COMP 3400 Mainframe Administration[1]

Christian Grothoff
christian@grothoff.org

http://grothoff.org/christian/

[1]These slides are based in part on materials provided by IBM's Academic Initiative.

README

http://grothoff.org/christian/teaching/2009/3400/

Today

- Administrative Issues
- Introduction to the Mainframe
- Mainframe Hardware

Academic dishonesty

- Webpage says what is allowed.
- If in doubt, ask first.
- Any violation will be reported to the dean.

Expectations

- Read the indicated chapters of the textbook – not every detail is covered in class, but it may still be helpful in exams!
- Study additional material (software documentation, other books, additional textbook chapters) as needed.
- Deliver tested, documented, packaged and working versions of projects on time using subversion.
- Answer questions on quizzes and final exam.

Workload

- 10h/week of studying is probably the minimum needed to pass.
- You may need 20h, 30h or 40h/week; nevertheless, make sure you get enough sleep.
- Projects range in difficulty from hard to impossible.

You can only do well if you treat this as **fun**.

Course Overview

- zSeries hardware (z9, z10), z/VM, z/OS, applications
- We will cover many topics, some somewhat superficially
- Content is centered around IBM hardware and software
- You will need to learn many acronyms

Grading

- 25% Project (individual or teams of two or three students)
- 25% Quizzes (many)
- 50% Final (IBM z/OS Mastery Exam)

Labs will **not** be part of your grade!

Projects

- All projects will require you to study extensive materials not covered in class
- Your TA and instructor have never tried to do these; in fact, you may be the first ever to try
- You can and should still ask for help with them
- Different projects require different skills; pick what you can do and like to do!

Project Choices

- Port libgcrypt to z/OS
- Data Set Explorer
- System Report
- Need for Speed: DB2
- Simple Event Notification

You will be expected to write a report and demo your project in class 18 or 19.

Lectures and Quizzes

- I will try to keep lectures to 1h 50 minutes, but they may be shorter or longer depending on the material and/or student questions.
- Most lectures will have a quiz at the beginning; the quizzes are largely on the previous lecture but can cover all previous lectures.
- Quizzes cover mostly material covered in the lecture but can also cover material only in the textbooks.
- You will have about 1 minute/question or 30s for true/false questions.
- I do not take attendance, but if you miss lectures without official prior excuse you will get zero points on the quiz.

Labs

- There will be several labs and textbook exercises assigned throughout the quarter
- You can work on those individually or in groups
- The labs will not be graded
- You need to do the labs to have a chance to successfully do the projects!

Employment

- If the university hears about specific opportunities, we will pass them on using the `mfa-scholars` mailinglist
- If you pass the final exam, your resume will get listed in an IBM database for employers looking for mainframe administrators
- There is a significant number of open positions worldwide
- We plan on having a job-fair / industry round on campus sometime during the quarter.

Jobs around Mainframes

- Application Developer
- System Programmer
- System Administrator
- System Operator
- Production Control Analyst

Questions

?

What are Mainframes Used for Today?

- Banking (ATM, credit cards, stock exchange)
- Insurance (claims processing)
- Travel (air line reservation systems, car rental)
- Manufacturing (inventory control, scheduling)
- Government (Tax processing, license management)

Why use Mainframes?

Why use Mainframes?

availability

Why use Mainframes?

- Reliability, availability and serviceability (RAS)
- Security
- Scalability, partitioning / virtualization
- Centralized control, workload management
- Continuing compatibility
- Evolving architecture

7th April 1964, Poughkeepsie NY

- A new generation of electronic computing equipment was introduced that day by IBM.
- IBM Board Chairman Thomas J. Watson Jr. called the event the most important product announcement in the company's history.
- The new equipment is known as the IBM System/360.
- S/370, S/390, zSeries/Systemz are S/360's successors

S/360 Evolution

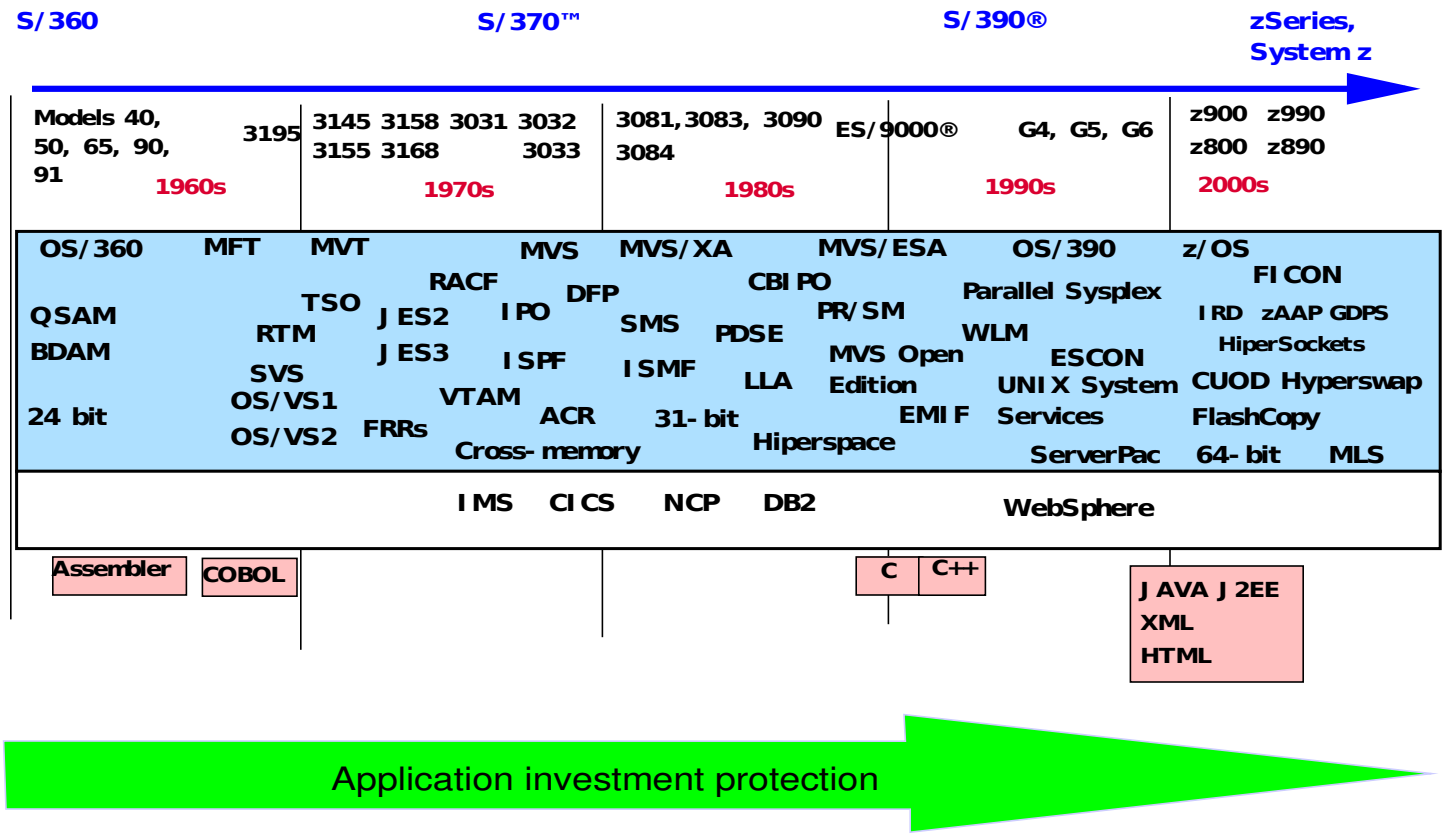

Mainframes in our midst

Mainframes:

- Are prevalent, yet hidden from public eye
- Not often publicized – stable, reliable, dependable

$\Rightarrow$ Not in the news!

What is a mainframe?

A mainframe is a computing system that businesses use to host the commercial databases, transaction servers, and applications that require a greater degree of security and availability than is commonly found on smaller-scale machines.

