Introduction to the new mainframe

Chapter 2 Hardware systems and LPARs

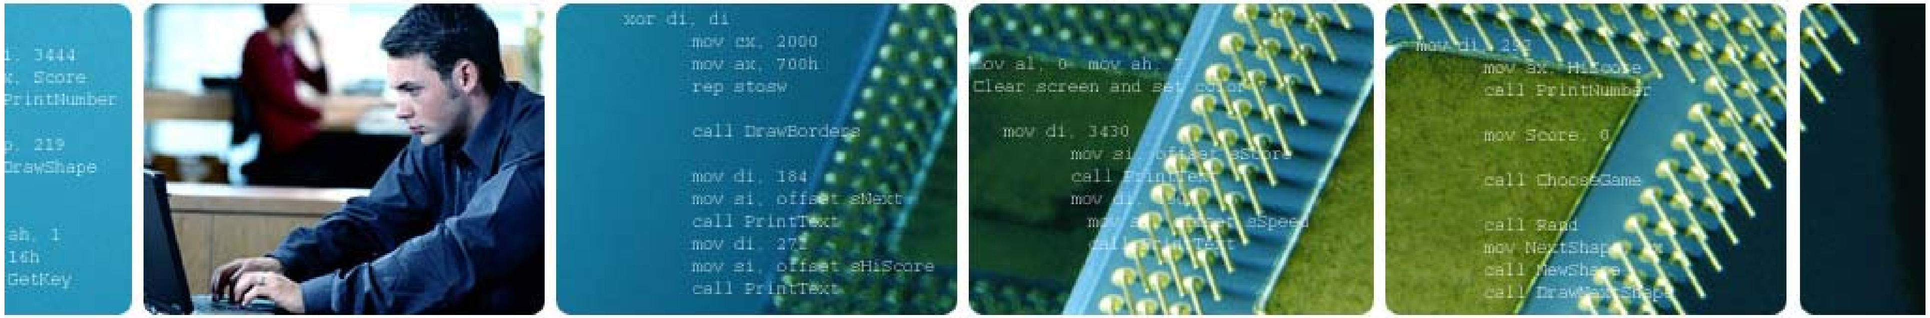

Objectives

In this chapter you will learn:

- About S/360 and zSeries hardware design
- About processing units and disk hardware
- How mainframes differ from PC systems in data encoding
- About some typical hardware configurations

Introduction

Here we look at the hardware in a complete system although the emphasis is on the processor "box"

Terminology is not straightforward

- **Ever since "boxes" became multi-engined, the terms system, processor, and CPU have become muddled**

Terminology:

