

iSeries Logical Partitioning

Logical Partition (LPAR)



LPARs operate independently!

||

iSeries Partition Characteristics

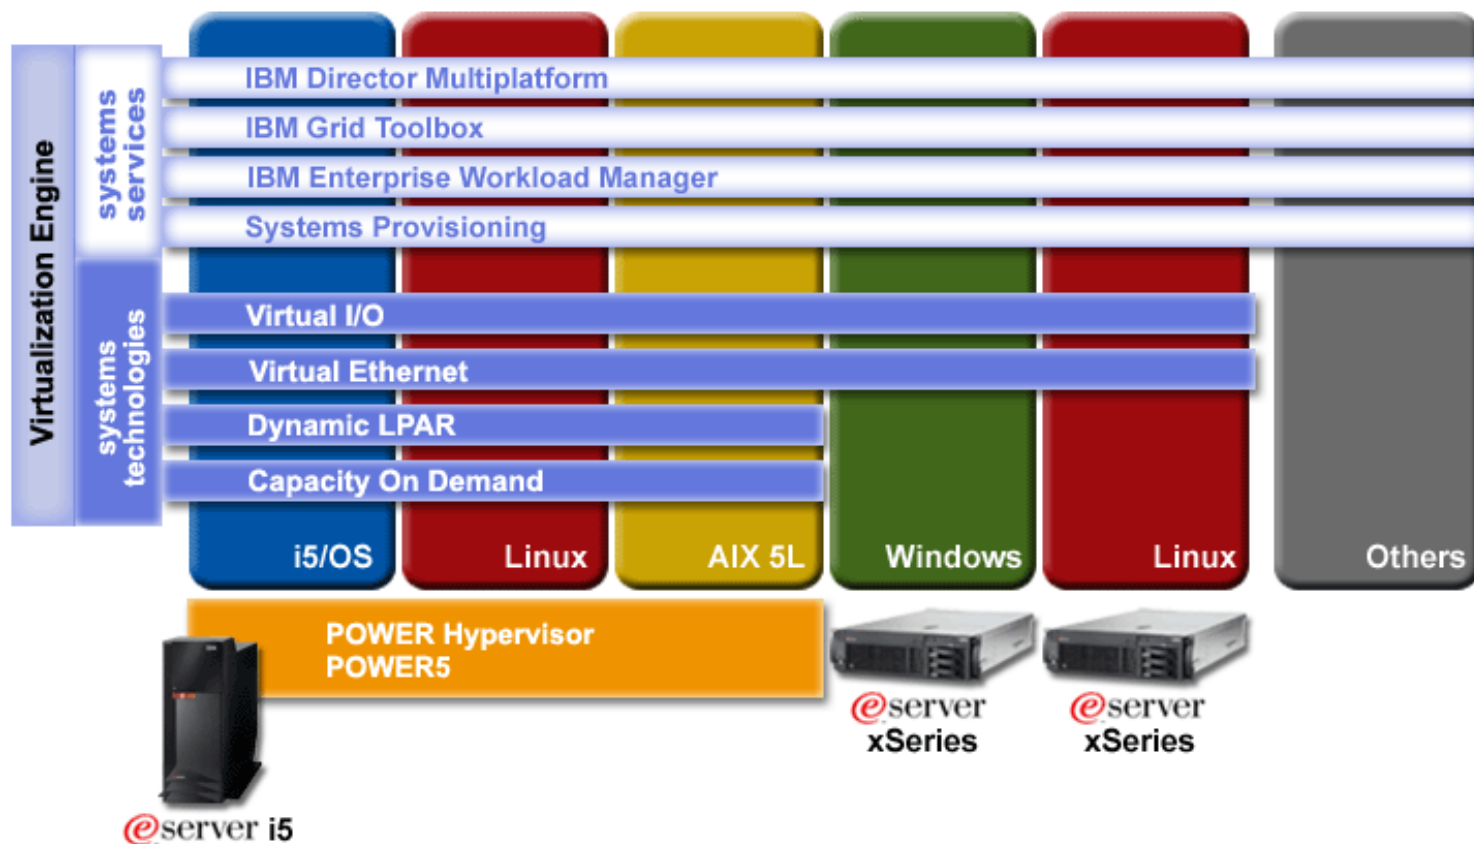
- Each partition has its own:
 - Operating system
 - License Internal Code (LIC) for i5/OS or Open Firmware for AIX 5L and Linux
 - Console
 - Load source disk or boot resources
 - ASPs and memory pools
 - And others expected in a standalone operating system environment:
 - Problem logs
 - Data (libraries, objects, filesystems)
 - Performance characteristics
 - Language feature codes
 - Network identity
 - Date and time

Benefits of Using Partitions

- Capacity management
 - Flexibility to allocate resources
- Consolidation
 - Consolidate hardware, floor space, software licenses, support contracts, in-house support/operations
 - Efficient use of resources
- Application isolation
 - Separate workloads
 - Guaranteed resources
 - Data integrity
- Merge production and test environments
 - Share same hardware
- Ability to have virtual Ethernet and virtual I/O devices

Virtual Resources

IBM Virtualization Engine



Flexibility and scalability

- Dynamic LPAR (CPU, memory, I/O)
- Virtual I/O
- Automatic CPU balancing
- Capacity on Demand


Enterprise-wide integration

- Cross-LPAR system technologies
- Cross-platform system services


Shared Processor

Understanding Shared Processors


To understand Processing Units – there are four main concepts



1. One single processor is equivalent to 1.00 Processing Units, 2 Processors = 2.00 Processing Units, etc. 0.5 processing units is NOT same as half a processor.




2. Shared processor pool. A processor must live in the shared processor pool (now the default) to become Processing units.



3. Virtual Processor – how many processors do you want the partition to be able to use (run jobs/threads on) simultaneously. It's also the number of processors that the operating system thinks it has to use.