Mainframe facts

- Most Fortune 1000 companies use a mainframe environment
- 60% of all data available on the Internet is stored on mainframe computers
- There are more CICS transactions processed daily than Web pages served

Typical Workloads

- Large-scale transaction processing (thousands of transactions per second)
- Support thousands of users and application programs
- Simultaneously accessing many resources (terabytes of information in databases)
- Large-bandwidth communications

Batch vs. Transactions

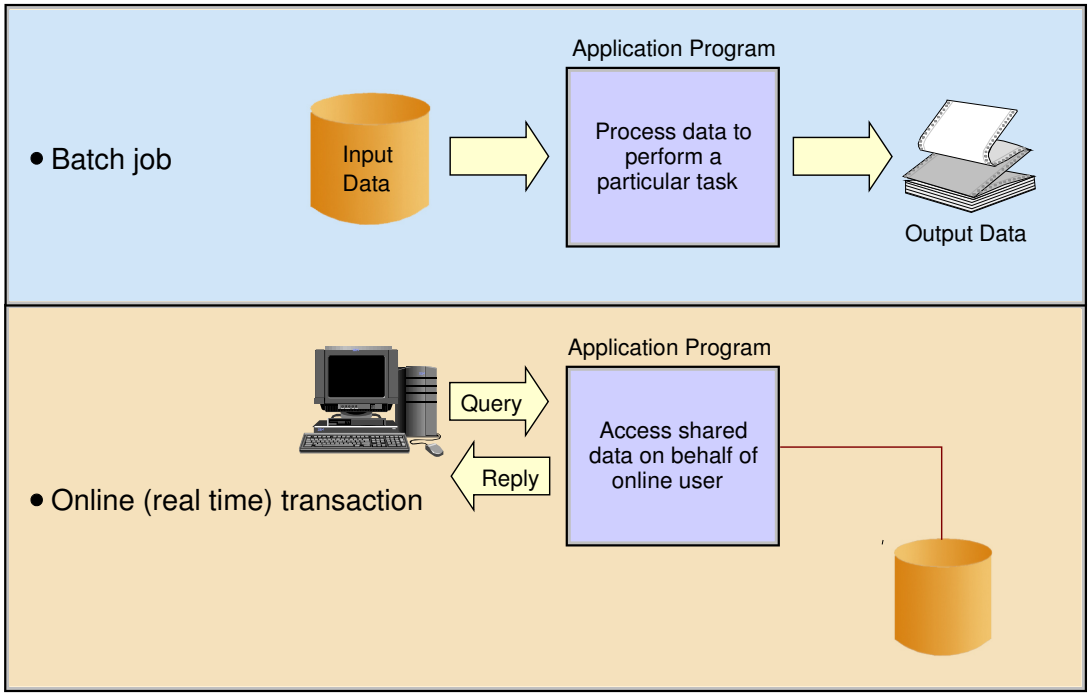

Typical batch use

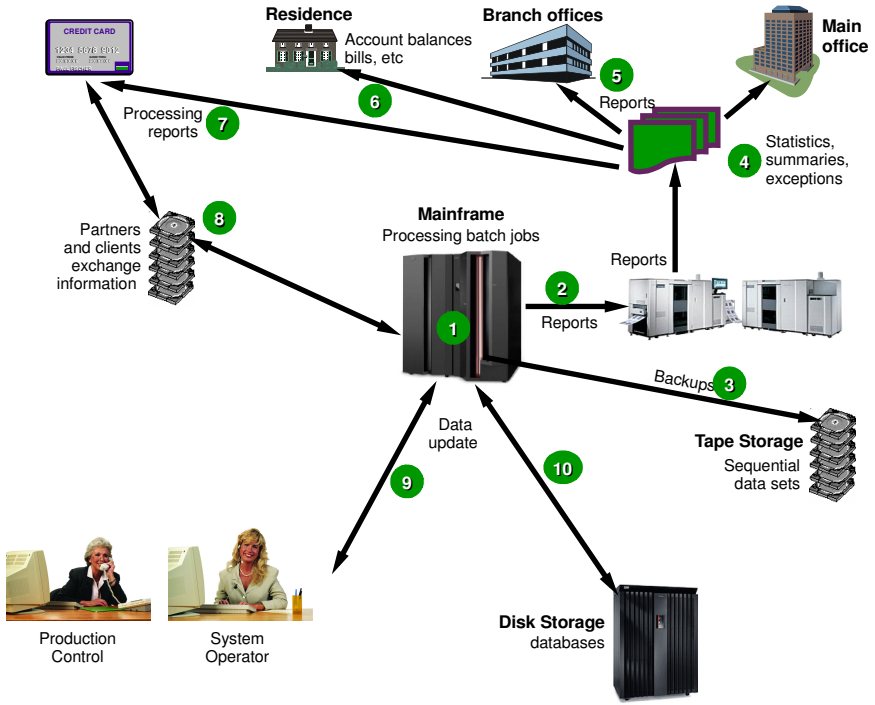

Typical online use

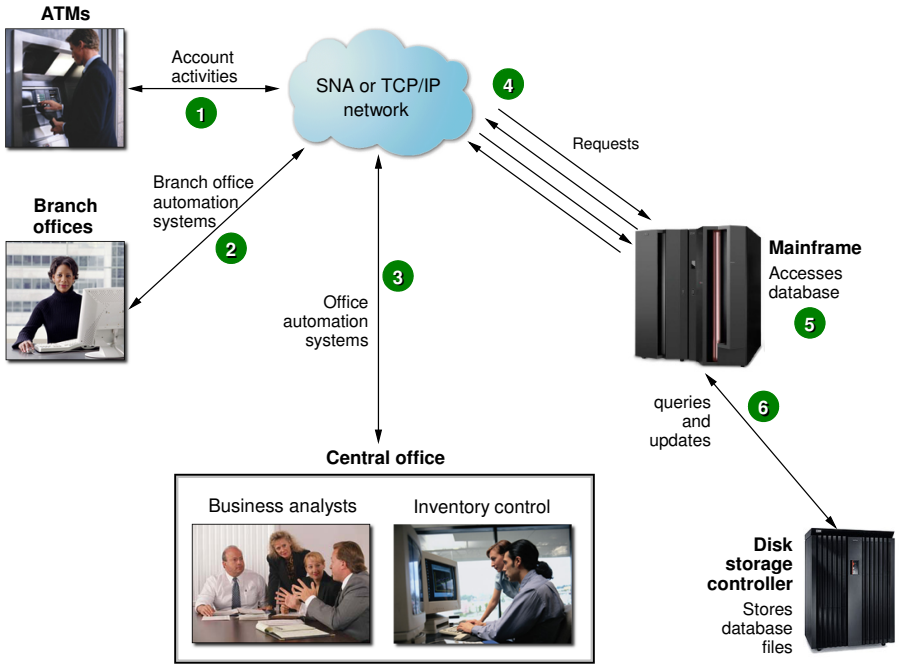

Mainframe Operating Systems

- OS/360, OS/VS, MVS, MVS/XA, MVS/ESA, OS/390, z/OS:
  Enterprise operating system providing a diverse application execution environment (PL/1 to Java)
- DOS/360, DOS/VS, VSE/SP, VSE/ESA, z/VSE:
  Batch and on-line transaction processing (CICS).
- ACP, TPF, z/TPF:
  High-performance transaction processing
- CP/67, VM/370, VM/SP, VM/XA, VM/ESA, z/VM:
  Virtualization for broad OS support and test environments

Questions

?

Mainframe Hardware

Terminology is not straightforward:

- Multiple physical "boxes", multiple processors running multiple cores running multiple operating systems (virtualization) means that the terms "system", "processor" and "CPU" are used for different things by different people

