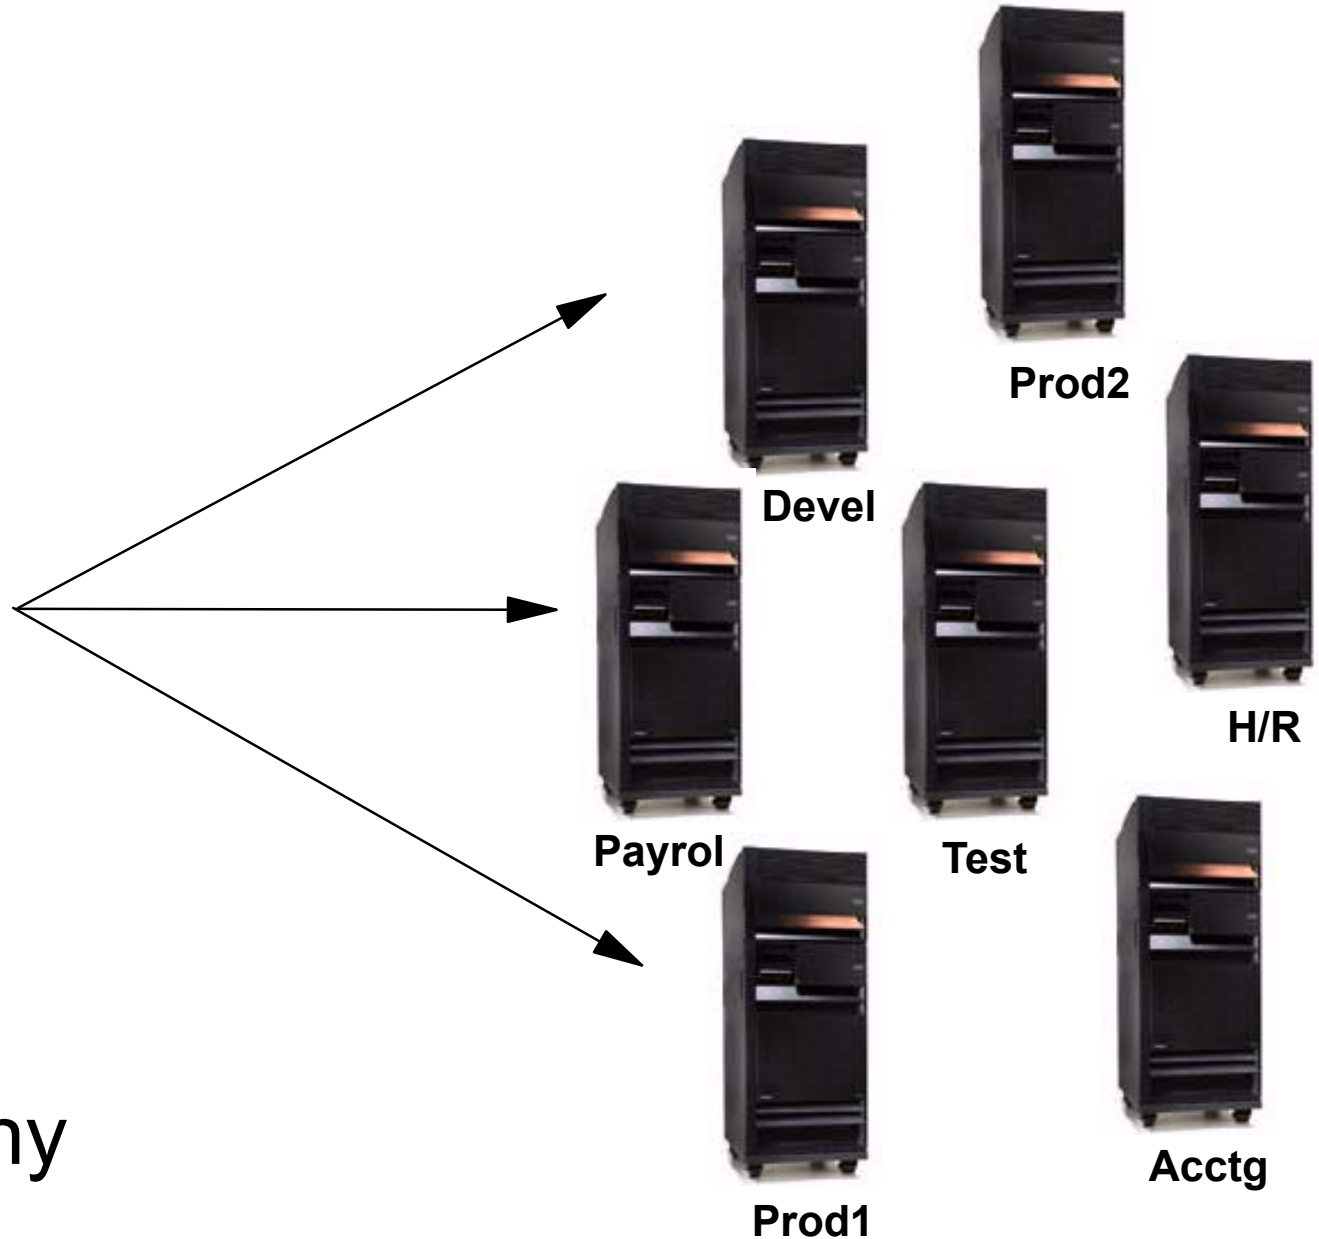


IBM **e**server
iSeries Technology Center

Introduction to Logical Partitioning aka LPAR



Serial # 100D3FL



- From one - many

What is it?

- ▶ LPAR = Logical (and physical) Partitioning
- ▶ Logical Partitioning = The allocation of system resources to create logically separate systems within the same physical footprint
- ▶ System resources include:
 - Processors
 - Main Storage
 - Interactive Performance
 - System buses
 - Disk
 - I/O controllers and devices

Terminology:

▶ Primary Partition

The partition which provides certain general functions on which all partitions are dependent.

▶ Secondary Partition

Any partition that is not the primary partition.


▶ Partition Identifier

A numeric value which represents a particular partition. Primary partition is always 0

▶ Hypervisor

A kernel which enables partitioning. Consists of a new layer of code (PLIC) along with parts of SLIC from the primary partition.