10 Milliseconds Time Slice

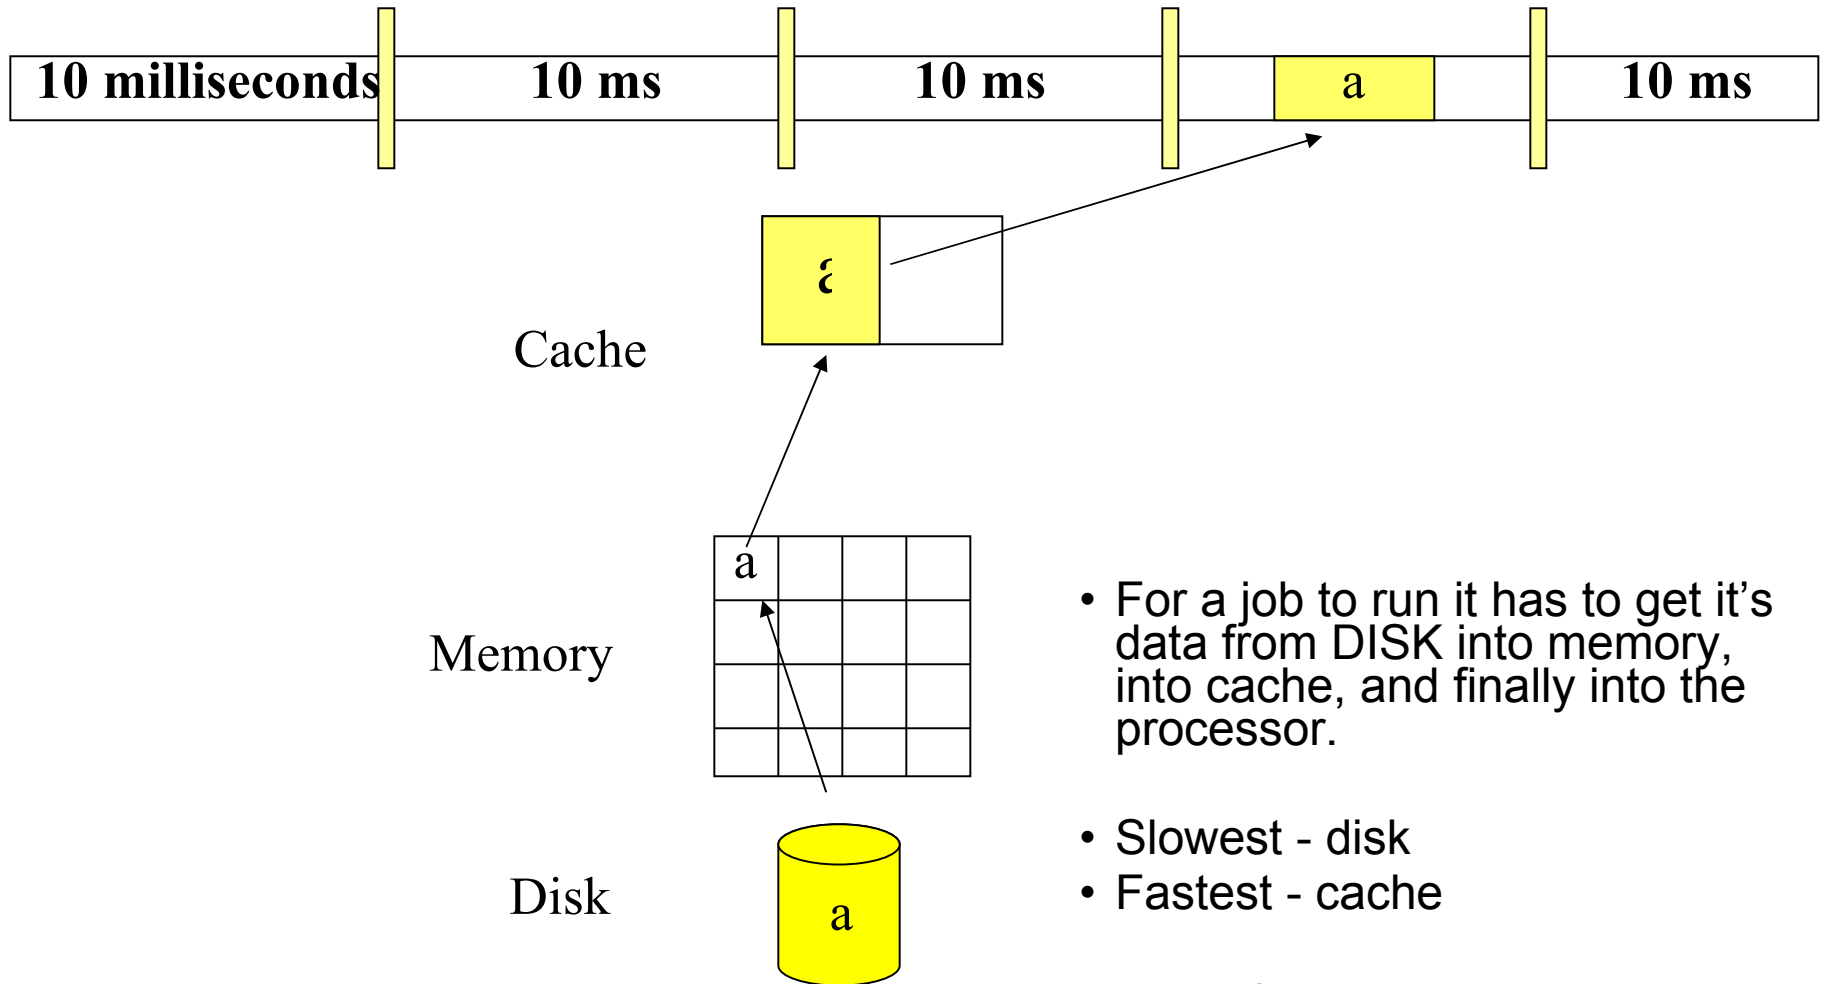
4. The iSeries processors run on 10 ms time slices



**Each Processor use is allocated within
a 10 ms Cycle**

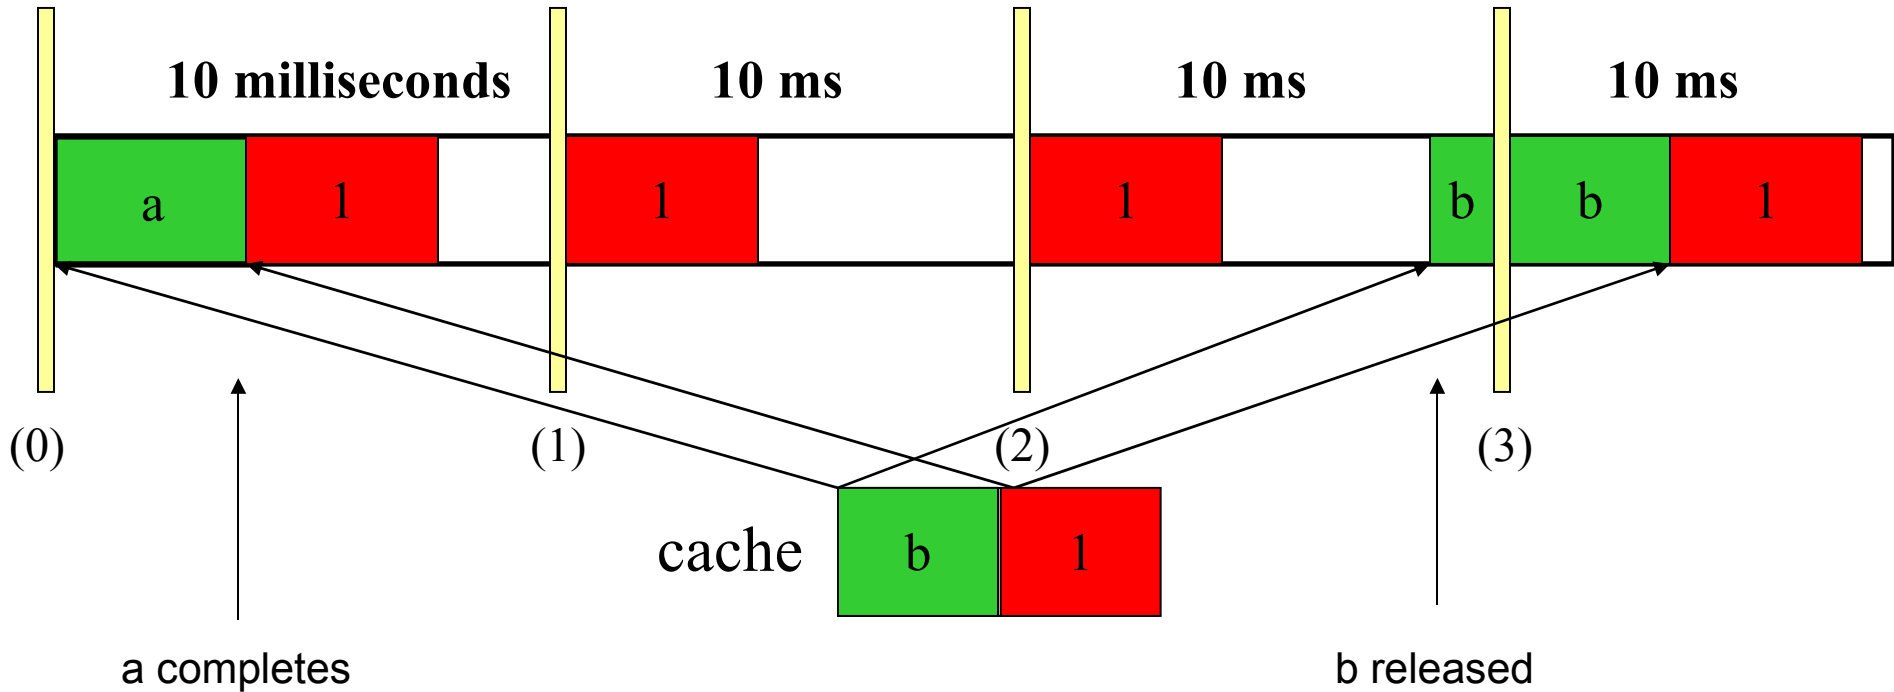
- A partition's use of a processor is limited to its allocation during each 10 ms cycle.
- For example, 80% of a processor (.80 Processing Units) yields up to 8 ms of processing time out of this 10 ms time slice. It also yields .8 X CPW rating of the processor.
- Every 10 ms this cycle repeats itself.

How Does a Job Get Into the Processor?



- For a job to run it has to get its data from DISK into memory, into cache, and finally into the processor.
- Slowest - disk
- Fastest - cache
- Even if the data is already in memory it has to be loaded into cache to run – this takes time.

