

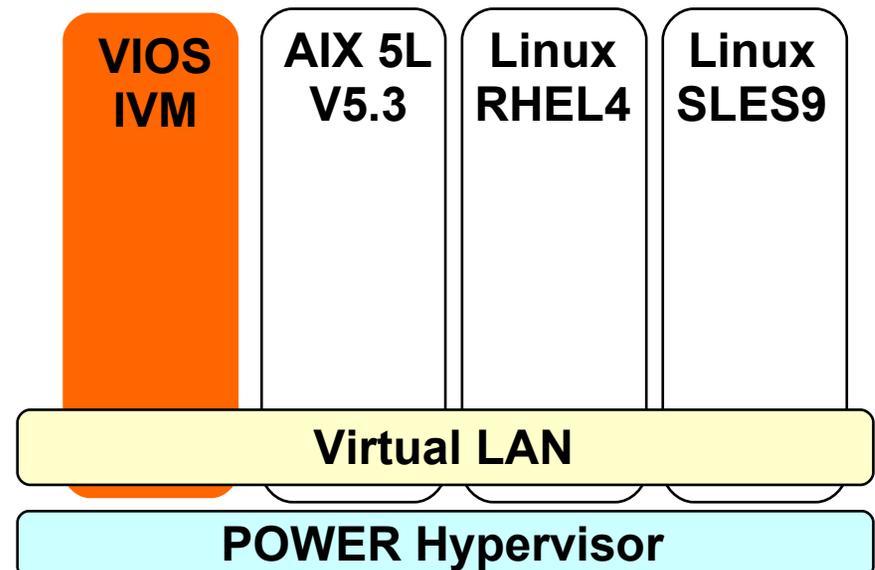
Systemverwaltung 2009

AIX / LPAR



Integrated Virtualization Manager (IVM) (1 of 2)

- Provides support for logical partitions without a physical HMC
- IVM software runs in Virtual I/O Server partition
- Virtual I/O Server partition owns all physical I/O
 - Client partitions can access virtual optical, disk, Ethernet, and console
- Configuration considerations
 - Maximum of two virtual Ethernet adapters per client partition
 - Maximum of four virtual LANs per managed system
 - Maximum of one virtual SCSI adapter per client partition
 - Limited DLPAR functionality
 - Subset of HMC service tools
 - One profile per partition



Integrated Virtualization Manager (IVM) (2 of 2)

Partition Management

- ◆ [View/Modify Partitions](#)
- ◆ [Create Partitions](#)
- ◆ [View/Modify System Properties](#)

Virtual Ethernet Management

- ◆ [View/Modify Virtual Ethernet](#)

Storage Management

- ◆ [View/Modify Devices](#)
- ◆ [Advanced View/Modify Devices](#)
- ◆ [Create Devices](#)

User Management

- ◆ [View/Modify User Accounts](#)
- ◆ [Create User Accounts](#)

Service Management

- ◆ [Service Focal Point](#)
 - ◆ [Manage Serviceable Events](#)
 - ◆ [Service Utilities](#)
 - ◆ [Create Serviceable Event](#)
 - ◆ [Manage Dumps](#)
 - ◆ [Collect VPD Information](#)
- ◆ [Updates](#)
- ◆ [Backup/Restore](#)
- ◆ [Application Logs](#)
- ◆ [Hardware Inventory](#)

View/Modify Partitions ?

To perform an action on a partition, first select the partition or partitions, and then select the task.

System Overview

Total system memory:	8 GB	Total processing units:	2.0
Memory available:	5.41 GB	Processing units available:	0.1
Reserved firmware memory:	416 MB	Processor pool utilization:	0.0%
System attention LED:	Inactive		

Partition Details

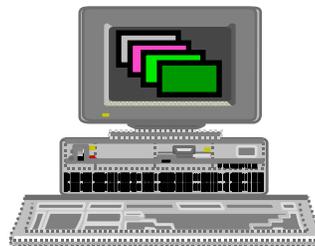
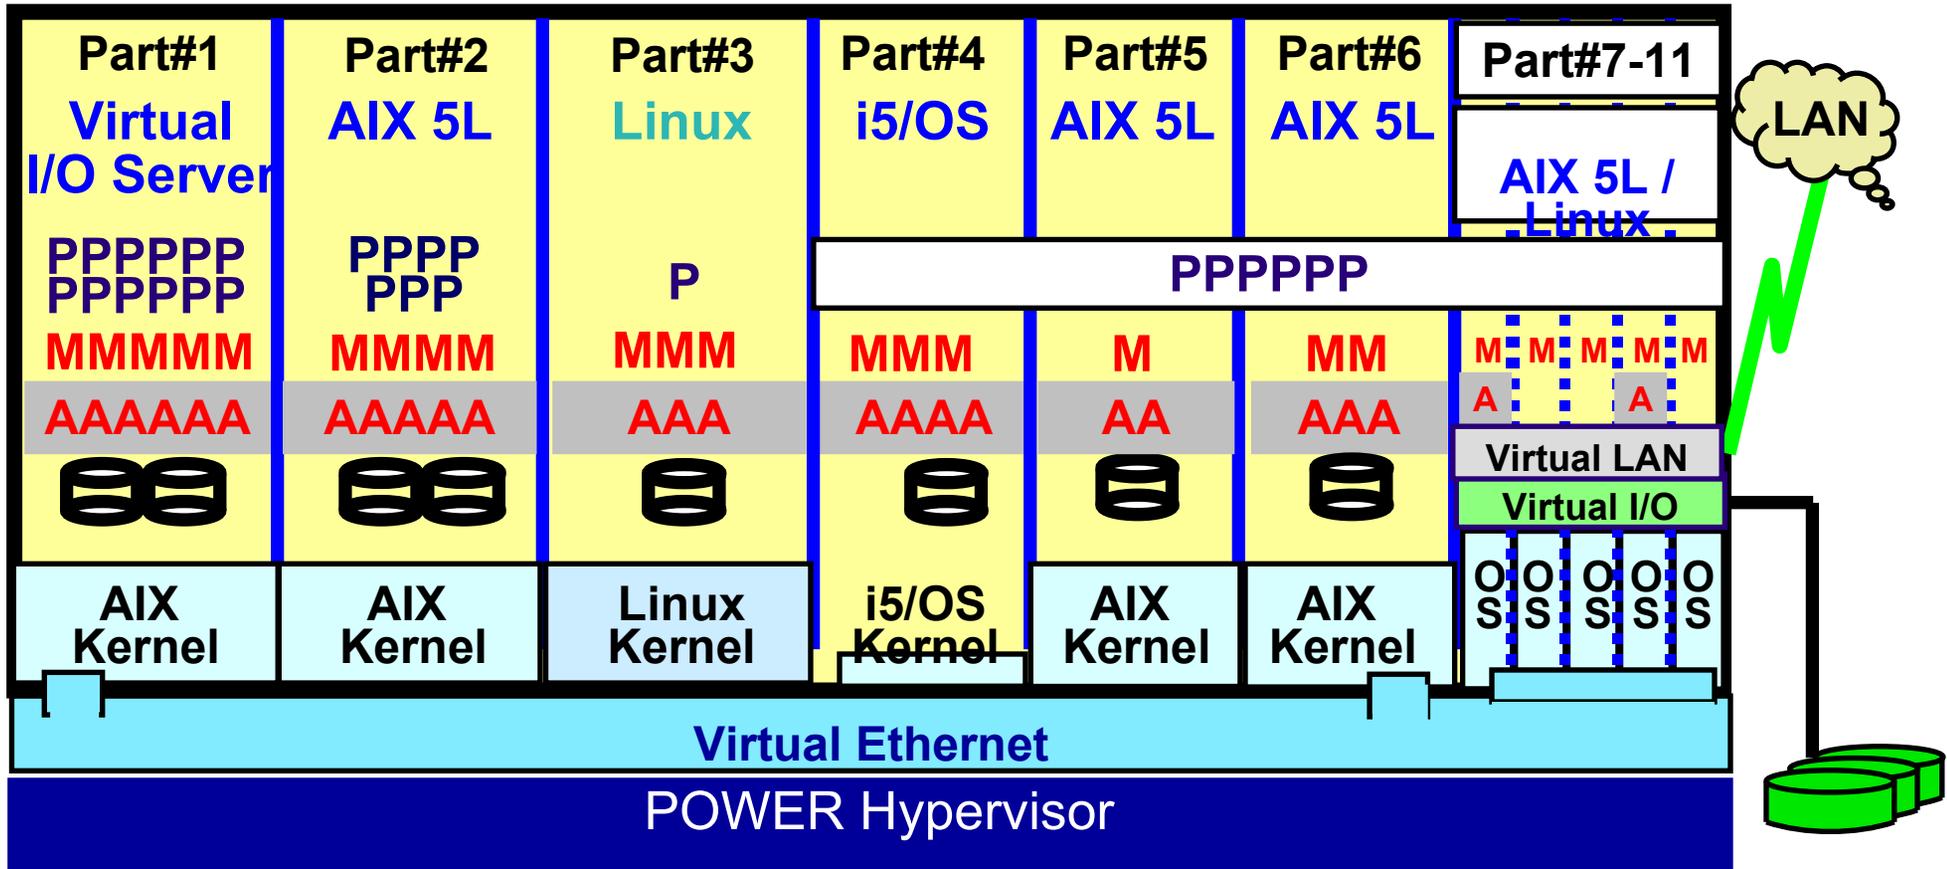
Select	ID ^	Name	State	Uptime	Memory	Processors	Processing Units	Processor Utilization	Reference Code
<input type="checkbox"/>	1	10-0044A	Running	0 Minutes	1 GB	2	1.5	0.0%	
<input type="checkbox"/>	2	zn65p2	Not Activated	0 Minutes	128 MB	1	0.1	0.0%	00000000
<input checked="" type="checkbox"/>	3	zn65p3	Open Firmware	0 Minutes	128 MB	1	0.1	0.0%	AA00E158
<input type="checkbox"/>	4	zn65p4	Not Activated	0 Minutes	128 MB	1	0.1	0.0%	00000000
<input type="checkbox"/>	5	zn65p5	Running	0 Minutes	832 MB	1	0.1	0.0%	CA00E1DC

Tasks

[Properties](#) |
 [Activate](#) |
 [Shutdown](#) |
 [Delete](#) |
 [Operator panel service functions](#) |
 [Reference Codes](#)

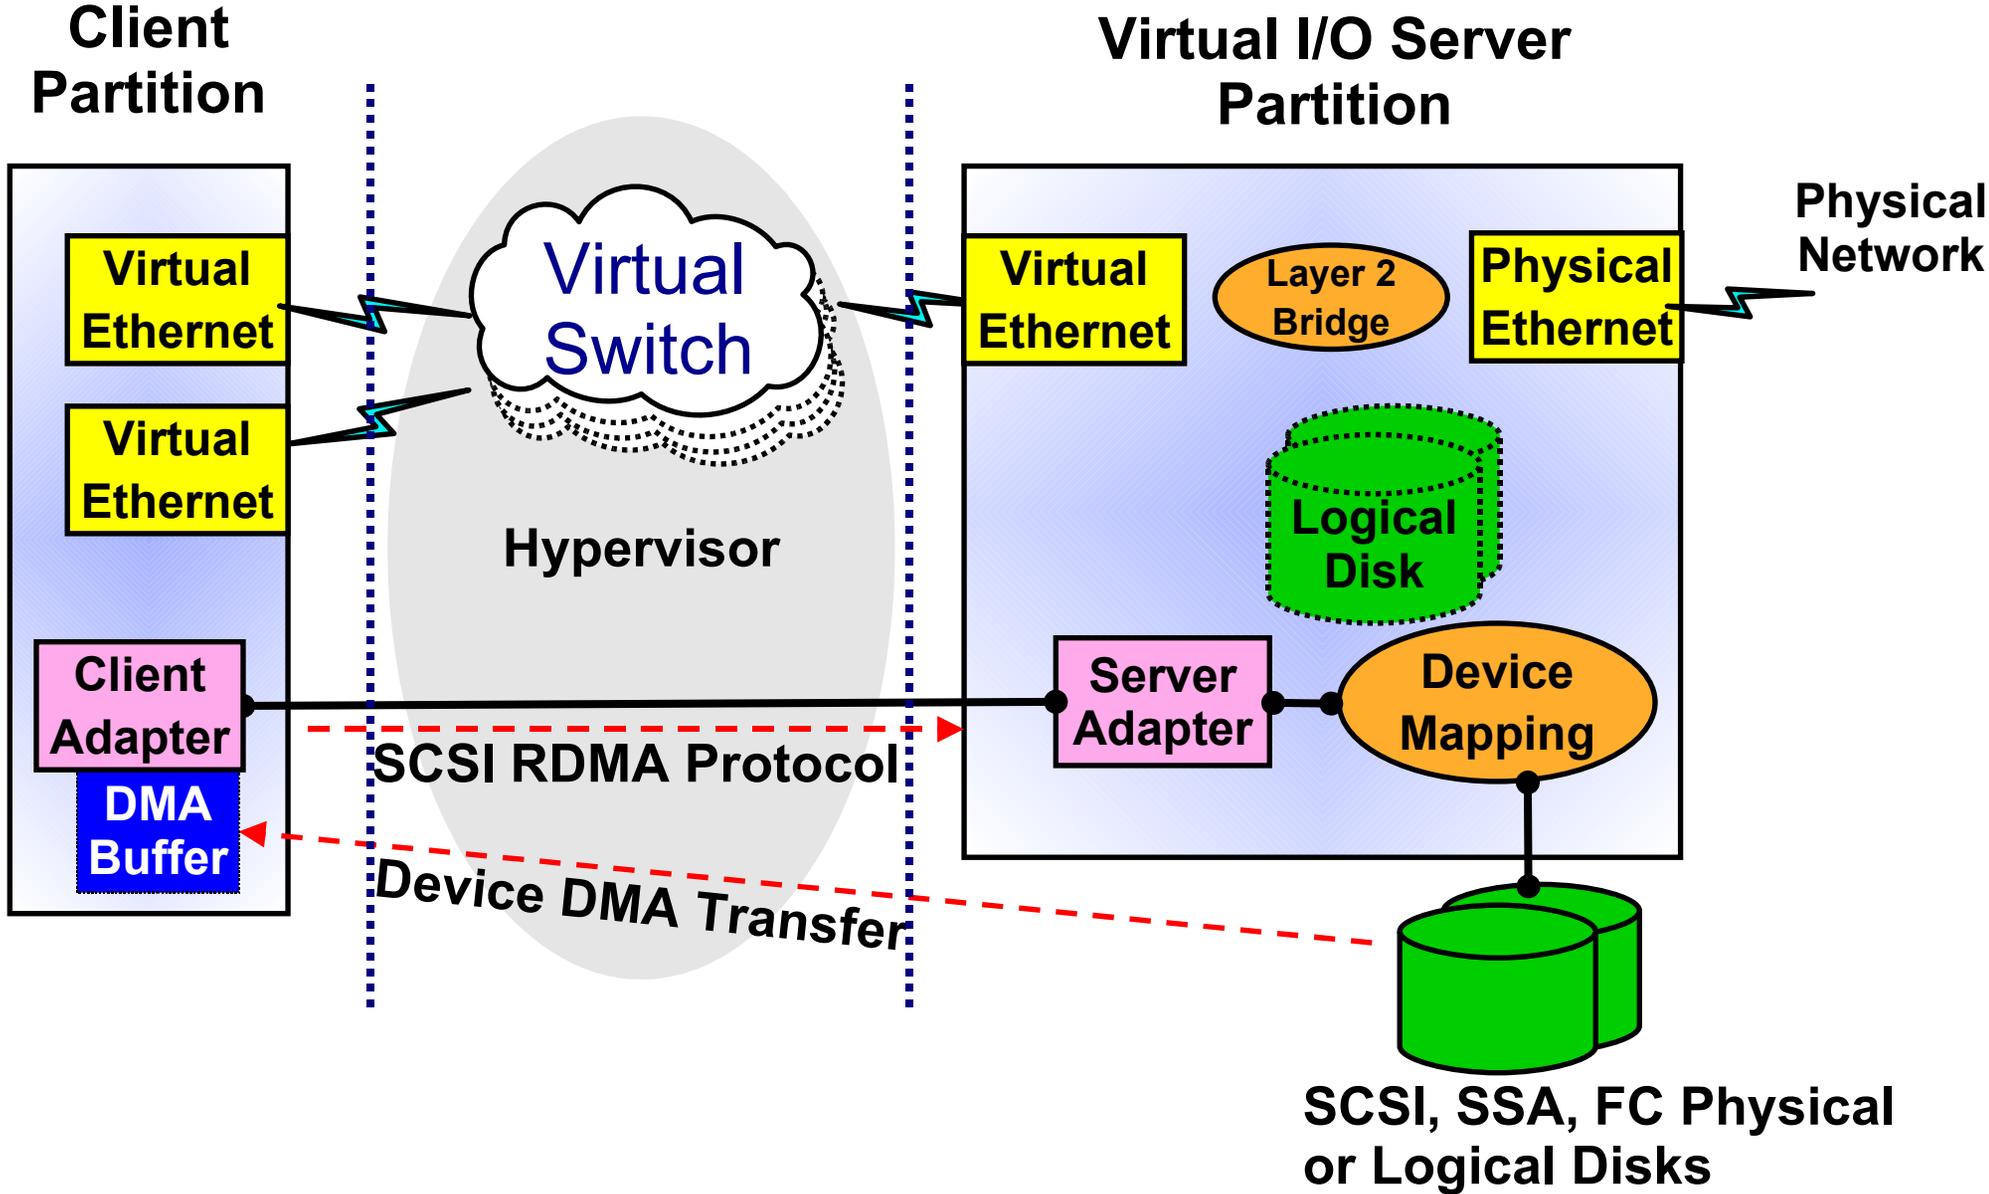
Done

Partitioning Concepts Summary



HMC

Virtual I/O Example



Logical Partition References

- Documentation

- <http://www.ibm.com/support/publications/us/library/>
- Software and hardware Information Centers available

- LPAR and Advanced POWER Virtualization information

- <http://www.ibm.com/systems/p/apv/index.html>

- IBM Redbooks

- <http://www.redbooks.ibm.com/>

- Technical Support Web sites

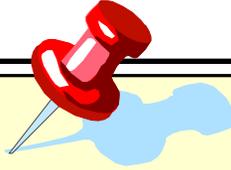
- <http://www.ibm.com/servers/eserver/support/unixservers/>
- <http://www14.software.ibm.com/webapp/set2/sas/f/hmc/>
- <http://techsupport.services.ibm.com/server/vios/>

- Linux Web site

- <http://www.ibm.com/systems/p/linux/>

Web links may change over time

Unit Summary



- **Logical partitions** are independent operating environments
- **Resources** are processors, memory, and I/O slots
 - Dedicated processors and shared processing units
 - Physical and virtual I/O
- The **POWER Hypervisor** is firmware that provides the isolation between partitions, virtual console support, and virtual memory management
- The **HMC** is used to manage partition configurations, consoles, and service tools
- With **dynamic logical partitioning**, changes can be made to a partition's resource allocation without a restart
- **Virtual I/O Server** software provides the ability to share SCSI devices and Ethernet adapters
- The **Advanced POWER Virtualization** feature provides Virtual I/O Server, Partition Load Manager, and Micro-Partitioning

Checkpoint (1 of 2)

- *True or False:* A partition is an independent operating environment.
- A partition is a *logical* partition if which **one** of the following is true?
 - a. Resource assignments are flexible.
 - b. Resources can be moved between partitions without a restart.
 - c. Partitioning is not dependent on physical system building blocks.
- 3. Which of the following are true for dynamic partition operations?
 - a. All AIX 5L partitions on POWER5 processor-based systems are capable of dynamic resource allocations.
 - b. Virtual I/O devices can be dynamically added or removed but not moved between partitions.
 - c. Partitions do not need to be restarted to add, remove, or move resources.
 - d. Applications may or may not be DLPAR-aware.
- 4. List the three types of resources that are configured in partitions.

Checkpoint Solution (1 of 2)

- *True or False:* A partition is an independent operating environment.
True
- A partition is a *logical* partition if which **one** of the following is true?
 - a. Resource assignments are flexible.
 - b. Resources can be moved between partitions without a restart.
 - c. Partitioning is not dependent on physical system building blocks.
- 3. Which of the following are true for dynamic partition operations?
 - a. All AIX 5L partitions on POWER5 processor-based systems are capable of dynamic resource allocations.
 - b. Virtual I/O devices can be dynamically added or removed but not moved between partitions.
 - c. Partitions do not need to be restarted to add, remove, or move resources.
 - d. Applications may or may not be DLPAR-aware.
- List the three types of resources that are configured in partitions.
Processors, Memory, I/O slots

Checkpoint Solution (2 of 2)

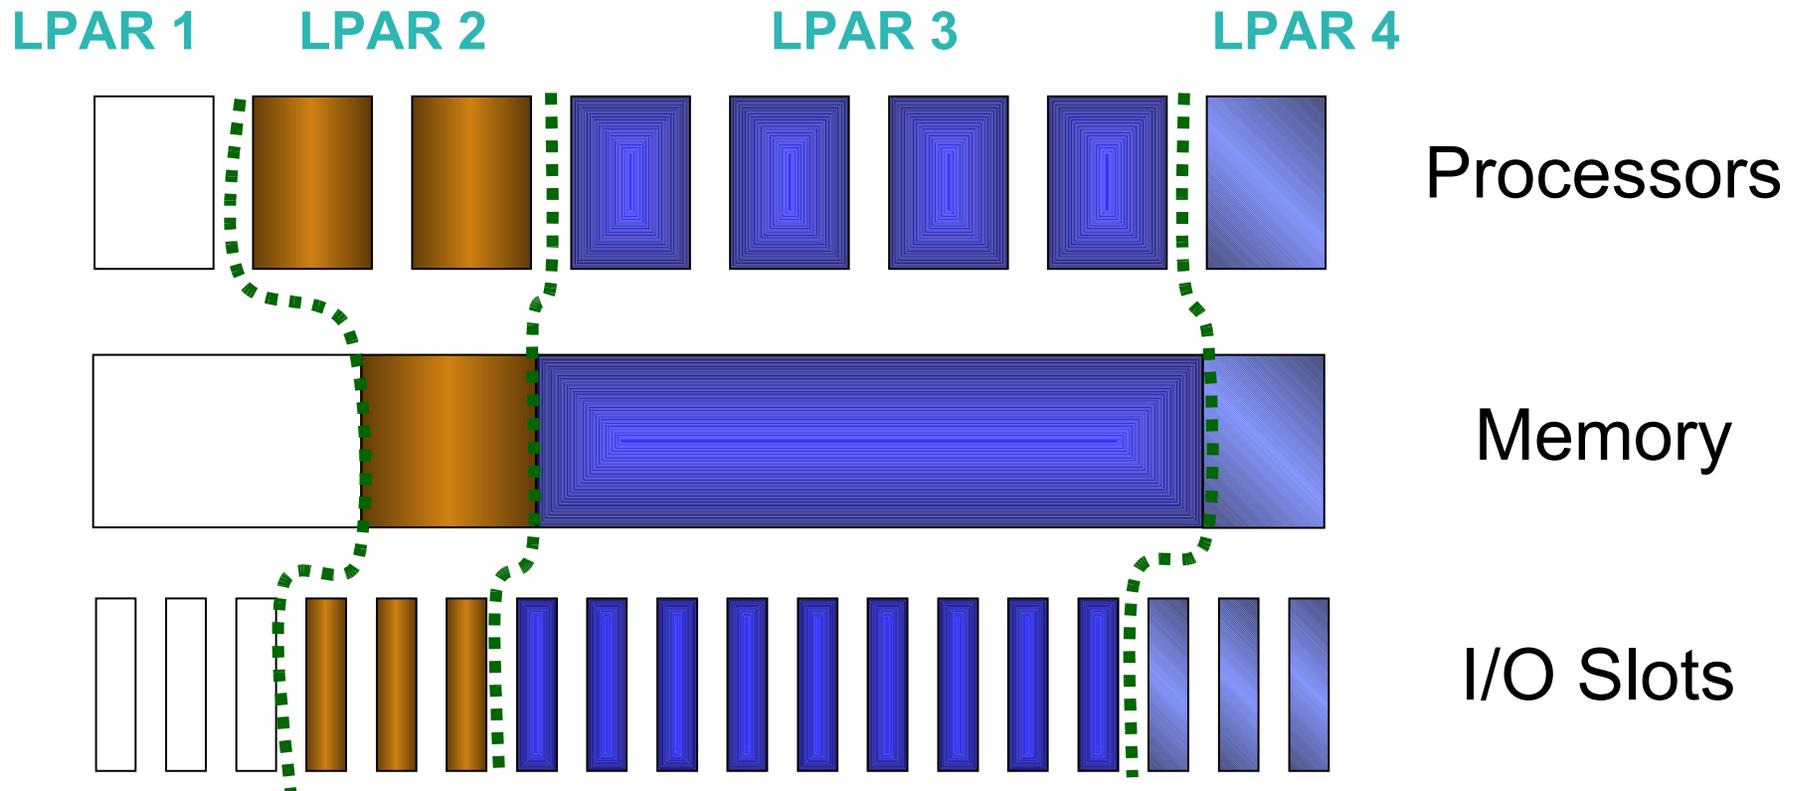
- What system option provides the ability to order additional resources that can be activated when you need them? [Capacity on Demand \(CoD\)](#)
- Which of the following are benefits of using partitions?
 - a. Better management of resources across operating environments
 - b. Isolate applications that were all running on one system
 - c. Better high availability solution
 - d. Manage operational costs by consolidating resources
- What is the system component that performs the resource allocation to partitions? [POWER Hypervisor](#)
- What is the system component needed to configure partitions, configure CoD, and provide access to virtual consoles? [HMC](#)
- What is the term used for the concept of allocating sub-processors to partitions? [Micro-Partitioning](#)

Checkpoint Solution (2 of 2)

- What system option provides the ability to order additional resources that can be activated when you need them? [Capacity on Demand \(CoD\)](#)
- Which of the following are benefits of using partitions?
 - a. Better management of resources across operating environments
 - b. Isolate applications that were all running on one system
 - c. Better high availability solution
 - d. Manage operational costs by consolidating resources
- What is the system component that performs the resource allocation to partitions? [POWER Hypervisor](#)
- What is the system component needed to configure partitions, configure CoD, and provide access to virtual consoles? [HMC](#)
- What is the term used for the concept of allocating sub-processors to partitions? [Micro-Partitioning](#)

Partition Resources

- **Resources** are allocated to partitions
 - Processors, memory, and I/O slots



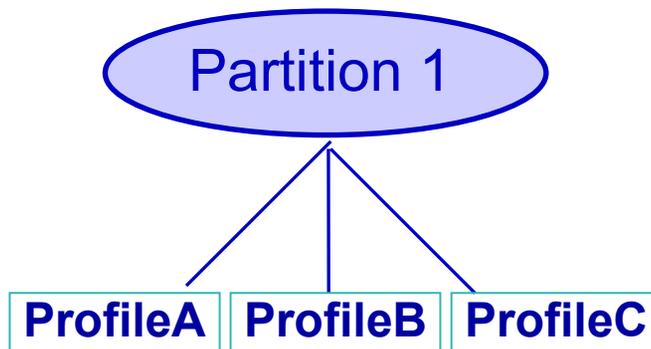
Dividing the System Resources (1 of 2)

- Minimum partition configuration
 - 0.1 processing units if shared; 1 processor if dedicated
 - 128 MB memory
 - Access to necessary I/O devices:
 - Adapter for boot disk
 - Network adapters
- Smallest granularity for allocating additional resources
 - 0.01 processing units if shared; 1 processor if dedicated
 - 1 LMB (16-256 MB) of memory
 - One I/O slot
- Maximum number of partitions depends on system model and available resources
 - Software maximum: 254 partitions

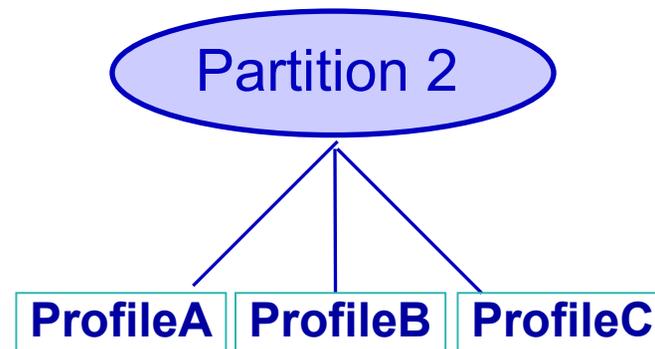
Dividing the System Resources (2 of 2)

- Profiles

- **Partition profiles** describe configurations for partitions
- **System profiles** are collections of partition profiles



Partition Profiles



Partition Profiles

Be sure to document the configuration



Memory Resources (1 of 2)

- Memory allocation
 - One LMB (16 MB to 256 MB sized logical memory blocks)
- For each partition, configure:
 - **Minimum**
 - Partition will not start if this amount is not available
 - Partition can be decreased to this amount if using dynamic LPAR
 - **Desired**
 - Partition will use up to this amount upon activation if available
 - **Maximum**
 - Partition can be increased to this amount if using dynamic LPAR
 - Used for sizing the page table for the partition

Memory Resources (2 of 2)

- Valid memory ranges are based on a ratio for minimum and maximum settings
 - Partitions configured with < 256 MB of memory
 - Maximum cannot be > 16 times actual memory size
 - Partitions configured with >= 256 MB of memory
 - Maximum cannot be > 64 times actual memory size

Minimum

128 MB

256 MB

512 MB

768 MB

1 GB

Maximum

2 GB

16 GB

32 GB

48 GB

64 GB

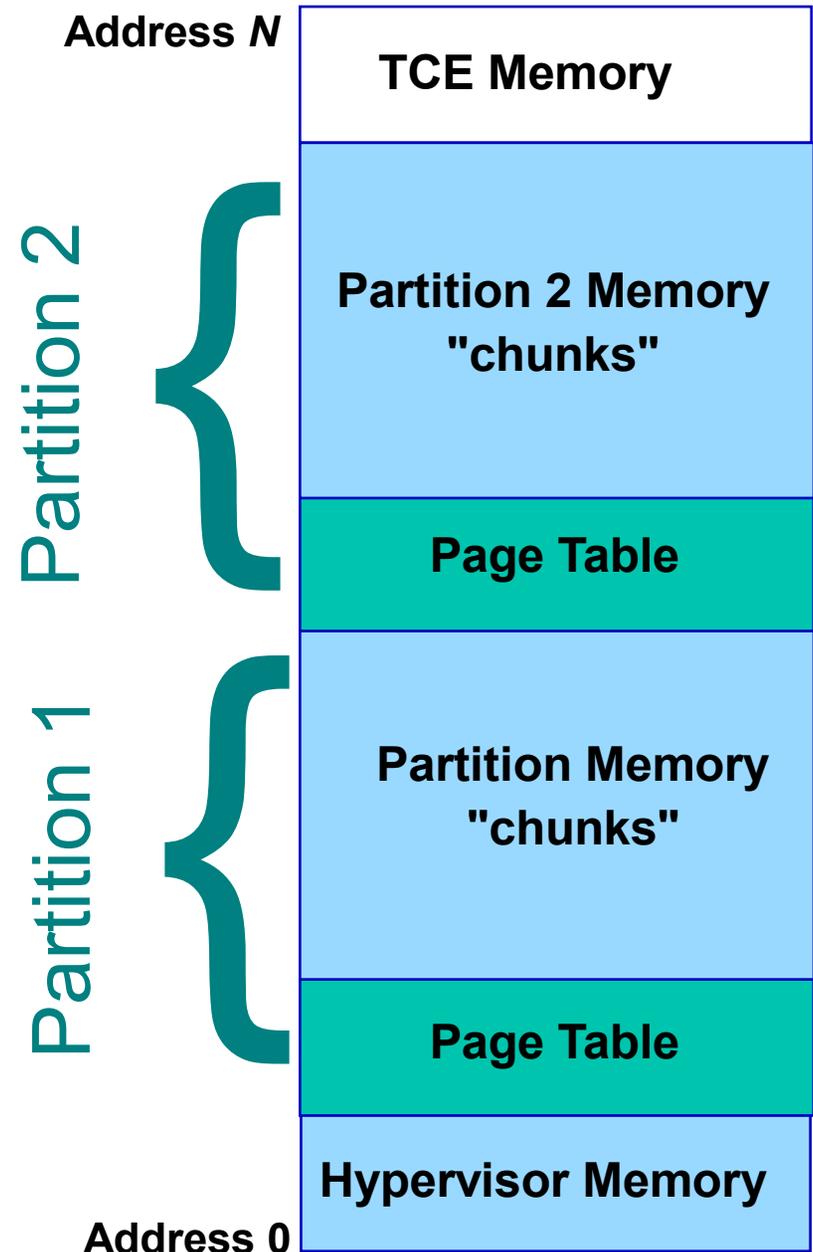
← **1:16 ratio**

← **1:64 ratio**

and so forth

Memory Usage (1 of 2)

- Power Hypervisor memory
 - Variable in size
- TCE memory
 - Used for DMA to I/O devices
- Page tables for each partition
 - Used for mapping partition pages
 - Size is 1/64th of partition maximum size
- Partition memory



Unit Summary

- **Partition concepts**
 - Minimum partition configuration
 - Processors, memory, I/O slots
 - Maximum number of partitions depends on model and amount of resources
 - Partition profiles contain resource configuration information
 - System profiles are lists of partitions and profiles to start in a particular order
- **Minimum, Maximum, and Desired** settings for memory and processors
- **Required** and **Desired** settings for I/O slots
- **Create Partition wizard** leads you through creating partitions and profiles
- **Activate** a partition to start it

I/O Resources

- A partition requires a minimum of:
 - One boot device in allocated I/O slot
 - One network adapter in allocated I/O slot
 - Other I/O adapters for storage and network access
- I/O slots can be *required* or *desired*
- Empty slots can be allocated to a partition
 - Useful if not able to allocate dynamically
- Allocation to partitions is by individual slots
 - Multiple devices connected to a single SCSI adapter are allocated together
- *Advanced POWER Virtualization* feature allows for virtual SCSI and Shared Ethernet configurations
 - Must install and configure Virtual I/O Server software

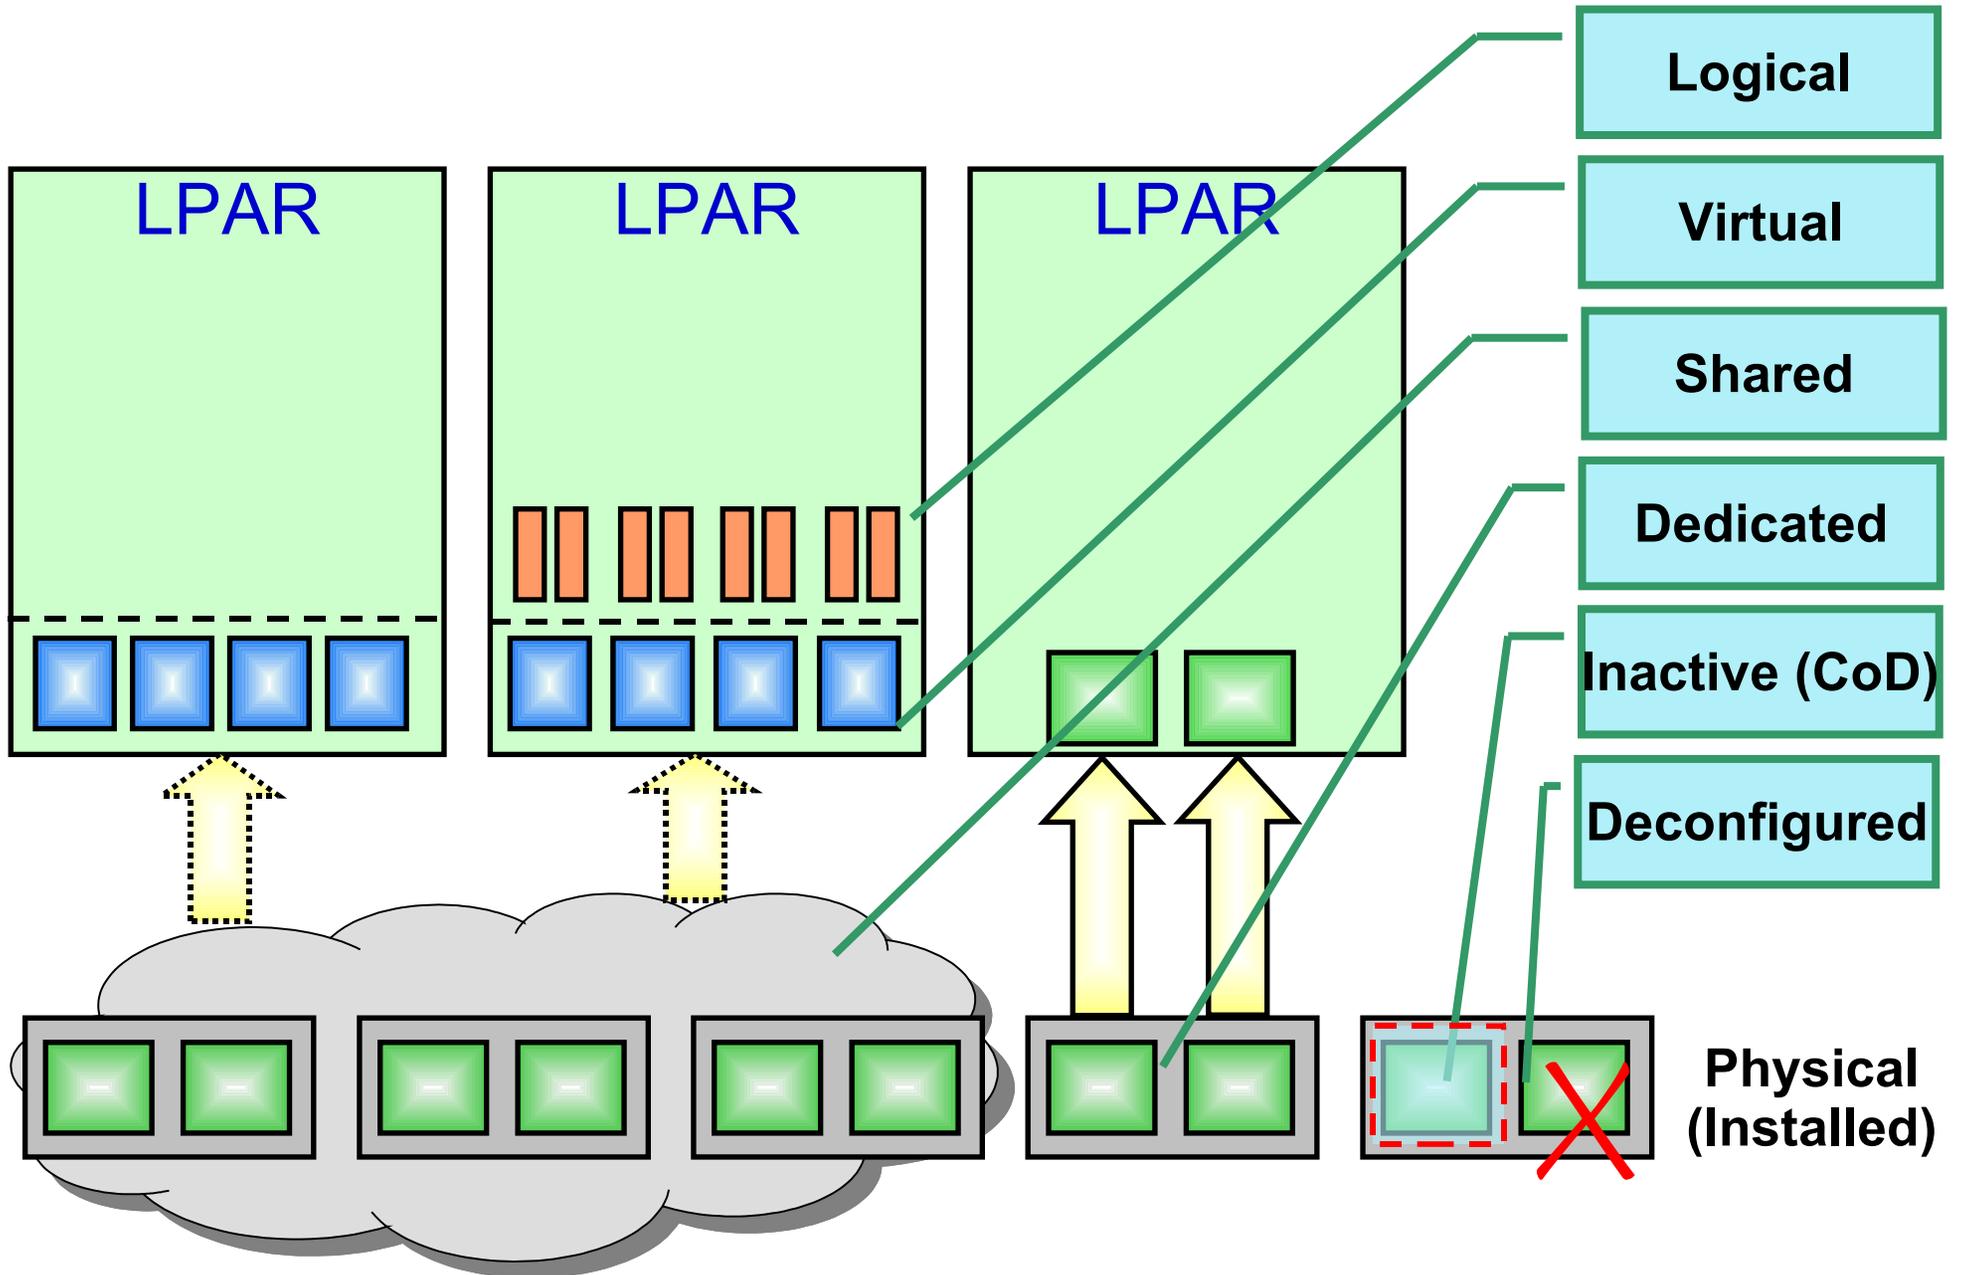
Checkpoint Solution (1 of 2)

1. Match the terms Minimum, Desired, and Maximum to the proper description:
 - This is the upper limit of processors or memory that cannot be exceeded when using dynamic operations. Maximum
 - This is the lower limit of processors or memory when using dynamic operations. Minimum
 - This is the amount of processors or memory that a partition receives if there are more than enough resources on the system when the partition is activated (starts). Desired
- *True or False:* The amount of *desired* processors must always be greater than or equal to the amount of *minimum* processors. True
- What is the minimum amount of memory for an AIX 5L partition? 128 MB or, if the LMB size of the partition is 256 MB, then the minimum would be 256 MB.
- *True or False:* Empty I/O slots can be allocated to partitions. True

Checkpoint (2 of 2)

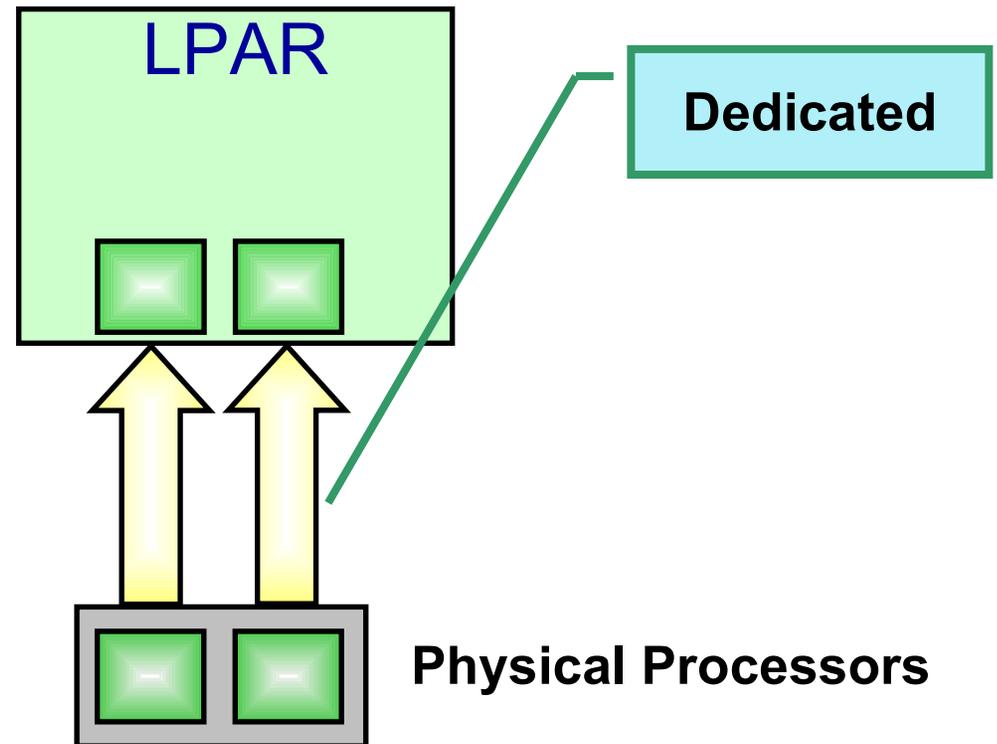
- What happens if the minimum (or required) amount of a resource is not available when a partition is activated?
- How many processors is a partition allocated if, when it starts, it has a *minimum* of 3 processors, *desires* 5 processors, and 4 processors are currently available on the system?
- 3. What happens if a partition is currently running with its 3 minimum processors, but it desires 5 processors and another processor becomes available after the partition is started?
- 4. What is the granularity of allocating processor resources when configuring partitions?
- 5. When incrementing the amount of memory for a partition profile, the minimum unit is how much?

Processor Concepts - Overview



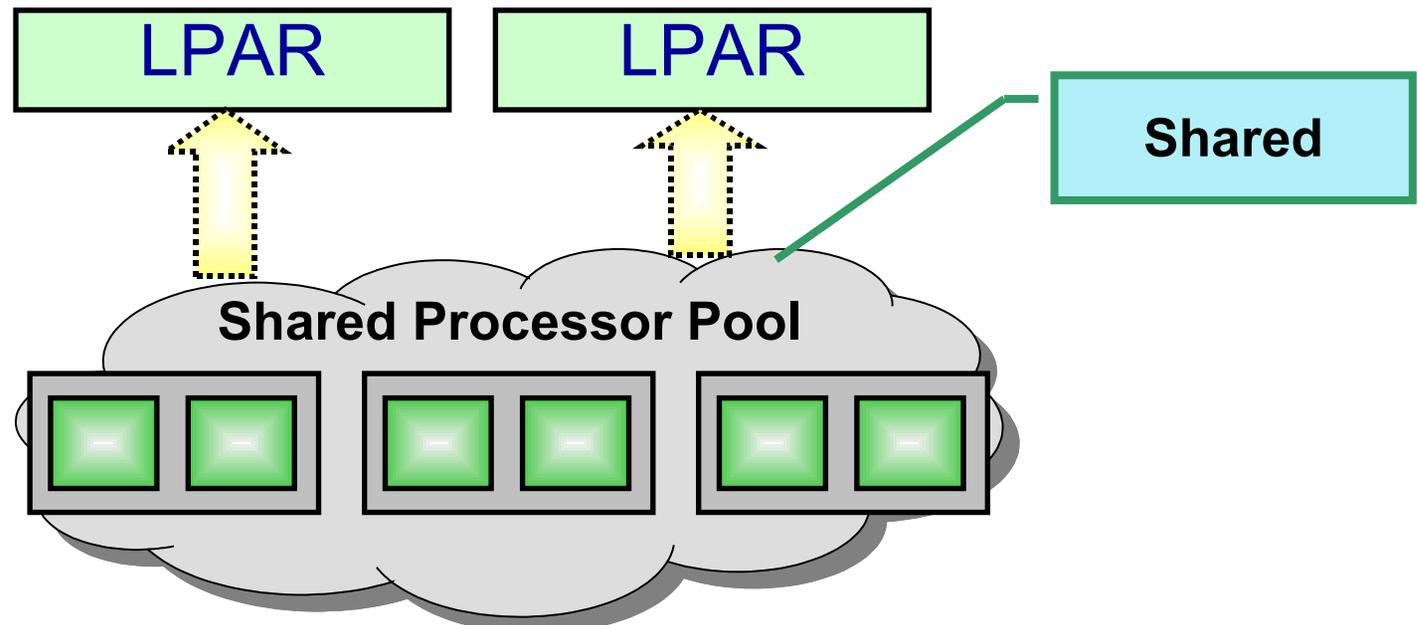
Dedicated Processors

- Allocated as whole processors to a specific partition
- Same physical processors are used for that partition while it is running
 - When partition is stopped, dedicated processors may or may not go to shared pool
- Processor affinity utilized for best performance



Shared Processors

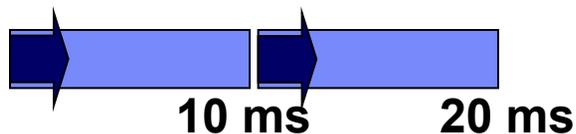
- Processor capacity assigned in *processing units* from the *shared processing pool*
 - Minimum per partition is 0.1 processing units
 - Additional capacity allocated in 0.01 processing units
- Partition's guaranteed amount is its *entitled capacity*
- Advantages:
 - Configuration flexibility
 - Excess capacity can be used by other partitions
- *Micro-Partition*: A partition utilizing shared processors



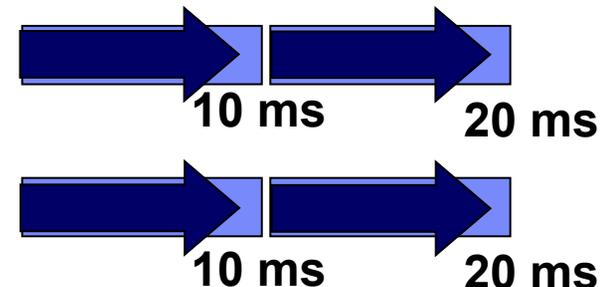
Shared Processor Pool (1 of 2)

- Each partition is configured with a percentage of execution dispatch time for each 10 ms time slice
- Examples:
 - A partition with 0.2 processing units is entitled to 20% capacity during each time slice
 - A partition with 1.8 processing units is entitled to 18 ms of processing time for each 10 ms time slice (using multiple processors)
- The Hypervisor dispatches excess idle time back to pool
- Processor affinity algorithm takes into account *hot* cache

0.2 = 20% of timeslice =
2 ms per timeslice

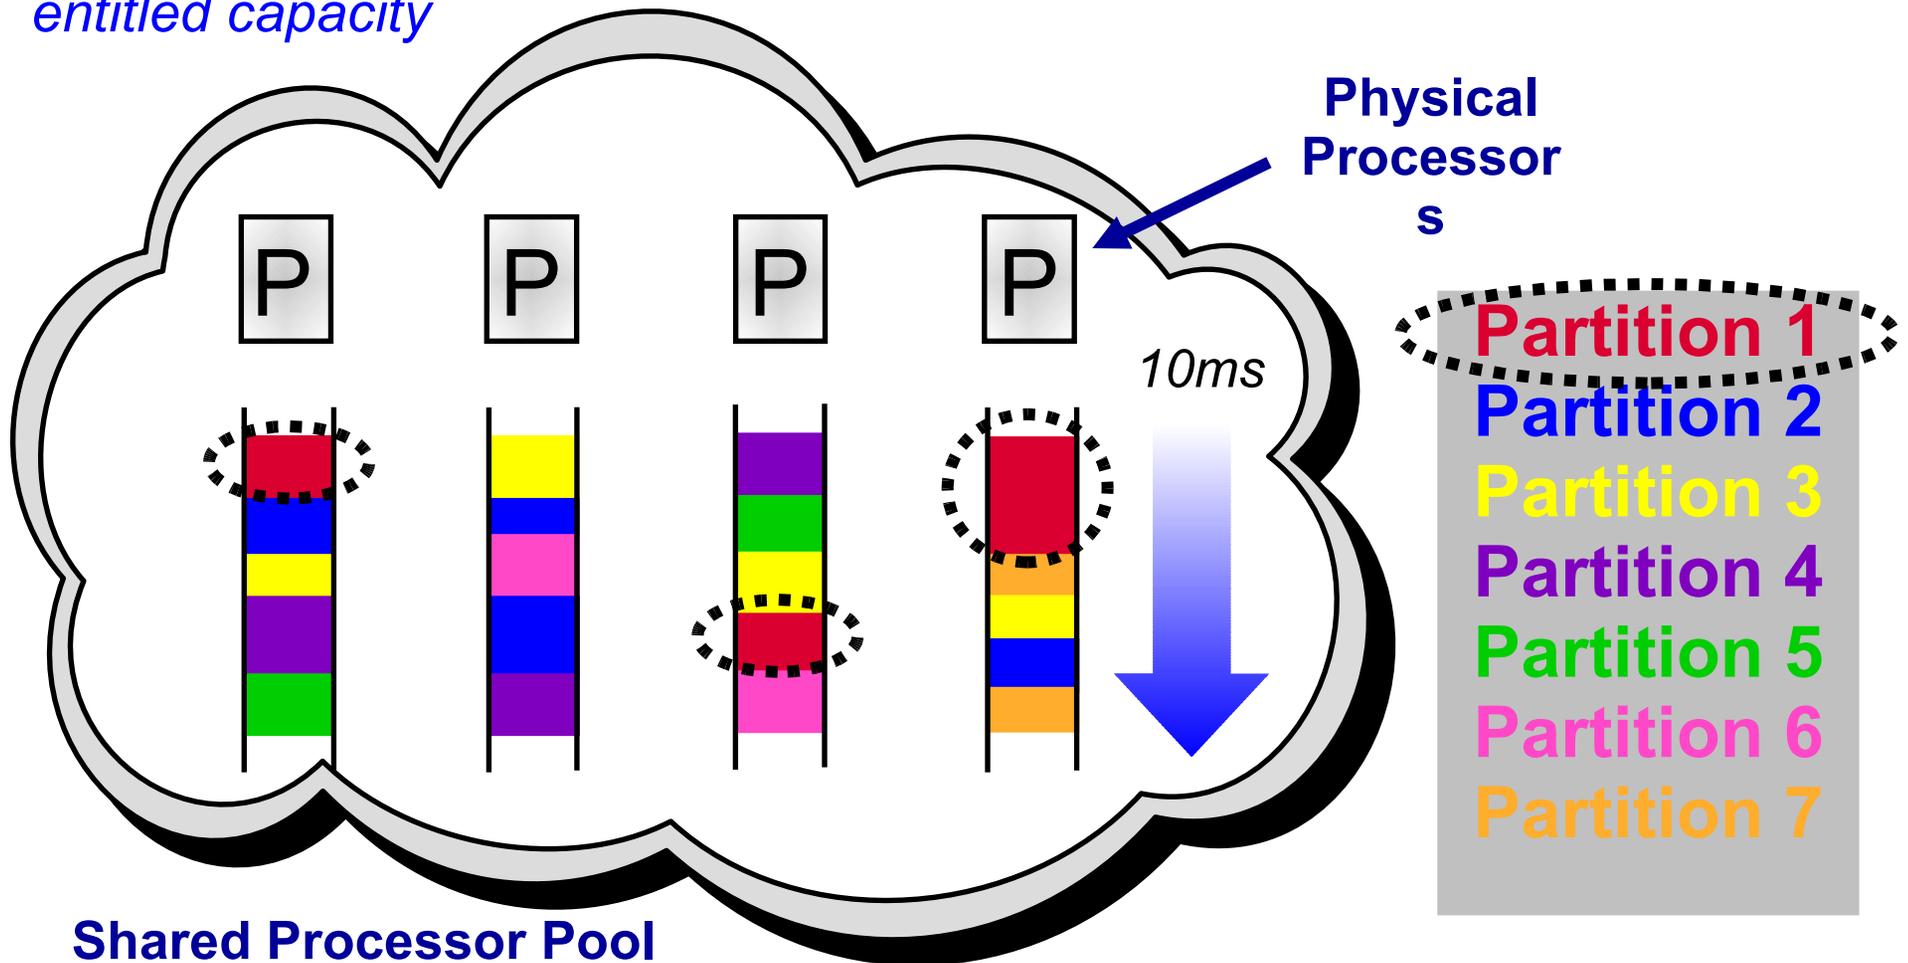


1.8 = 180% of timeslice =
18 ms per 10 ms timeslice



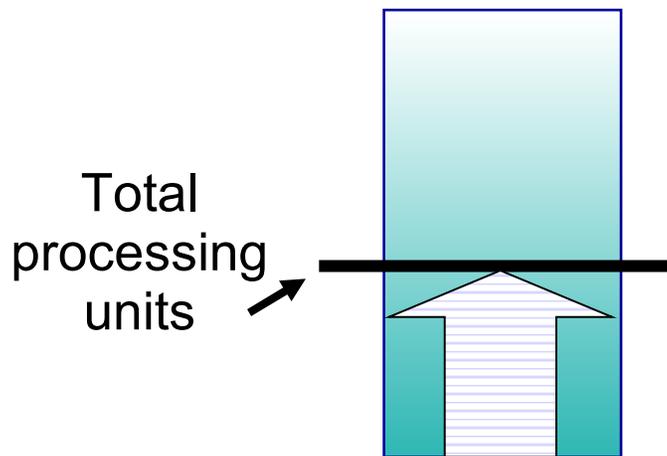
Shared Processor Pool (2 of 2)

- This example shows one 10 ms time slice, 7 running partitions, and 4 processors
 - A partition may run on multiple processors depending on interrupts and its entitled capacity



Capped Partitions

- Partitions with shared processors are either *capped* or *uncapped*
- *Capped*: Limited to the entitled capacity
- Example: 1.5 capped processing units means a partition can use *up to* 15 ms of execution time during each time slice, but no more than that

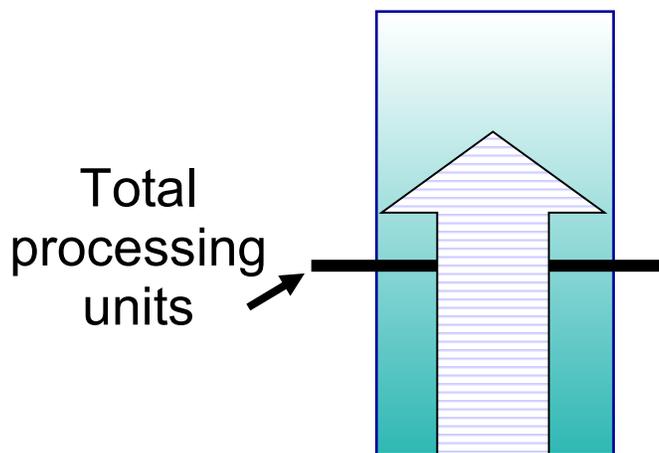


The partition *may* utilize CPU cycles up to its capacity – guaranteed

Excess cycles are ceded back to shared pool

Uncapped Partitions

- *Uncapped*: If a partition needs extra CPU cycles (more than its entitled capacity), it can utilize unused capacity in the shared pool
- Example: An uncapped partition with 1.5 processing units is guaranteed to be able to use 1.5 units, but may use more if necessary (and more is available)



All uncapped partitions share the unused capacity based on user-defined weighting

Weight scale is 0 to 255

Virtual Processors (2 of 3)

- By default, for each 1.00 of a processor, or part thereof, a virtual processor will be allocated
 - Example: 3.6 processing units would have 4 virtual processors
- Up to 10 virtual processors *can* be assigned per processing unit
 - Example: 3.6 processing units can have up to 36 virtual processors
- Both entitled capacity and number of virtual processors can be changed dynamically for tuning
 - Maximum virtual processors per partition is 64

Example:

Partition with 4.2 entitled capacity:

Minimum virtual processors = _____

Maximum virtual processors = _____

Virtual Processors (3 of 3)

- The number of virtual processors does not change the *entitled capacity*
- Example: A partition has 1.5 *capped* processing units
 - For each 10 ms time slice, the partition is only entitled to 15 ms of processing time
 - If 2 virtual processors
 - 7.5 ms runs on 2 physical processors
 - If 4 virtual processors
 - 3.75 ms runs on 4 physical processors
- For uncapped partitions, performance may be limited if the virtual processor setting is too low
 - Number of simultaneous physical processors is limited by the virtual processor setting
 - Each virtual processor is worth up to 10 ms of processing time

Unit Summary



- Processing modes are **Dedicated** and **Shared**
- Shared processors are allocated in **processing units** from one **shared processing pool**
- **Entitled capacity** is the amount of processing power a partition is guaranteed
- Entitled capacity is assigned as milliseconds within a **10 ms time slice**
- **Capped** partitions cannot exceed their entitled capacity
- **Uncapped** partitions may use excess capacity in the shared processing pool
- Priority is determined for uncapped partitions by their **weight value**
- The operating system in a partition views **virtual processors** as physical processors

Checkpoint (1 of 3)

1) Match the following processor terms to the statements that describe them:

Dedicated Shared Capped Uncapped Virtual Logical

- a. _____ These processors cannot be used in Micro-Partitions.
- b. _____ Partitions marked as this may use excess processing capacity in the shared pool.
- c. _____ There are two of these for each virtual processor if simultaneous multi-threading is enabled.
- d. _____ This type of processor must be configured in whole processor units.
- e. _____ These processors are configured in processing units as small as one hundredth of a processor.
- f. _____ By default, there is one of these for each whole processing unit, or part thereof.
- g. _____ Partitions marked as this may use up to their entitled capacity but not more.

2. Can a partition have both dedicated and shared processors?

Checkpoint Solution (1 of 3)

1) Match the following processor terms to the statements that describe them:

Dedicated **Shared** **Capped** **Uncapped** **Virtual** **Logical**

- a. **Dedicated** These processors cannot be used in Micro-Partitions.
- b. **Uncapped** Partitions marked as this may use excess processing capacity in the shared pool.
- c. **Logical** There are two of these for each virtual processor if simultaneous multi-threading is enabled.
- d. **Dedicated** This type of processor must be configured in whole processor units.
- e. **Shared** These processors are configured in processing units as small as one hundredth of a processor.
- f. **Virtual** By default, there is one of these for each whole processing unit, or part thereof.
- g. **Capped** Partitions marked as this may use up to their entitled capacity but not more.

2. Can a partition have both dedicated and shared processors?

No, a partition must use one or the other.

Checkpoint (2 of 3)

3. *True or False:* Dedicated processors are returned to the shared processor pool if the dedicated partition becomes inactive.

4. If a partition has 2.5 processing units, what is the *minimum* number of virtual processors it must have?
 - a. 1
 - b. 3
 - c. No minimum

5. If a partition has 2.5 processing units, what is the *maximum* number of virtual processors it can have?
 - a. 25
 - b. 30
 - c. Total number of physical processors X 10
 - d. No maximum

Checkpoint Solution (2 of 3)

3. *True or False:* Dedicated processors are returned to the shared processor pool if the dedicated partition becomes inactive.
Trick question! The answer is maybe and it depends on how the dedicated partition is configured.
4. If a partition has 2.5 processing units, what is the *minimum* number of virtual processors it must have?
a. 1
b. 3
c. No minimum
5. If a partition has 2.5 processing units, what is the *maximum* number of virtual processors it can have?
a. 25 **(Max can be no more than 10 times processing units.)**
b. 30
c. Total number of physical processors X 10
d. No maximum

6. What is the minimum number of processing units in a partition?

Checkpoint (3 of 3)

What is the smallest granularity of processing units when allocating additional processing units to a partition?

- What is the maximum amount of processing units that can be allocated to a partition?
- If there are multiple uncapped partitions running, how are excess shared processor pool resources divided between the partitions?
- *True or False:* If a partition is allocated 2.5 processing units, this means that it can use up to 25 ms of processing time for every 10 ms of clock time.
- What is the maximum number of virtual processors that can be configured for an individual partition?

- What is the minimum number of processing units in a partition?

0.1 processing units

Checkpoint Solution (3 of 3)

What is the smallest granularity of processing units when allocating additional processing units to a partition?

0.01 processing units

- What is the maximum amount of processing units that can be allocated to a partition? All available processing units
- If there are multiple uncapped partitions running, how are excess shared processor pool resources divided between the partitions?
The uncapped weight configuration value is used to allocate excess resources.
- *True or False:* If a partition is allocated 2.5 processing units, this means that it can use up to 25 ms of processing time for every 10 ms of clock time. True
- What is the maximum number of virtual processors that can be configured for an individual partition? Up to ten times the amount of processing units.

Virtual Ethernet Overview (1 of 3)

- Virtual Ethernet:

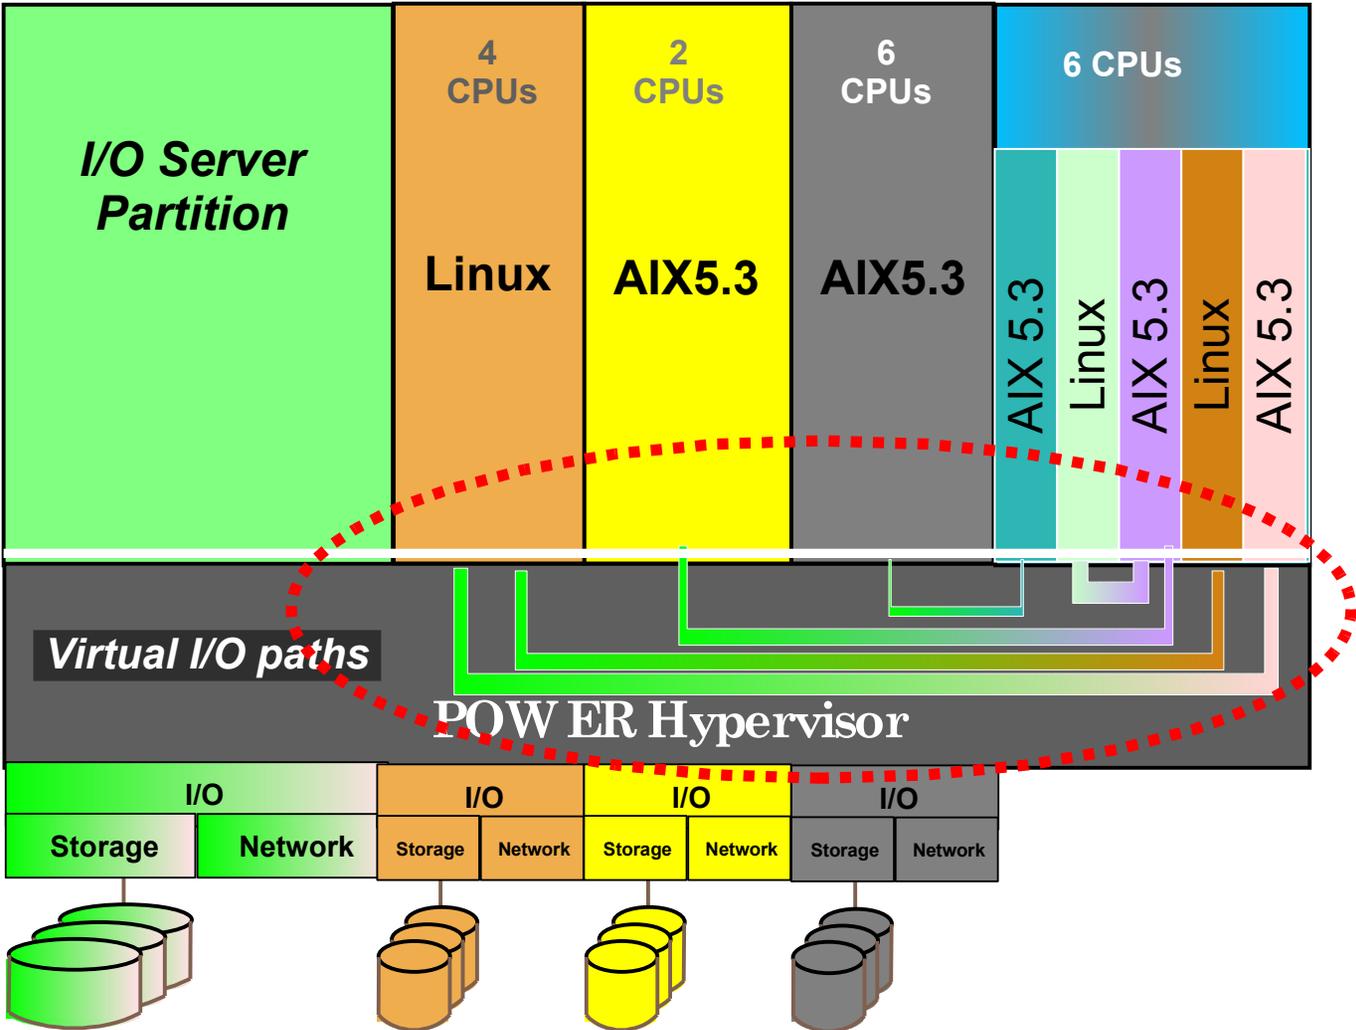
- Enables inter-partition communication without the need for physical network adapters assigned to each partition
- Adapters are virtual and if communication is only between partitions on same managed system, no physical Ethernet adapter is needed
- Requires a POWER5 processor-based system and HMC for configuration
- Partitions can run either AIX 5L V5.3 or Linux
- Does not require the purchase of any additional features or software such as the Advanced POWER Virtualization feature

Virtual Ethernet Overview (2 of 3)

- Up to 256 virtual Ethernet adapters supported per partition
- Support for large MTUs (up to 65394) on AIX 5L
- Supports both IPv4 and IPv6 standards
- Virtual Ethernet adapter with only a PVID is configured just like physical Ethernet adapter
- Virtual Ethernet adapter with multiple VIDs is configured as a VLAN device with tagging

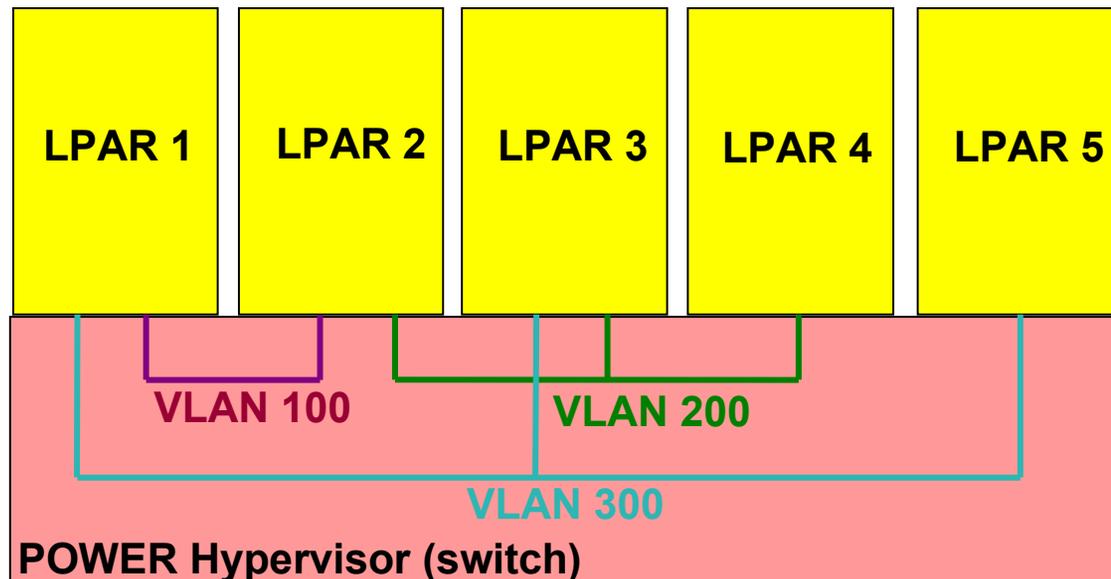
Virtual Ethernet Overview (3 of 3)

- Allows partitions to communicate without using a physical Ethernet adapter

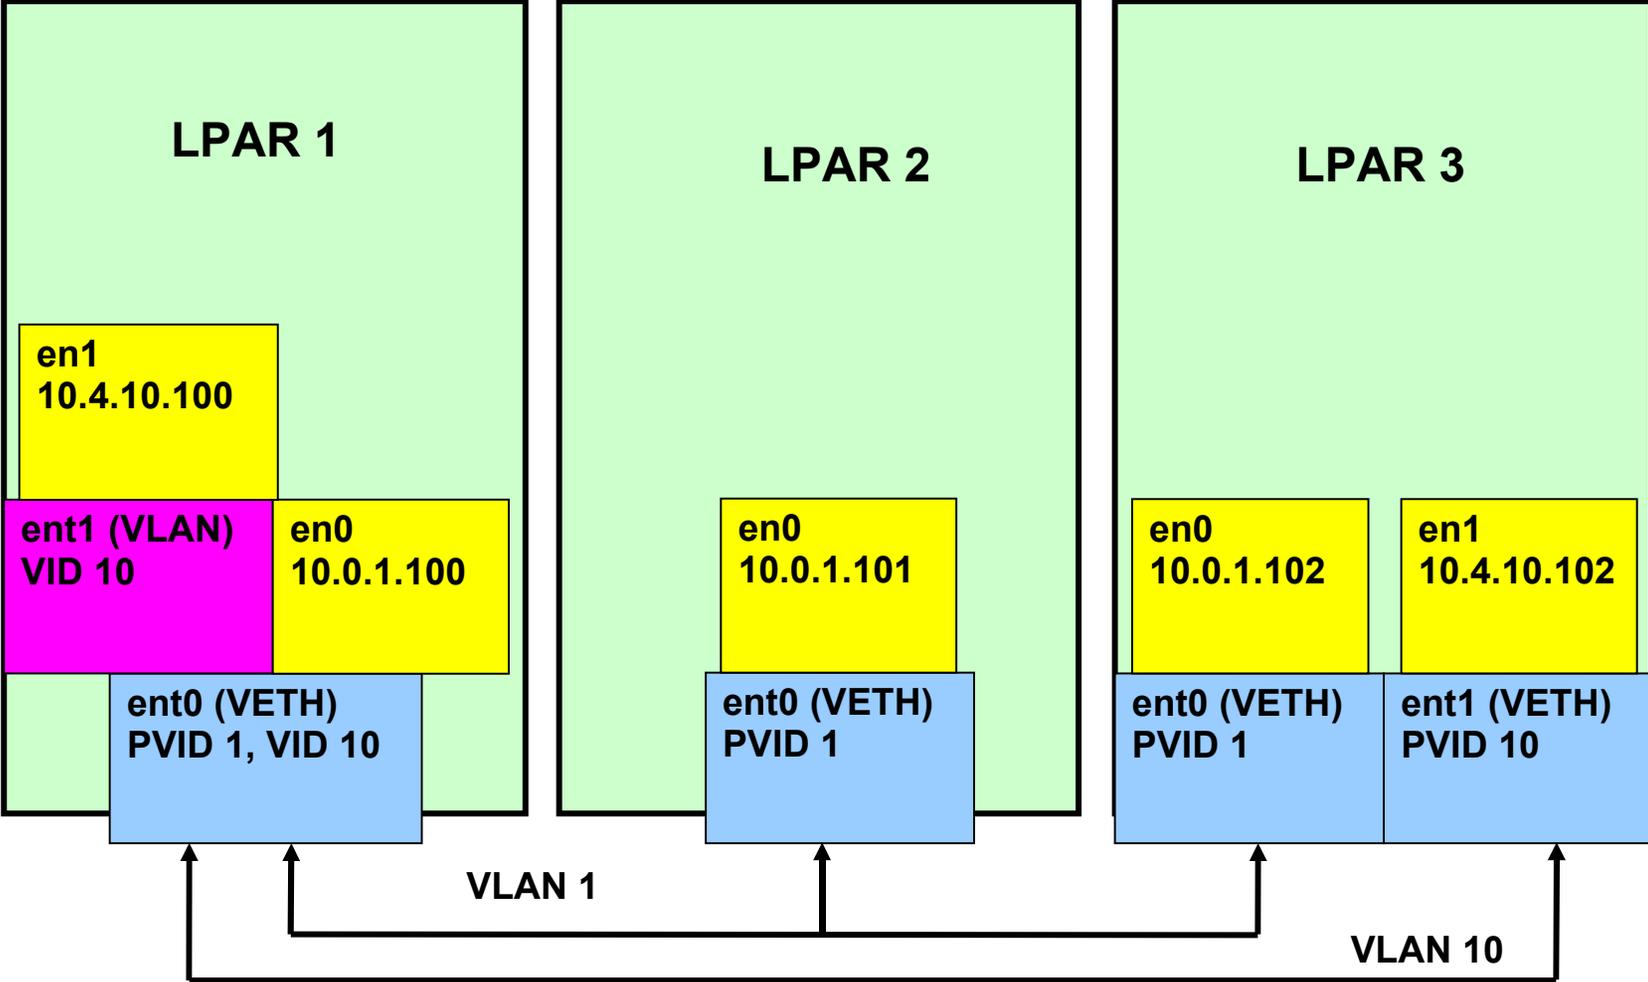


Virtual Ethernet Adapters

- Virtual Ethernet adapters are required to access the Hypervisor Ethernet switch
- Use HMC to create virtual adapters
 - Created in a partition's profile or dynamically
- HMC generates MAC address
 - Local bit to prevent conflict with physical Ethernet
 - Based on machine serial number for uniqueness
- Access to 20 VLAN (VID) networks per adapter (plus the PVID)



VLAN Configuration Examples

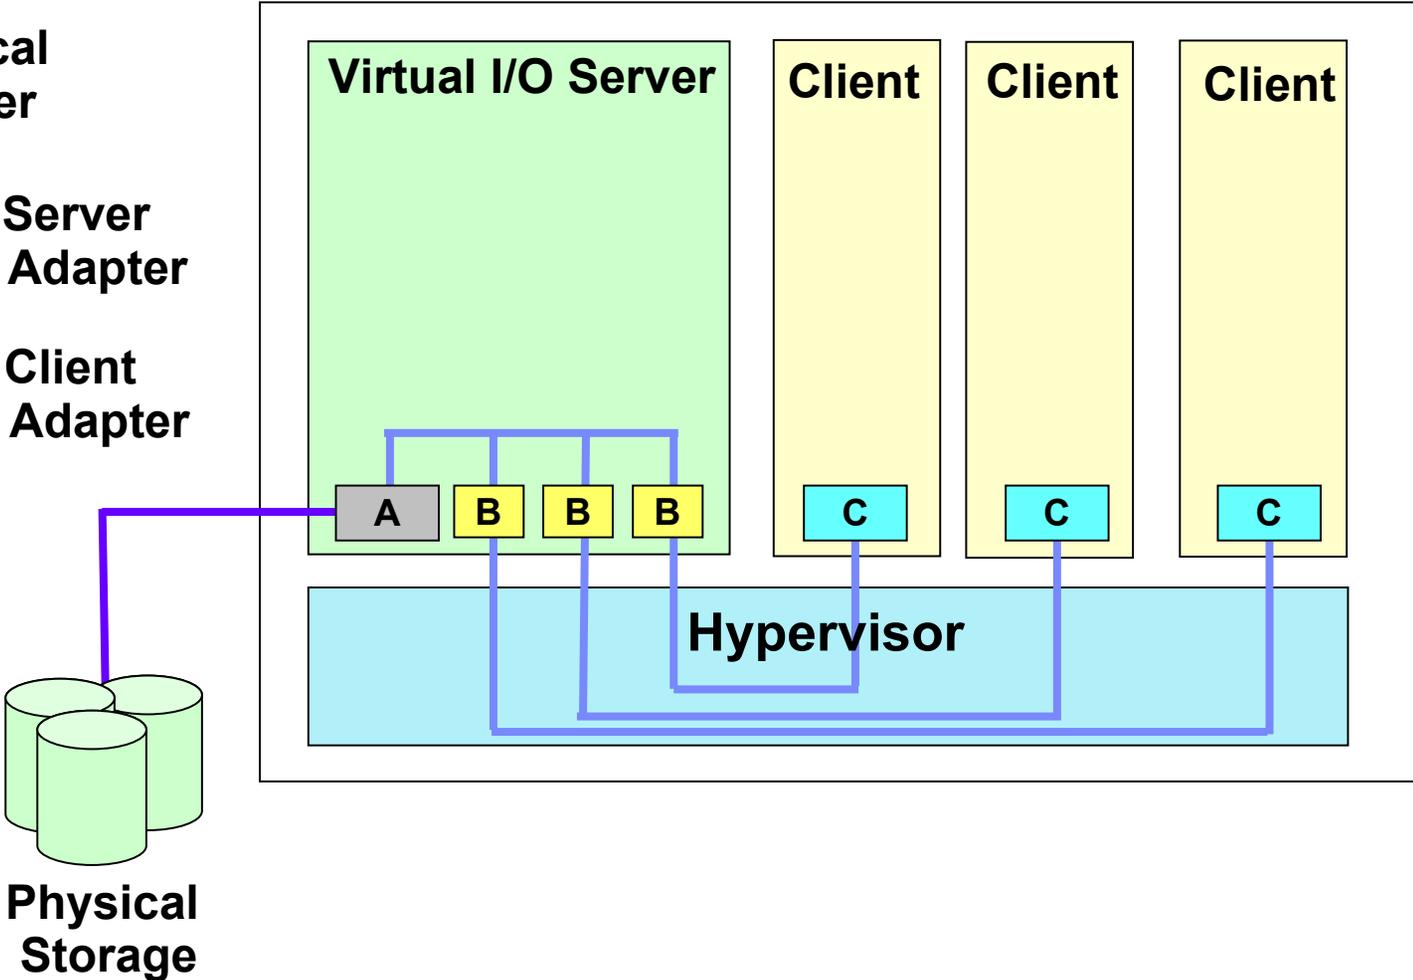


Virtual SCSI Overview (1 of 2)

- Implemented as client/server relationship
- Virtual I/O server owns the physical resource
- Client partition sees standard SCSI devices
- Current device types supported:
 - Disk (backed by physical volume)
 - Disk (backed by logical volume)
 - Optical (DVD-ROM, DVD-RAM, and CD-ROM)
- Inter-partition communication provided by Hypervisor
- Client partition can boot from virtual SCSI devices

Virtual SCSI Overview (2 of 2)

- A** Physical Adapter
- B** VSCSI Server Virtual Adapter
- C** VSCSI Client Virtual Adapter

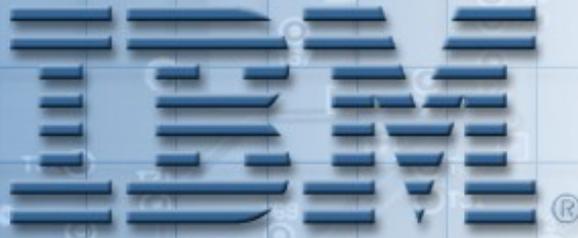


Checkpoint (2 of 2)

- *True or False:* A partition may boot from virtual SCSI disk.
- *True or False:* A partition may not use both virtual SCSI disk and virtual Ethernet at the same time.
- When should the *Access external network* setting of a virtual Ethernet adapter be used?

Checkpoint Solution (2 of 2)

- *True or False: A partition may boot from virtual SCSI disk.*
True.
- *True or False: A partition may not use both virtual SCSI disk and virtual Ethernet at the same time.*
False. A partition can use both virtual SCSI disk and virtual Ethernet at the same time.
- When should the *Access external network* setting of a virtual Ethernet adapter be used? The setting should be used on the virtual adapter associated with the shared Ethernet adapter used by the virtual I/O server partition to connect internal VLANs to the outside network.



Systemverwaltung 2009

AIX / Hardware



Location Codes

- Types of location codes
 - Physical location codes refer to a specific component
 - A physical location code is a sequence of location labels that, when followed in order, leads to the resource
 - Operating system location codes also refer to components and use a different convention
 - Example for AIX 5L, hdisk0: 05-08-00-3,0
- Virtual devices have codes as if there were on another physical planar

Physical Location Code Convention

- Physical location code format
 - **U**ttt.**m**mm.**s**sssss-additional device information

U ↑ **m** ↑ **s** ↑
Unit Type . Model . Serial number

- Example physical location code
 - **U787A.001.DNZ0713-P1-C3**
 - *P1* specifies Planar 1
 - *C3* specifies Card slot 3
 - After a card slot, a T followed by a number specifies a particular port
 - If no card slot, a T followed by a number specifies an integrated slot
 - An L followed by a number specifies a logical path identifier specific to the protocol being used (for example, SCSI identifiers)

AIX 5L Location Codes

- Use `lscfg` to see physical location codes

- For `hdisk0`

```
lscfg | grep hdisk0
hdisk0 U787A.001.DNZ00G0-P1-C1-T1-L3-L0
```

- For Ethernet adapter

```
lscfg | grep en
ent0 U787A.001.DNZ00G0-P1-C3-T1
```



- Use `lsdev` to see AIX 5L location code

- For disks

```
lsdev -Cc disk
hdisk0 Available 05-08-00-3,0
hdisk1 Available 05-08-00-4,0
```



- For Ethernet adapter

```
lsdev | grep ent
ent0 Available 03-08
```