Confusing terminology

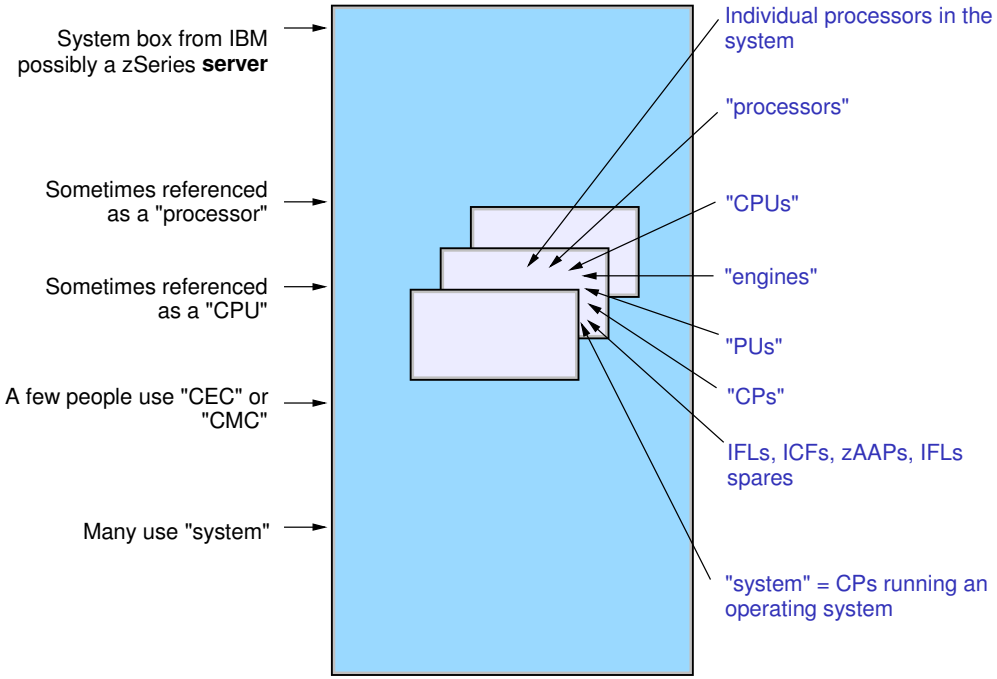

Conceptual S/360

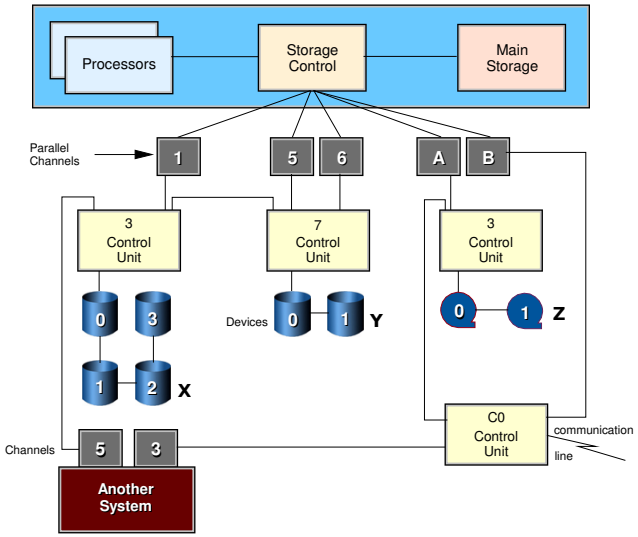

Parallel Channels on the S/360

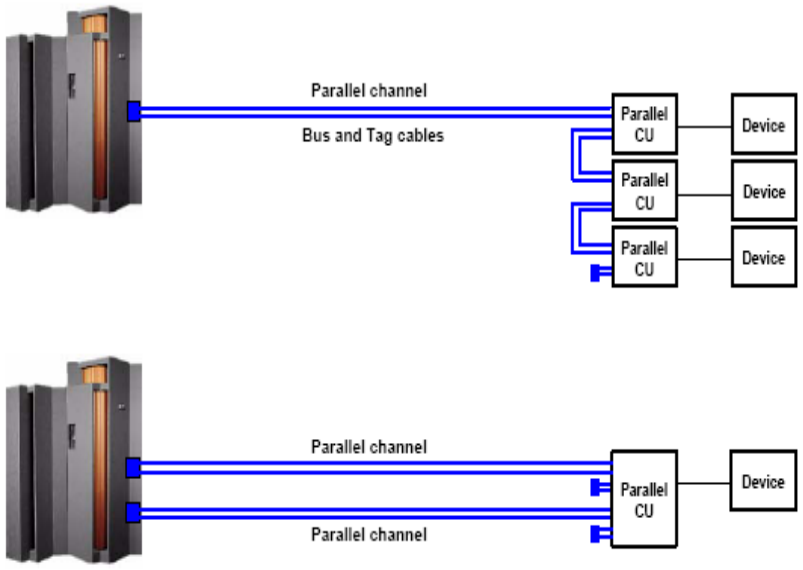

The maximum data rate of the parallel channel was 4.5 MB, the maximum distance was 122 meters (400 ft).

ESCON Connectivity

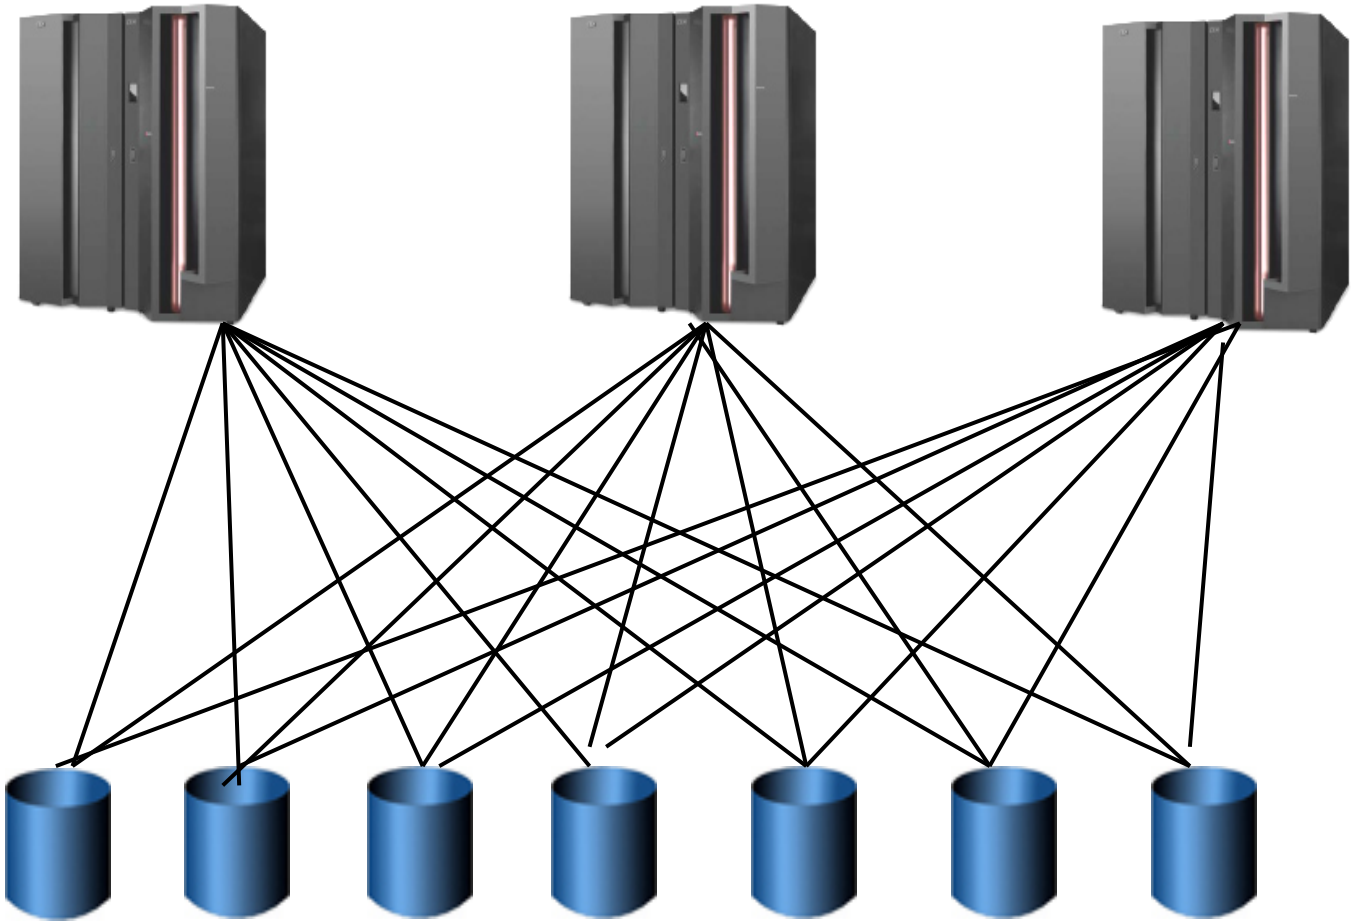

ESCON Directors

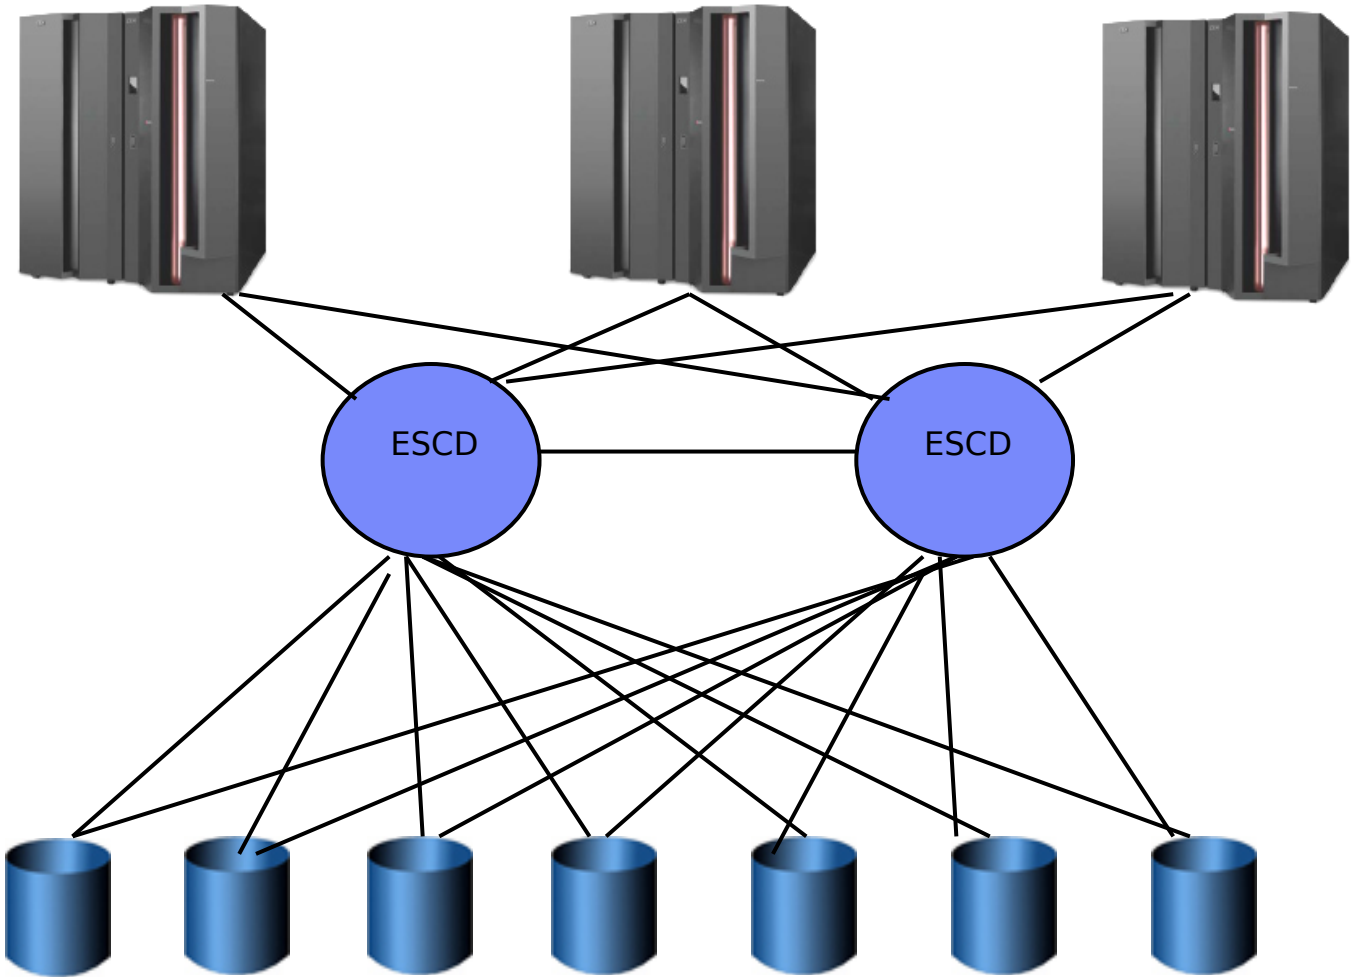

ESCON vs. FICON

	ESCON	FICON
Speed	20 MB/s	400-800 MB/s
Parallelism	1 request	64 requests
Target Control Units	1 target	64 targets
Cascading switches	no	yes
Cable type	electric	fiber optic

FICON

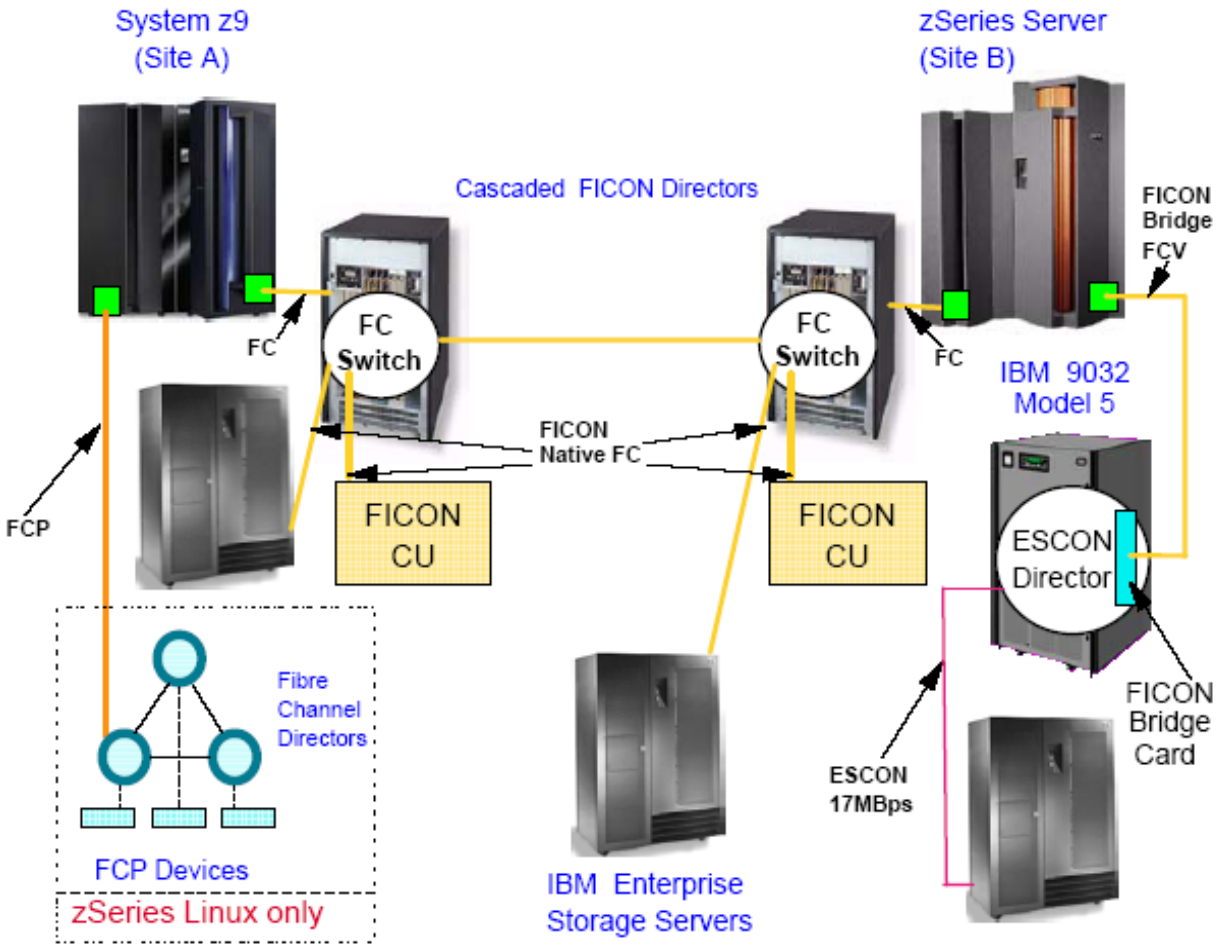
System z9
(Site A)
zSeries Server
(Site B)
Cascaded FICON Directors
FICON
Bridge
FCV
FC
FC
Switch
FC
Switch
FC
IBM 9032
Model 5
FICON
Native FC
FCP
FICON
CU
FICON
CU
ESCON
Director
Fibre
Channel
Directors
FICON
Bridge
Card
ESCON
17MBps
FCP Devices
zSeries Linux only
IBM Enterprise
Storage Servers

IBM 3390

Current mainframes use 3390 disk devices (DASDs).

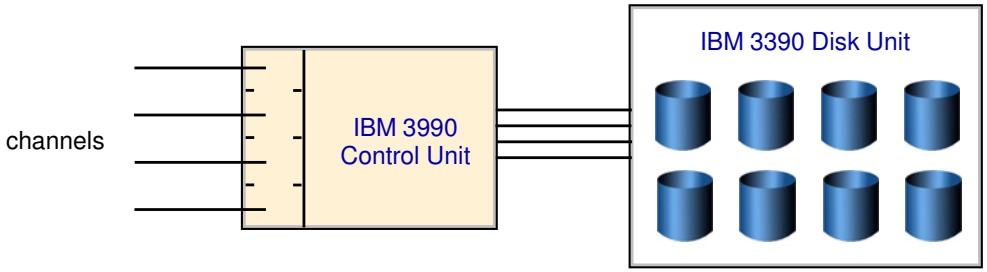

Modern 3390 Devices

- Physical disks are commodity SCSI-type units
- Usual configuration 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 extend the original 3390 architecture with new features such as improved copying mechanisms and parallel access to improve performance
- To the operating system, they still appear as traditional 3390 devices

