



IBM eServerJ iSeriesJ

# iSeries Migration Topics

ITSO Forum February 2003

ITSO iSeries Technical Forum  
GP06

F03GP06.PRZ

© 2000-2003 IBM Corporation

02/21/03

# Redbooks



- This ITSO presentation was produced by:
  - ▶ Nick Harris - ITSO Rochester
  - ▶ Louis Cuypers - IBM Belgium
  - ▶ Jos Vermaere - IBM Belgium
  - ▶ Subramanian Sivaramakrishnan - IBM India
  - ▶ Dale Barrick - IBM USA

# Objectives

- During this presentation we will cover the last iSeries announcements and how they affect upgrade and migration.
- There are new PCI-X towers and expansion units. We show hw these attach and combine with new and existing systems and towers.
- Capacity Upgrade on Demand, no pain processor upgrades. We will show how these new announcements work on iSeries.
- We will also cover some iSeries advanced topics like upgrades with LPAR, Linux, Integrated Environments and SAN.



# Agenda

- **Server considerations**
- **Main Storage**
- **Expansion and IO**
- **High Speed Loop**
- **Capacity Upgrade on Demand**
- **Upgrade scenarios**
- **Advanced topics**
  - ▶ **LPAR**
  - ▶ **High Availability**
  - ▶ **iASP**
  - ▶ **Linux**
  - ▶ **Integrated Windows Environment**
  - ▶ **SAN and External devices**
- **Software consideratons**



# Agenda

- **Server considerations**
- Main Storage
- Expansion and IO
- High Speed Loop
- Capacity Upgrade on Demand
- Upgrade scenarios
- LPAR
- Linux
- Integrated Windows Environment
- SAN and External devices
- Software considerations



## Model considerations

- Model 890
  - ▶ Existing 890 available until May 2003
- Model 870
  - ▶ Same as 890 with less in it .....
  - ▶ has #0126 reduction feature
- Model 825
  - ▶ 5 disks per Expansion tray
  - ▶ Note for upgrade Model 820 has 6 per tray
- Model 810
  - ▶ HSL not HSL-2
- Model 800
  - ▶ Upgrade from 800 (300) to 800 (950 CPW)
    - Do not get additional Twinax, memory and disk
    - Websphere express not required
  - ▶ HSL not HSL-2



# General server considerations

- General
  - ▶ HSL- HSL-2 cables (conversion)
  - ▶ RAID to RAID
  - ▶ 35GB vs 70GB drives
  - ▶ New models require DVD drive
  - ▶ Buy enterprise upfront - less expensive than upgrade
  - ