

# Module 21: Windows 2000

- **History**
- **Design Principles**
- **System Components**
- **Environmental Subsystems**
- **File system**
- **Networking**
- **Programmer Interface**

# Windows 2000

## ■ Features:

- ☞ 32-bit OS
- ☞ Preemptive OS
- ☞ multitasking
- ☞ OS for Intel Pentium CPUs, only.
- ☞ micro-kernel architecture

## ■ Key goals for the system:

- ☞ portability
- ☞ security
- ☞ POSIX compliance
- ☞ multiprocessor support
- ☞ extensibility
- ☞ international support
- ☞ compatibility with MS-DOS and MS-Windows applications.

# Windows 2000

- **Available in 4 versions:**

- 👉 **Professional**

- 👉 **Server**

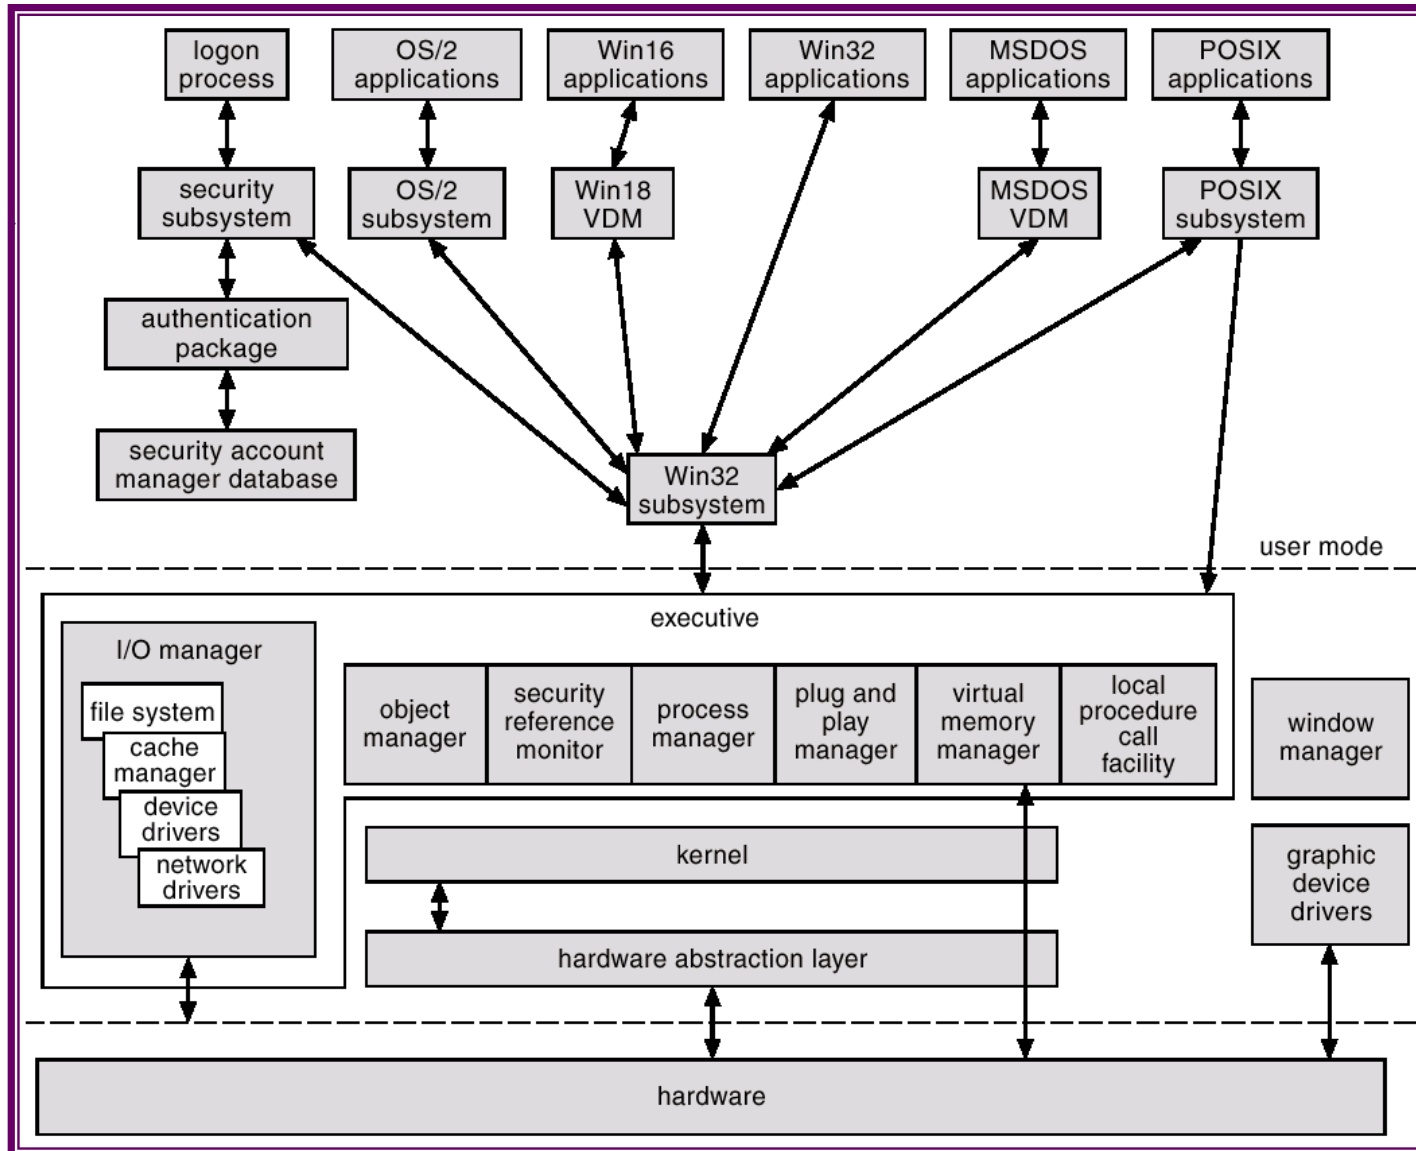
- 👉 **Advanced Server**

- 👉 **Data Center Server**

- 📄 (up 32 CPU, up to 64GB RAM).

- In 1996, more **NT server licenses** were **sold** than **UNIX licenses**


# Depiction of 2000 Architecture




# System Components — Kernel

- **Foundation for the executive and the subsystems.**
- **Never paged out of memory; execution is never preempted.**
- **Four main responsibilities:**
  - ☞ **thread scheduling**
  - ☞ **interrupt and exception handling**
  - ☞ **low-level processor synchronization**
  - ☞ **recovery after a power failure**

# One-to-one Model



# Many-to-Many Model



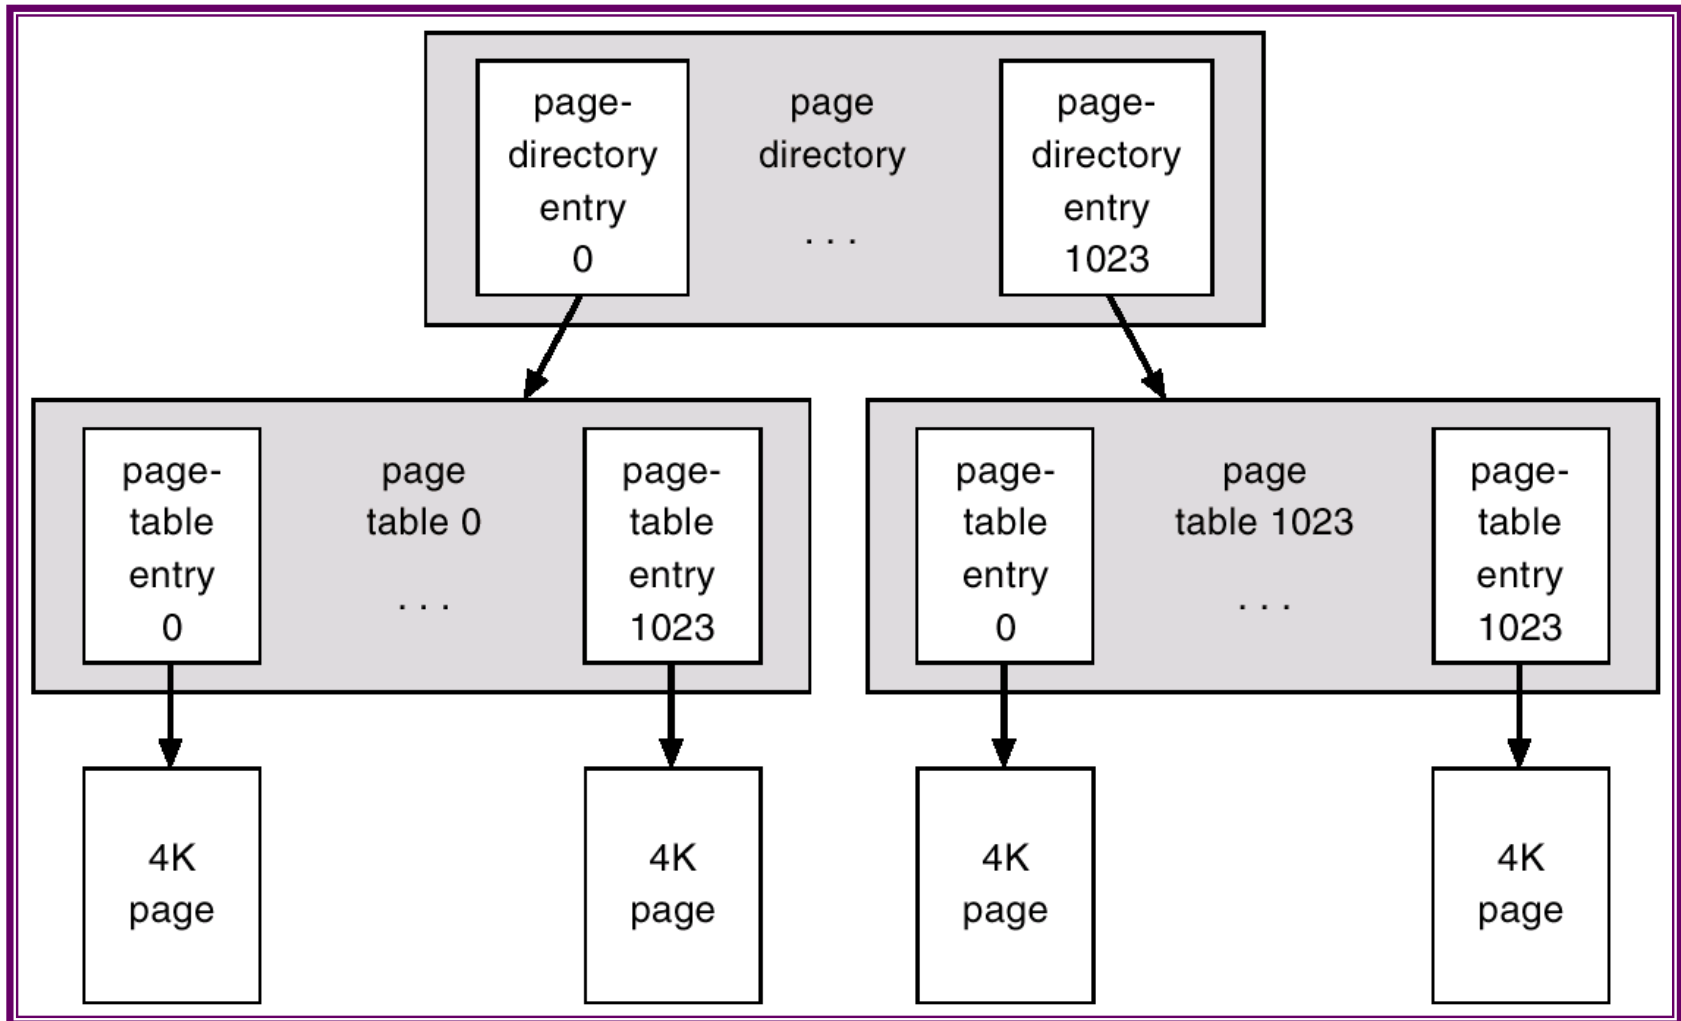
# Kernel — Scheduling

- The dispatcher uses a 32-level priority scheme
  - ☞ to determine the order of thread execution.
- Priorities are divided into two classes..
  - ☞ **1. real-time class**
    - ☞ contains threads with priorities ranging from 16 to 31.
  - ☞ **2. variable class**
    - ☞ contains threads having priorities from 0 to 15.
- Characteristics of 2000's priority strategy.
  - ☞ Trends to give very good response times
    - ☞ to interactive threads
    - ☞ that are using the mouse and windows.
  - ☞ Enables I/O-bound threads to keep the I/O devices busy.
  - ☞ Complete-bound threads
    - ☞ soak up the spare CPU cycles
    - ☞ in the background.


# Executive — Virtual Memory Manager

- The **design** of the VM manager assumes
- **that the underlying hardware supports**
  - ☞ virtual to physical mapping a paging mechanism,
  - ☞ transparent cache coherence on multiprocessor systems,
  - ☞ virtual addressing aliasing.
- The **VM manager in 2000 uses**
  - ☞ a page-based management scheme
  - ☞ with a **page size of 4 KB**.
- The **2000 VM manager uses a two step process**
- **to allocate memory**
  - ☞ **first step reserves a portion** of the process's address space.
  - ☞ **second step commits the allocation**
    - ☞ by assigning space in the **2000 paging file**

# Virtual-Memory Layout



# Virtual-to-Physical Address Translation



- 10 bits for page directory entry,
- 10 bits for page table entry, and
- 12 bits for byte offset in page.

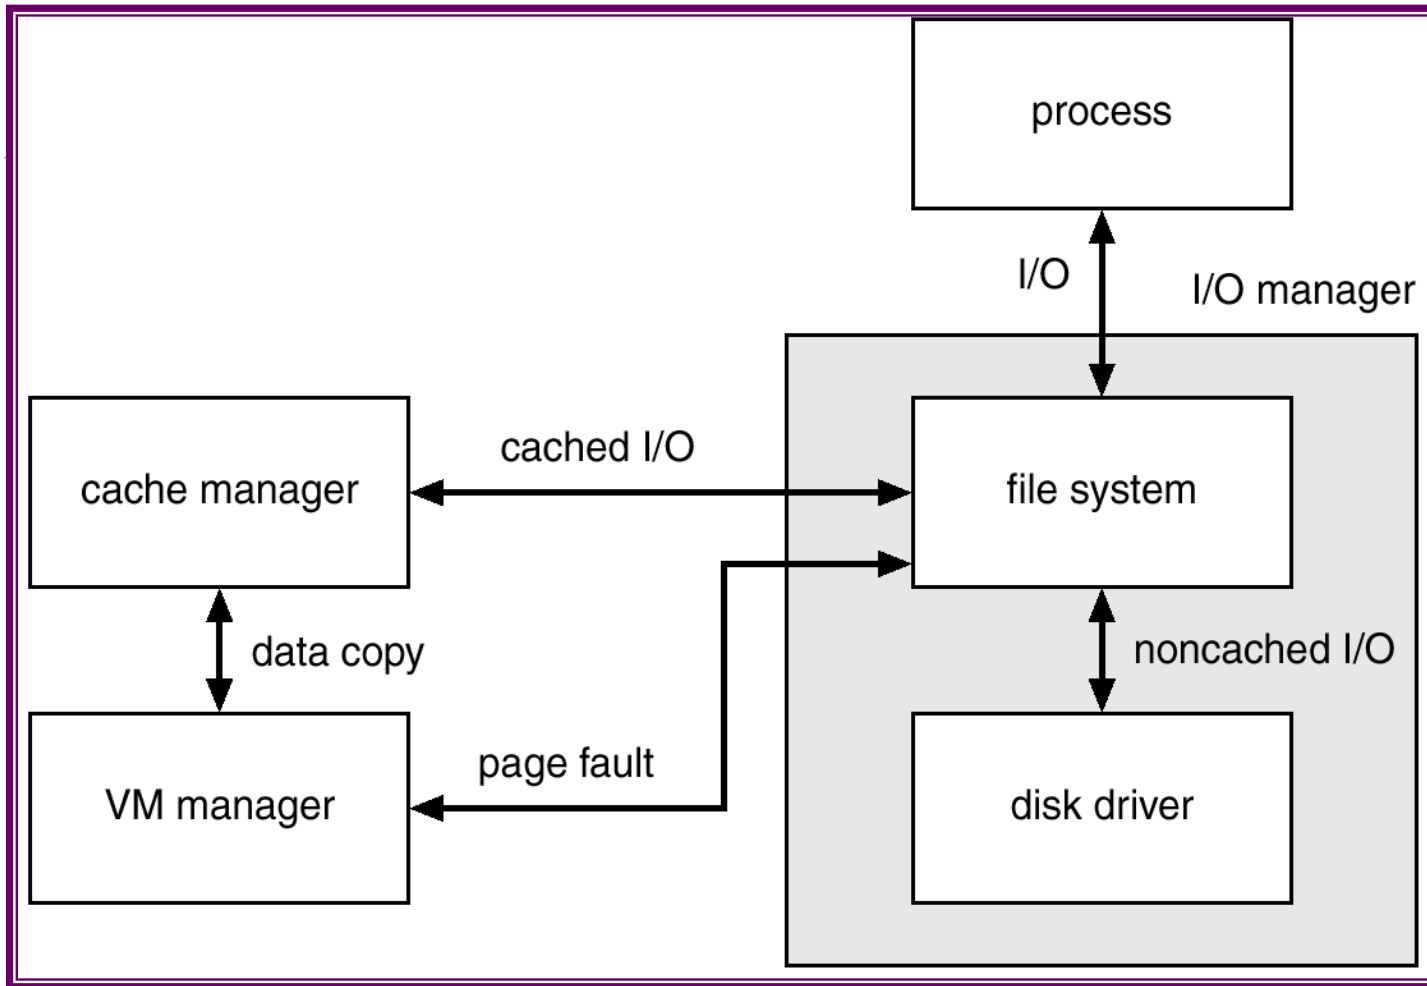
# Executive — Local Procedure Call Facility

- The **LPC passes** requests and results
  - ☞ between client and server processes
  - ☞ within a single machine.
- In particular,
  - ☞ it is used to request services
  - ☞ from the various 2000 subsystems.
- When a **LPC channel** is created,
- one of **3 types** of message passing techniques must be specified.
  - ☞ **1. First type is suitable for small messages**
    - ☞ up to 256 bytes;
    - ☞ port's message queue is used as intermediate storage, and
    - ☞ the messages are copied from one process to the other.
  - ☞ **2. Second type avoids copying large messages**
    - ☞ by pointing to a shared memory section object
    - ☞ created for the channel.
  - ☞ **3. Third method, called quick LPC was used by**
    - ☞ **graphical display portions of the Win32 subsystem.**

# Executive — I/O Manager

- The **I/O manager** is responsible for
  - ☞ **file systems**
  - ☞ **cache management**
  - ☞ **device drivers**
  - ☞ **network drivers**
- **Keeps track of which installable file systems are loaded, and**
  - ☞ **manages buffers for I/O requests.**
- **Works with VM Manager**
  - ☞ **to provide memory-mapped file I/O.**
- **Controls the 2000 cache manager,**
  - ☞ **which handles caching for the entire I/O system.**
- **Supports**
  - ☞ **both synchronous and asynchronous operations**
  - ☞ **provides time outs for drivers, and**
  - ☞ **has mechanisms for one driver to call another.**

# File I/O



# Executive – Plug-and-Play Manager

- **Plug-and-Play (PnP) manager**
- **is used**
- **to recognize and**
- **adapt to changes**
- **in the hardware configuration.**
  
- **When new devices are added**
  - ☞ **(for example, PCI or USB),**
  - ☞ **the PnP manager loads the appropriate driver.**
  
- **The manager also keeps**
  - ☞ **track of the resources**
  - ☞ **used by each device.**

# File System

- **The fundamental structure of the 2000 file system (NTFS) is a volume.**
  - ☞ **Created by the 2000 disk administrator utility.**
  - ☞ **Based on a logical disk partition.**
  - ☞ **May occupy a**
    - 📄 **portions of a disk**
    - 📄 **an entire disk or**
    - 📄 **span across several disks**
- **All metadata,**
  - ☞ **such as information about the volume,**
  - ☞ **is stored in a regular file. (original concept)**
- **NTFS uses clusters as the underlying unit of disk allocation.**
  - ☞ **A cluster is a number of disk sectors that is a power of two.**
  - ☞ **Because the cluster size is smaller than for the 16-bit FAT file system,**
    - 📄 **the amount of internal fragmentation is reduced.**

# File System — Recovery

- **All file system data structure updates (metadata)**
  - ☞ are performed
  - ☞ inside transactions
  - ☞ that are logged.
- **Before a data structure is altered,**
  - ☞ the transaction writes a **log record**
  - ☞ that contains redo and undo information.
- **After the data structure has been changed,**
  - ☞ a **commit record** is written to the log
  - ☞ to signify that the **transaction succeeded**.
- **After a crash,**
  - ☞ the file system data structures
  - ☞ can be restored to a consistent state
  - ☞ by processing the log records.

# File System — Compression

- To **compress a file**,
- **NTFS divides the file's data**
  - ☞ into compression units,
  - ☞ which are blocks
    - 📄 of 16 contiguous clusters.
  
- For **sparse files**,
- **NTFS uses another technique to save space.**
  - ☞ Clusters that contain all zeros
    - 📄 are not actually allocated
    - 📄 or stored on disk.
  
  - ☞ Instead, gaps are left
    - 📄 in the sequence of virtual cluster numbers
    - 📄 stored in the MFT entry for the file.
  
  - ☞ When reading a file,
    - 📄 if a gap in the virtual cluster numbers is found,
    - 📄 NTFS just zero-fills that portion of the caller's buffer.

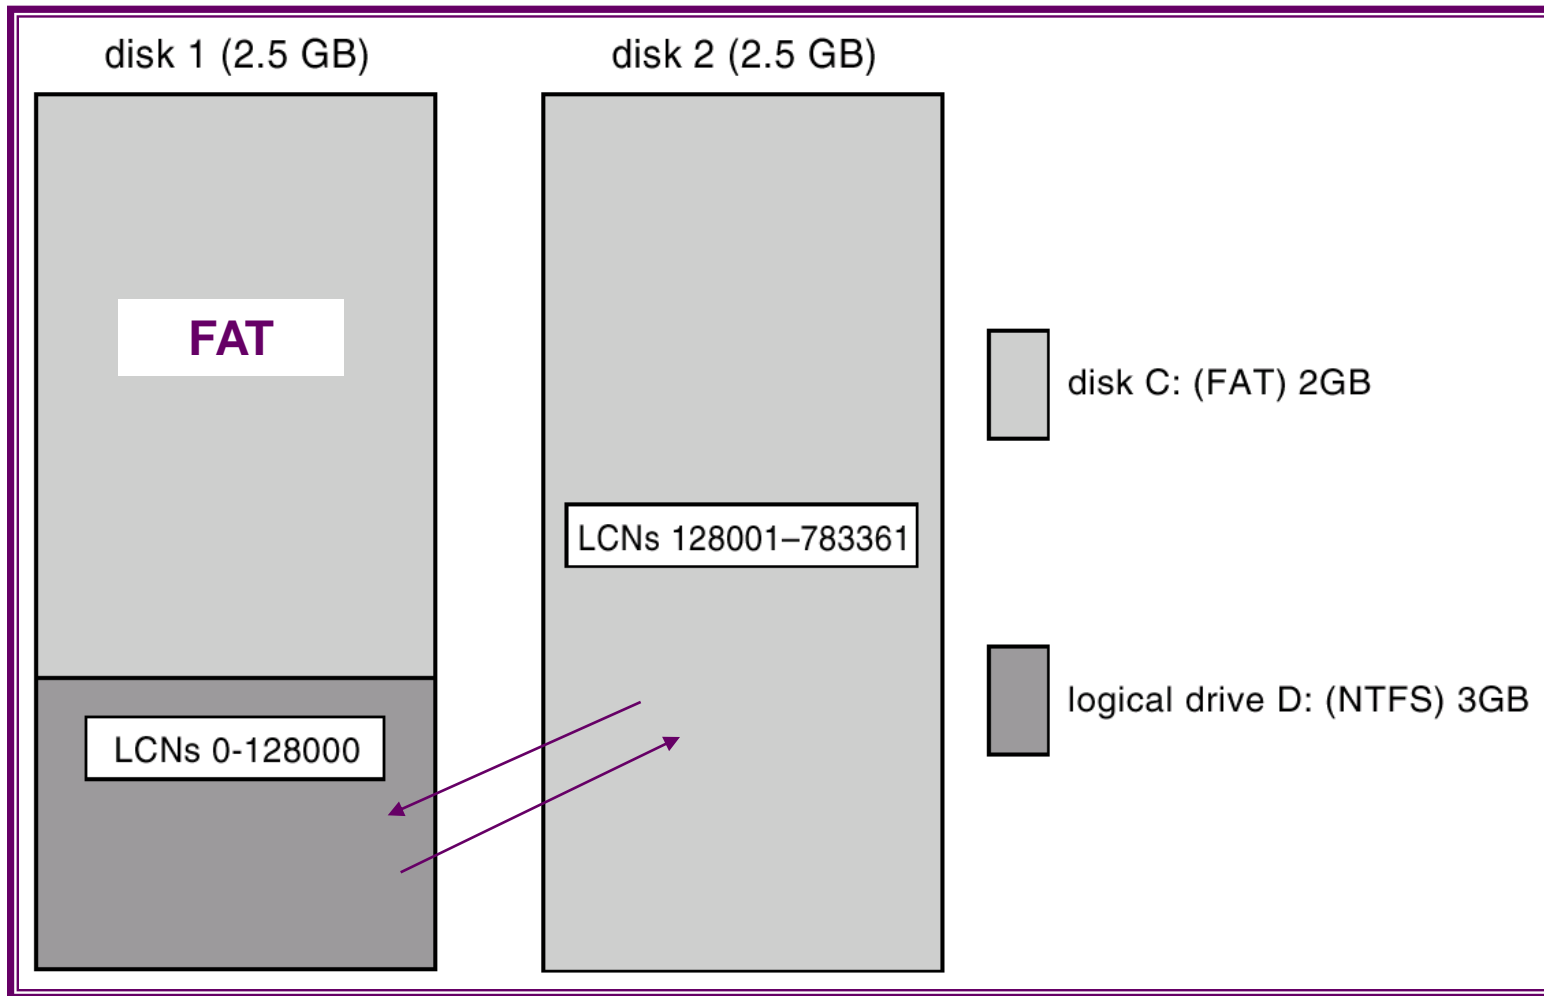
# File System — Reparse Points

- A reparse point returns
  - ☞ an error code when accessed.
- The reparse data
  - ☞ tells the I/O manager
  - ☞ what to do next.
- Reparse points
  - ☞ can be used to provide
  - ☞ the functionality of **UNIX mounts**
- Reparse points
  - ☞ can also be used
  - ☞ to access files
  - ☞ that have been moved to offline storage

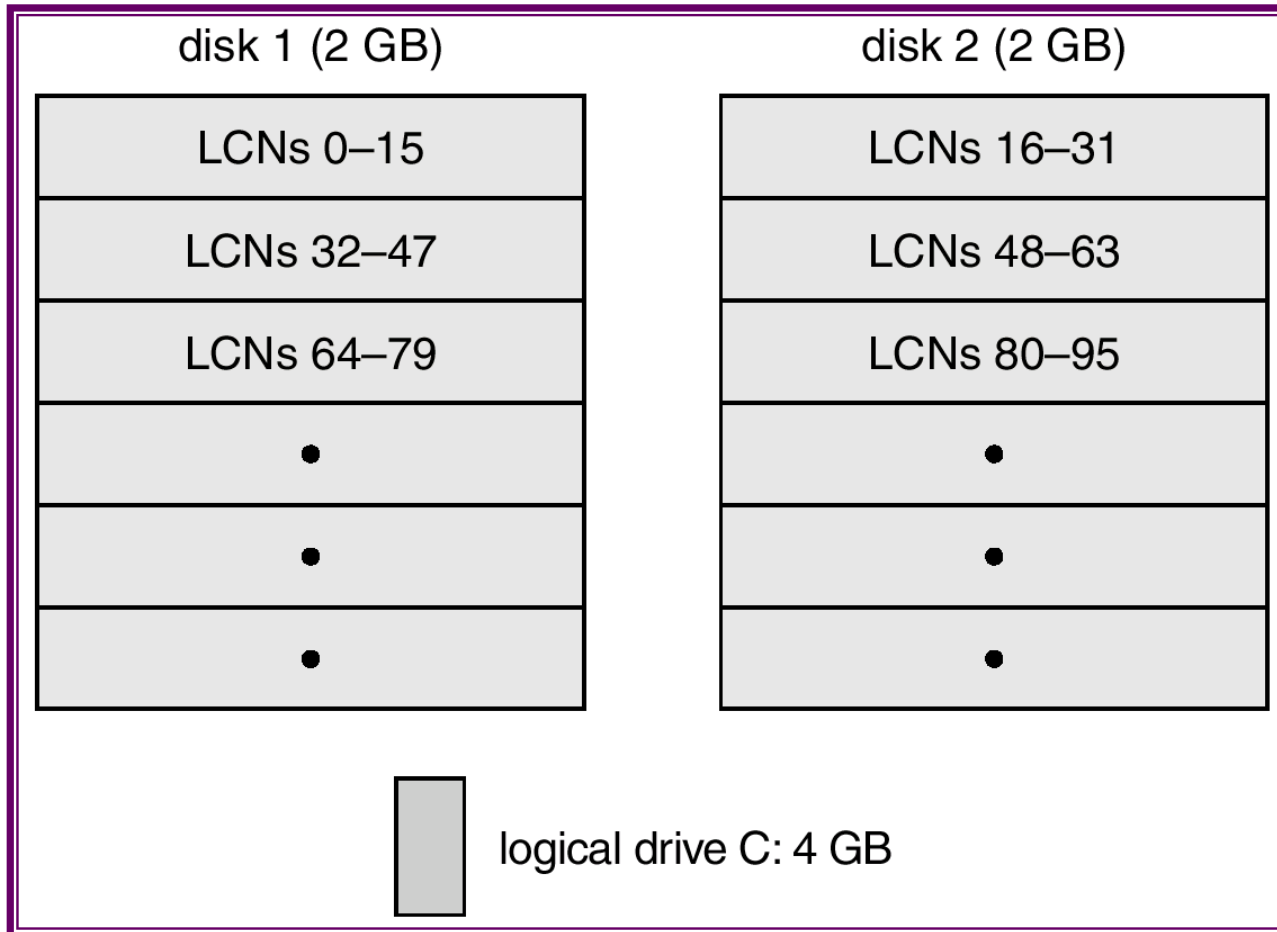
# Volume Management and Fault Tolerance

- **FtDisk,**
  - the **fault tolerant disk driver for 2000,**
  - provides several ways
  - to **combine multiple SCSI disk drives**
  - **into one logical volume**
- **Logically concatenate multiple disks**
  - ☞ to form a large logical volume,
  - ☞ a volume set


# Volume Set On Two Drives



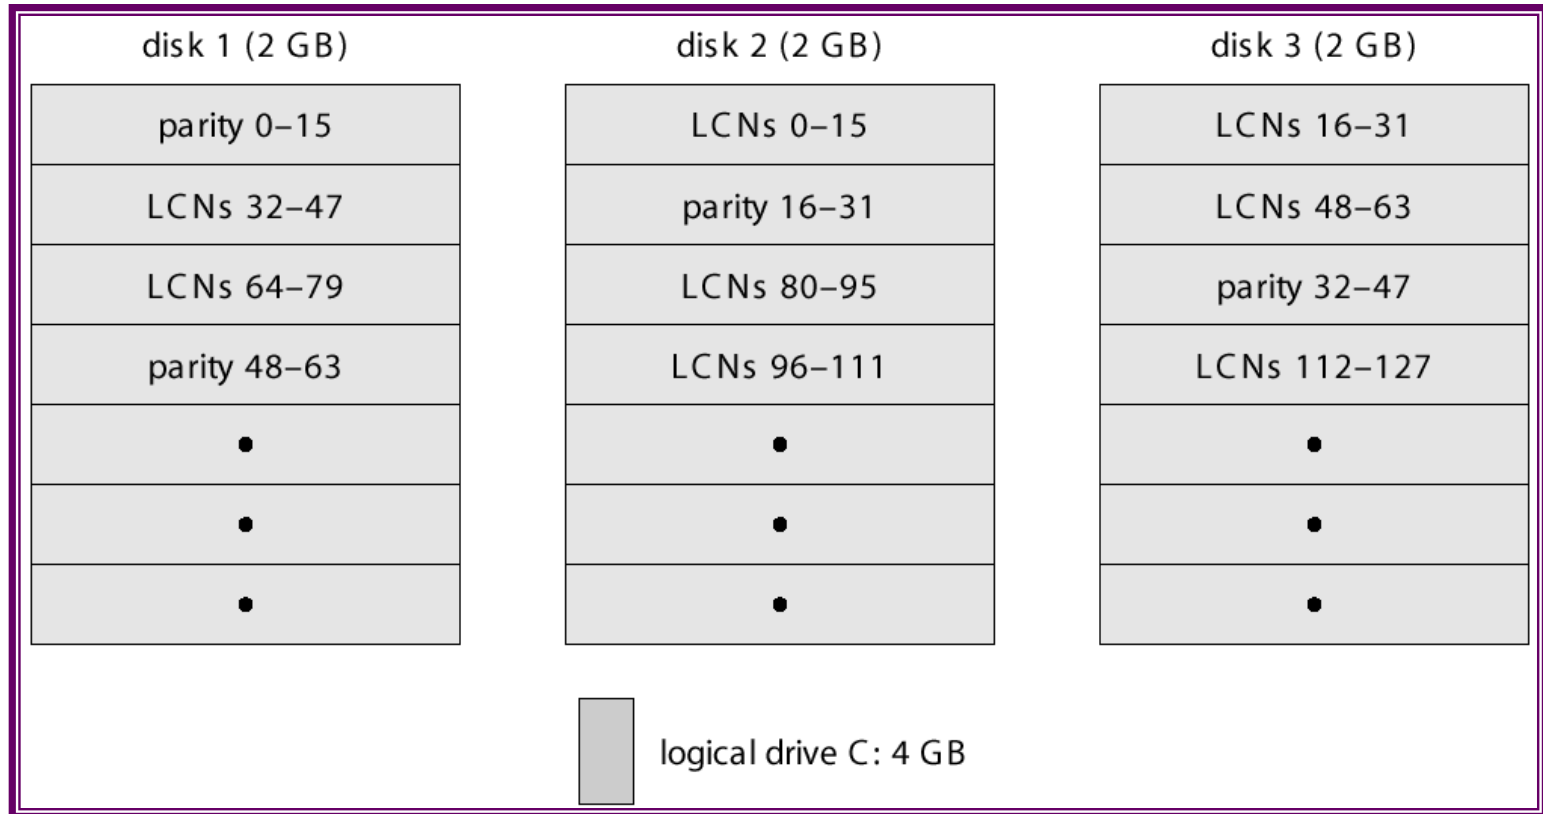
# Stripe Set on Two Drives – RAID 0



# Mirror Set on Two Drives – RAID 1



# Stripe Set With Parity on Three Drives-RAID 5



# Networking

- 2000 supports both:
  - ☞ 1. peer-to-peer networking;
  - ☞ 2. client/server networking;
  - ☞ it also has facilities for network management.
- To describe networking in 2000, we refer to two of the internal networking interfaces:
  - ☞ 1. NDIS (Network Device Interface Specification):
    - ☞ separates network adapters
    - ☞ from the transport protocols
    - ☞ so that either can be changed without affecting the other.
  - ☞ 2. TDI (Transport Driver Interface):
    - ☞ enables any session layer component
    - ☞ to use
    - ☞ any available transport mechanism.
- 2000 implements
  - ☞ transport protocols as drivers
  - ☞ that can be loaded and unloaded from the system dynamically.

# Networking — Protocols

- The **server message block protocol (SMB)**
- is used
- to send I/O requests over the network.
- It has **4 message types**:
  - **1. Session control** (start and end connection to shared resource to server)
  - **2. File** (use for file access at the server)
  - **3. Printer** (send data to a remote print queue)
  - **4. Message** (used to communicate with another station)

# Networking — Protocols (Cont.)

- **NetBEUI (NetBIOS Extended User Interface):**
  - ☞ default protocol for
    - 📄 Win95 peer networking
    - 📄 Windows for Workgroups
  - ☞ used when 2000 wants to share resources with these networks.
- **2000 uses the TCP/IP Internet protocol**
  - ☞ to connect to
  - ☞ a wide variety of operating systems and
  - ☞ hardware platforms.
- **PPTP (Point-to-Point Tunneling Protocol) is used**
  - ☞ to communicate
  - ☞ between Remote Access Server modules
  - ☞ running on 2000 machines
  - ☞ that are connected over the Internet.
- The **2000 NWLink protocol**
  - ☞ connects the NetBIOS to Novell NetWare networks.

# Networking — Protocols (Cont.)

- The **Data Link Control protocol (DLC)** is used to access
  - ☞ **IBM mainframes** and
  - ☞ **HP printers** that are **directly connected to the network.**
  
- **2000 systems**
  - ☞ **can communicate with Macintosh computers**
  - ☞ **via the Apple Talk protocol**
  
- **if an 2000 Server on the network**
  - ☞ **is running the Windows 2000 Services**
  - ☞ **for Macintosh package.**

# Networking — Domains

- **NT uses the concept of a domain**
  - ☞ to manage
  - ☞ global access rights within groups.
  
- **A domain is a group of machines**
  - ☞ running NT server
  - ☞ that share a **common security policy** and **user database**.
  
- **2000 provides three models of setting up trust relationships.**
  - ☞ 1. One way, A trusts B
  
  - ☞ 2. Two way, transitive,
    - ☞ A trusts B, B trusts C so A, B, C trust each other
  
  - ☞ 3. Crosslink –
    - ☞ allows authentication
    - ☞ to bypass hierarchy
    - ☞ to cut down on authentication traffic.

# Name Resolution in TCP/IP Networks

- 2000 provides **several methods** of **name resolution**:
  - ☞ **Windows Internet Name Service (WINS)**
  - ☞ **broadcast name resolution**
  - ☞ **domain name system (DNS)**
  - ☞ **a host file**
    - 📄 an **LMHOSTS** file
- **WINS** consists of **two or more WINS servers**
  - ☞ that maintain
    - 📄 a **dynamic database** of **name to IP address bindings**
  - ☞ **and**
  - ☞ **client software** to **query the servers**.
- **WINS** uses
- **the Dynamic Host Configuration Protocol (DHCP)**,
  - ☞ **DHCP** automatically updates
    - 📄 **address configurations** in the **WINS** database
  - ☞ **without user or administrator intervention**