Identifying I/O Devices

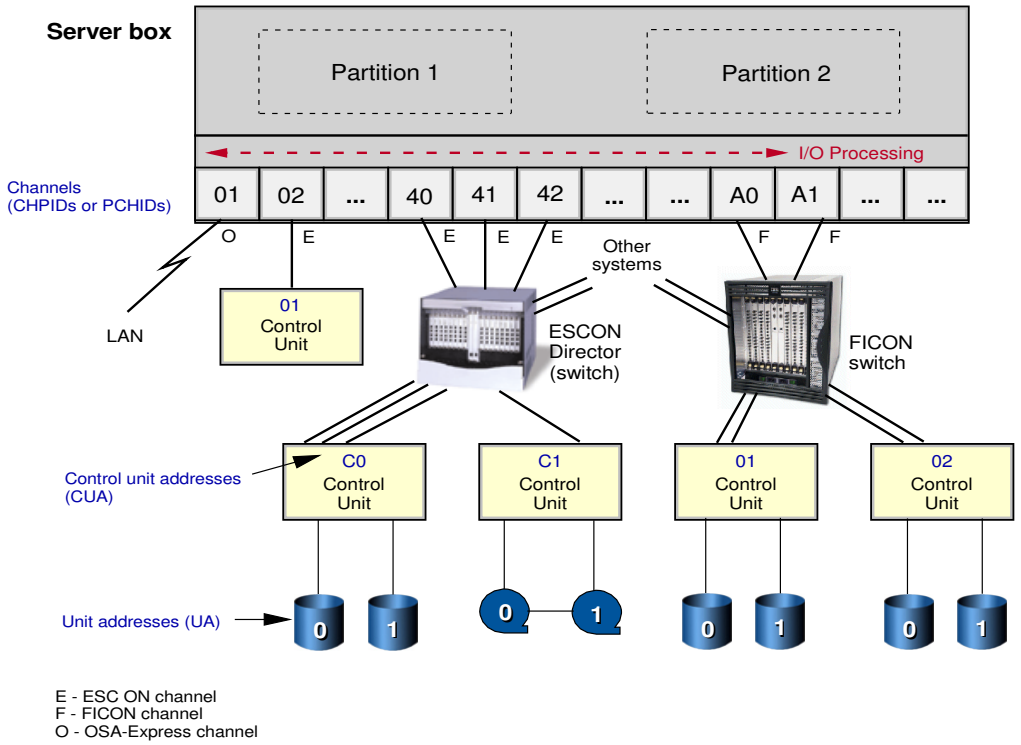

Identifying I/O Devices

- CHPID addresses are two hex digits; LPARs can share CHPIDs
- I/O subsystem layer exists between OS and the CHPIDs, configured using control file `IOCDS`
- Device numbers are three or four hex digits assigned by the system programmer when creating the `IODF` and `IOCDS` configurations

Logical Partitions (LPARs)

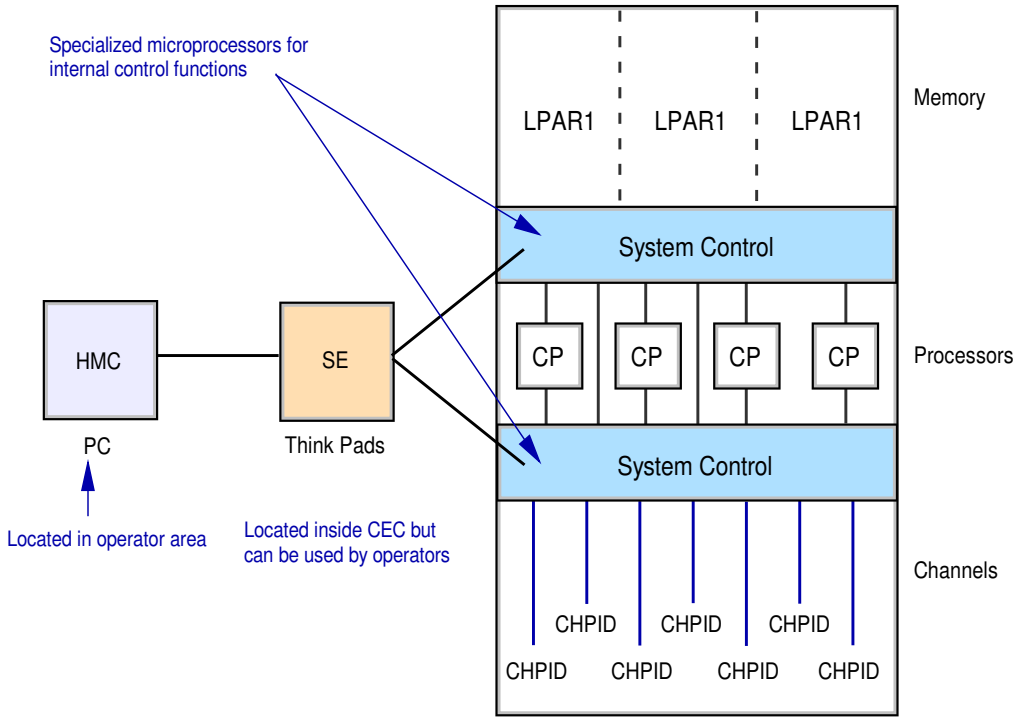

Partitioning

- The first MVS systems allowed physical partitioning
- Modern mainframes allow logical partitioning: Virtual machines in hardware!
- Initial limit was 15 LPARs, modern machines allow 60 LPARs
- Each LPAR is assigned memory, I/O and processing power (in `IOCDS`)
- Changing LPAR resource assignments usually requries power-on reset (POR)!

LPAR Characteristics

- Each LPAR runs its own operating system
- Devices can be shared across several LPARs (through the Multiple Image Facility, MIF)
- Processors can be dedicated or shared
- LPARs are independent and are the equivalent of a separate mainframe for most practical purposes

Processing Units

- Central Processor (CP) – a full z/OS processor
- System Assistance Processor (SAP) – used for I/O subsystem, each machine has at least one
- Integrated Facility for Linux (IFL) – Special processor for Linux
- zAAP – Special processor for Java (on z/OS)
- zIIP – Special processor for DB2
- Spares

The actual hardware for all of these processors is always the same!

Shared CPs

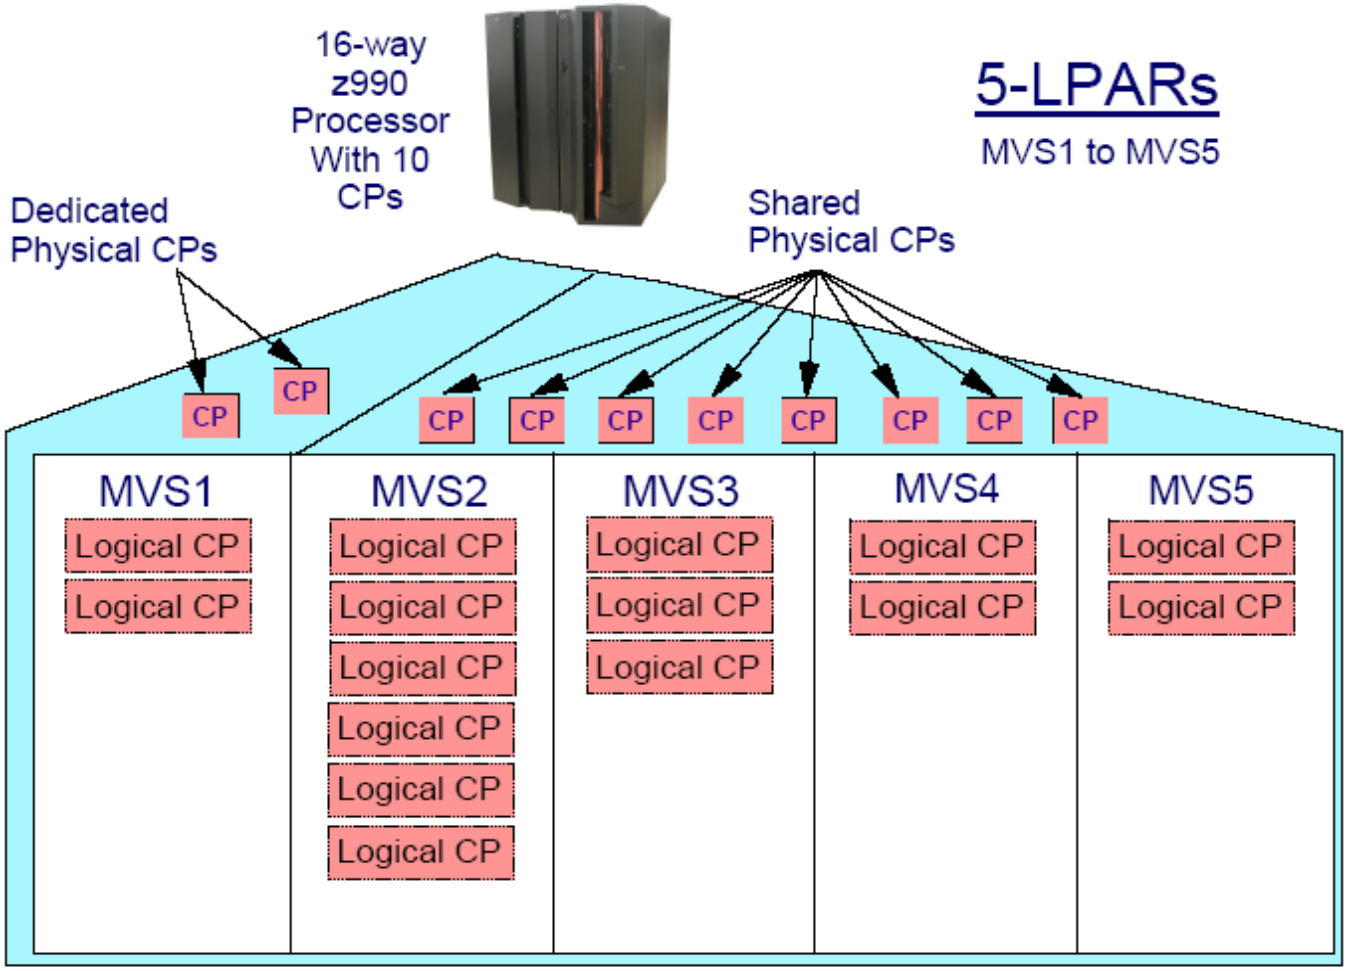
16-way
z990
Processor
With 10
CPs
5-LPARs
MVS1 to MVS5
Dedicated
Physical CPs
Shared
Physical CPs
CP
CP
CP
CP
CP
CP
CP
CP
CP
CP
MVS1
Logical CP
Logical CP
MVS2
Logical CP
Logical CP
Logical CP
Logical CP
Logical CP
Logical CP
MVS3
Logical CP
Logical CP
Logical CP
MVS4
Logical CP
Logical CP
MVS5
Logical CP
Logical CP