What minimum resources does an AS/400 need to be able to?



```
Sign On
System . . . . . : TSCLPAR3
Subsystem . . . . : QINTER
Display . . . . . : DSP28


User . . . . . _____
Password . . . . . _____
Program/procedure . . . . . _____
Menu . . . . . _____
Current library . . . . . _____

RELEASE: V05R02M00
DRIVER: 20N0802

USE OF THIS SYSTEM IS FOR IBM MANAGEMENT APPROVED PURPOSES ONLY.
USE IS SUBJECT TO AUDIT AT ANY TIME BY IBM MANAGEMENT.

(C) COPYRIGHT IBM CORP. 1980, 2002.
```

What minimum resources does an AS/400 need to be able to?



```

Sign On
System . . . . . : TSCLPAR3
Subsystem . . . . : QINTER
Display . . . . . : DSP28


User . . . . . : _____
Password . . . . . : _____
Program/procedure . . . . . : _____
Menu . . . . . : _____
Current library . . . . . : _____

RELEASE: V05R02M00
DRIVER: 20N0802




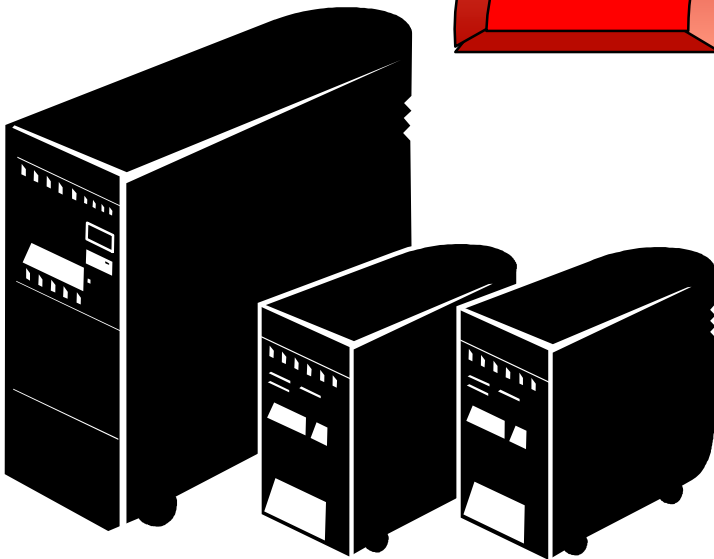
USE OF THIS SYSTEM IS FOR IBM MANAGEMENT APPROVED PURPOSES ONLY.
USE IS SUBJECT TO AUDIT AT ANY TIME BY IBM MANAGEMENT.

(C) COPYRIGHT IBM CORP. 1980, 2002.
    
```

1. Processor(s)
2. Memory
3. Interactive (?)
4. IOP
5. Console - IOA
6. Disk Drive(s) - IOA
7. LIC/OS
- LPPs - Apps (?)
8. CD/DVD/Tape



Logical Partitioning





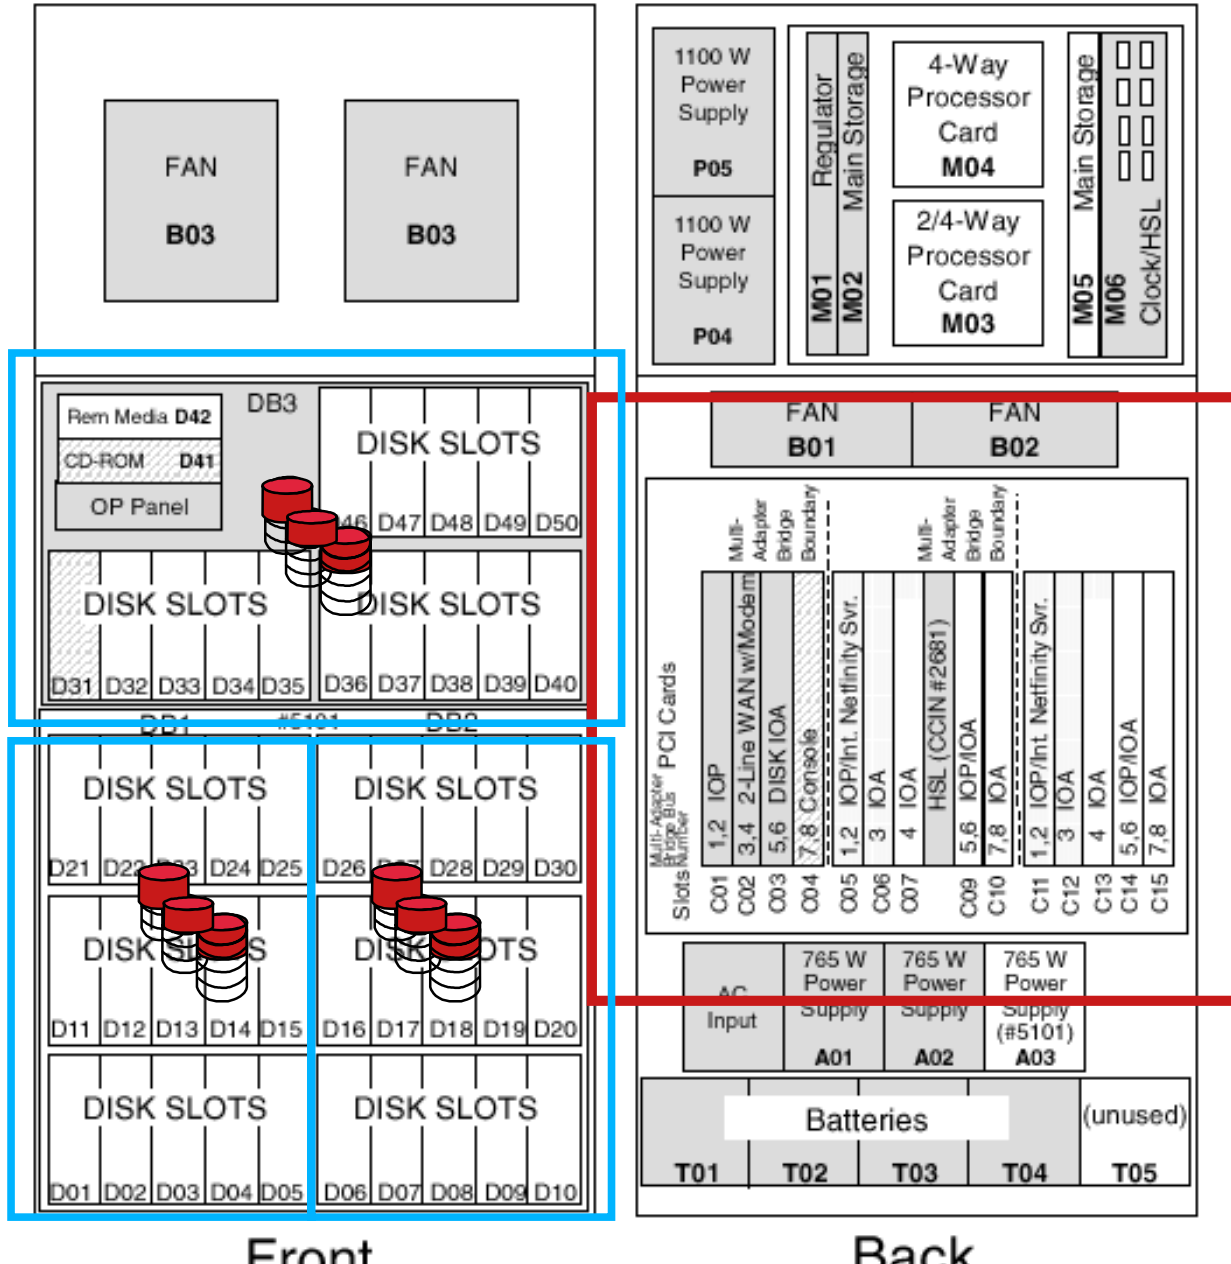
LPAR: Introduction

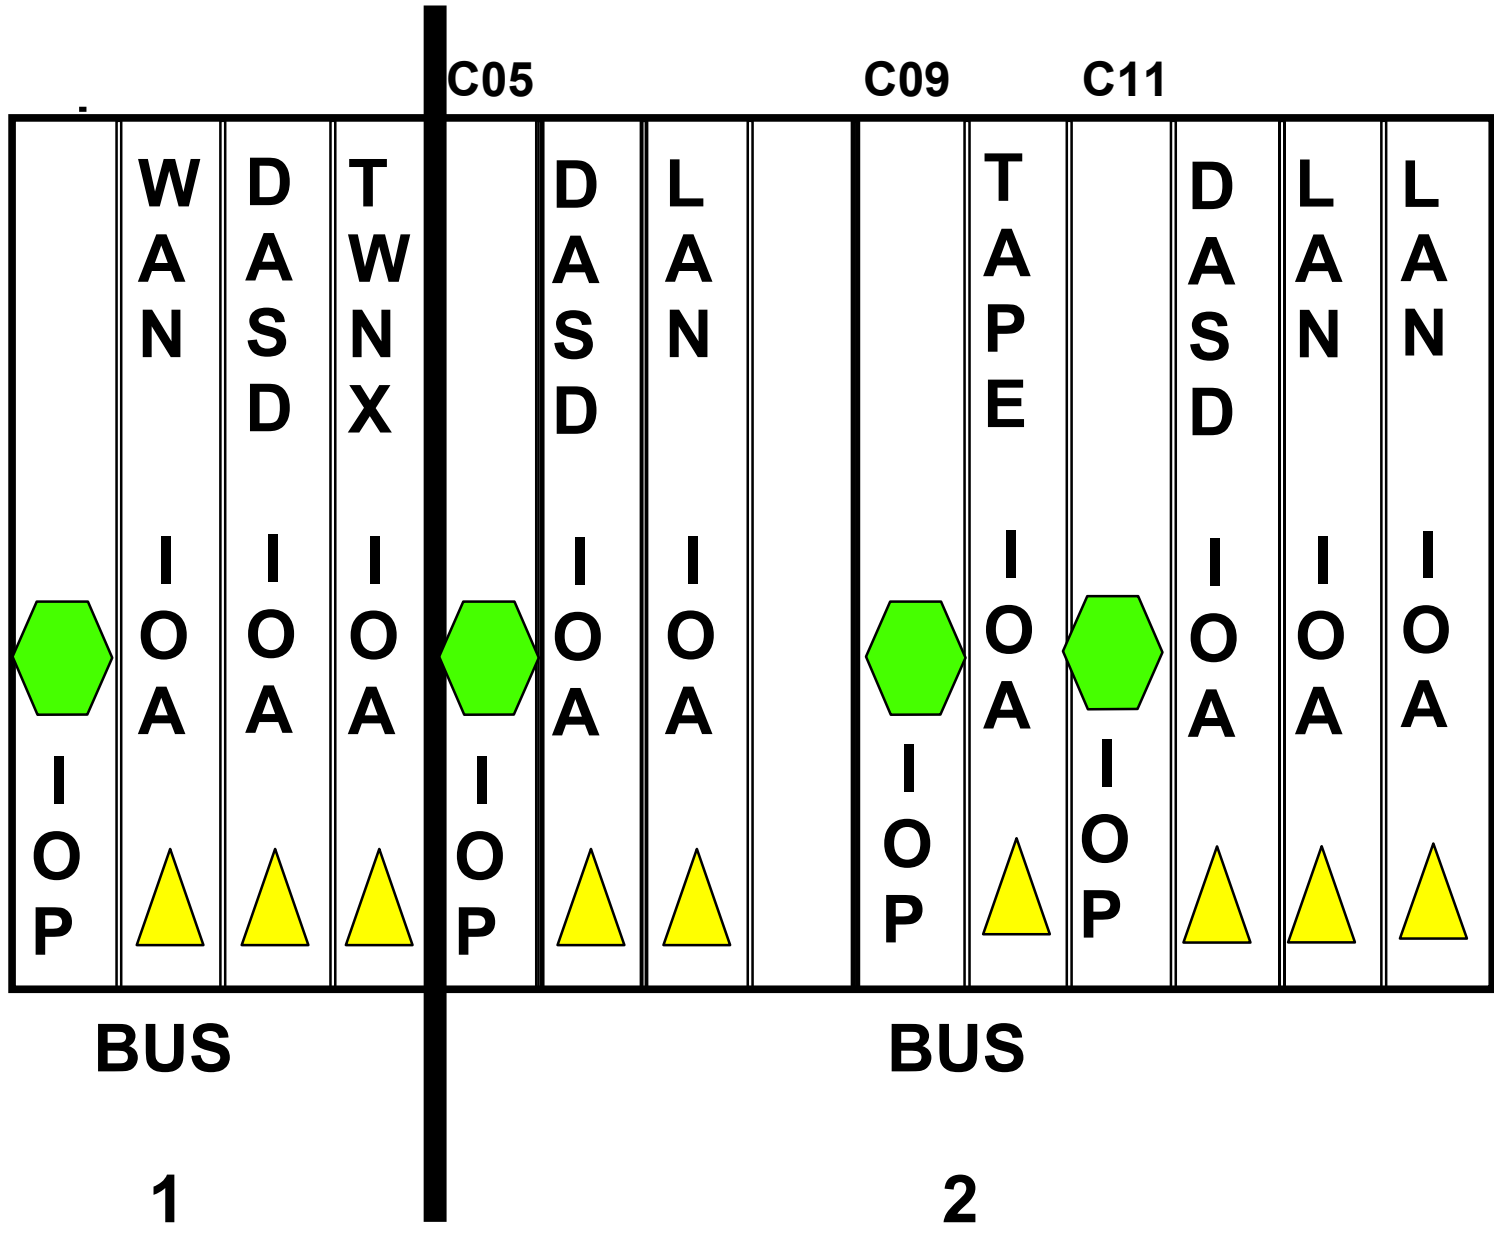
Model 830 System Unit (n - Way)

#2400, #2402, #2403 Processors

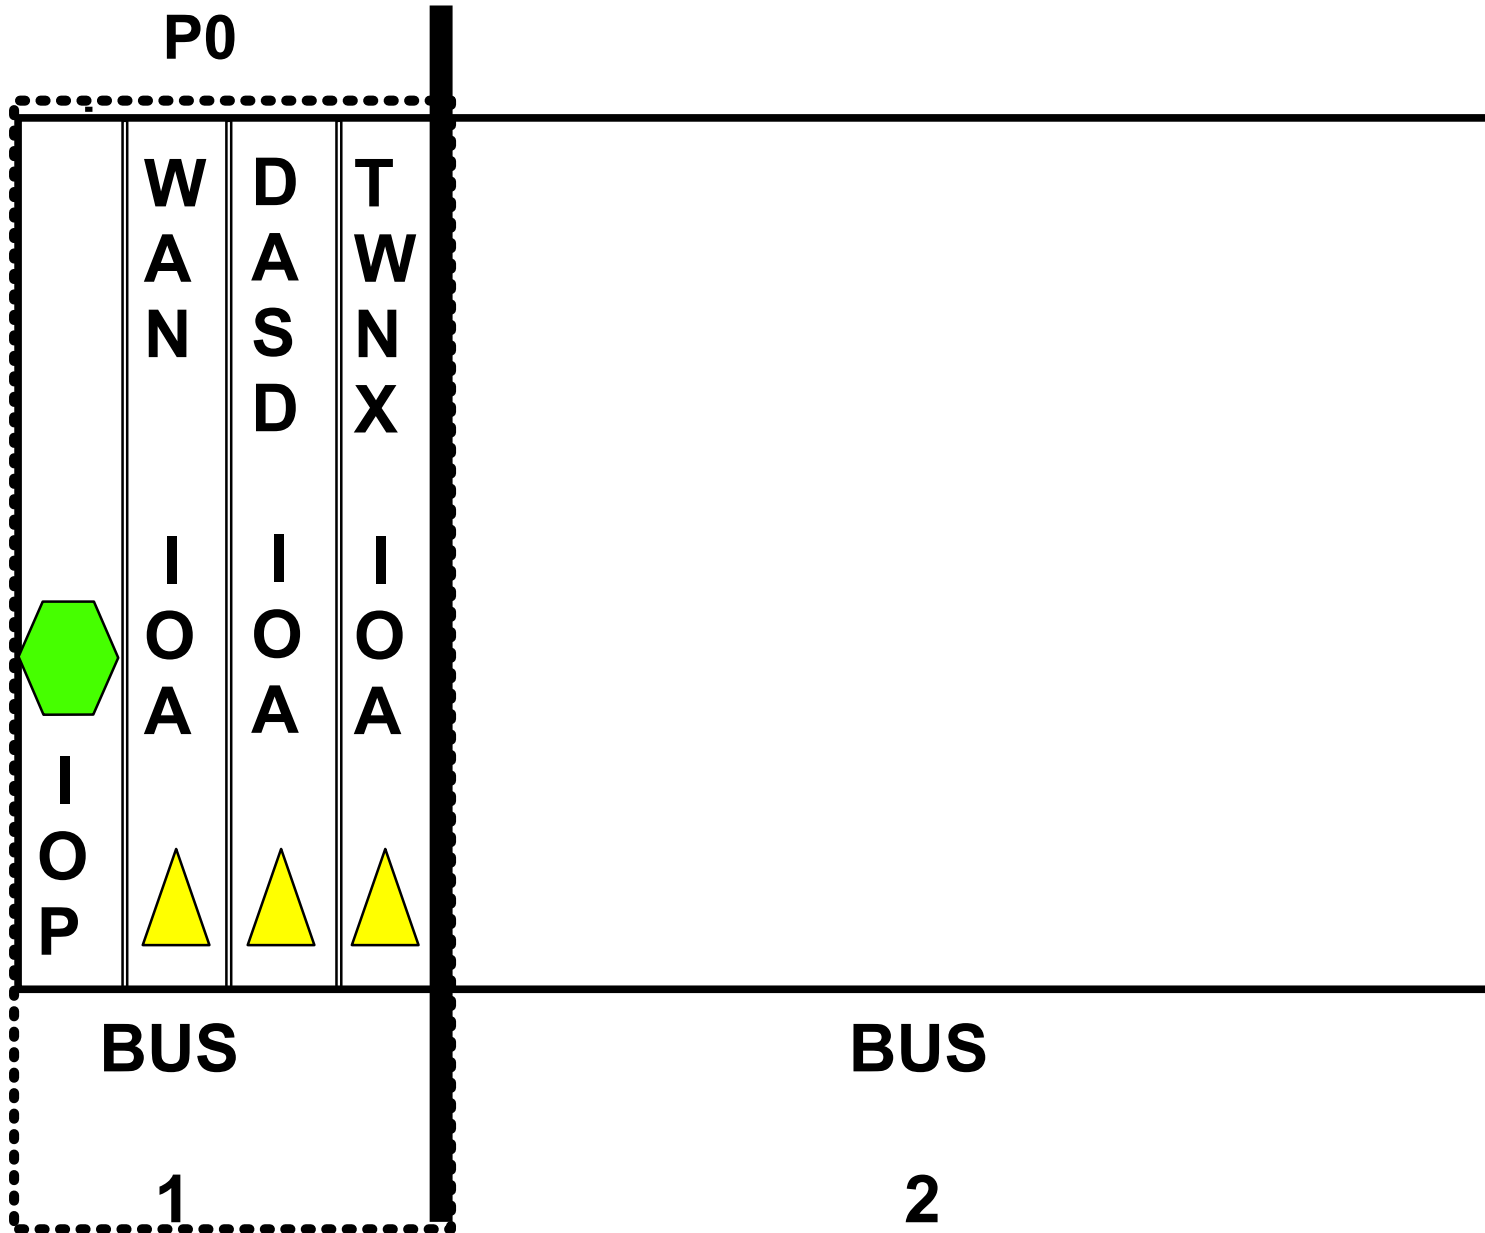


How to get
3 partitions
out of a single
830

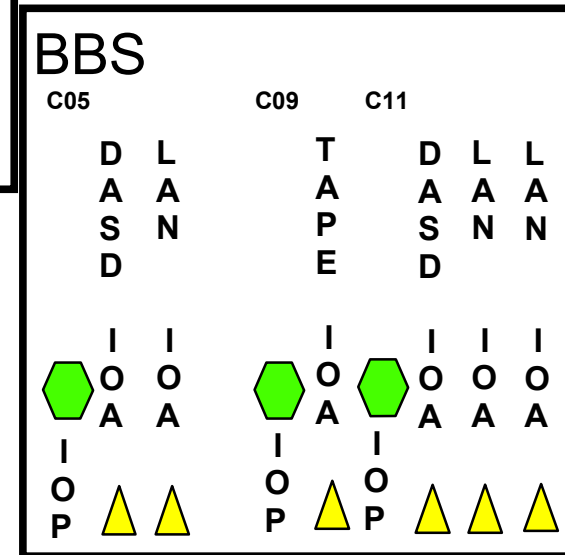


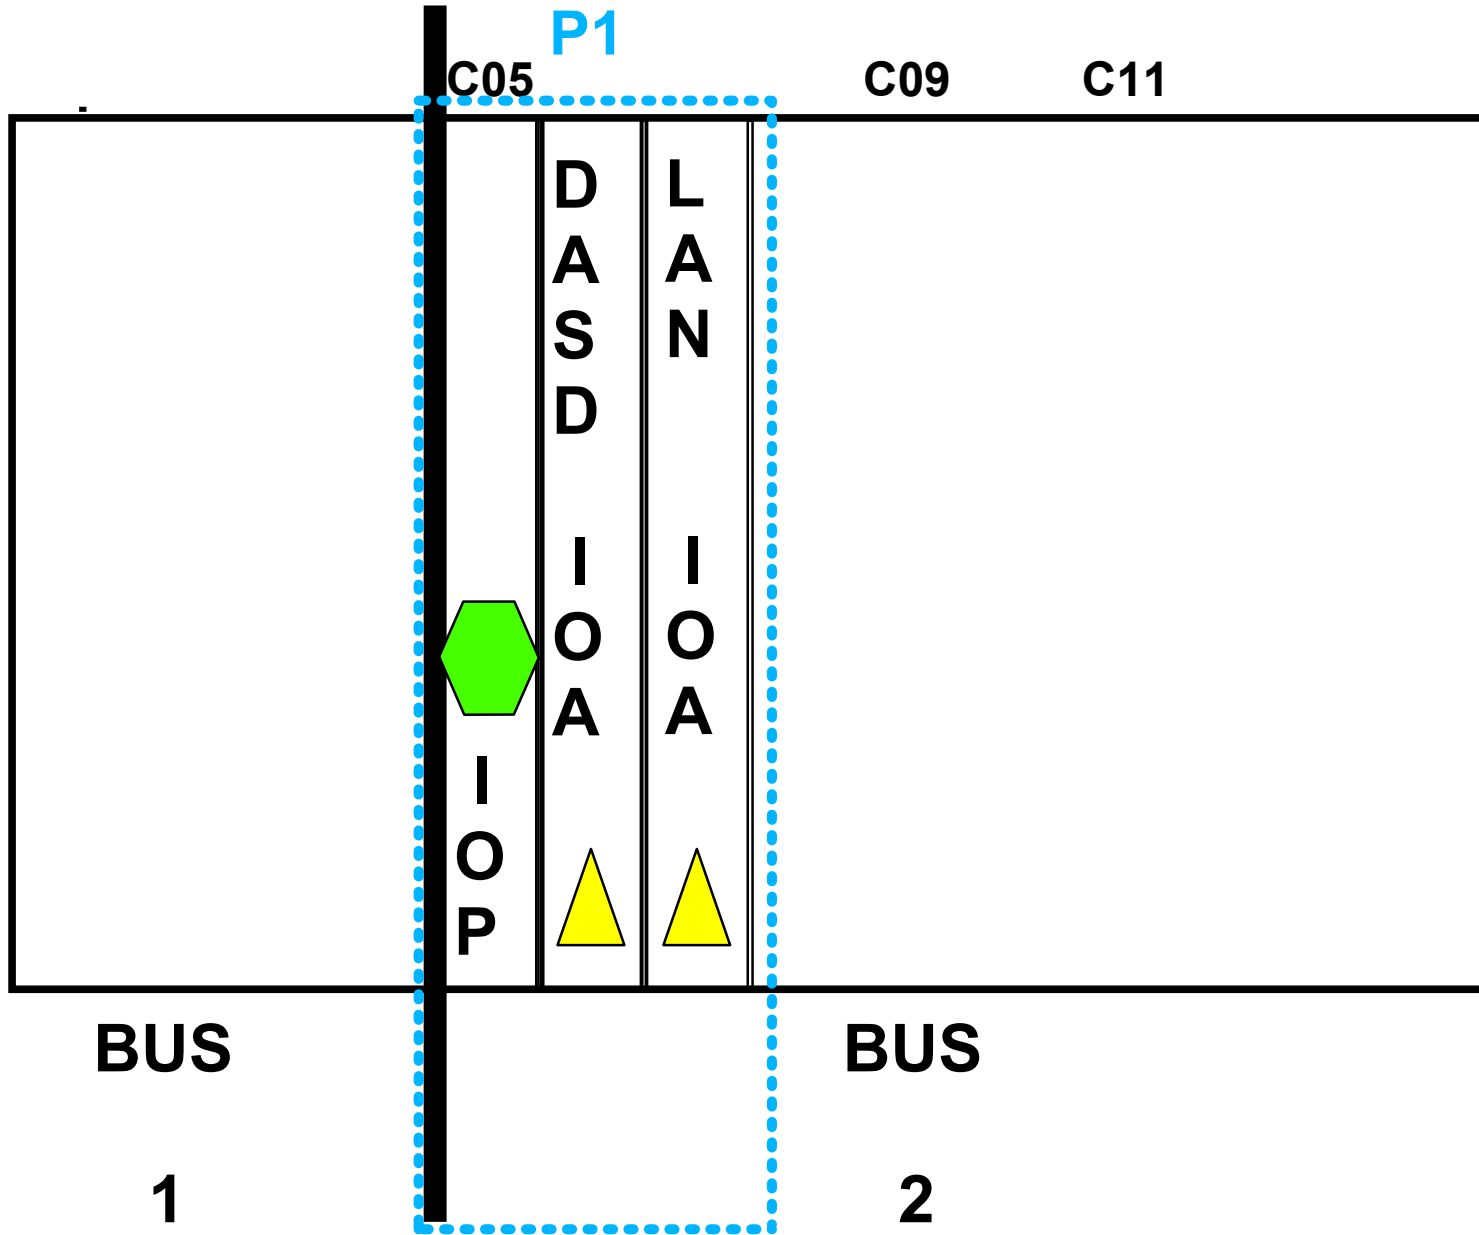


• Here is how I would position my IOPs and IOAs

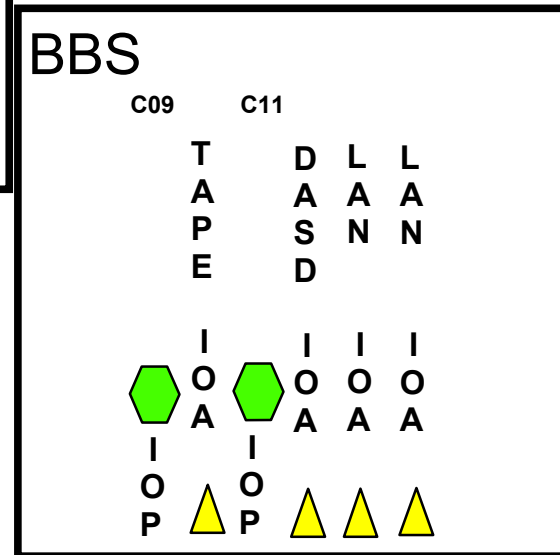


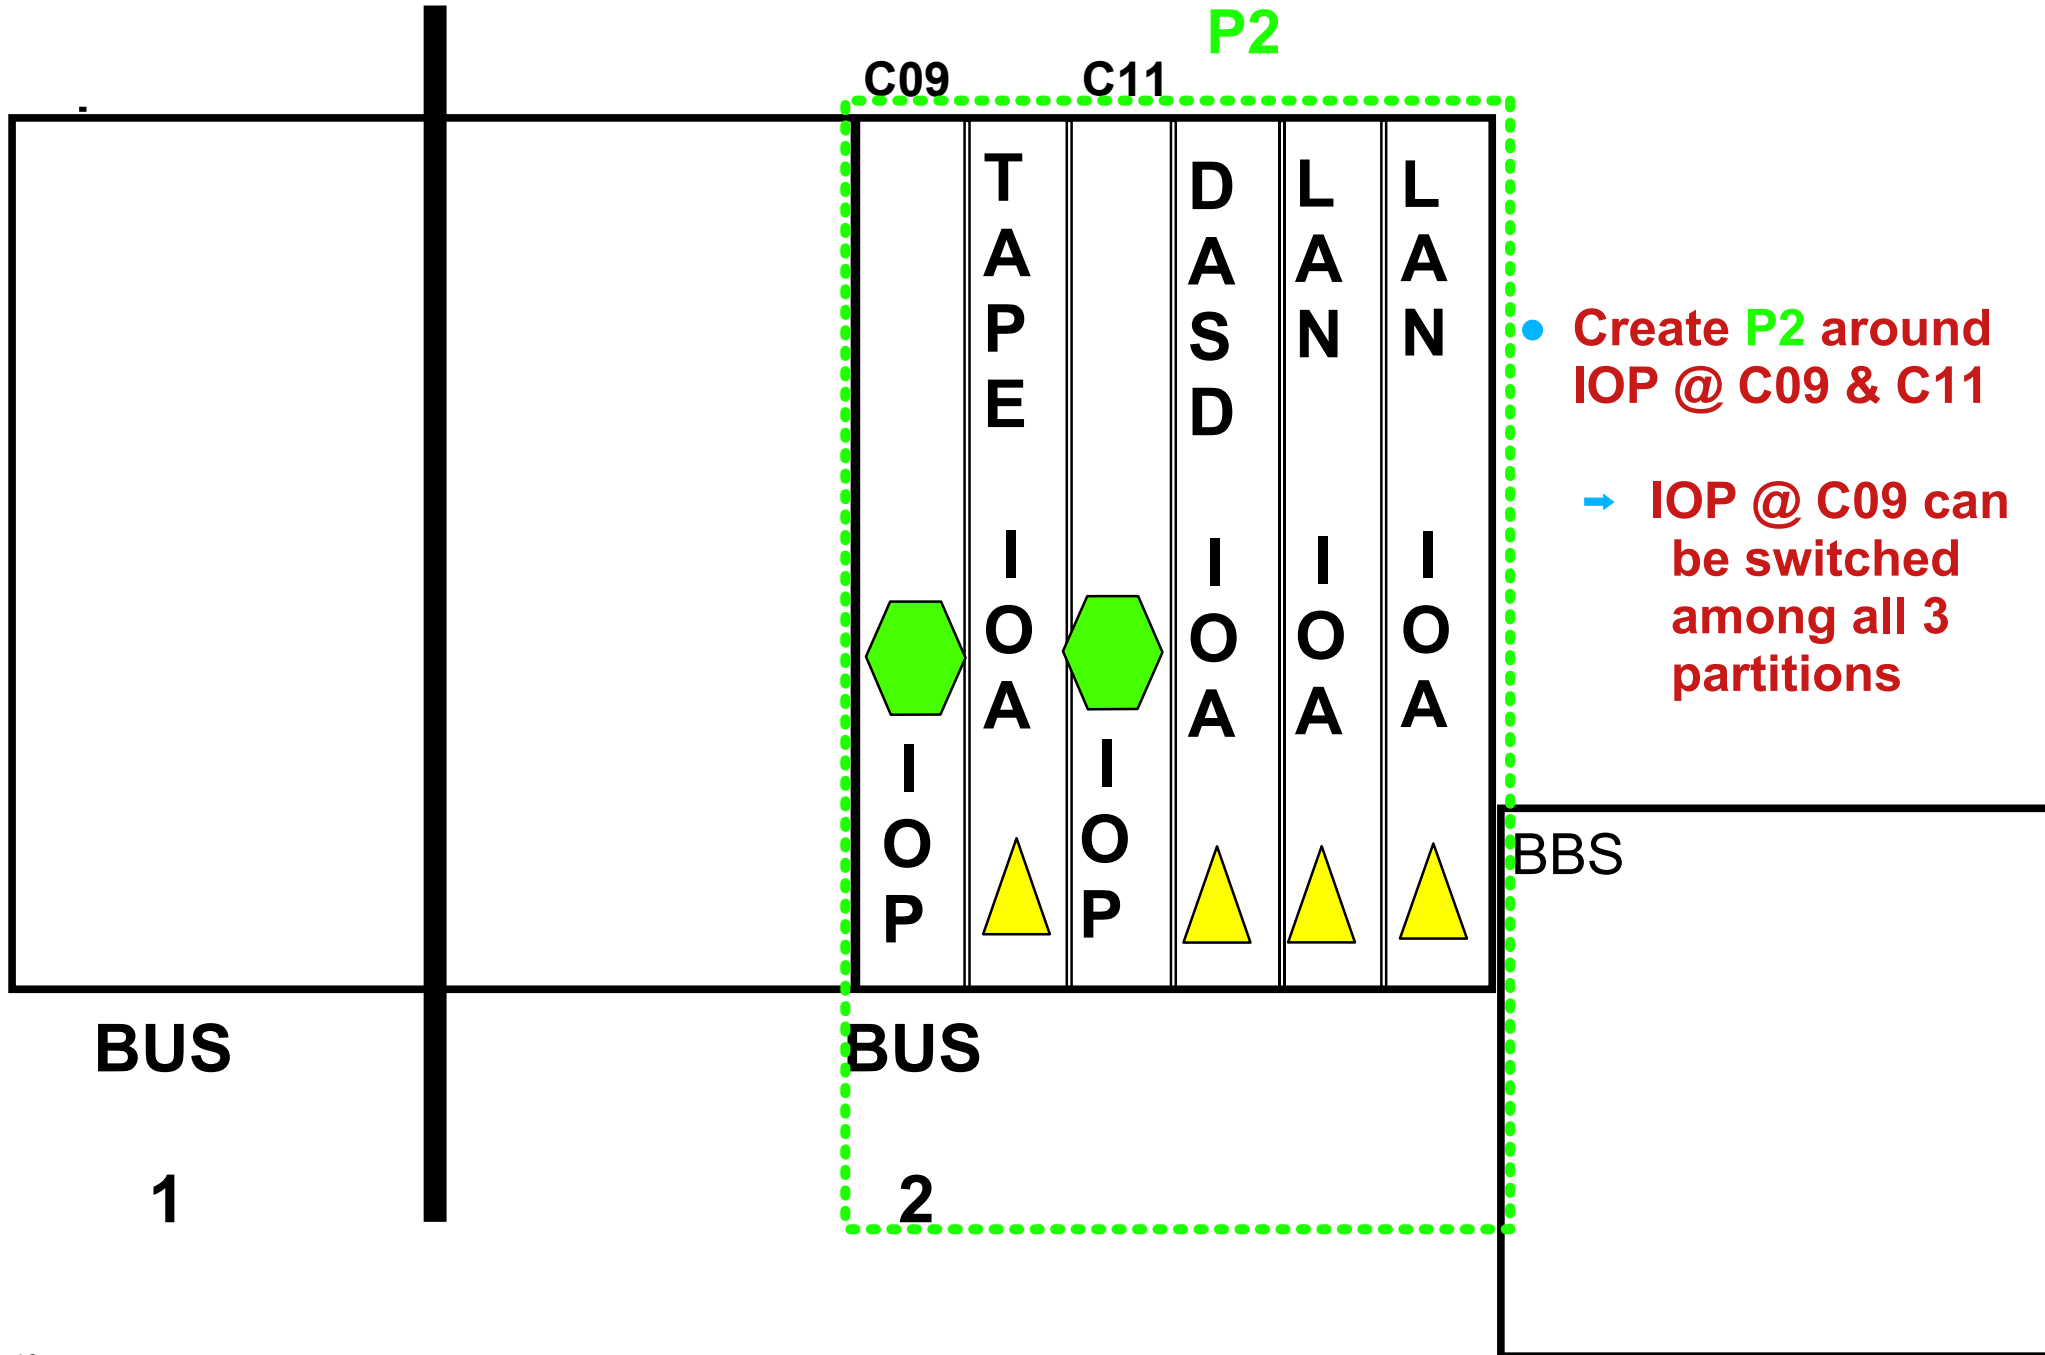
- Share Bus 2
- Deallocate (logically remove) IOPs @ C05, C09 and C11
- A defacto creation of P0 around Bus 1

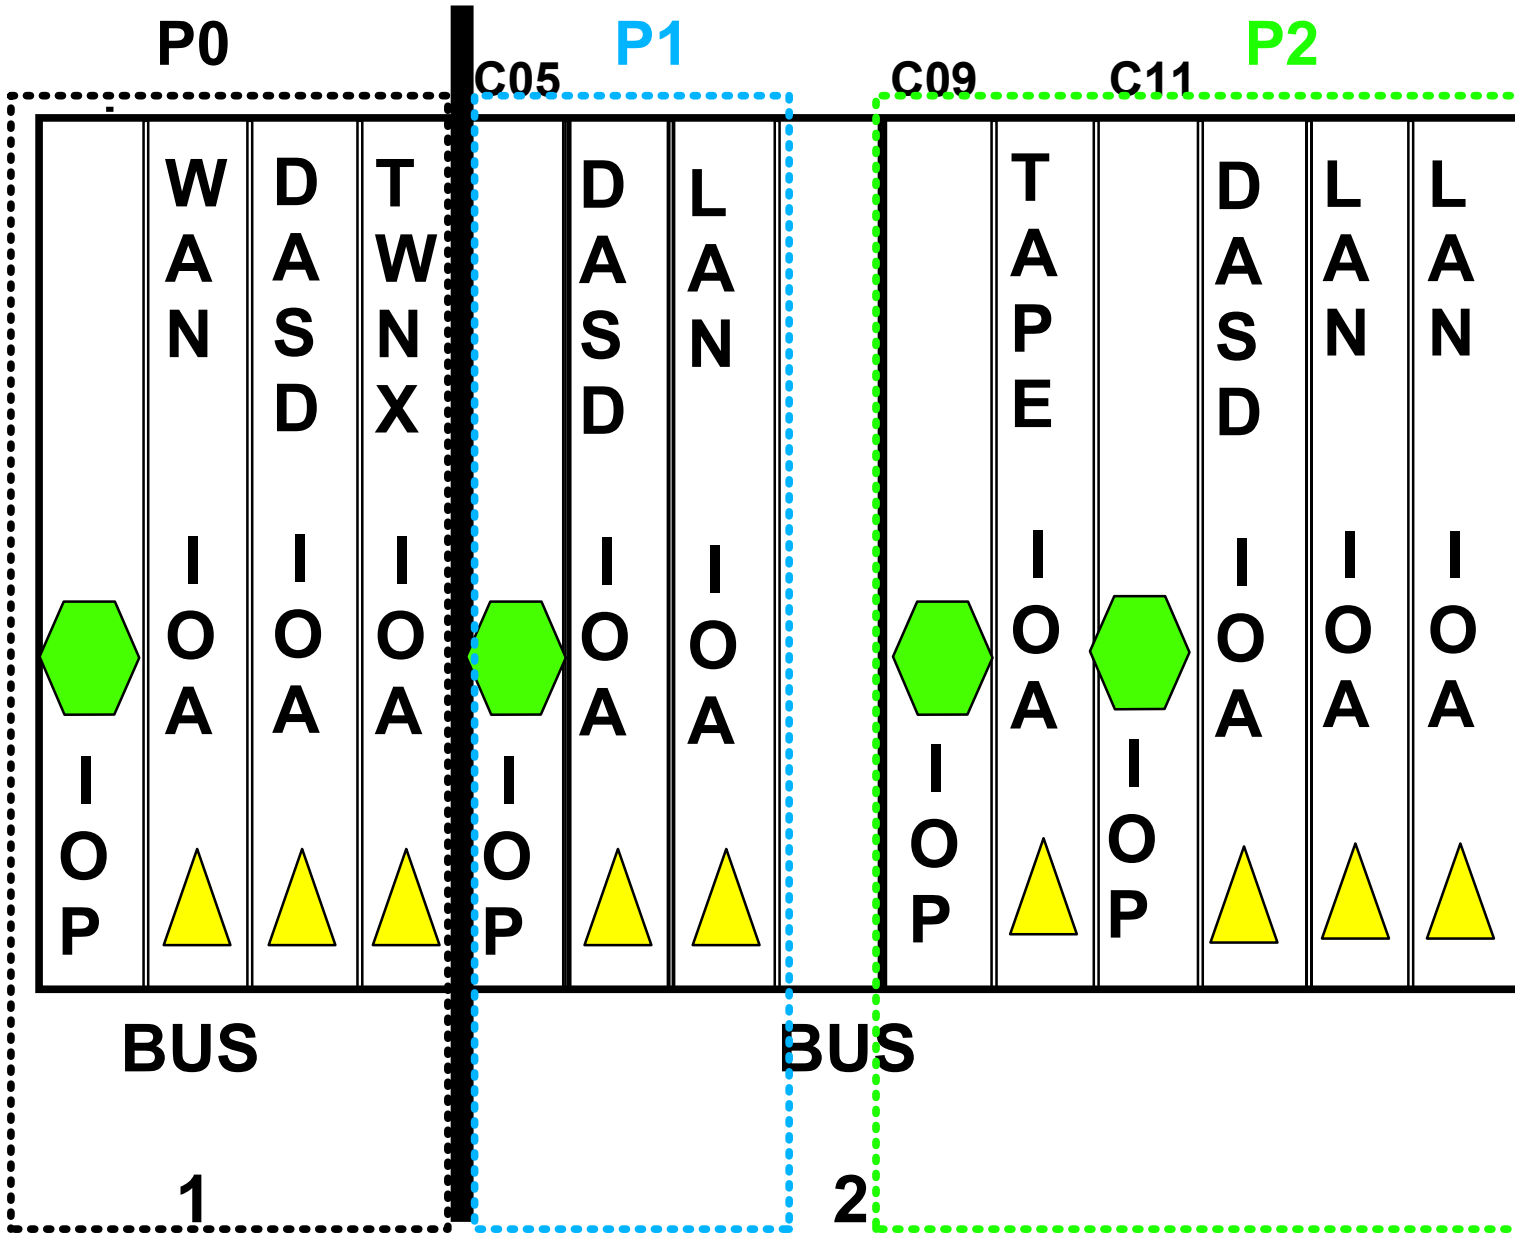




- Create P1 around IOP @ C05



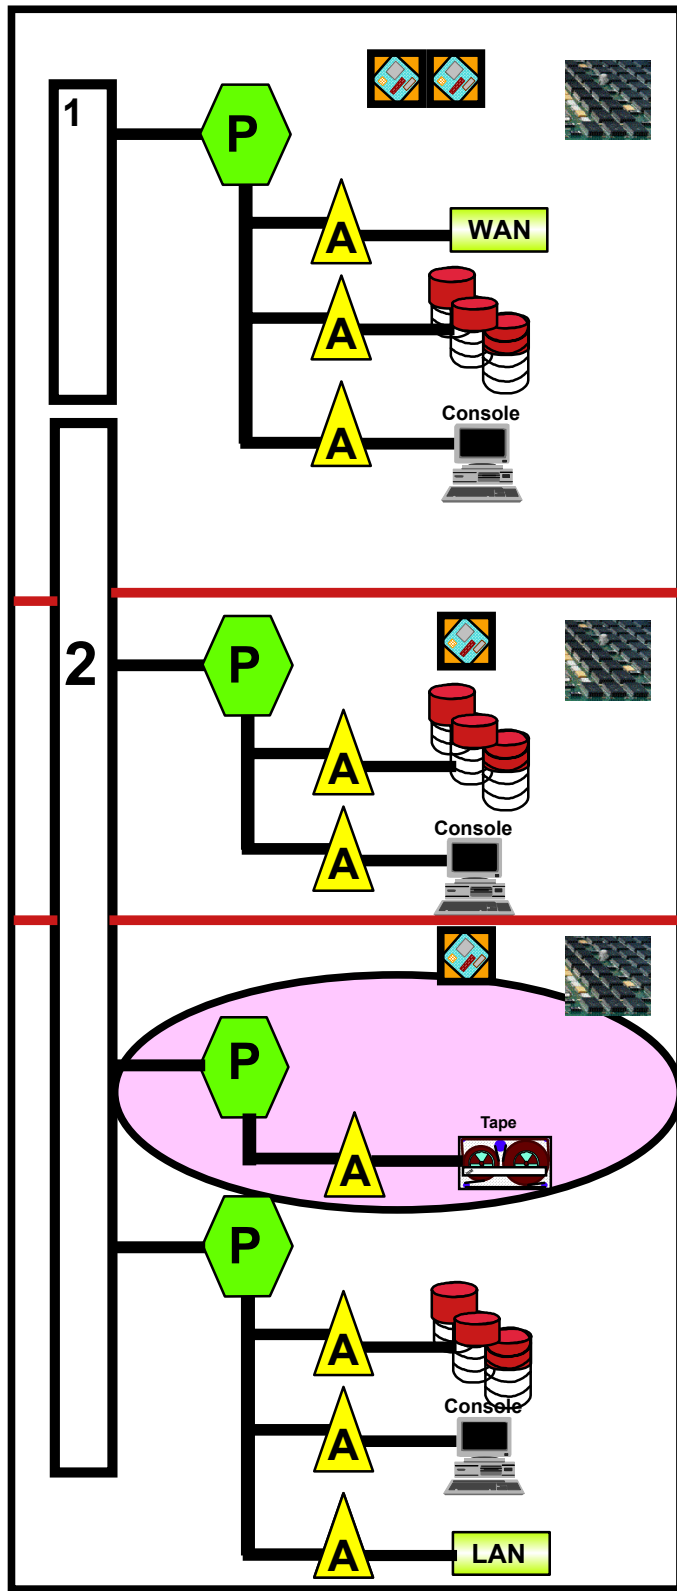




• Physically this now looks like:



• Logically it now now looks like:



P0 Owns Bus 1 Dedicated
Owns Bus 2 Shared

P1 Uses Bus 2 Shared
Owns IOP @ C05

P2 Uses Bus 2 Shared
Owns IOP @ C07 - 4 now
Owns IOP @ C09

Partition Characteristics

Each logical partition operates as a complete system with its own:

- SLIC and OS/400
- Processors, main storage, and interactive performance
- Load source disk and console
- 64 bit address space (single level storage)
- ASPs and memory pools
- Problem logs, LIC logs, PAL entries
- Data (libraries, objects, file systems)
- Performance characteristics
- Language feature code(s)

Logical partitions are supported only on the following iSeries systems with the ability to support the required I/O resources and at least V4R4.

- Models 8xx
- Models 7xx
- Models 6xx
- Models Sxx
- Models 270 - some

Minimum Requirements



- ▶ Version 4.5 - **unsupported** - or V5.1/2 for 8xx/270 *
- ▶ Version 4 Release 5 - **unsupported** - for 7xx/6xx/Sxx
- ▶ Minimum processor requirements
 - 0.10 per partition - if applicable
 - 0.25 Primary - if applicable
 - 1 per partition - if not applicable
- ▶ Minimum memory requirements
 - 64MB (Linux)
 - 128MB secondary (OS)
 - 512MB Better for primary
- ▶ I/O requirements
 - Load source disk IOP and disk
 - Console IOP and console
 - Available Alt IPL device

* Check Website

Software Licensing


- ▶ OS/400 is licensed to the physical machine (regardless of the number of partitions).
- ▶ New i8XX systems = 1 OS/400 License per active processor
- ▶ LPs are licensed by processor group - one license per physical machine (regardless of the number of partitions).
- ▶ User based licensing is checked at the partition level in V4R4. Beginning with V4R5, the check will be performed across all partitions.
- ▶ License keys need to be entered for all partitions. There is no automatic transfer of license information between partitions.

Note: Check with your software vendor for specific licensing questions.

Benefits of LPAR to the customer

- ▶ Reduced Cost
 - Hardware
 - Maintenance
 - Software
 - Linux Enablement
 - Operations Console via LAN
- ▶ High-speed Inter- partition Communications
 - Virtual OptiConnect
 - Virtual LAN
 - HSL
- ▶ Dynamic Reassignment - No IPL - of Resources
 - IOPs
 - Processors
 - Main Storage
 - Interactive Performance
 - Busses



Internal communications



- **Virtual OptiConnect is an any-to-any protocol. Each partition needs LPP**
- **Virtual Ethernet is 1Gb ethernet protocol, no soft/hardware necessary**
- **HSL Opticonnect allows partitions to use HSL for system-to-system or partition-to-system-communication**

What LPAR is NOT:

- ▶ NOT intended to contain hardware failures to a single partition or provide hardware fault recovery
 - Certain hardware failures will cause a failure of the entire system (all partitions).
 - ◆ Processor failure
 - ◆ Non recoverable Memory failure (can not recover the data in the memory which has failed, i.e., whole card)
- ▶ Primary partition failure = probable system failure
 - A failure in the Hypervisor will cause a failure of the entire system (all partitions).
 - If the primary partition fails for any reason, all secondary partitions may also fail.



- Server Consolidation
- Seperate production and test environments
- Departmental Systems
- Release Migration
- Diverse Workloads
- Integrated Cluster
- Granular Fail-over
- 24X7 Partition Availability
- On-site D/R test

Why partition a system?

Simplicity

Money - and the saving thereof - not limited to:

- Floor space
- Software Licenses
- Support Contracts
- In house support/operations
- Hardware

24/7 Availability w/o additional system

Separate workloads i.e. development/production

Partial processors

Communication w/o additional hardware or software

Etc.

Trademarks and Disclaimers

© IBM Corporation 1994-2003. All rights reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

AS/400	IBM
AS/400e	IBM (logo)
eServer	iSeries
	OS/400

Lotus and SmartSuite are trademarks of Lotus Development Corporation and/or IBM Corporation in the United States, other countries, or both.

MMX, Pentium, and ProShare are trademarks or registered trademarks of Intel Corporation in the United States, other countries, or both.

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

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

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

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

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

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

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information in this presentation concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information in this presentation addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.