Example of Two Partitions Sharing a Processor (“capped”)



Partition **dog** jobs **a,b,c** allocated .6 Processing Units

Partition **cat** jobs **1,2,3** allocated .4 Processing Units

Potential Shared Processor Penalty

- There is a potential for a performance penalty (from 0 to 10%) when using shared processors, due to:
 - Increasing the possibility that jobs won't complete, and
 - Having to be redispach and potentially reload cache, and
 - Increasing the chance of a cache miss
- Reduce the chance for processor and memory affinity
- The POWER Hypervisor overhead of:
 - Managing multiple processors
 - Tracking each partitions use of its allotted milliseconds
 - Managing time slices from multiple partitions
- All of the above are affected by how you allocate your virtual processors – next couple of foils

Desired Minimum/Maximum Processing Units

Create Logical Partition Profile - Processing Settings

Specify the desired, minimum, and maximum processing settings in the fields below.

Total usable processing units: 2.00

Desired processing units:

Minimum processing units:

Maximum processing units:

- How about 0.2 Processing Units
- Minimum of .1
- Maximum of 2.00
- Select Advanced

Capped Partitions

Advanced Processing Settings

Sharing modes

You must specify a processing sharing mode for this partition profile.

Capped
The processor usage never exceeds the assigned processing capacity.

Uncapped
Processing capacity may be exceeded when the shared processing pool has spare processing power.

Weight : 128

Virtual processors

The default virtual processor settings have been filled in for you. You may change the default settings below.

Minimum processing units required for each virtual processor : 0.10

Desired number of virtual processors : 2

Minimum number of virtual processors : 1

Maximum number of virtual processors : 20

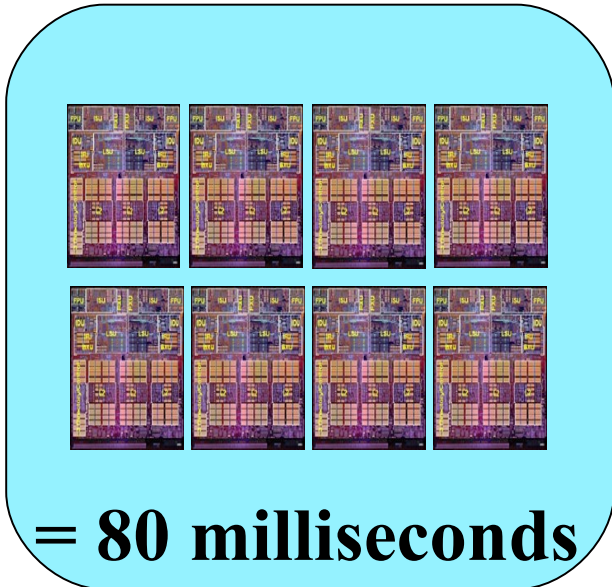
OK Cancel Help ?

- First time through we select Capped
- You can't have less than .10 processing units per virtual processor
- I'll allocate two virtuals for my .2 PUs
- What's a virtual processor?

Introduction to Virtual Processors




For every 10 milliseconds of wall clock time each processor in the shared pool is capable of 10 milliseconds of processing time

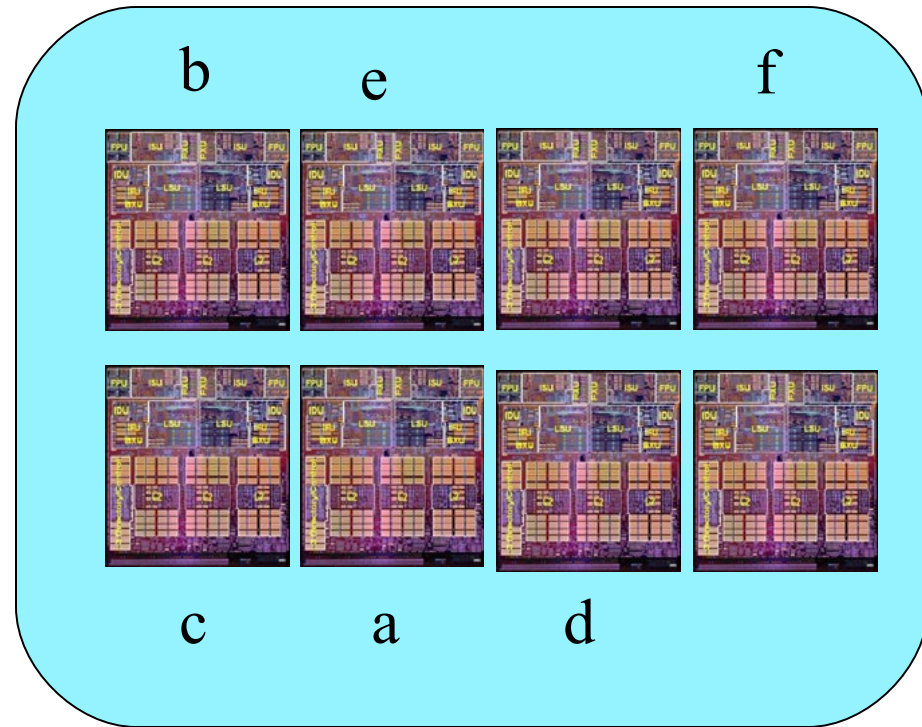


- If you give Partition x .5 processing units it could use (up to) 5 milliseconds of processing time – capped. (more on capped soon)
- But you have **ABSOLUTELY** no control over which processors your jobs/threads run on
- All you **CAN** control is how many of the processors in the pool, your jobs/threads do run on (potentially) simultaneously, via Virtual Processors

Virtual Processors - Capped



P1 1.5 processing unit default of **2** virtual processors – max of **15 milliseconds** – capped. Each job potentially could get 7.5 milliseconds ($15/2 = 7.5$)

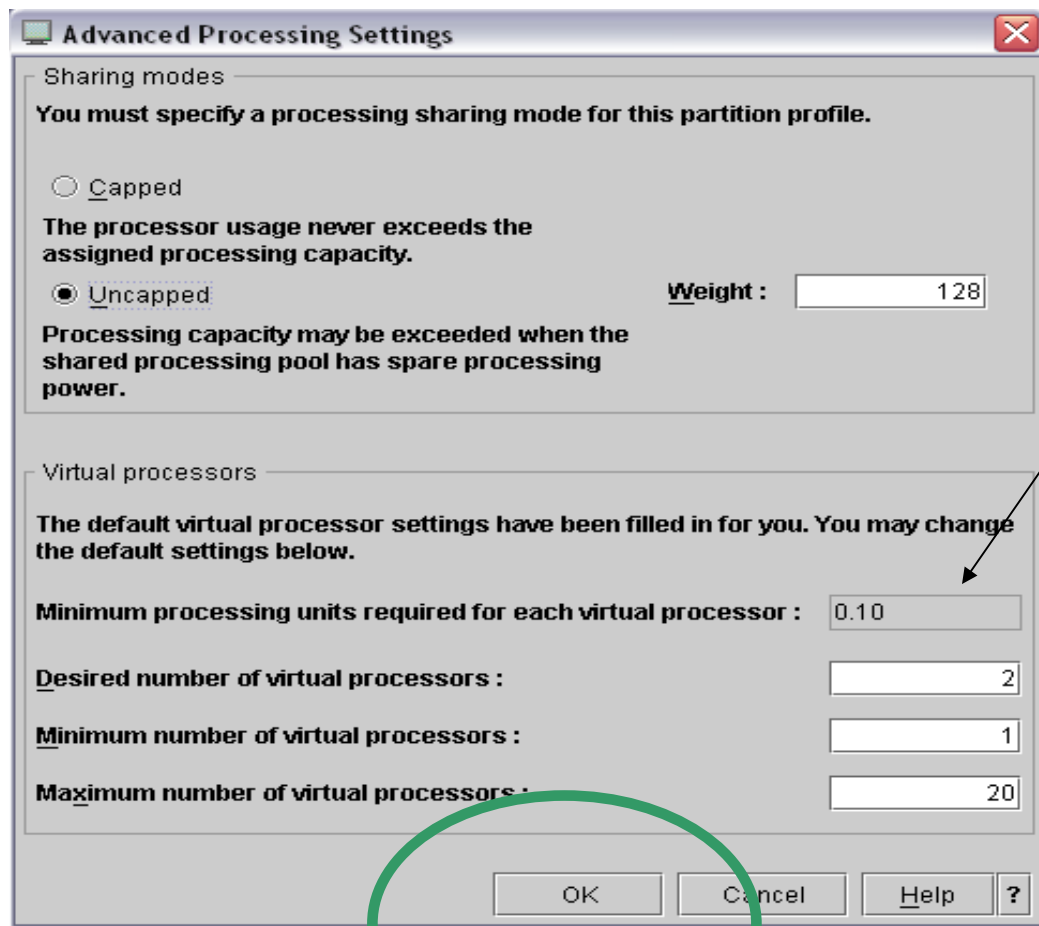


P2 1.5 processing unit but using **6** virtual processors – max of **15 milliseconds** – capped. But if all 6 Jobs ran at same time each may get no more than 2.5 milliseconds per job. ($15/6 = 2.5$)

Uncapped - Introduction


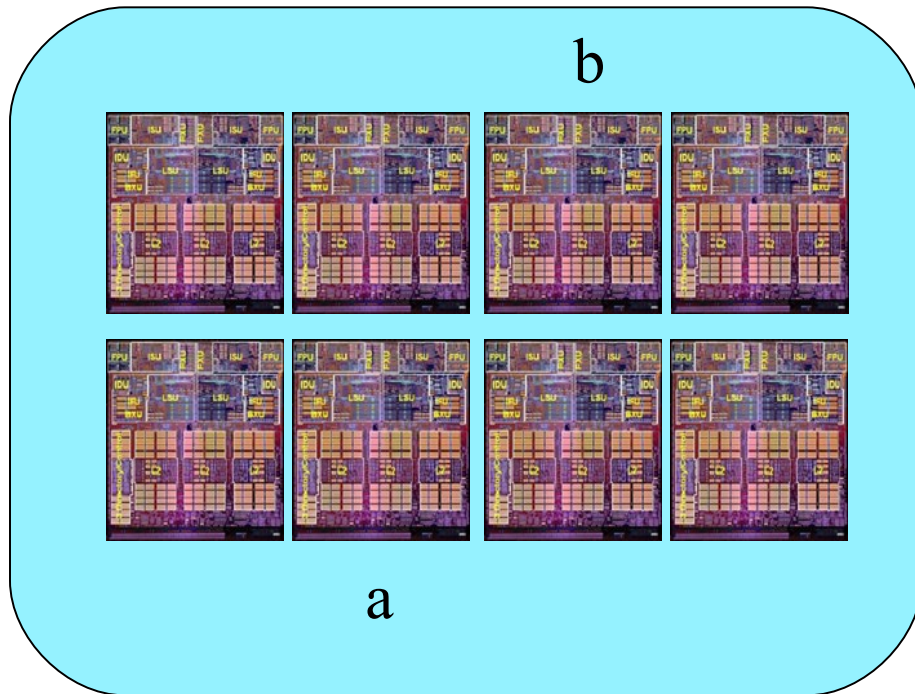
- As of IBM eServer i5/OS, and POWER5-based servers, it is now possible, by using the uncapped mode, to use more milliseconds than are allocated to a partition.
- An uncapped partition can use excess pool Processing Units.
- But even an uncapped partition could still be limited by setting the number of virtual processors too low. The number of processors it can use simultaneously is still limited by the number of virtual processors assigned.
- As of Power5 it's also possible to allocate more virtual processors than there are processors in the pool, since the actual number of processors in the pool is a 'floating' number. However, you still cannot allocate less than 1 ms (.10 PUs) per processor per job (virtual processor). For example, .5 PUs and 6 virtuals is a dog that doesn't hunt. $5 \text{ (milliseconds)}/6 \text{ (jobs)} < 1 \text{ milliseconds per job}$.

Uncapped - Configuring



- This time we deal with uncapped
- You can't have any less than .10 processing units per virtual processor
- Allocate two virtuals for my .2 PUs
- Select OK

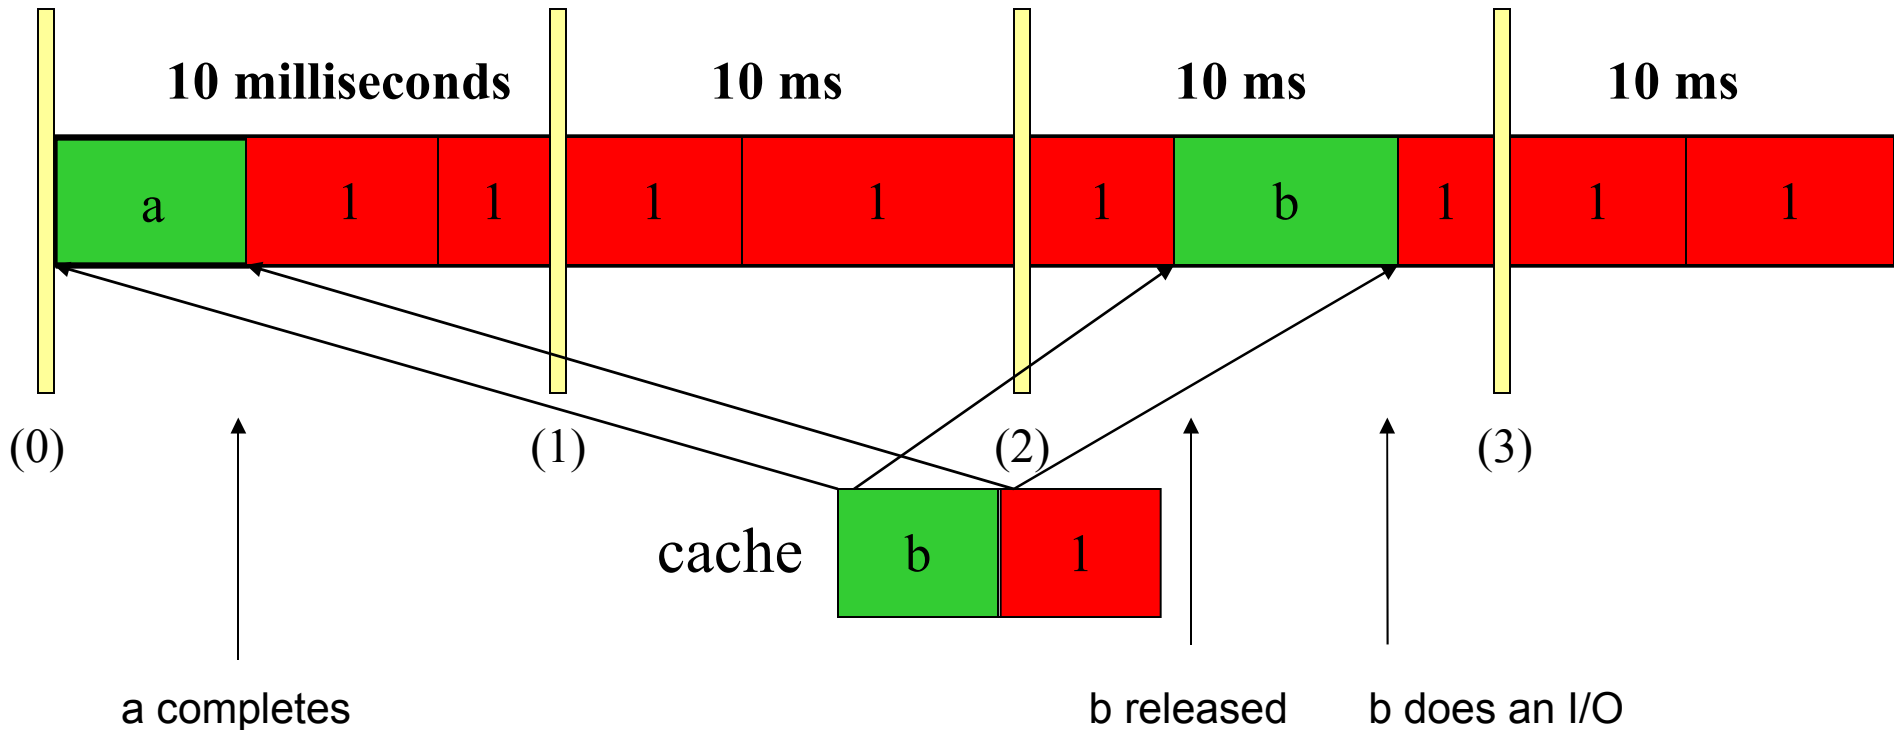
Virtual Processors (Limited) – Uncapped



P1 1.5 processing unit default of **2** virtual processors – max of **20 milliseconds** – uncapped because you are limited to only use 2 Processors simultaneously

P2 1.5 processing unit **6** virtual processors – max of **60 milliseconds** – uncapped

Example of Two Partitions Sharing a Processor (“uncapped”)




Partition **dog** jobs **a,b,c** allocated .6 Processing Units


Partition **cat** jobs **1,2,3** allocated .4 Processing Units

iSeries “uncapped shared pool” with CoD


LPAR 1 (uncapped)
 Optimized capacity of 0.4 processors for LPAR1




LPAR 2 (uncapped)
 Optimized capacity of an entire processor for LPAR 2




LPAR 3 (uncapped)
 Optimized capacity of 0.6 processors for LPAR3



- Ressources can be requested by any partition
- Unused resources can be released
- Priorities can be assigned
- Unused resources CPUs/MEM will automatically be used to solve failures in a running operating environment




Virtual Processors (Unlimited) - Uncapped



P2 1.5 processing unit with 15 Virtual Processors (maximum allowed) gives Partition – potentially – **ALL 80 milliseconds** of Processing time – for every 10 physical milliseconds BUT only as long as the other shared processor partitions DON'T have jobs ready to run!



Floating Processors

- You have eight processors on your system. Seven are in the pool and one partition uses a dedicated processor.



- Dedicated partitions can allow its processors to be used for uncapped capacity (returned to the shared pool) when the partition is powered off, or a processor is removed from the partition. This is the default.

Dedicated or Shared / Capped or Uncapped?




- The best performance may well be achieved by using dedicated processors. However, dedicated processors cannot utilize excess capacity.
- For both capped and uncapped partitions, setting the virtual processor number too high may degrade performance.
- Shared uncapped allows use of excess capacity of the processors in the shared pool. Setting virtual processors too low limits the amount of excess you can use. Setting too high may negatively impact performance.
- Also be aware for uncapped partitions the operating system sees the number of desired virtual processors as equal to the number of physical processors, you need an OS license (i5/OS, Linux and AIX 5L) for the lesser of the number of virtual processors or the number of processors in the shared pool.
- So what could be recommended? The right answer depends on workload.

LINUX and AIX Integration

AIX 5L and Linux in LPAR on eServer i5

- ❑ Mission-critical UNIX/Linux on eServer i5
 - 64-bit or 32-bit AIX 5L and 64-bit Linux kernel running on POWER5
 - Consolidate i5/OS and AIX/Linux storage
 - Utilizes IBM Virtualization Engine
- ❑ POWER5 LPAR Enhancements
 - Micro-partitioning
 - Up to 254 LPARs
 - Automatic CPU balancing
 - Virtual SCSI and Ethernet
- ❑ Across eServer i5 and p5
 - AIX 5L 5.3: Up to 10 LPARs/processor, virtual I/O
- Common Linux distributions for eServer i5 and p5
 - " RHEL 3 AS QU3 from Red Hat, Inc.
 - " SLES 9 from Novell, Inc. (SUSE part of Novell)
- ❑ Enterprise Edition (depends on model)
 - Extra Processor Activated
 - Service and Education Voucher



Create a New Partition

A - MY520

File Edit View Communication Actions Window Help

CREATE NEW PARTITION

SYSTEM: Z2238P1

COMPLETE BLANKS, PRESS ENTER.

PARTITION IDENTIFIER AND NAME 2 Linux1

NUMBER OF AVAILABLE SYSTEM PROCESSORS : 1

NUMBER OF PARTITION PROCESSORS 1

MINIMUM / MAXIMUM NUMBER OF PROCESSORS 1 / 1

USE SHARED PROCESSOR POOL 1 1=YES, 2=NO

SHARED PROCESSOR POOL UNITS 10

MINIMUM / MAXIMUM PROCESSOR POOL UNITS 10 / 1 . 00

UNCAPPED PROCESSING 1 1=YES, 2=NO

UNCAPPED PROCESSING WEIGHT 128 0, 64, 128, 255

SIZE OF AVAILABLE MEMORY (MB) : 1024

SIZE OF PARTITION MEMORY (MB) 256

MINIMUM / MAXIMUM SIZE OF MEMORY (MB) 128 / 512

ENABLE WORKLOAD MANAGEMENT 2 1=YES, 2=NO

VIRTUAL ETHERNET IDENTIFIERS (1=YES, 2=NO)

1 2 3 4

1 2 2 2

F3=EXIT F12=CANCEL

MA a 21/010

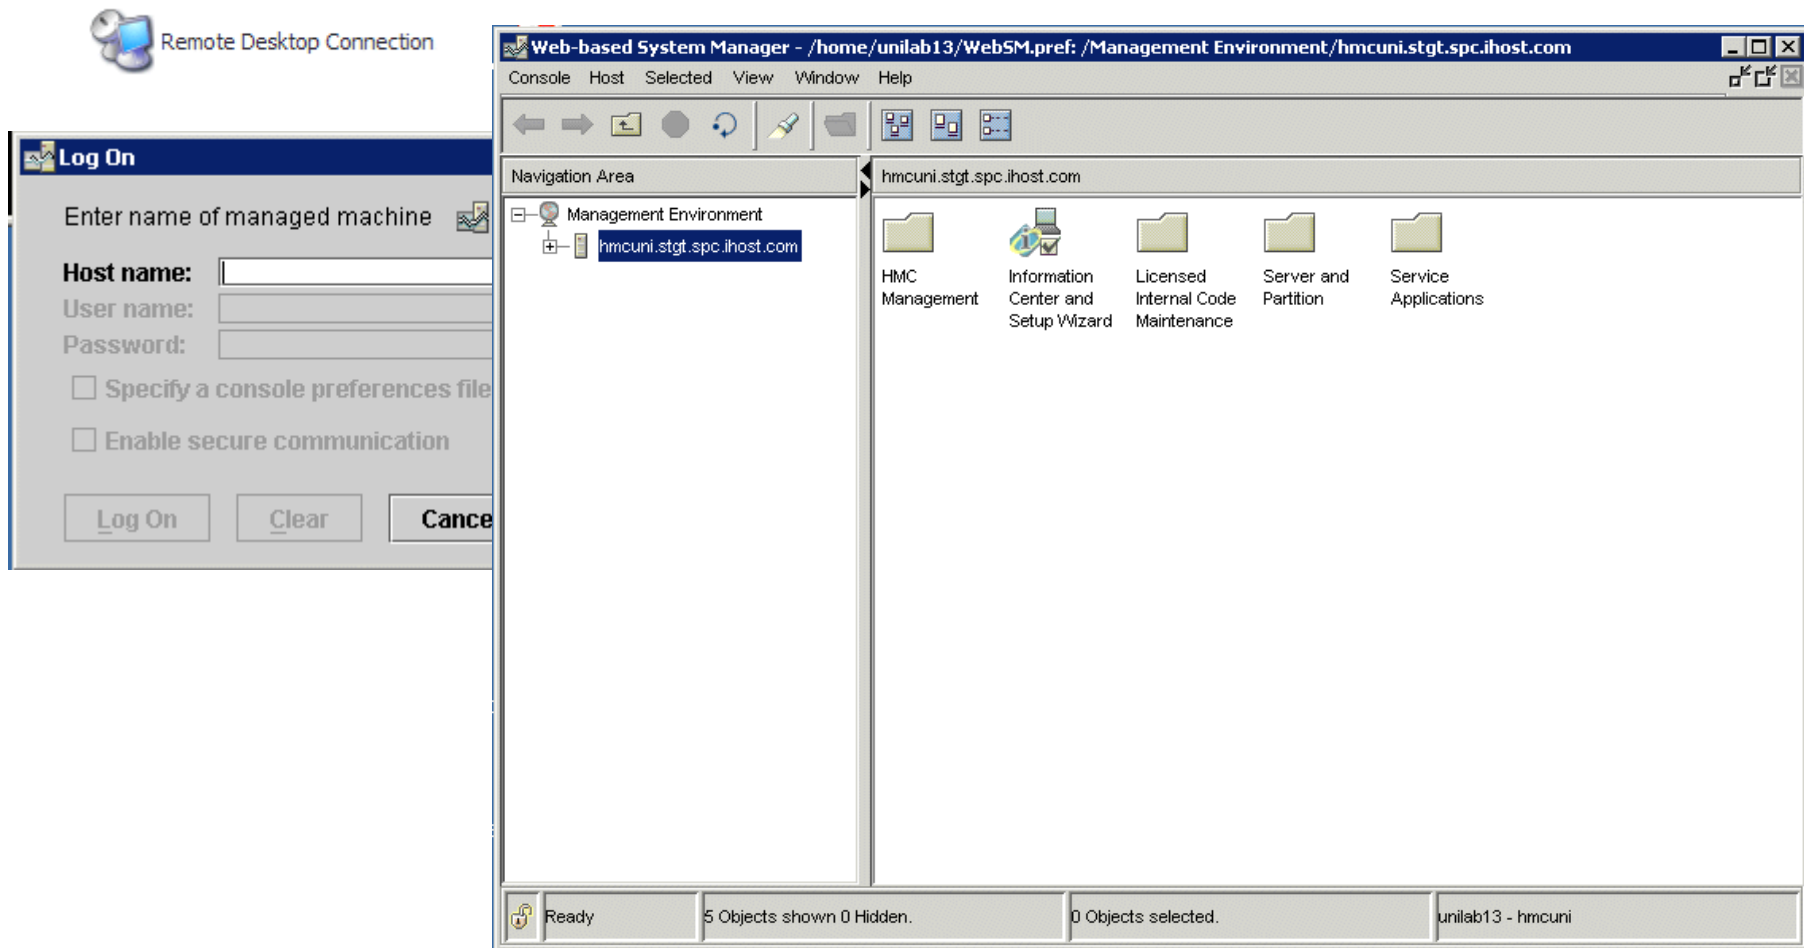
Lab Preparation

Lab Exercise Unit 13 Overview

- Use WebSM to remotely connect to the HMC
- Create a partition and a partition profile for Linux
- Create virtual adapters on the Linux partition and the i5/OS partition
- Install SUSE Linux Enterprise Server 9 from an image catalog
- Set up network communications through virtual Ethernet
- Optionally, install Platform Enablement Library

Using WebSM

- Connect to a Windows 2003 server on an IXS via Remote Desktop Connection
- From the remote desktop, start WebSM and log in to the HMC



© Copyright IBM Corporation 2005

Material may not be reproduced in whole or in part without the prior written permission of IBM.

Create a Partition

- Create a logical partition and the required partition profile using WebSM

Create Logical Partition Wizard

This wizard helps you create a new logical partition and a default profile for it. You can use the partition properties or profile properties to make changes after you complete this wizard.

Ensure you have your logical partition planning information before you use this wizard. You may also find it helpful to be familiar with logical partition concepts. Click Help for more information.

To create a partition, complete the following information:

System name : Server-9406-550-SN65DAA3E Uni

Partition ID :

Partition name :

Partition environment :

AIX or Linux

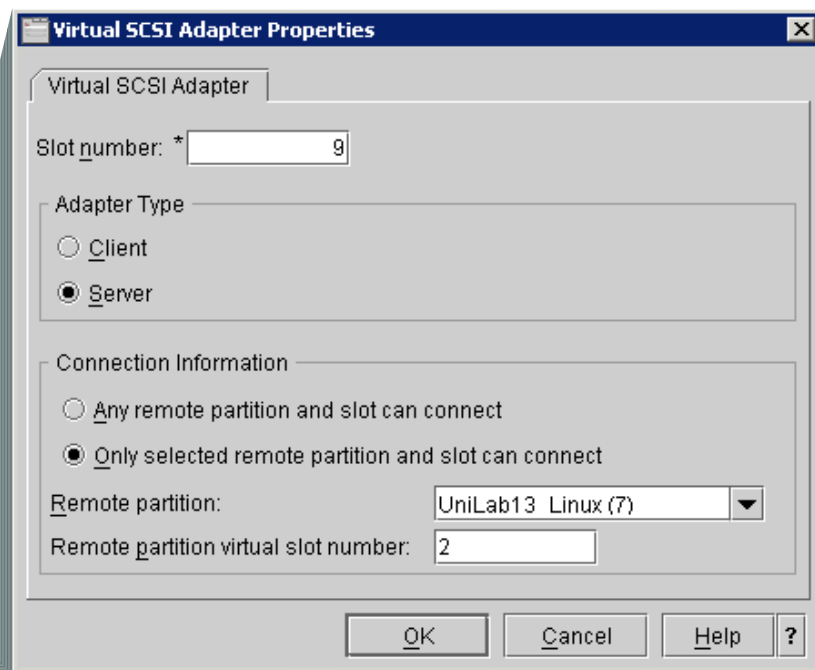
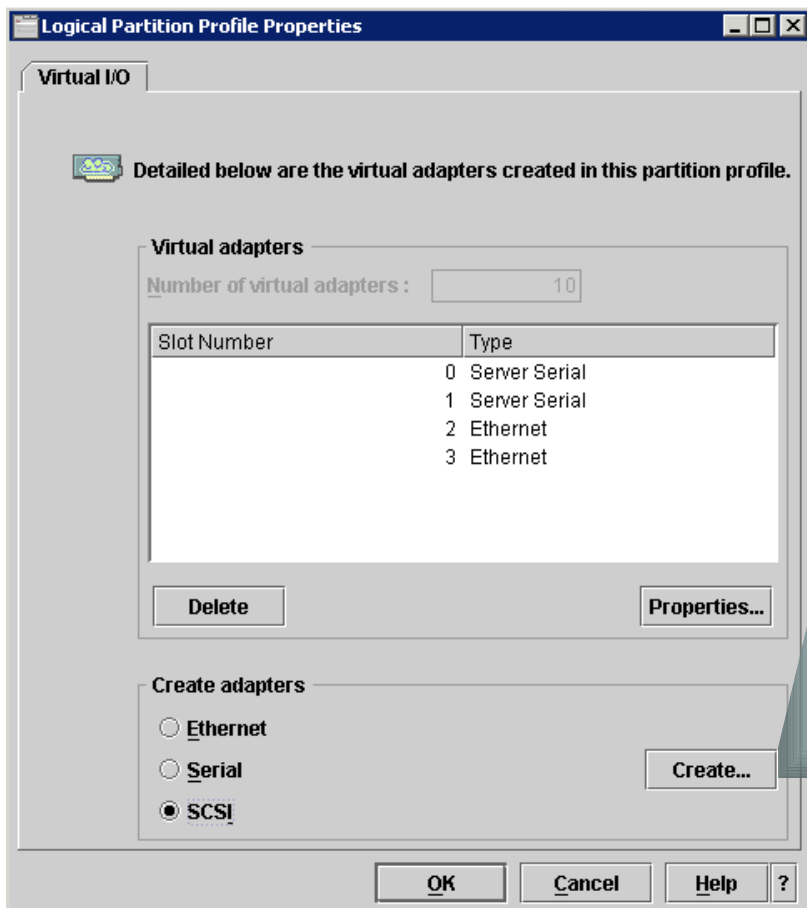
i5/OS

Virtual I/O server

Help ? < Back Next > Finish Cancel

Adding SCSI Adapter via DLPAR

- Create a virtual SCSI client adapter via the partition creation wizard
- Create a virtual SCSI server adapter via Dynamic LPAR on the i5/OS partition – Does not require any restart



Installing Linux from Virtual Media (1)

- Linux CDs are loaded in to an image catalog in i5/OS
- Virtual CD/DVD optical drive used for installation

```
Session A - [24 x 80]
File Edit View Communication Actions Window Help
Work with Image Catalog Entries
System: I5OSP5
Image Catalog . . . . . : SUSE_SLES
Image Catalog Status . . . . . : Not ready
Virtual optical device . . . . . :
Directory . . . . . : /imagecatalog/sles9

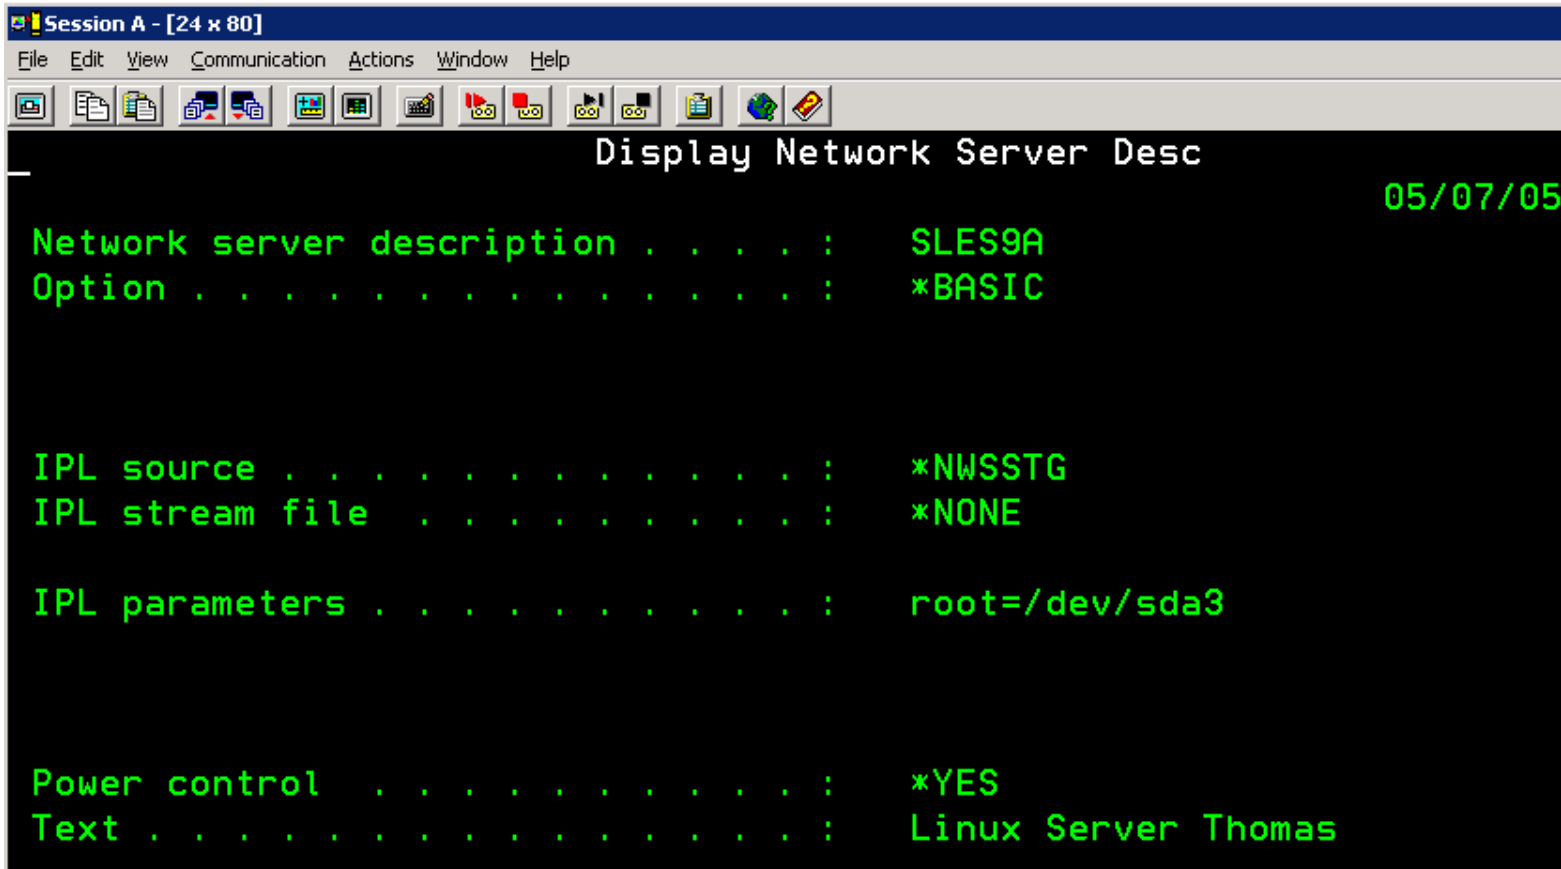
Type options, press Enter.
 1=Add   2=Change   4=Remove   6=Mount   8=Load   9=Unload
10=Initialize volume   12=Work with volume

  Opt  Index  Status  Image File Name
  ---  ---
  *AVAIL
  ---   1  Mounted  SU90.001
  ---   2  Loaded   SU90.C01
  ---   3  Loaded   SU90.C02
  ---   4  Loaded   SU90.C03
  ---   5  Loaded   SU90.C04

Bottom
F3=Exit   F5=Refresh   F6=Load image catalog   F7=Verify image catalog
F8=Reorder by index   F12=Cancel   F24=More keys
```

Installing Linux from Virtual Media (3)

- Creating Network Server Description (NWSD) and network Server Storage (NWSSTG)
- Set one-time boot parameter for installation



```
Session A - [24 x 80]
File Edit View Communication Actions Window Help
Network Server Desc
05/07/05
Network server description . . . . . : SLES9A
Option . . . . . : *BASIC

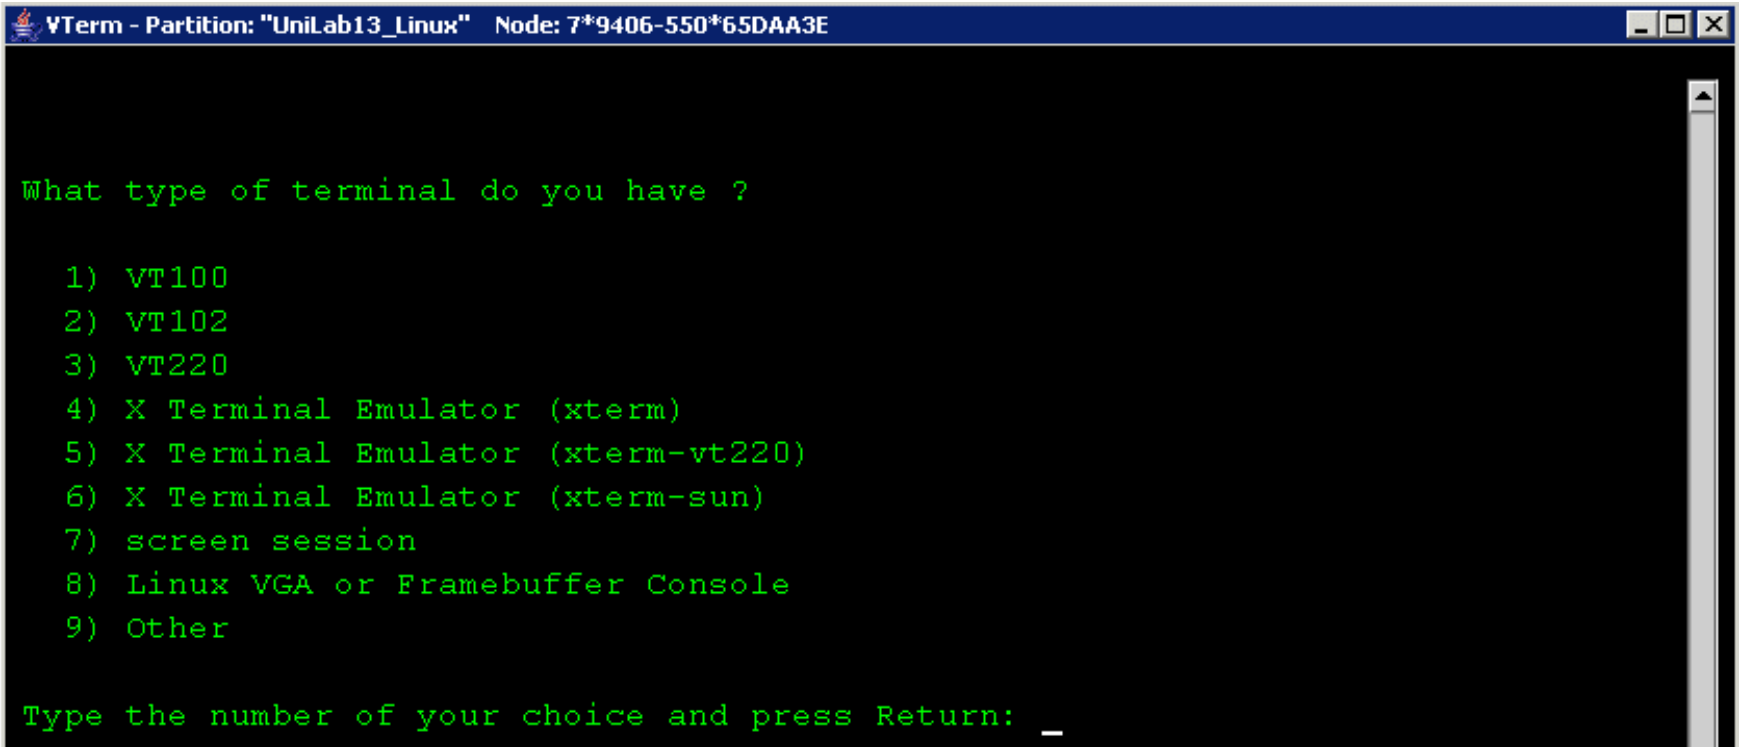
IPL source . . . . . : *NWSSTG
IPL stream file . . . . . : *NONE

IPL parameters . . . . . : root=/dev/sda3

Power control . . . . . : *YES
Text . . . . . : Linux Server Thomas
```

Installing Linux from Virtual Media (2)

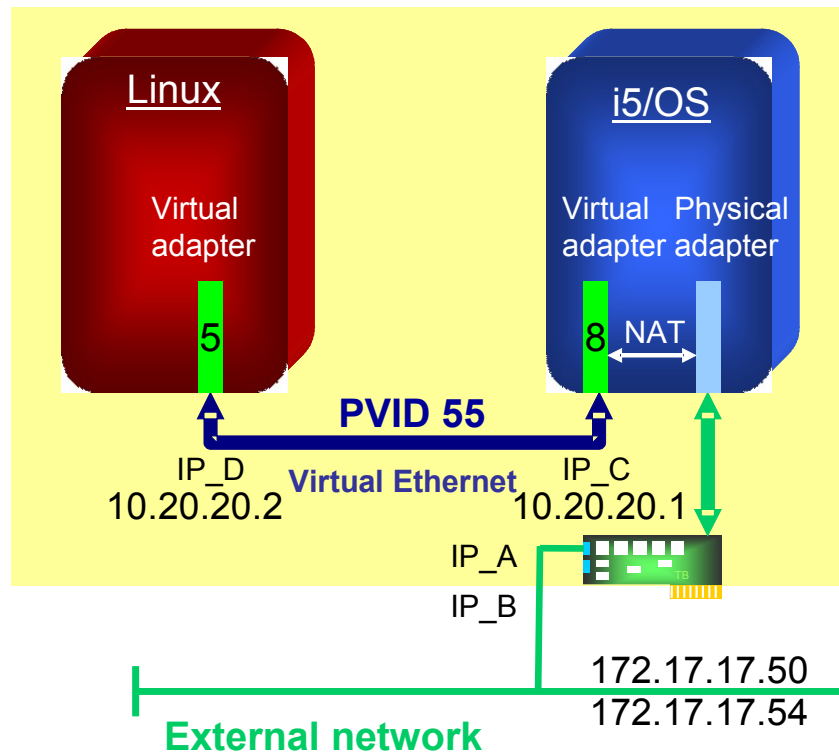
- Console session started from WebSM (virtual terminal)
- Installation from console

A screenshot of a VTerm window. The title bar reads "VTerm - Partition: 'UniLab13_Linux' Node: 7*9406-550*65DAA3E". The terminal content is as follows:

```
What type of terminal do you have ?  
  
1) VT100  
2) VT102  
3) VT220  
4) X Terminal Emulator (xterm)  
5) X Terminal Emulator (xterm-vt220)  
6) X Terminal Emulator (xterm-sun)  
7) screen session  
8) Linux VGA or Framebuffer Console  
9) Other  
  
Type the number of your choice and press Return: _
```

Setting up IP Networking (1)

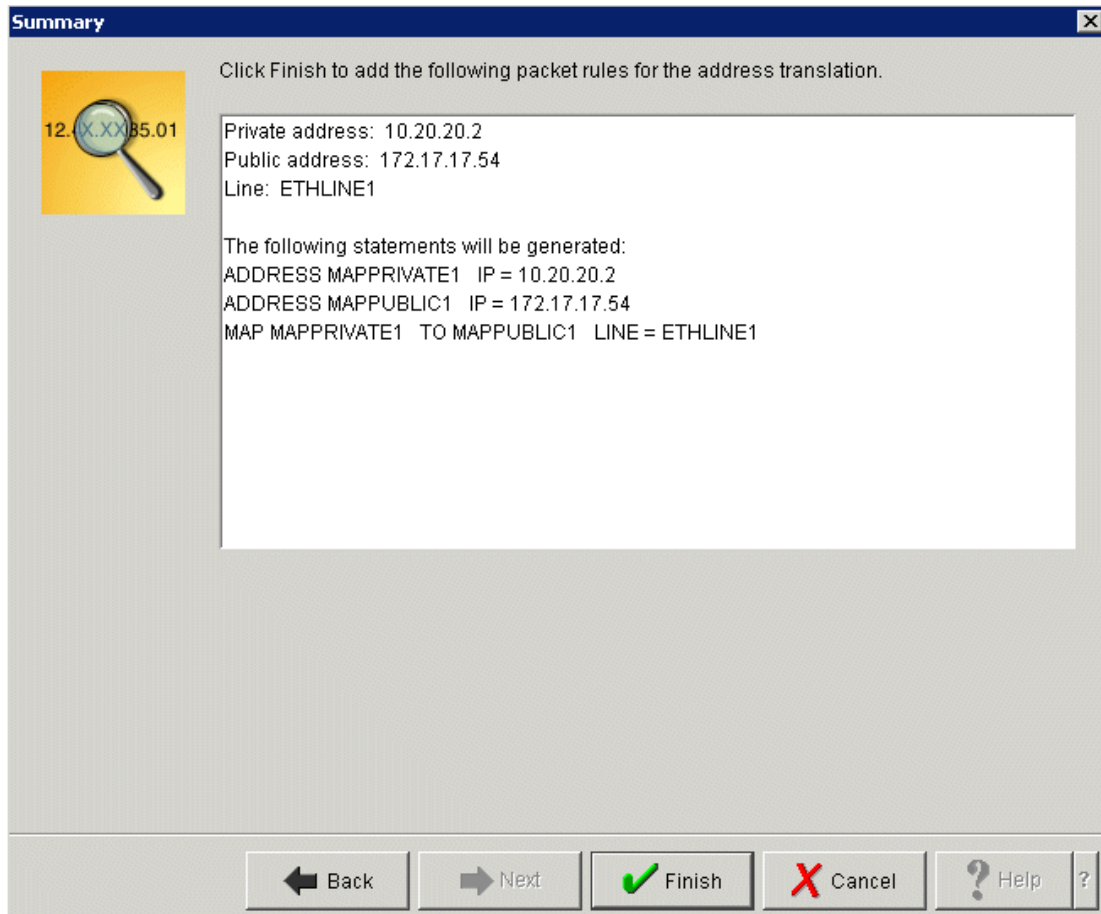
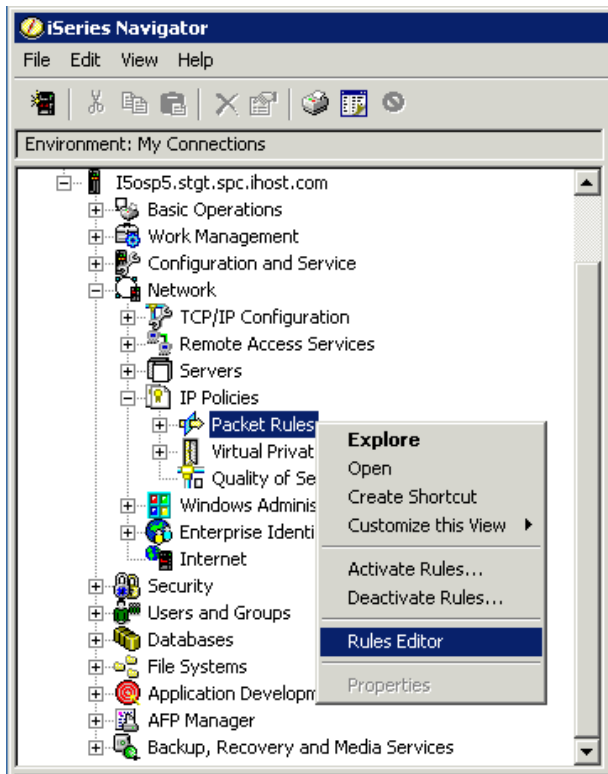
- Virtual Ethernet adapter created via partition creation wizard
- Virtual Ethernet adapter for i5/OS added via DLPAR
- External access to Linux partition through i5/OS partition



172.17.17.54 translated to 10.20.20.2

Setting up IP Networking (2)

- Communications to a partition that is only reachable via virtual Ethernet can be achieved by:
 - Using Network Address Translation (used in this lab)
 - Transparent Subnetting (a.k.a Proxy ARP)
 - IP Routing
 - Virtual I/O Server with Shared Ethernet adapter



© Copyright IBM Corporation 2005

Material may not be reproduced in whole or in part without the prior written permission of IBM.

Installing the Platform Enablement Library (optional)

- To enable full Linux integration in a hosting i5/OS environment, additional software needs to be installed under Linux
- Enables full shutdown support for Linux when NWSD is varied off