Sharing Example

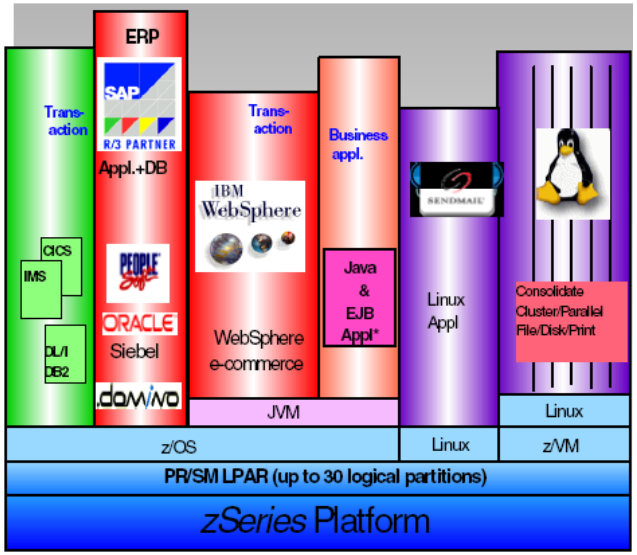

z9-109 – Under the covers

More Terms

- MCM: multi-chip module – glass ceramic containing multiple CPs (5 for z10) and I/O support elements (2 for z10)
- MCM and memory are together in a $book$ (together with cooling)
- Up to four $books$ are in a $CEC\ cage$
- CEC: Central Electrical Complex

z10 MCM

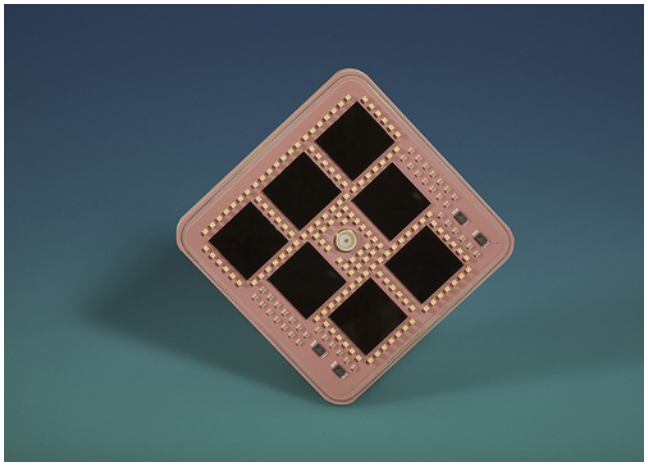

Clustering

- Clustering is done in several forms:
  - Basic shared DASD
  - CTC/GRS rings
  - Basic and Parallel sysplex
- The term *image* describes a single (z/OS) system, which might be standalone or an LPAR on a large box

Basic shared DASD

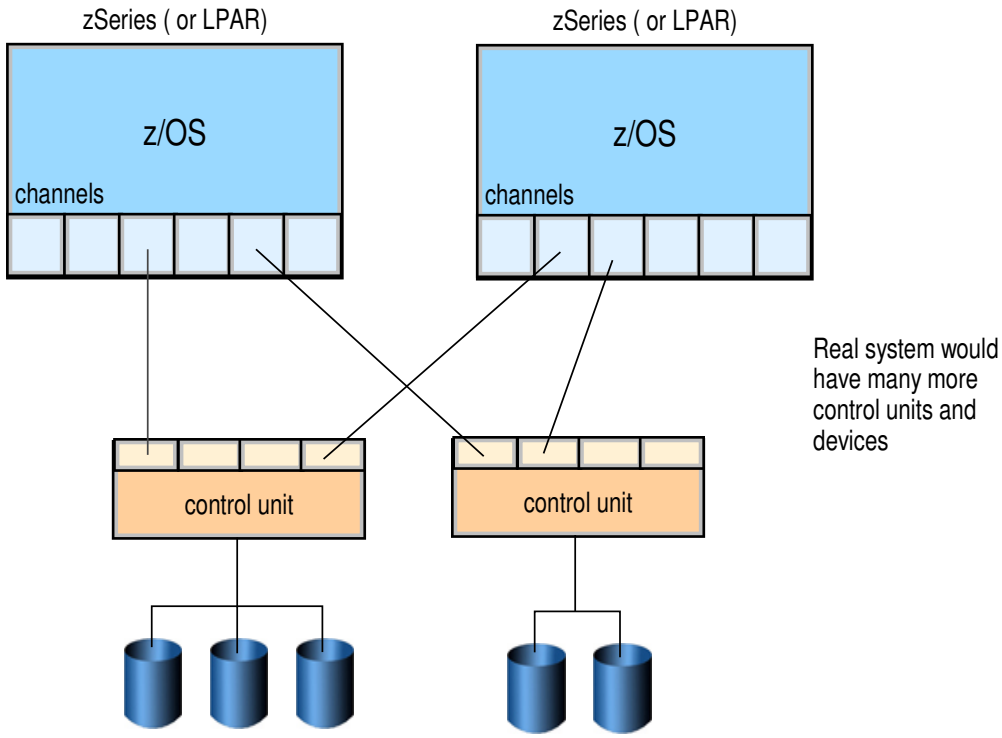

RESERVE and *RELEASE* commands limit access to an entire DASD for the duration of an update.

CTC and GRS

- A Channel-to-Channel (CTC) connection between systems can be used to exchange locking information to avoid the $RESERVE/RELEASE$ overheads.
- Job queue information and security controls can also be transmitted using the CTC ring.
- The z/OS component managing the CTC ring is called the "global resource serialization" (GRS) function.

A CTC Ring (or GRS Ring)