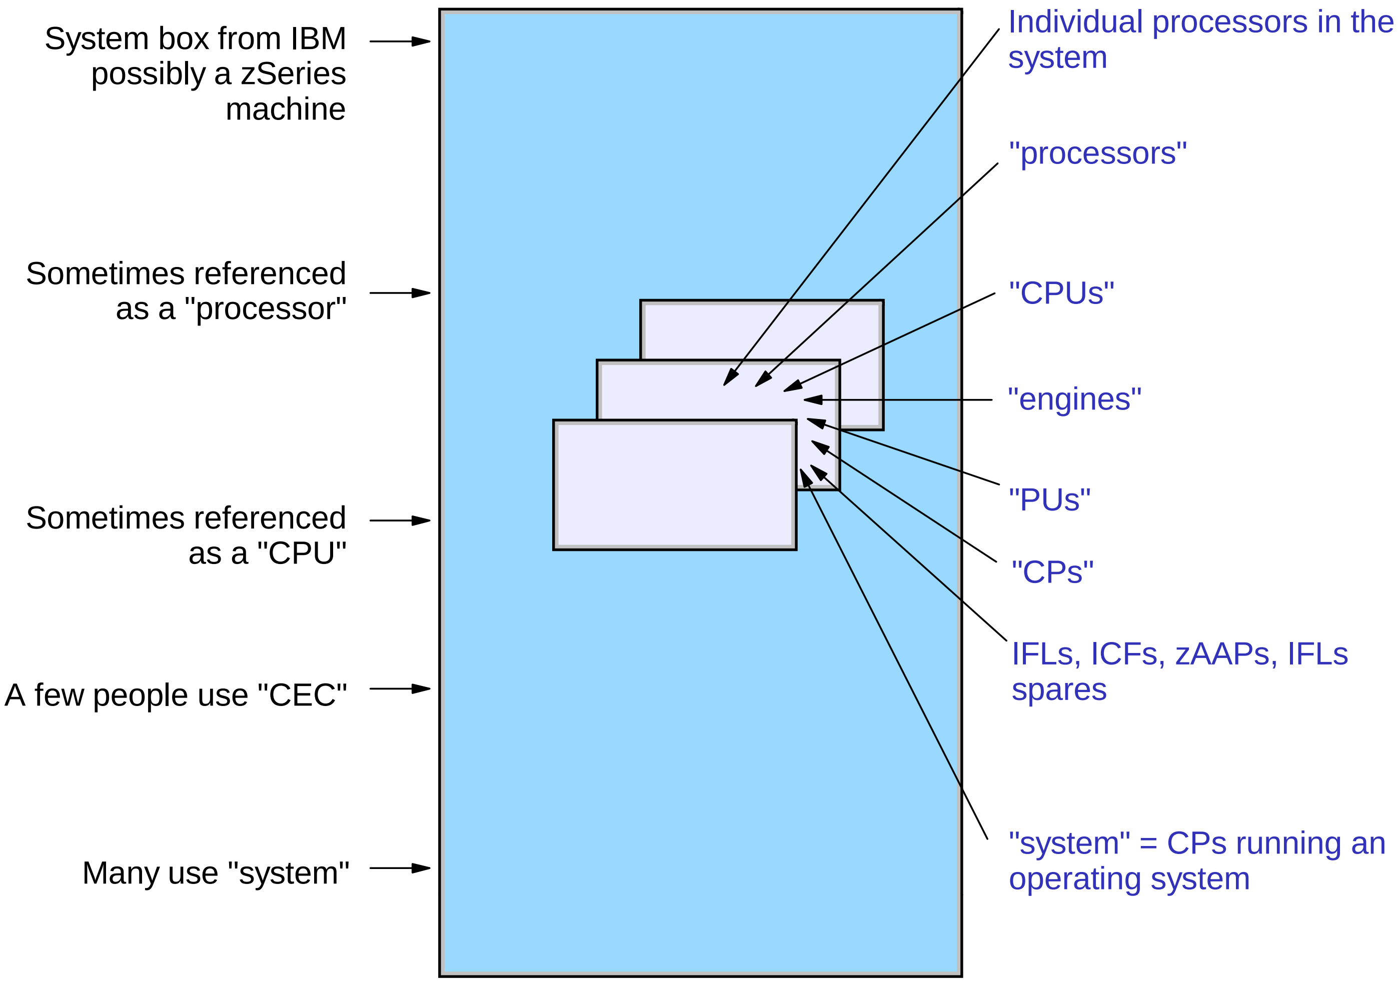

Conceptual S/360

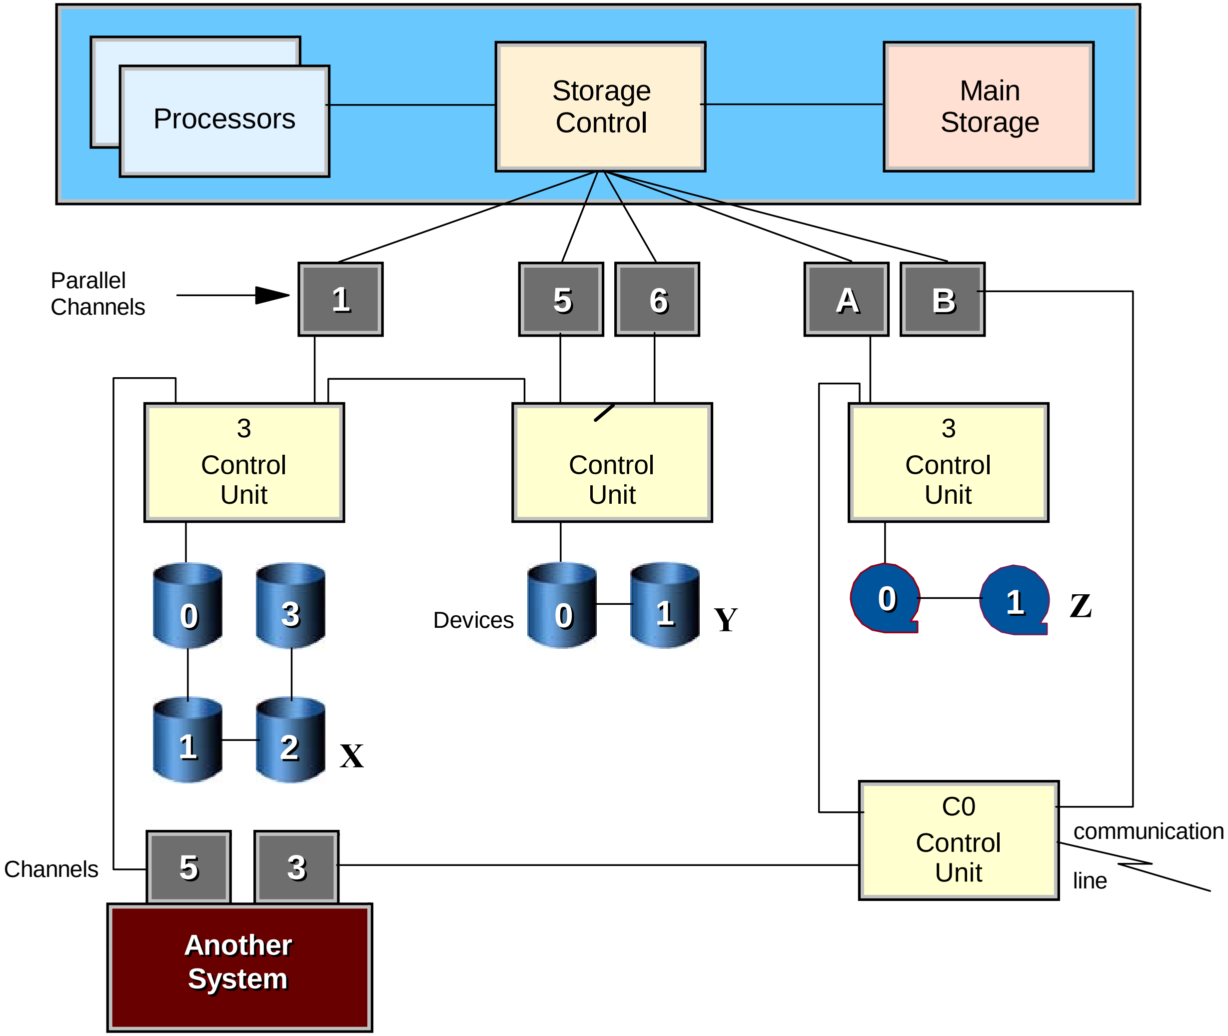

Early system design

- **System/360 was designed in the early 1960s**
- **The central processor box contains the processors, memory, control circuits and channel interfaces**
  - Early systems had up to 16 channels whereas modern systems have over 1000
- **Channels connect to control units**
- **Control units connect to devices such as disk drives, tape drives and communication interfaces**

Device address

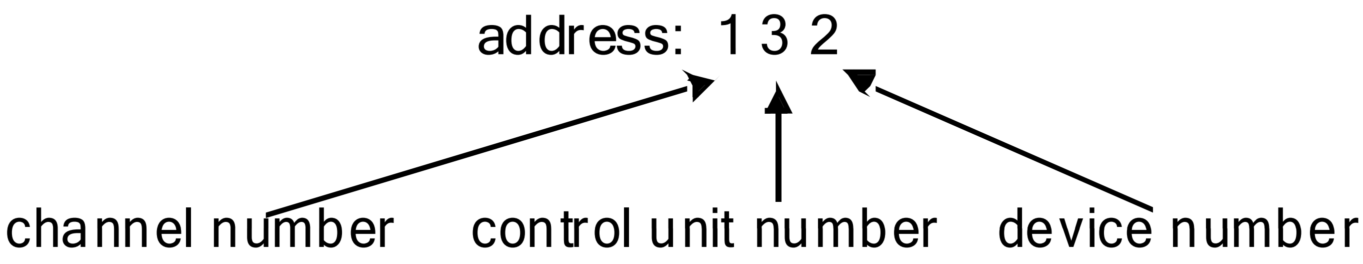

In the early design the device address was physically related to the hardware architecture

Parallel channels had large diameter heavy copper “bus and tag” cables

Current design

Current CEC designs are considerably more complex then the early S/360 design

This new complexity includes

- I/O connectivity
- I/O operation
- Partitioning of the system

I/O Connectivity

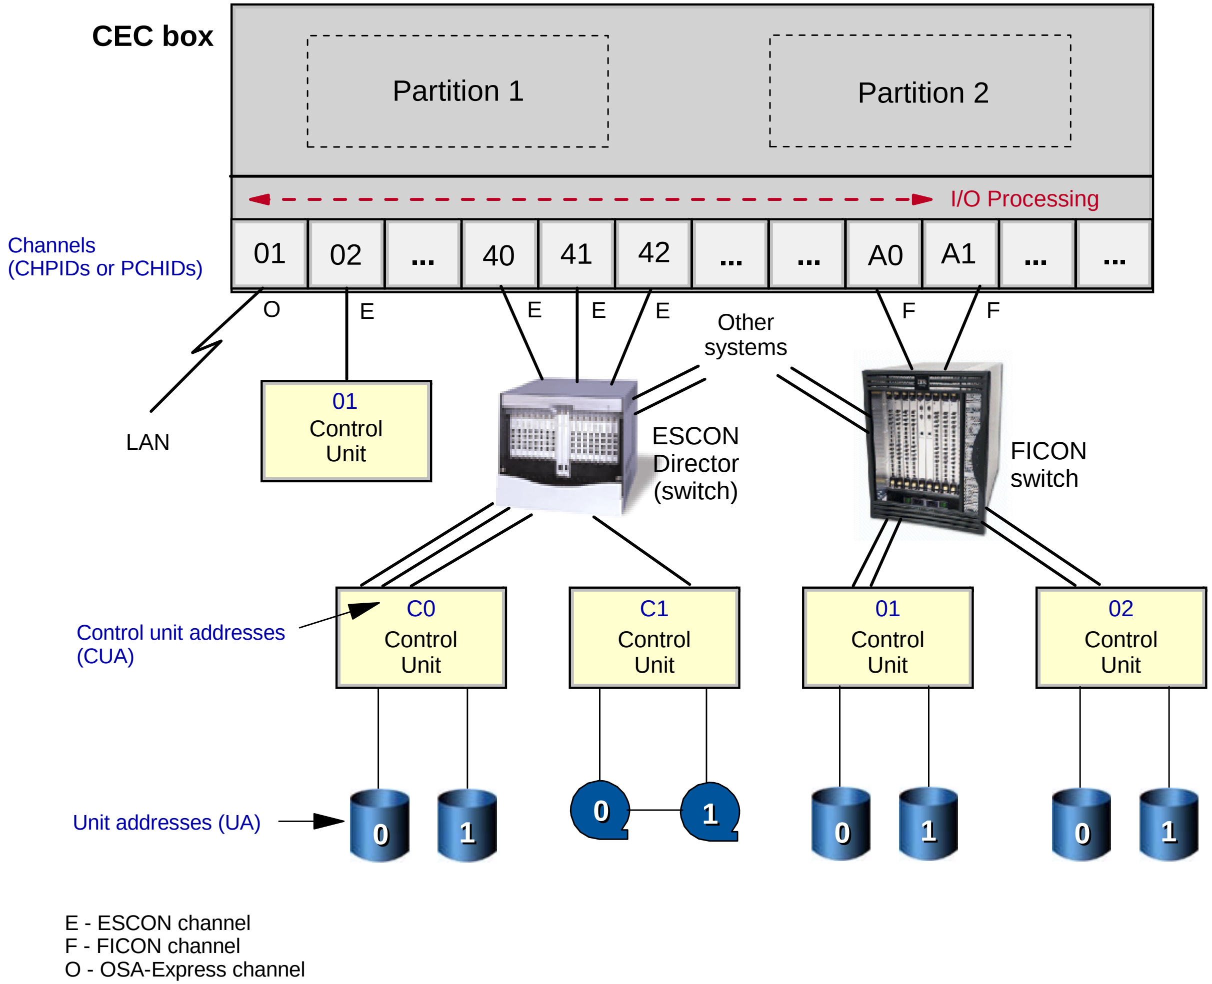

I/O Connectivity

- ESCON and FICON channels
- Switches to connect peripheral devices to more than one CEC
- CHPID addresses are two hex digits
- Multiple partitions can share CHPIDs
- I/O subsystem layer exists between the operating system and the CHPIDs

I/O Connectivity

- I/O control layer uses a control file IOCDS that translates physical I/O addresses into devices numbers that are used by z/OS
- Device numbers are assigned by the system programmer when creating the IODF and IOCDS and are arbitrary (but not random!)
- On modern machines they are three or four hex digits

I/O Connectivity

- External device label
- Four hex digits in range 0000-FFFF
- Arbitrarily assigned by sysprog
- Used in JCL, commands, messages, EREP

