a scalability advantage that reduces edge device complexity.

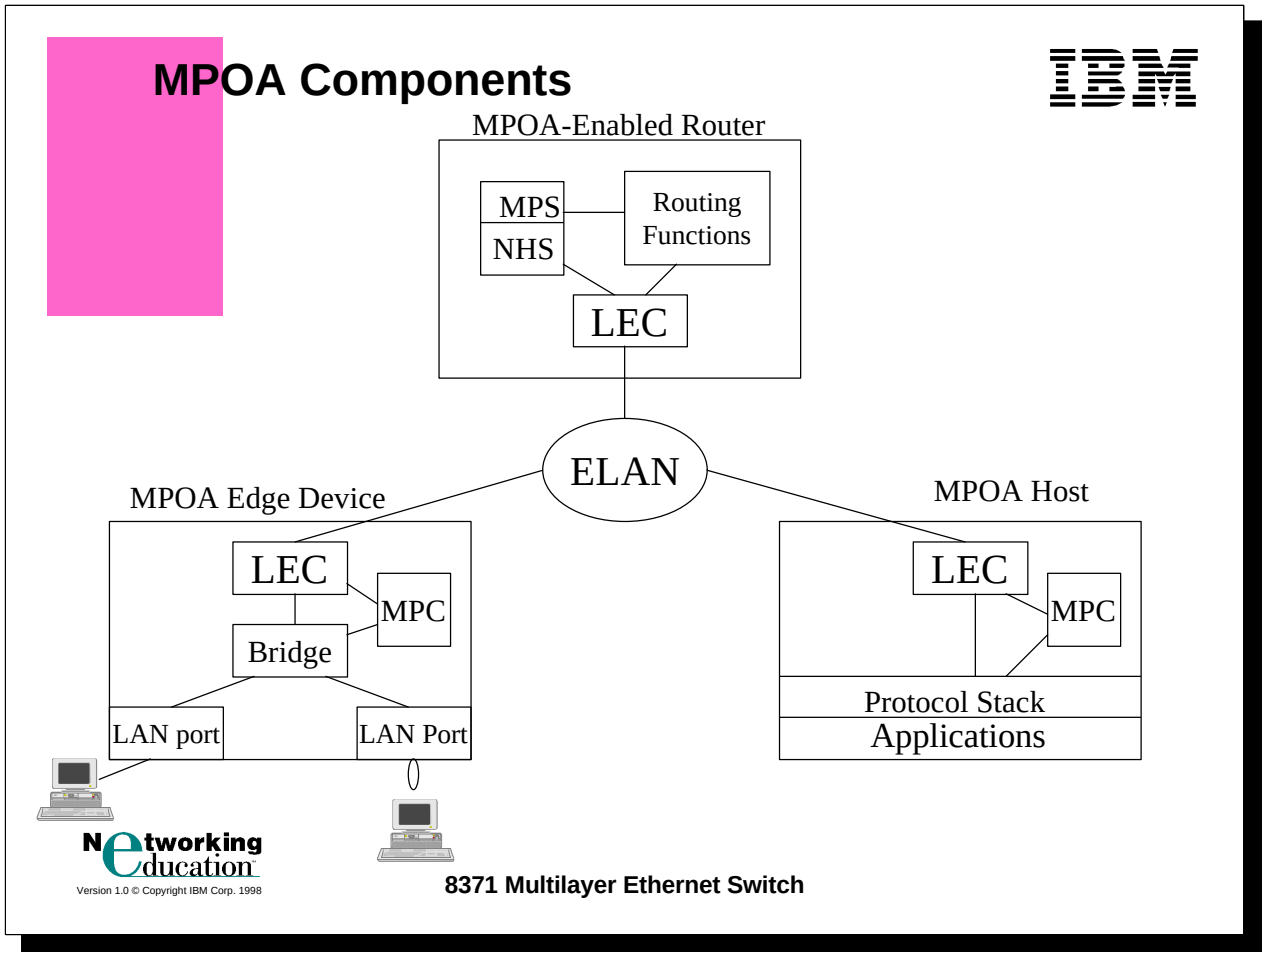

Notes:

MPOA is able to operate in a relatively simple manner by building upon standard networking technologies such as bridging, LAN Emulation and Next Hop Resolution Protocol (NHRP). As illustrated above, there are three basic types of MPOA-enabled devices. MPSs are co-located with router functions and a NHRP Server (NHS), while MPCs reside in *MPOA Hosts* or *MPOA Edge Devices. All of the devices include a LAN Emulation Client that provides default-path interconnection.* MPOA Edge Devices like the 8371, also include a bridge component for representing stations residing on Ethernet or Token-Ring segments.

The functions performed by a MPC in a MPOA Host are very similar to those performed by a MPC in a MPOA Edge Device, with the main difference being that the MPC in an edge device can represent several hosts. The MPC's primary objective is the same in both cases: establishing *shortcut VCCs* and forwarding intersubnet traffic over these VCCs to improve system performance.

MPOA relies on LAN Emulation for three important functions: auto-configuration, dynamic device discovery, and intranet/default-routed-path connectivity. Auto-configuration and dynamic device discovery are features designed to simplify the management of MPOA systems.

1. Auto-configuration allows MPOA configuration parameters to be stored and distributed from a centralized repository, namely the LECS. MPOA devices can then obtain their configuration parameters from the LECS while they are being initialized, minimizing configuration at individual devices and helping to create a *plug-and-play* environment.

(continued on next page)

MPOA - (continued)

2. Dynamic device discovery is the other piece of the plug-and-play story. MPOA devices are not configured with information about neighboring MPOA entities. Instead, MPOA devices dynamically learn about neighboring components through the discovery protocol. The discovery protocol mechanism finds the required information by attaching information to LANE control messages. This requires LANE V2. Fortunately, not all of the devices on the Emulated LAN need to upgrade to LANE V2, only the LECs in MPOA devices and the LANE Service components.

3. When there is no shortcut available, intrasubnet traffic is bridged to the MPS router over an Emulated LAN. Since most MPOA edge devices include LAN switching hardware capabilities, intrasubnet traffic is handled very efficiently with end-to-end switching. This use of bridging, coupled with dynamic device discovery, enables the MPC to be independent of router technology and redundancy protocols, while maintaining the change management benefits provided by Virtual LANs (e.g., allowing a station to be moved from a segment behind one MPC to a segment behind another MPC without any reconfiguration).

Notes:

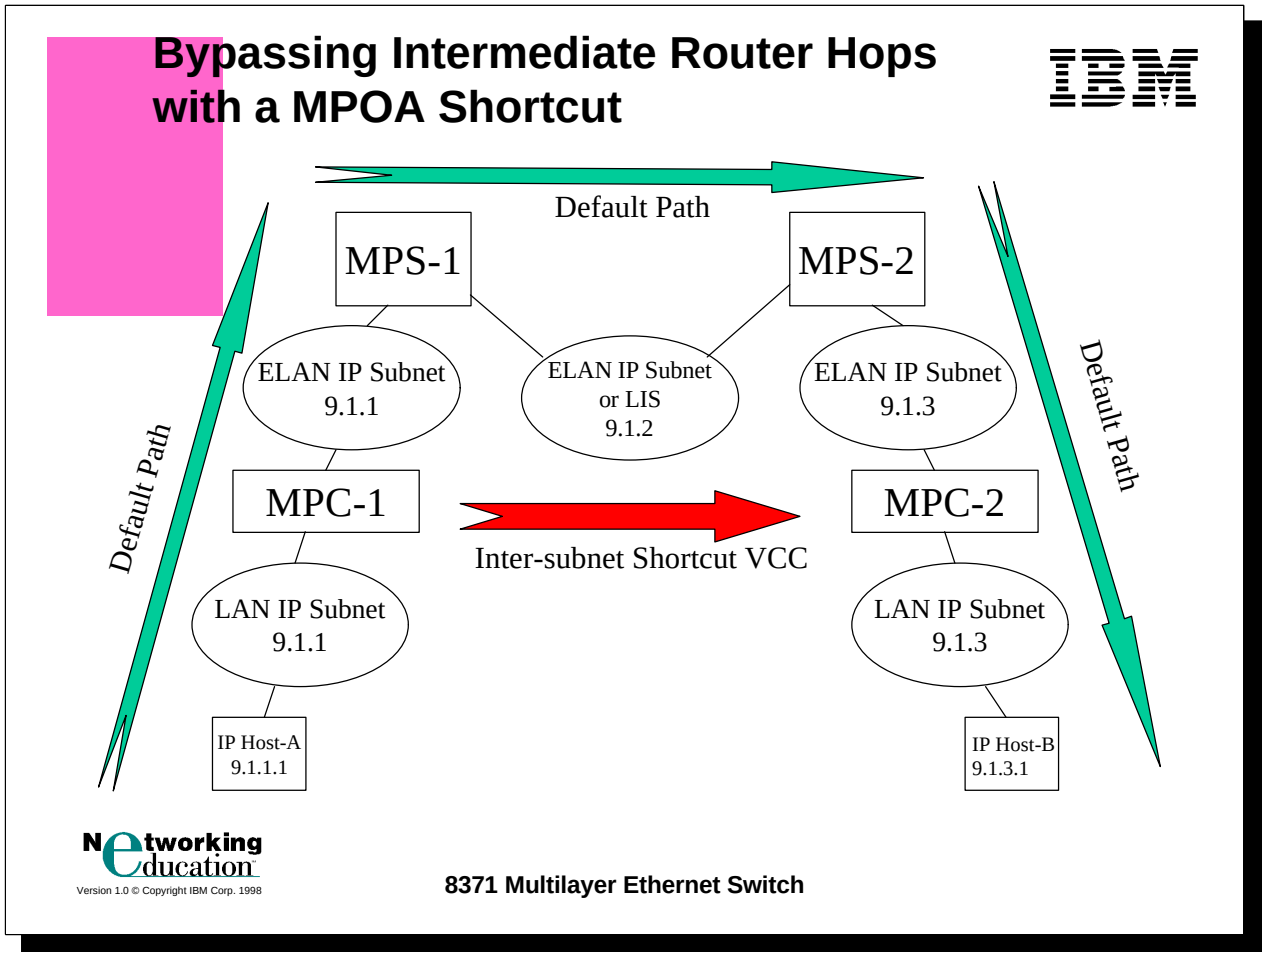

Notes:

Establishing Shortcuts:

The MPC is responsible for initiating shortcut establishments. The steps are as follows:

1. The MPC discovers the MAC address of the MPS routers along with the corresponding Control ATM address.

2. The MPC monitors the flow of traffic destined for these MPS MAC addresses. When the rate of traffic to a particular internetworking-layer destination address exceeds the configured threshold (10 frames/sec), the MPC initiates shortcut establishment by sending a MPOA Resolution Request to the associated MPS.

3. Assume that IP Host-A is transmitting to IP Host-B and that MPC-1 has determined that the flow exceeds the threshold for establishing a shortcut.

4. MPC1 then sends a MPOA Resolution Request for Host-Bs IP address to MPS-1.

5. Upon receiving the request, MPS-1 converts the MPOA Resolution Request to a NHRP Resolution Request, and forwards the converted request along the routed path to MPS-2.

6. MPS-2 responds to the NHRP Resolution Request because it serves IP Host-B (I.e., a MPS serves destinations that reside on subnets local to the MPS).

(continued on next page)

MPOA Shortcuts (continued)

7. Before responding to the resolution request, MPS-2 imposes an egress cache entry on MPC-2. MPS-2 imposes the cache entry because it recognizes, through the discovery protocol, that the destination MAC address is associated with a MPC.

8. MPC-2 records the cache entry and responds with a MPOA Cache Imposition Reply that includes an ATM address that can be used to set up a shortcut VCC for traffic destined to IP Host-B.

9. MPS-2 inserts the ATM address into the NHRP Resolution Reply that is sent back to MPS1.

10. MPS-1 transforms the NHRP Resolution Reply into a MPOA Resolution Reply that is sent to MPC-1.

11. MPC-1 then uses the ATM address to set up a shortcut VCC to MPC-2

12. After the VCC is established, traffic destined for IP Host-B is transmitted over the shortcut VCC, bypassing the intermediate routing hops at MPS-1 and MPS-2.

13. When MPC-2 receives the frames over the shortcut VCC, it delivers the frames to IP Host-B as if they were received via the routed path from MPS-2.

Notes:

Router Accelerator - Standalone 8371

- **What is a Router accelerator**
  - Network device placed between a router and it's associated LAN interfaces
  - Distributes layer-3 forwarding functions out to LAN switches
  - Similar to MPOA but it's for Ethernet - not based on LANE
  - Learns topology information by snooping on contents of selected packets - uses ***Self Learning IP***
- **Identifies routers attached to ports on switch**
  - Builds IP forwarding table
  - Offloads the IP forwarding functions from the attached routers
- **Realize same benefits in LAN environment as MPOA does in ATM environment - router offload**

8371 Multilayer Ethernet Switch

Notes:

The Self Learning IP feature is similar to MPOA in a number of ways. Both are Router Accelerator techniques that enable scalable solutions by distributing layer-3 forwarding functions out to LAN switches. Self Learning IP also relies on many of the same implementation mechanisms as MPOA.

One way that the two features differ is that Self Learning IP is for non-ATM environments (i.e., Ethernet). Another difference is that MPOA has standardized protocols for discovering routers and resolving shortcut routes to Layer-3 destinations, while Self Learning IP must improvise and learn this same topology information by snooping on the contents of selected packets.

As a result of its snooping activities, Self Learning IP is able to identify adjacent routers (i.e., router attached to ports of the 8371 switch) and build an IP forwarding table. With this database, Self Learning IP is able to identify packets destined for routers and offload the IP forwarding function from these routers.

NOTE: The Self Learning IP and MPOA Client functions are mutually exclusive in that only one of the functions may be enabled at any given time. On the standalone 8371 Switch, Self Learning IP has an automatic configuration facility that is invoked when no configuration records exist for either Self Learning IP or the MPOA Client and there is no ATM feature card installed. When these conditions are satisfied, Self Learning IP is automatically enabled.

Self Learning IP

- **Auto-configuration**
  - Self Learning IP snoops IP address information from ARP and ICMP requests
  - No device configuration needed - enable or disable
  - No configuration of existing network equipment needed
  - Self Learning IP takes time to learn the IP hosts since it uses the ARP process, so the benefits will take effect over time
- **Hardware based switching/bridging**
  - Layer-3 traffic can be forwarded at wire-speed over Fast Ethernet LANs

8371 Multilayer Ethernet Switch

Notes:

The Self Learning IP function of the 8371 builds a IP forwarding table with the following information:

- IP address of the station
- MAC address of the station
- interface through which the station can be accessed
- type of LAN encapsulation used by the station (e.g., DIX or LLC/SNAP)
- MAC address of the station's default router
- a timout value indicating when the entry is to be aged-out (e.g, if it's not being used)

Self Learning IP Example 1

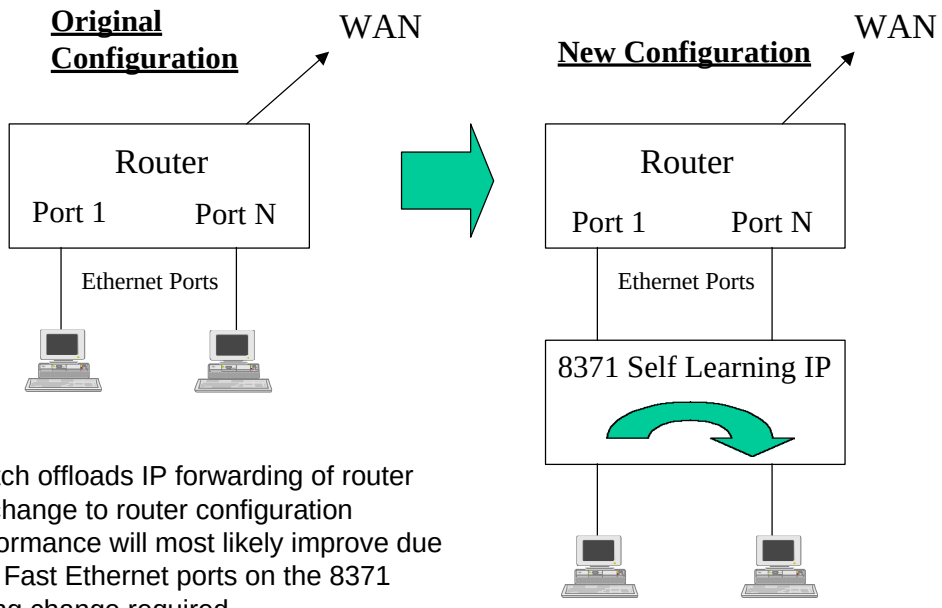

- Switch offloads IP forwarding of router
- No change to router configuration
- Performance will most likely improve due to the Fast Ethernet ports on the 8371
- Wiring change required

8371 Multilayer Ethernet Switch

Notes:

This figure represents one of the configurations where Self Learning IP is useful. In this configuration, stations are moved from an Ethernet port on the router to an ethernet port on the 8371 Switch. With this simple wiring change, The Self Learning IP function of the 8371 switch can begin to offload IP forwarding functions from the router.

See the next slide for details about how it works.

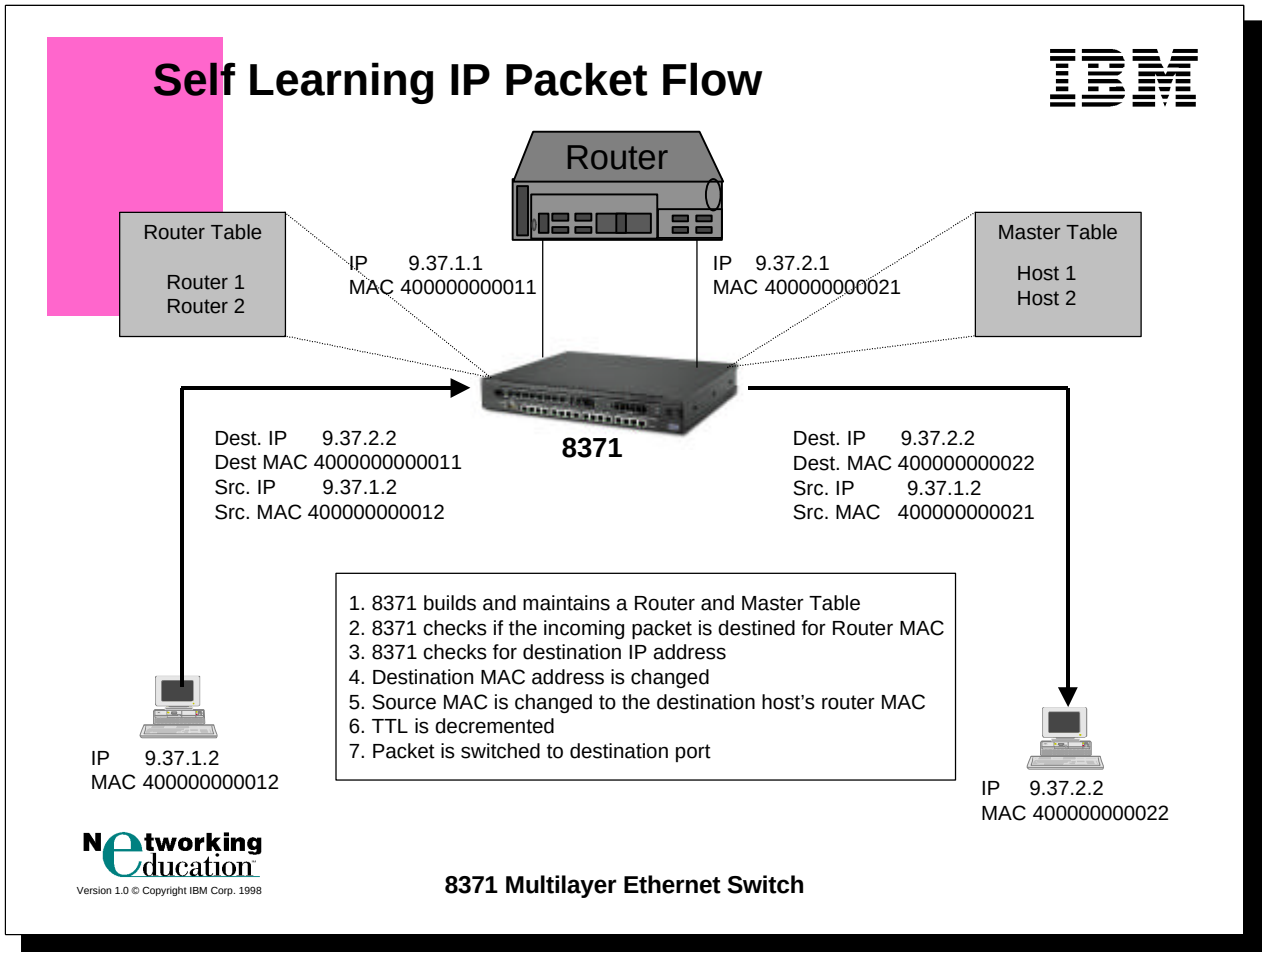

Notes:

This chart describes the detail operation of the previous chart.

The seven steps described on this chart are taken by the 8371 to switch a packet across 8371 Switch ports rather than sending the packets to the router. This is similar to the IP cut-through that is done by NHRP and MPOA. The packet headers are changed so that the receiving workstation believes that the packet came from the router. Neither the source nor the destination workstations are aware of the presence of the switch and neither is the router. There are no configuration changes needed at the router. The 8371 Switch needs no configuration either in this environment. This becomes a "plug-and-play" installation.

Self Learning IP Example 2

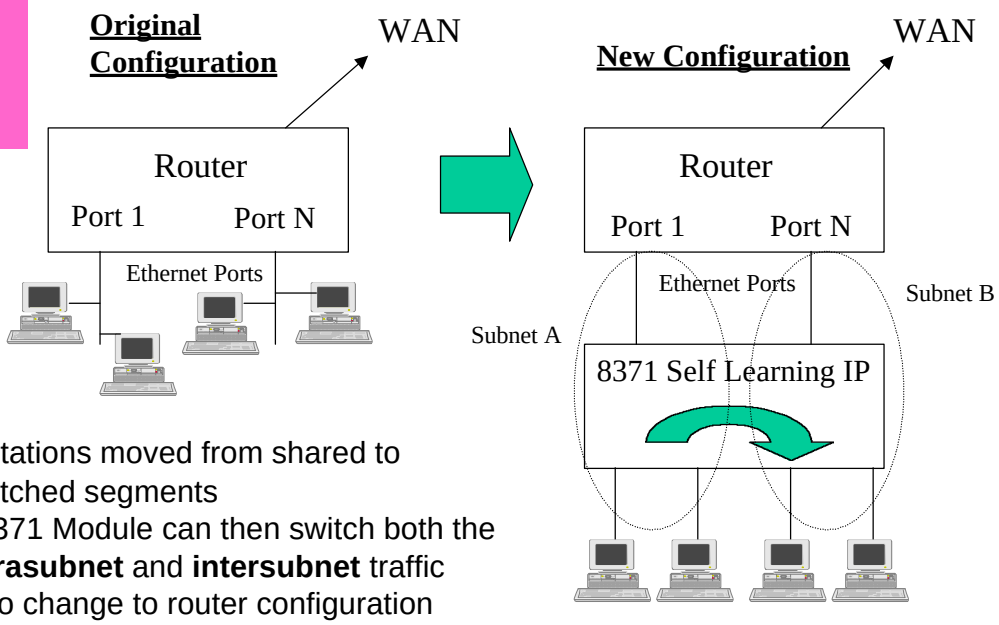

- Stations moved from shared to switched segments
- 8371 Module can then switch both the **intrasubnet** and **intersubnet** traffic
- No change to router configuration

8371 Multilayer Ethernet Switch

Notes:

The benefits provided by Self Learning IP can also be realized in conjunction with the advantages availed by micro-segmentation as depicted in the chart above. In this example, stations are moved from shared router segments to dedicated 8371 Switch segments. The 8371 Switch can then switch both the intra-subnet traffic and the inter-subnet traffic. This change could realize significant improvements in performance from three factors:

- Dedicated switch ports rather than shared
- Router bypass by the switch
- Fast Ethernet ports on the 8371 (most likely replaced 10 Mbps ports on the router)

Self Learning IP Example 3

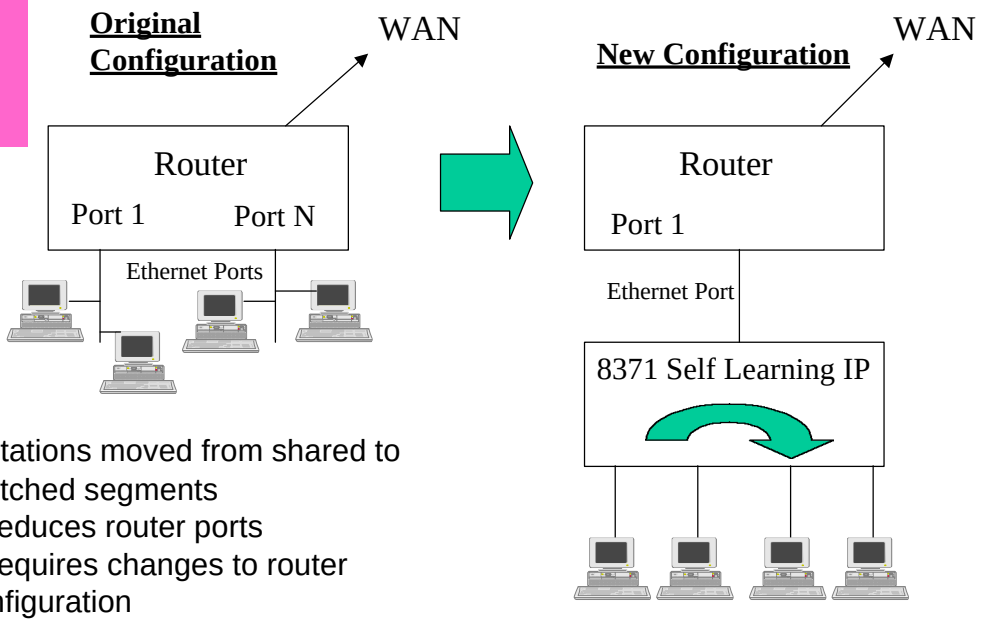

- Stations moved from shared to switched segments
- Reduces router ports
- Requires changes to router configuration

Networking education
Version 1.0 © Copyright IBM Corp. 1998

8371 Multilayer Ethernet Switch

Notes:

This chart shows another configuration where Self Learning IP is advantageous. Here, stations are moved from an ethernet port on the router to an ethernet port on the 8371 Switch, but the router is only connected to the 8371 Switch with a single ethernet port, thereby reducing the number of expensive ports required on the router. Setting up such a network would, however, require changes to the router's configuration.

Configuration & Console Interfaces

- **Uses Common Code Configurator**
  - Command Line
  - Built-in Web Browser interfaces
  - Standard RS-232 ASCII Interface
  - Telnet
- **Configurator easier to use due to auto-type feature**
- **8371 MLS intended for "Plug-and-Play" operation**
  - Only Set IP Address is required configuration step
- **Un-needed functions are removed**
  - Device ADD and Device DELETE
  - ADD and REMOVE for LECs (replaced by Assign LEC)
  - Source Route bridging, tunneling, RFC1483 bridging
- **Automatically configured:**
  - Network interfaces
  - Transparent Bridge
  - All Ethernet interfaces and all LECs are configured as bridge ports
  - One Ethernet LEC is enabled - other 23 disabled
  - All LECs, Ethernet interfaces and ATM interface given unique MACs
  - MPOA Client auto-configured and enabled

8371 Multilayer Ethernet Switch

Notes:

There is pre-configuration done for certain solutions available in the 8371. If you want to stray from that configuration (e.g, ATM uplink UFC in slot 2), then you will need to re-configure. See the next chart.

If you need to get into the configurator, it will look familiar to the 22xx Router or MSS administrator since it is the standard Common Code Configurator (i.e., Talk 5 and Talk 6). The configurator has been improved through the use of the type-ahead feature that works similar to the 8265 where the administrator types a minimum number of characters and presses the space bar. The configurator fills in the rest of the command. Also, the tab key will cycle through the available parameter options.

Some of the functions have been removed (listed above) since they are not needed by the 8371.

Each of the LECs are configured as follows:

ATM Interface:	36
ELAN Name:	ELANj (where j = 1-24)
LECS Auto-Config:	Yes
MAC Address/ESI:	burned-in MAC address for net I (where I = 40-63)
Selector:	0

There is a pool of unique MAC addresses in the 8371 to satisfy current and future requirements.

ATM Redundancy

- Two ATM uplinks ports can be configured for redundancy and additional bandwidth
- Spanning Tree blocks one Uplink/LEC if going to the same ELAN
- If Primary Link fails - Backup link is brought up automatically after Spanning Tree reconfigures
- After failed link is brought back online, Backup link will remain the Primary link until re-boot
  - Spanning Tree will then block initial Primary link

8371 Multilayer Ethernet Switch

Notes:

The following scenario shows the status of the LECs and Port Interfaces in a Redundant scenario. Port Interface 36 is assigned to LEC Interface 40 and Port Interface 37 is assigned to LEC Interface 41.

Initially:
Interface 36 (Slot 1, Port 1) - up
Interface 37 (Slot 1, Port 2) - down (blocked)
LEC Interface 40 - up
LEC Interface 41 - down (blocked)

Disconnect port 36:
Interface 36 - down
Interface 37 - up
Interface 40 - down
Interface 41 - up

Reconnect port 36:
Interface 36 - down (blocked)
Interface 37 - up
Interface 40 - down (blocked)
Interface 41 - up

ATM Traffic Shaping and QOS

- **QOS established on a LEC basis**
- **QOS support same as in MSS 1.1**
- **Best effort LECs use PCR (Peek Cell Rate) of 155 Mbps**
  - If Call Setup is rejected due to PCR, the call will be re-tried with a lower PCR
- **Reserved Bandwidth connections are treated as CBR (Constant Bit Rate)**
  - Traffic transmitted onto QOS VCCs is shaped according to SCR (Sustained Cell Rate)
  - No PCR used
  - No Maximum Burst Size used
  - QOS setup applies to both Backward and Forward VCCs
  - If Call Setup fails, LEC falls back to Best Effort with the configured PCR

8371 Multilayer Ethernet Switch

Notes:

When a Best Effort connection is rejected due to the Peak Cell Rate, the call may be automatically retried with a lower PCR. Retries are performed under the following conditions:

1. If the rejected PCR is greater than 100 Mbps, the call is retried with a PCR of 100 Mbps
2. Otherwise, if the rejected PCR is greater than 25 Mbps, the call is retried with a PCR of 25 Mbps.

VLANs, LECs and Filters

- **Separation of upstream ATM traffic onto multiple LECs possible via filters**
- **Filters included:**
  - IP by subnet
  - IPX by network
  - NetBIOS
  - IP Multicast
  - Sliding Window by offset and mask
- **Include or exclude protocols or specific traffic by LEC Interface**

8371 Multilayer Ethernet Switch

Notes:

Other Supported and Non-supported Function (vs. Common Code)

- **Protocols supported:**
  - ASRT (Adaptive Source Route Bridging)
  - SR-TB Translational Bridging - to send frames to T-R MPCs
  - MPOA (client only)
  - HST (IP Host)
  - SNMP
- **Not supported:**
  - *QCONFIG* - Quick Configurator
  - Routing
  - Classical IP
  - LAN Emulation Services
  - Self Learning IP on Blades
  - Token-Ring LECs
  - RFC1483 Bridging

8371 Multilayer Ethernet Switch

Notes:

8371 Network Management

- **SNMP Agent Supported by:**
  - Nways Workgroup Manager for Windows NT V1.1.2
  - Nways Manager for AIX V1.1.2 - Campus LAN Manager
- **RMON Groups 1, 2, 3 and 9 of RMON I MIB**
  - Group 1 - Statistics
  - Group 2 - History
  - Group 3 - Alarms
  - Group 9 - Events
- **13 MIBs supported**
  - 6 IETF MIBs
  - 2 ATM Forum MIBs
  - 5 Enterprise Specific MIBs

8371 Multilayer Ethernet Switch

Notes:

MIBs:

- IETF MIBs
  - MIB II specified in RFC 1213
  - MIB II Extension specified in RFC 1573b
  - Ethernet MIB specified in RFC 1650
  - AToM MIB specified in RFC 1695
  - RMON MIB specified in RFC 1757
  - Bridge MIB specified in RFC 1493
- ATM Forum Standard MIBs
  - LAN Emulation Client MIB
  - MPOA Client MIB
- Enterprise Specific MIBs
  - IBM LAN Emulation Extension MIB
  - IBM Common Routing MIB
  - IBM Netview 6000 MIB
  - IBM CPU Utilization MIB
  - Proteon MIB