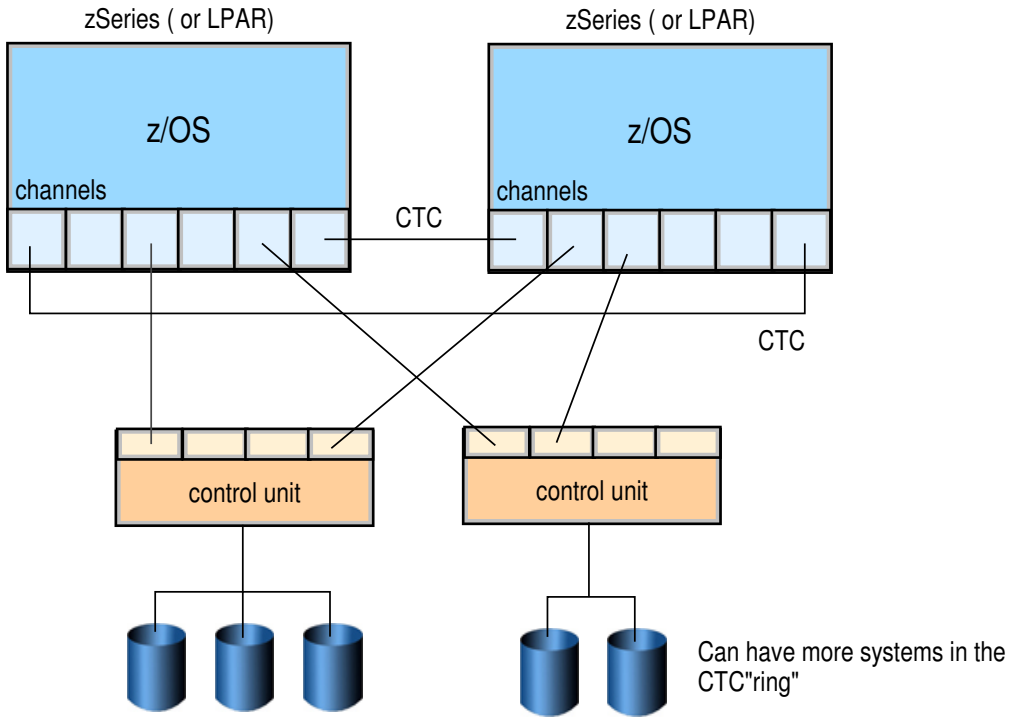

Parallel Sysplex

- Extension of the CTC ring that uses a dedicated Coupling Facility (CF) to store data for GRS.
- The Coupling Facility can also be used to share application data (request queues, DB2 tables, locks)
- Up to 32 systems in a Parallel Sysplex can appear as a single system and be managed as such (single system image)

Why use a Parallel Sysplex?

- Even higher availability
- Greater capacity, easier to grow and upgrade
- Workload balancing across systems
- Systems are still individual

$\Rightarrow$ Software and hardware upgrades can be rolled through one system at a time

Parallel Sysplex

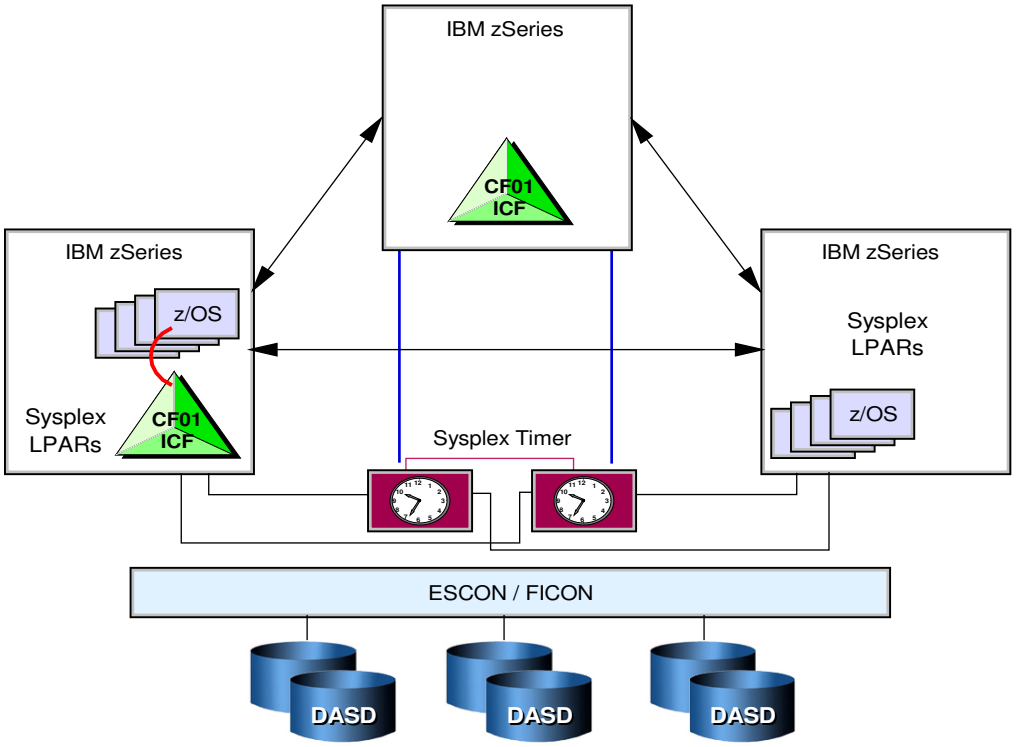

The Coupling Facility

The CF essentially helps with communication:

Cache structure Provides buffer invalidation to ensure consistency of cached data. The cache structure also provides a high-speed buffer for storing shared data.

List structure Enables sharing of data organized in lists, for example shared work queues and shared status information.

Lock structure Provides shared and exclusive locking capability for serialization of shared resources

XCF and XES are the z/OS components managing z/OS–CF interactions.

CF Illustrated

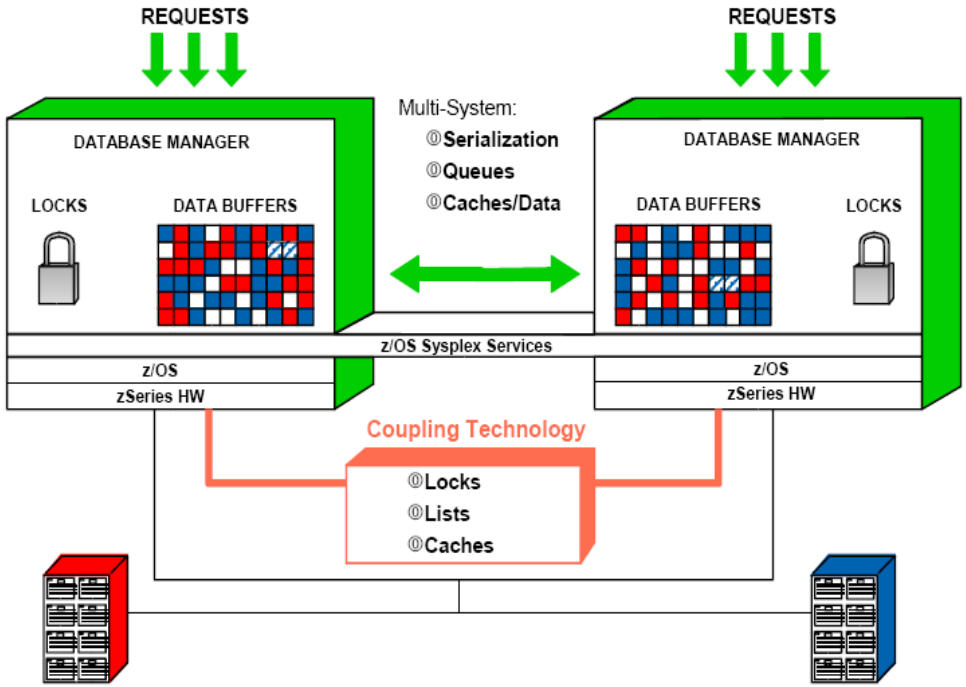

CF Users

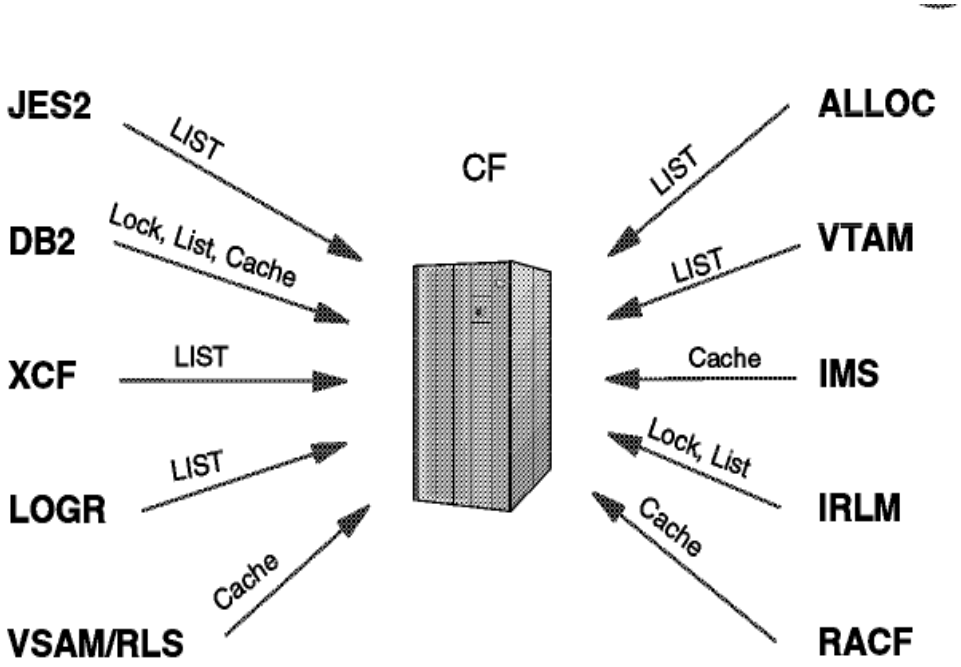
JES2
LIST
DB2
Lock, List, Cache
XCF
LIST
LOGR
LIST
VSAM/RLS
Cache
CF
ALLOC
LIST
VTAM
LIST
IMS
Cache
IRLM
Lock, List
RACF
Cache

GDPS

A geographically dispersed parallel sysplex (GDPS) further improves availability.