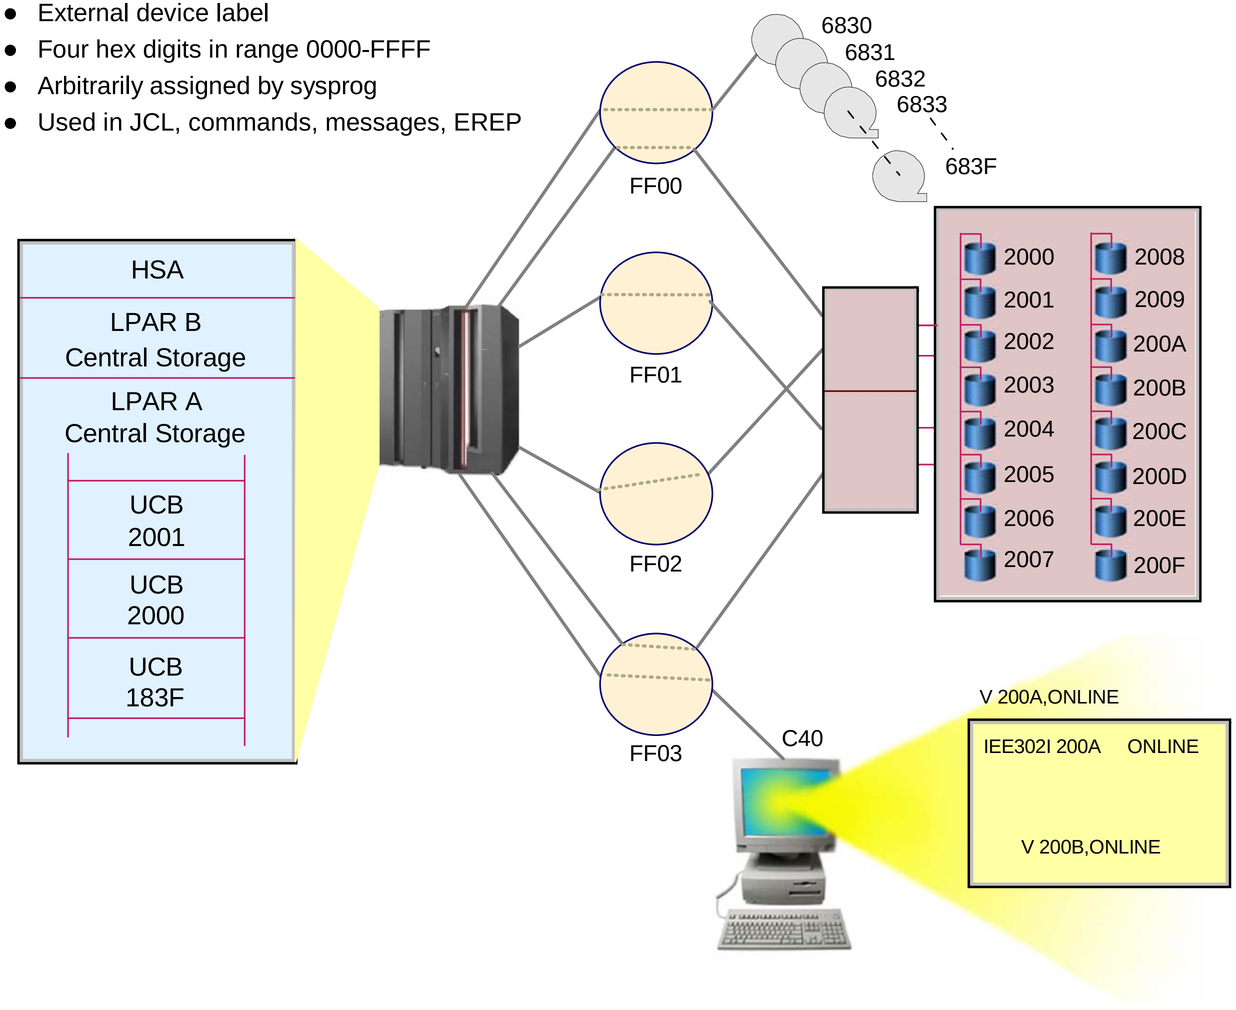

System Control and Partitioning

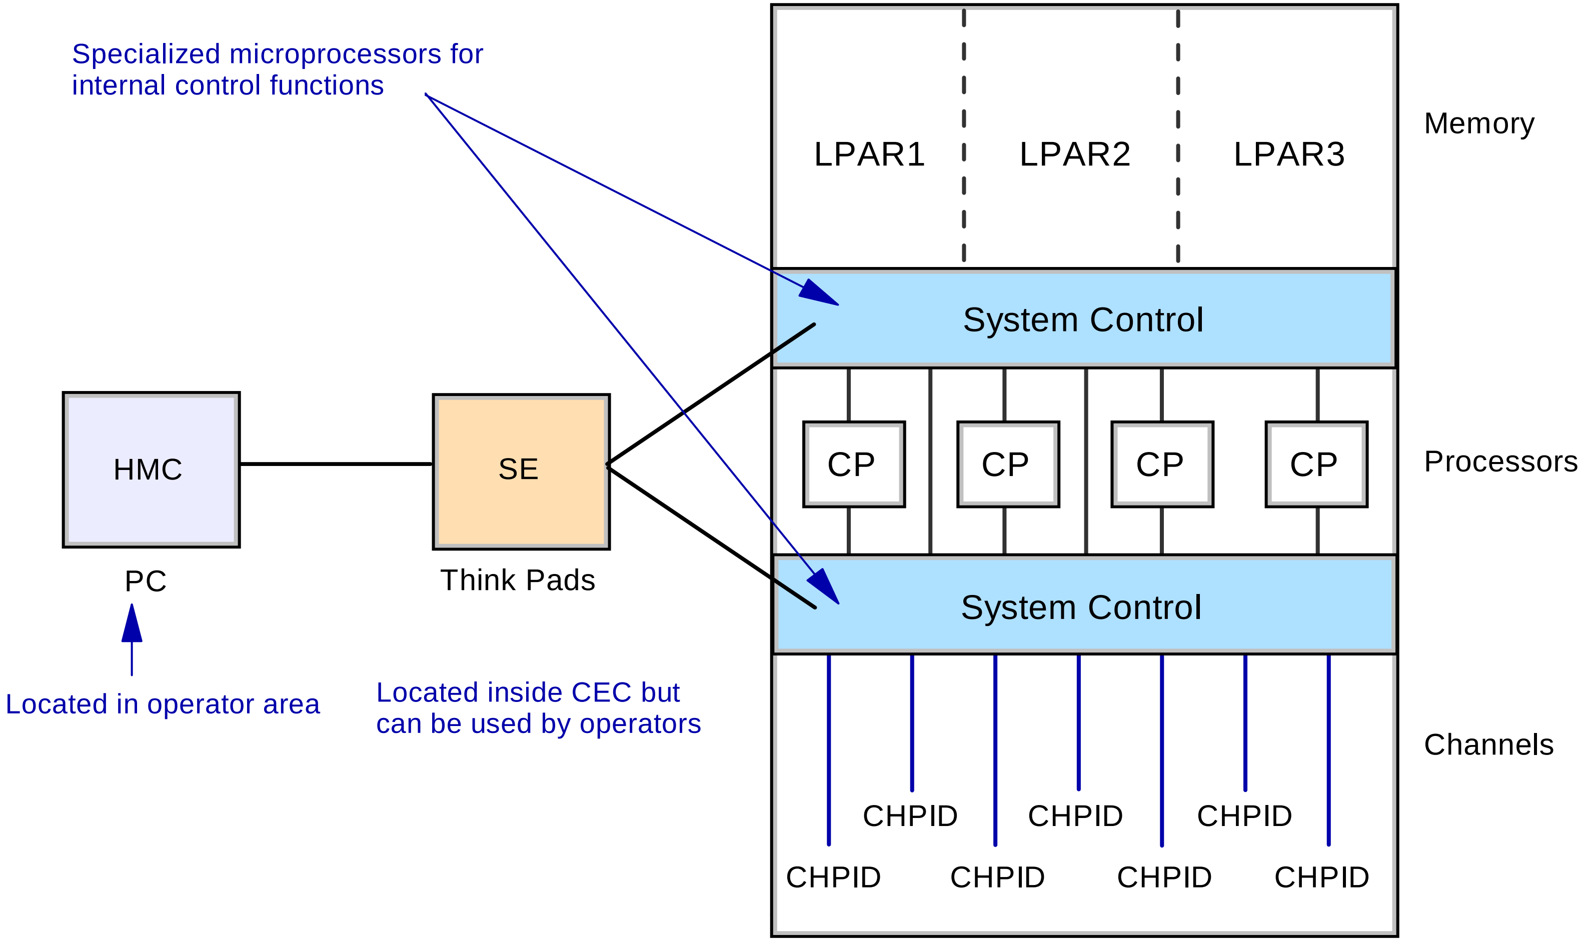

System Control and Partitioning

- **Among the system control function is the ability to partition the system into logical partitions (LPARs)**
- **Initial limit was 15 LPARs but newer machines allow 30 partitions**
- **Practical considerations can limit the number to less than this as each LPAR needs memory, I/O, and processing power**

LPARs

System administrators assign:

- **Memory**
- **Processors**
- **CHPIDs either dedicated or shared**

This is done partly in the IOCDS and partly in a system profile on the Support Element (SE) in the CEC. This is normally updated through the HMC.

Changing the system profile and IOCDS will usually require a power-on reset (POR) but some changes are dynamic

Characteristics of LPARs

- **LPARs are the equivalent of a separate mainframe for most practical purposes**
- **Each LPAR runs its own operating system**
- **Devices can be shared across several LPARs**
- **Processors can be dedicated or shared**
- **When shared each LPAR is assigned a number of logical processors (up to the maximum number of physical processors) and a weighting**
- **Each LPAR is independent**

Consolidation of mainframes

- **Many installations used to have several boxes**
- **A single larger machines running many LPARs is often more cost effective, as software licenses for multiple small machines can total more than those for a single large one**

Processing Units

Modern processor offer many specialist processor types.

Most of the additional types are not counted as full processors for software charging

- Central Processor (CP) - A full z/OS processor
- System Assistance Processor (SAP) - Used for the I/O subsystem – each machine has at least one
- Integrated Facility for Linux (IFL) - Special processor for Linux - optional
- zAAP – Used for JAVA code – optional
- zIIP – Used for Dataprocessing code – optional
- Integrated Coupling Facility (ICF) - For coupling facilities
- Spares

SAPs, IFLs, ICFs, zAAPs and zIIPS are not counted in the model number and or against software costs

Capacity on Demand

Various forms of Capacity on Demand exist

Additional processing power to meet unexpected growth or sudden demand peaks

Disk Devices

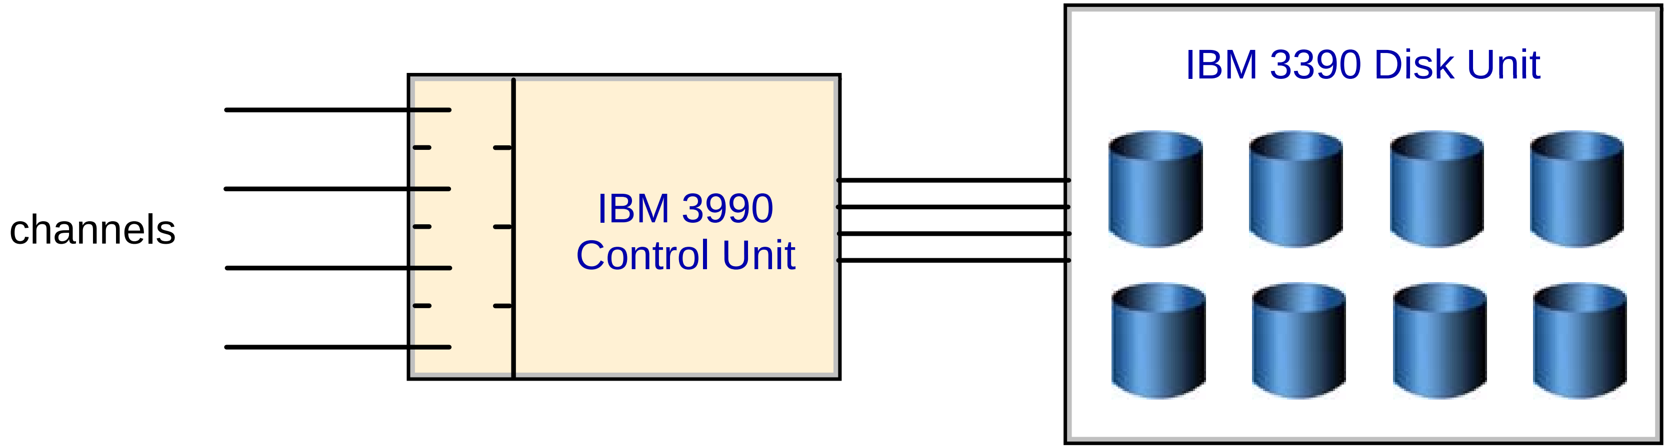

Current mainframes use 3390 disk devices

The original configuration was simple with a controller connected to the processor and strings of devices attached to the back end

Current 3390 Implementation

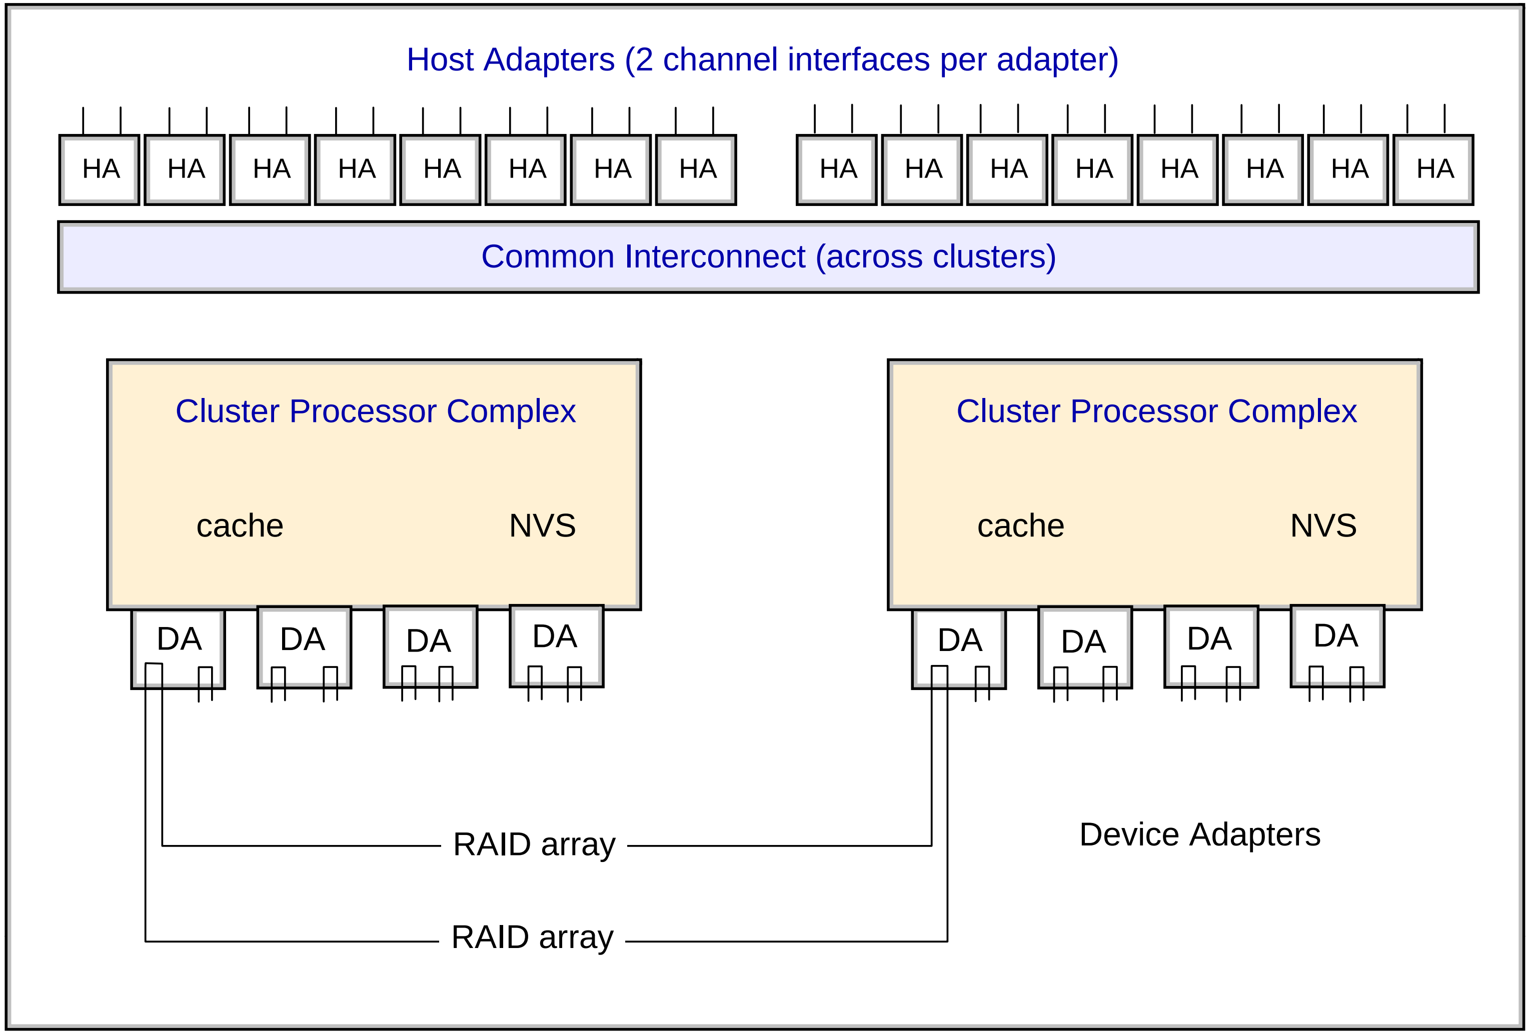

Modern 3390 devices

- The 2105 Enterprise Storage Server just shown is very sophisticated
- It emulates a large number of control units and 3390 disks. It can also be partitioned and connect to UNIX and other systems as SCSI devices.
- There are 11 TB of disk space up to 32 channel interfaces, 16GB cache and 284MB of non-volatile memory

Modern 3390 Devices

- **The physical disks are commodity SCSI- type units**
- **Many configurations are possible but usually it is RAID-5 arrays with hot spares**
- **Almost every part has a fallback or spare and the control units are emulated by 4 RISC processors in two complexes.**

Modern 3390 Devices

- **The 2105 offers FlashCopy, Extended Remote Copy, Concurrent Copy, Parallel Access Volumes, Multiple Allegiance**
- **This is a huge extension of the original 3390 architecture and offers a massive performance boost.**
- **To the z/OS operating system these disks just appear as traditional 3390 devices so maintaining backward compatibility**

EBCDIC

- **The IBM S/360 through to the latest zSeries machines use the Extended Binary Coded Decimal Interchange character set for most purposes**
- **This was developed before ASCII and is also an 8 bit character set**
- **z/OS Web Server stores ASCII data as most browsers run on PCs which expect ASCII data**
- **UNICODE is used for JAVA on the latest machines**

Clustering

Clustering has been done for many years in several forms

- **Basic shared DASD**
- **CTC/GRS rings**
- **Basic and Parallel sysplex**

Image is used to describe a single z/OS system, which might be standalone or an LPAR on a large box

Basic shared DASD

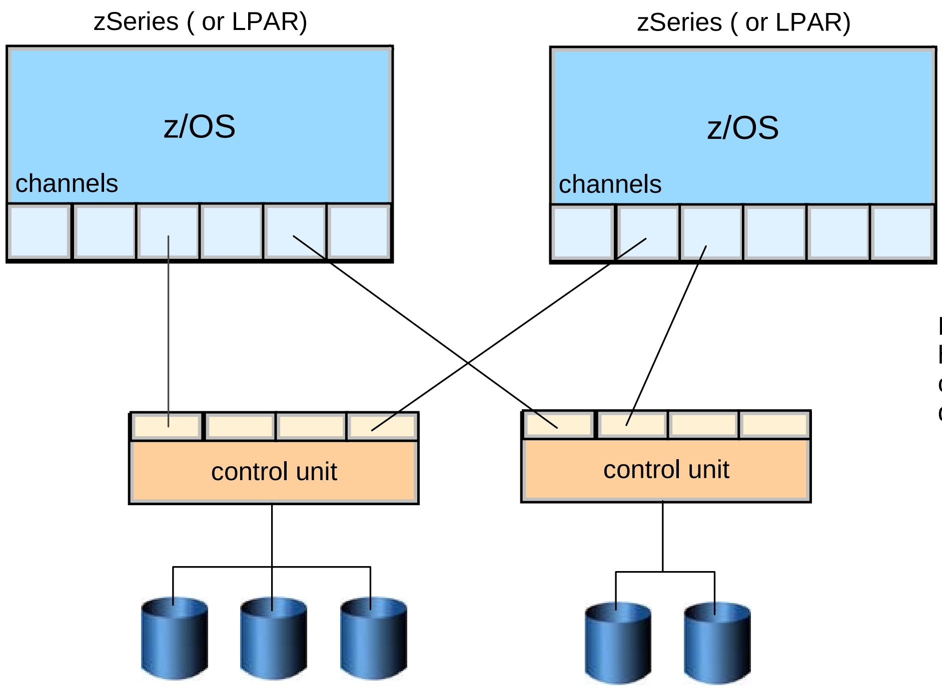

Real system would have many more control units and devices

Limited capability

Reserve and release against a whole disk

Limits access to that disk for the duration of the update

CTC Ring

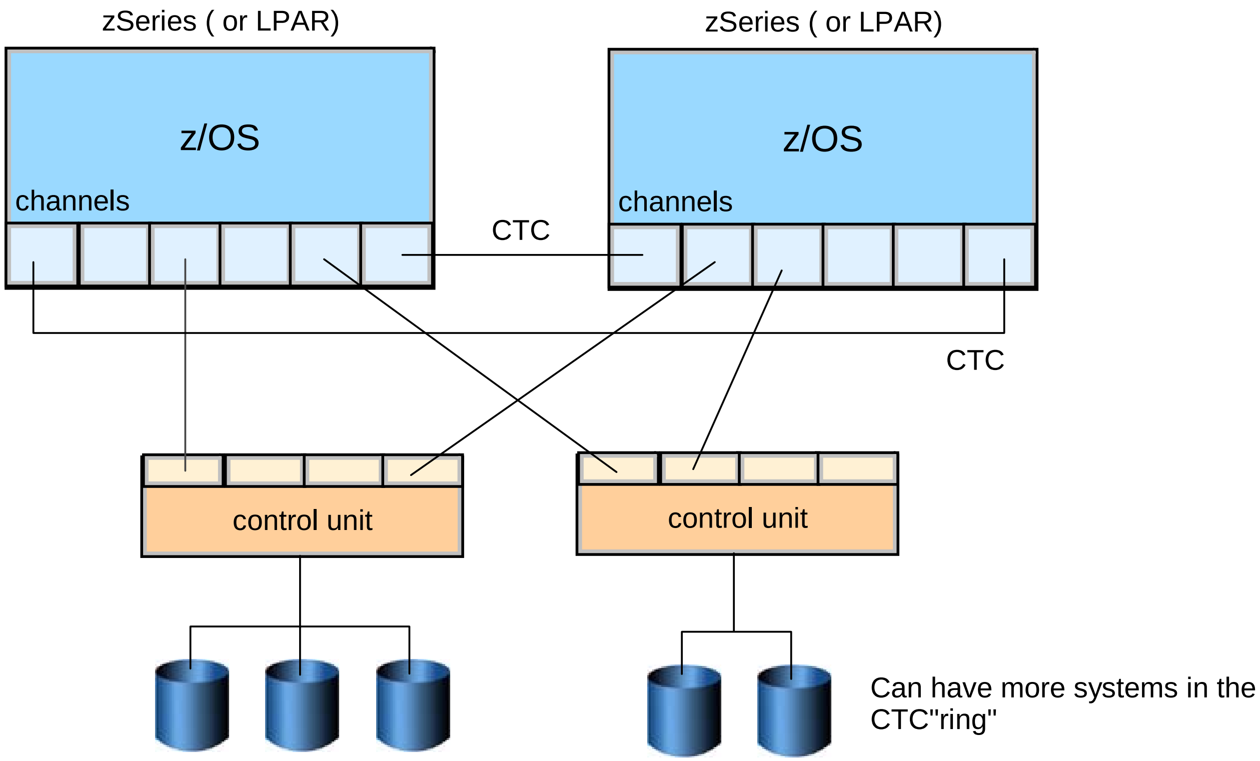

Global Resource Sharing (GRS) used to pass information between systems via the CTC ring

Request ENQueue on a data set, update, the DEQueue

Loosely coupled system

Parallel Sysplex

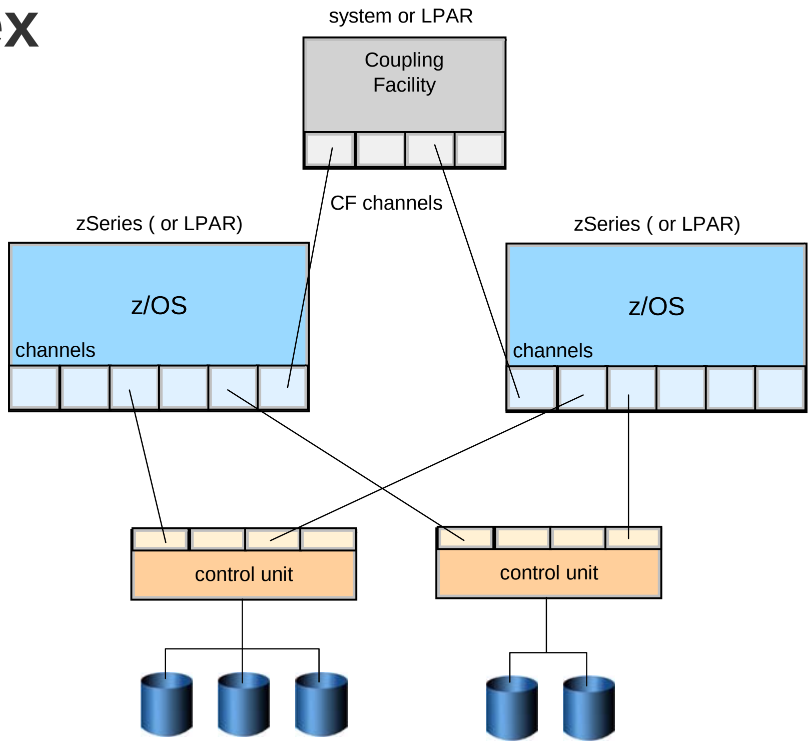

- This extension of the CTC ring uses a dedicated Coupling Facility to store ENQ data for GRS
- This is much faster
- The CF can also be used to share application data such as DB2 tables
- Can appear as a single system

Typical Systems

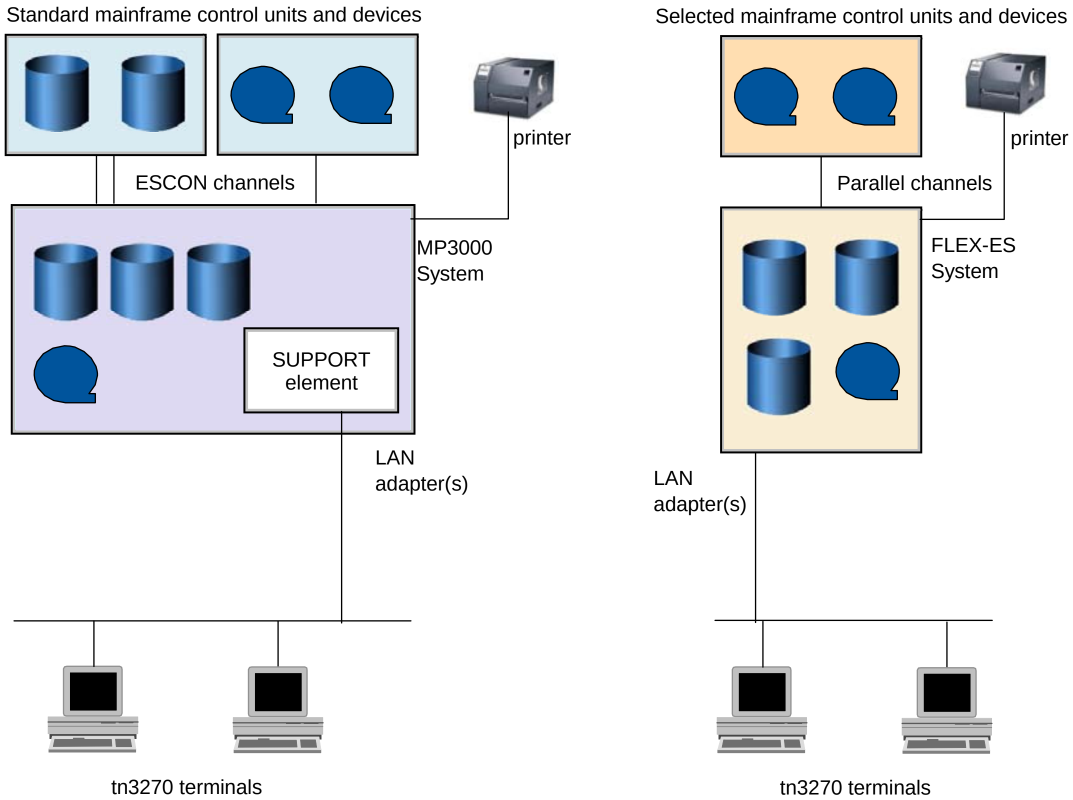

This shows two very small systems

- On the left is a Multiprise 3000, which was designed for small installations with internal disk drives
- On the right is a FLEX-ES emulation system, which runs on a PC running Linux or UNIX

Medium Mainframe Configuration

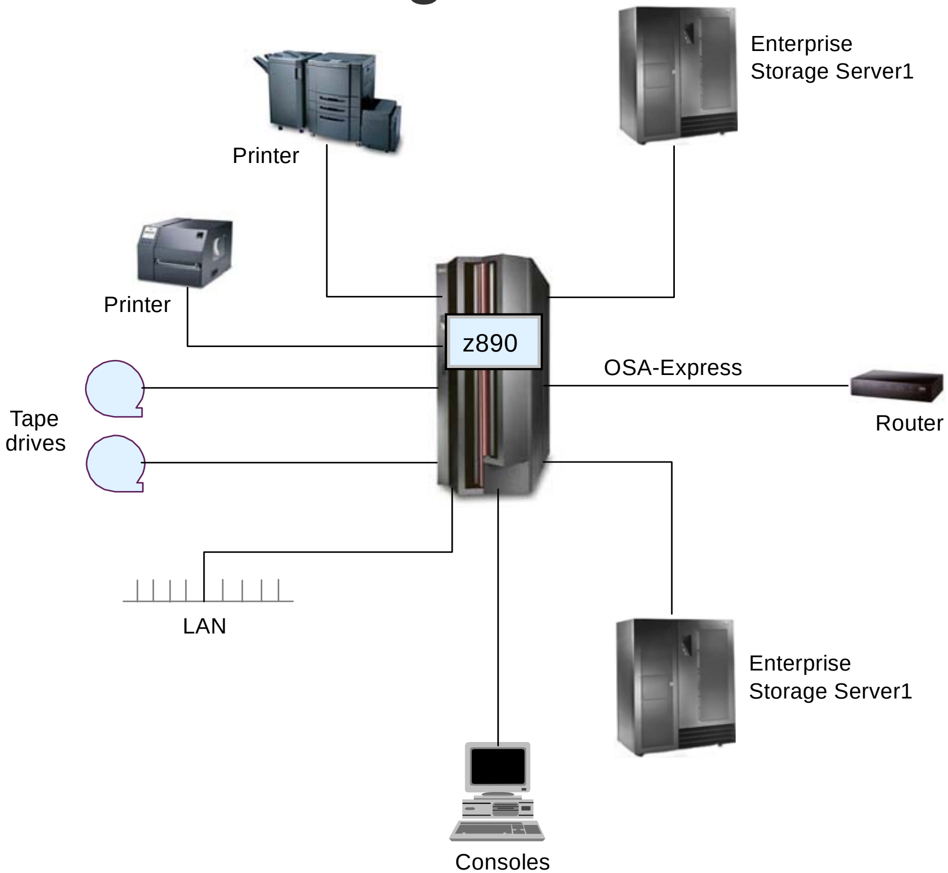

This is a completely modern implementation with no older devices

A z890 is featured with two external disk controllers, tape drives, printers, LAN attachments and consoles

Medium Single System

This system could be running several LPARs

- A production z/OS with interactive applications
- A second production LPAR devoted to batch
- A test z/OS LPAR for testing new software
- One or more LPARs running Linux, perhaps running web applications

Moderately Large Configuration

Two machines – a new z990 and an older 9672, Parallel Sysplex with coupling facility

Multiple ESS and older DASD connected via a switch 3745 communications controllers for an SNA network 3490E tape drives retained for compatibility OSA Express connections to various LANs Consoles

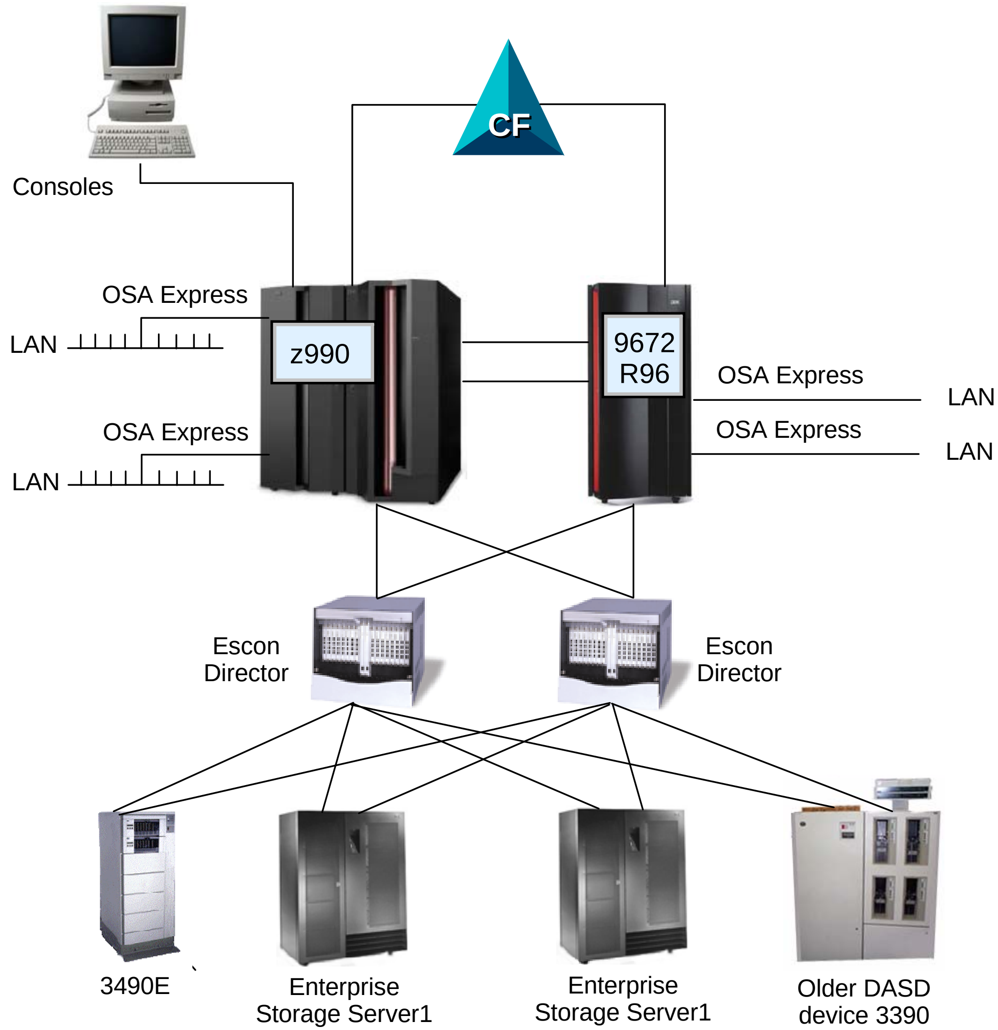

Larger system

- This is the sort of mixture of device generations that might be found in any enterprise
- New devices are brought in but older ones are still viable until they reach the end of their life
- z/OS will usually run on older devices until an architectural change forces their withdrawal
- This sort of change is usually well publicized and several years notice is given of any incompatibility which will usually coincide with the machine being at end of life

Summary

- Terminology is important
- The classic S/360 design is important as all later designs have enhanced it. The concepts are still relevant
- New processor types are now available to reduce software costs
- EBCDIC character set
- Clustering techniques and parallel sysplex

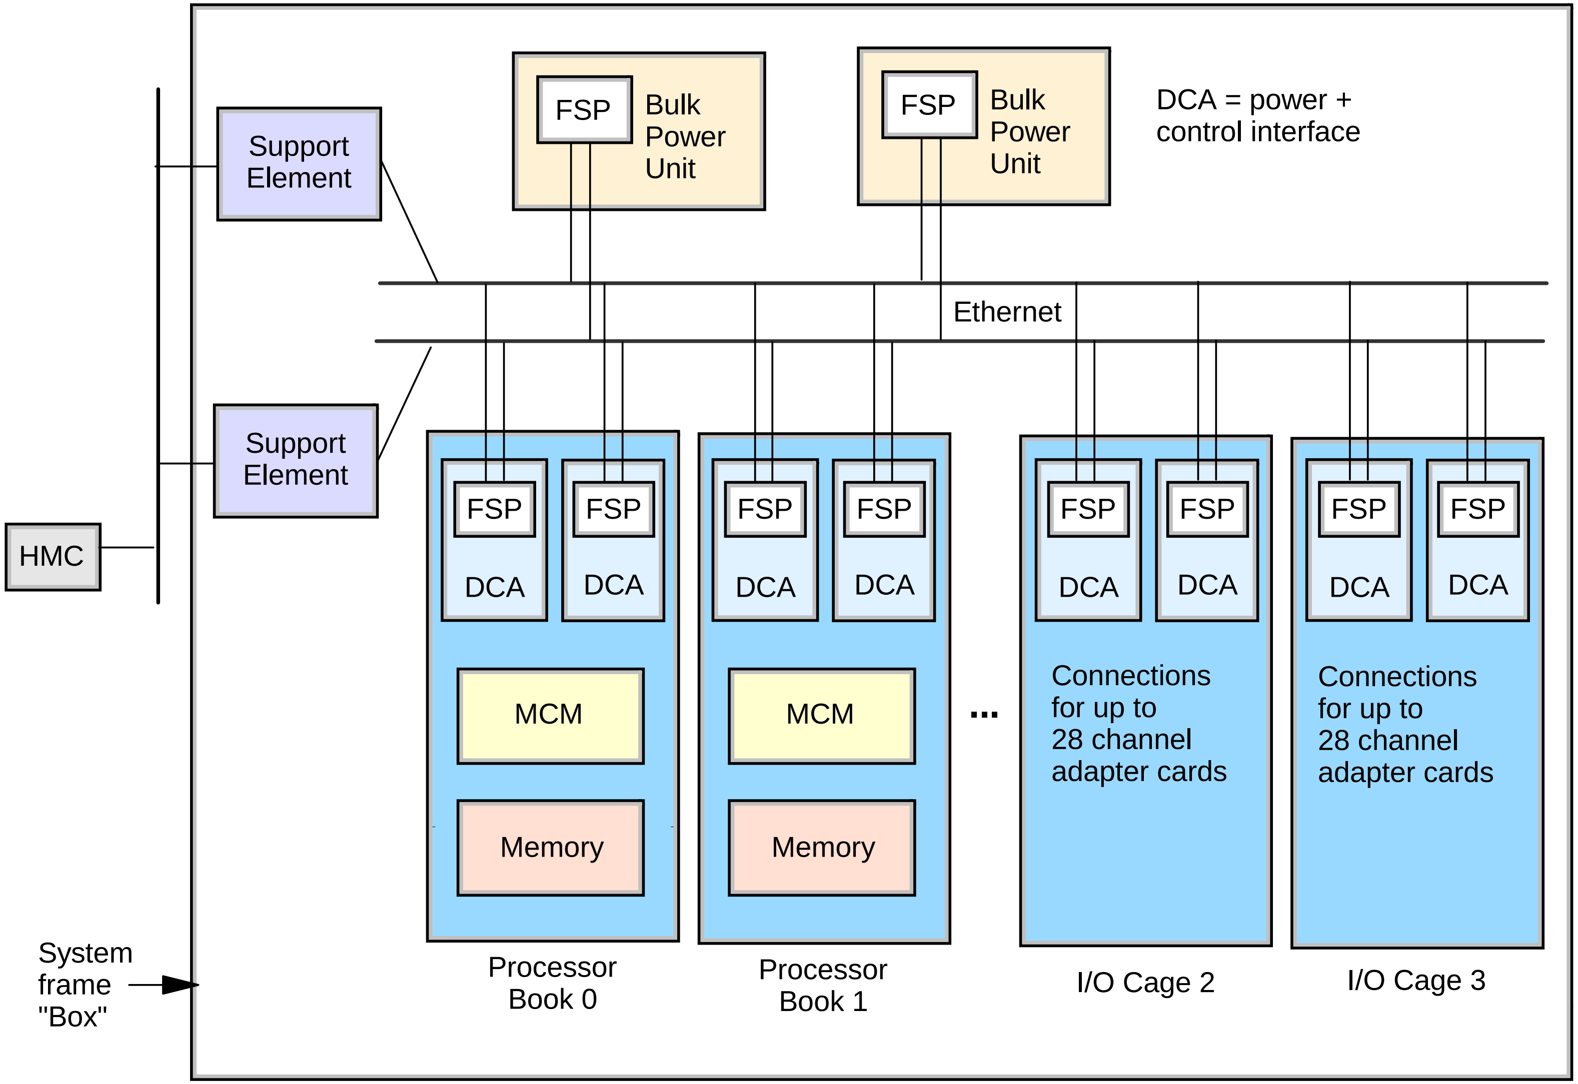
Support Element
FSP
Bulk Power Unit
FSP
Bulk Power Unit
DCA = power + control interface
Ethernet
Support Element
HMC
FSP
DCA
FSP
DCA
MCM
Memory
FSP
DCA
FSP
DCA
MCM
Memory
...
FSP
DCA
FSP
DCA
Connections for up to 28 channel adapter cards
FSP
DCA
FSP
DCA
Connections for up to 28 channel adapter cards
System frame "Box"
Processor Book 0
Processor Book 1
I/O Cage 2
I/O Cage 3