- Two sites up to 100 fiber kilometers apart maybe connected for synchronous updates
- Asynchronous techniques can be used for “arbitrary” distances

Global Mirroring

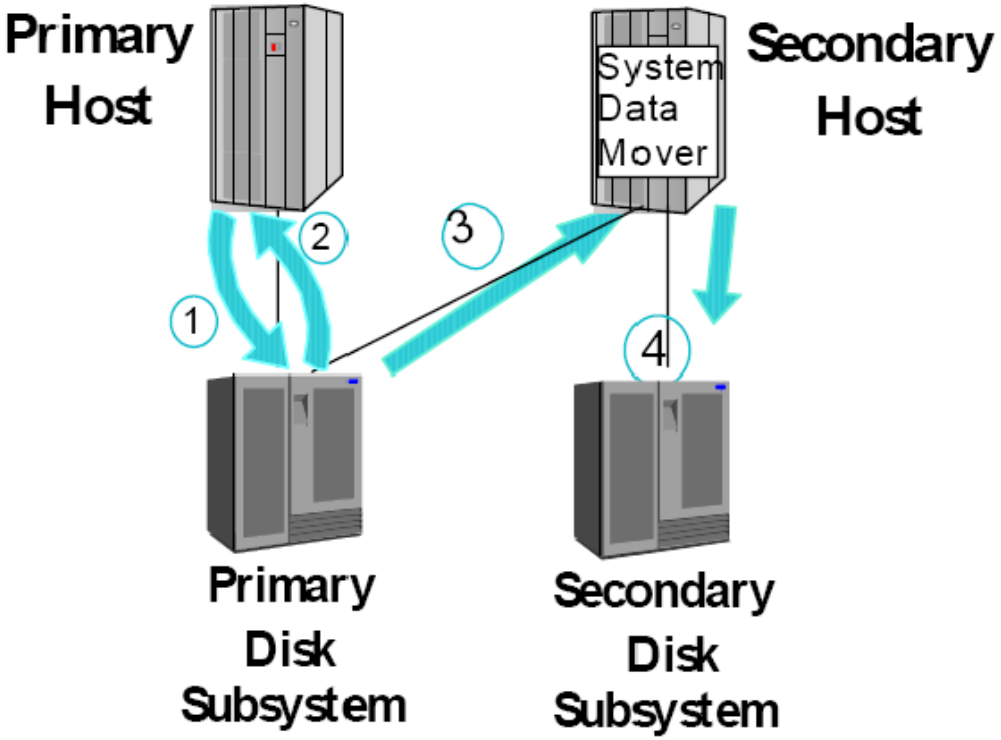

Capacity on Demand

- Additional processing power to meet unexpected growth or sudden demand peaks
- CBU – Capacity Back Up
- OOCUoD – On/Off Capacity Upgrade on Demand
- SubCapacity Licensing Charges
- LPAR CPU Management (IRD)

Harware Scalability

Number of z990 CPs	Base	Ratio
1 CP	z900-2C1	1.54 - 1.61
8 CPs	z900-2C8	1.52 - 1.56
16 CPs	z900-216	1.51 - 1.55
32 CPS	z900-216	2.46 - 2.98

Intelligent Resource Director (IRD)

- Integrates with Workload Management, Parallel Sysplex, PR/SM and the Channel Subsystem
- Moves physical resources to priority workloads

$\Rightarrow$ View a cluster of LPARs on a zSeries as a single pool of computing resources

Key Software Components

- Workload management (WLM)
- Automated Restart Manager (ARM)
- Sysplex failure manager (SFM)

A Typical Moderately Large System

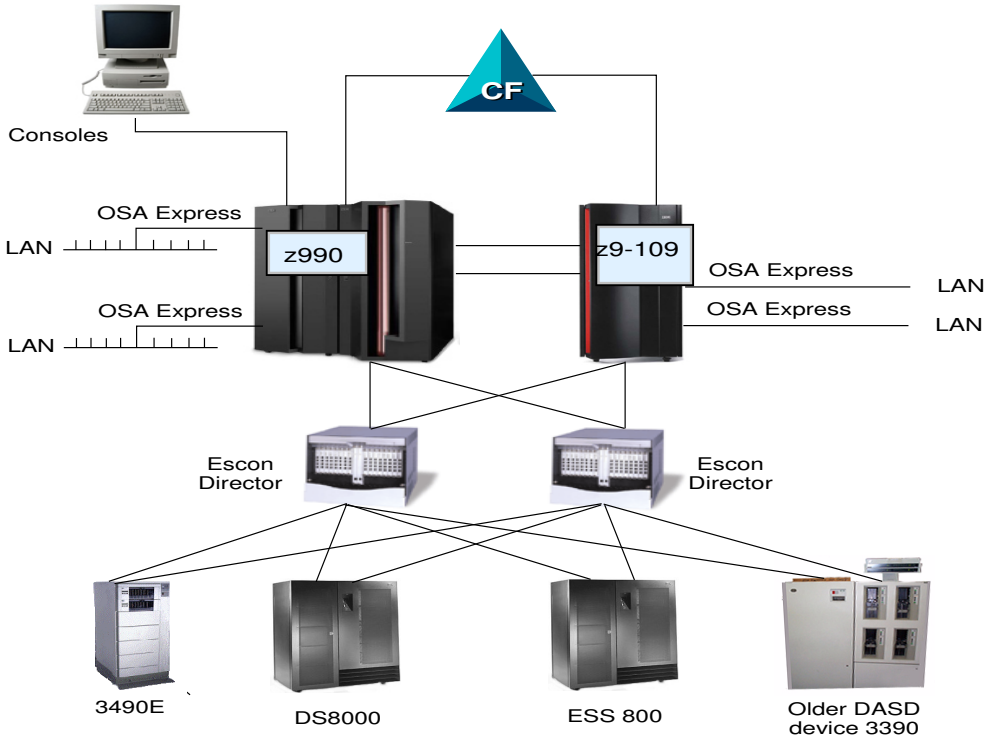

Homework Summary

Before the next lecture:

- Generate password with `htpasswd` and register account.
- Read the chapters in the z/OS and z/VM textbooks indicated on the webpage.
- Install TN3270 emulator software (or use department machines).

Questions

?

PRE-Resource Sharing

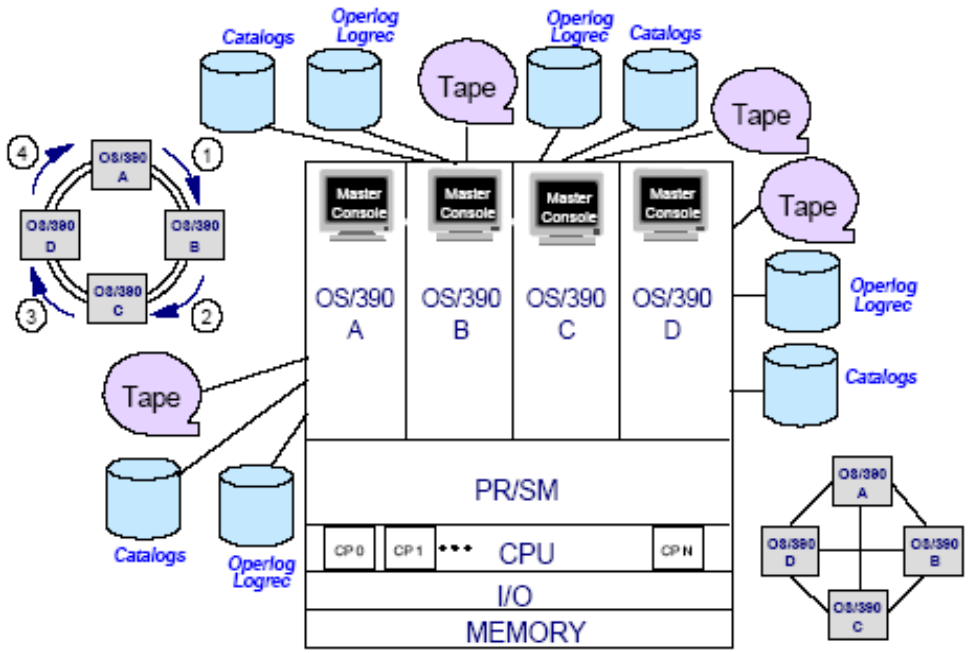

Resource Sharing Today

Even with a single CEC, a CF can still provide benefits: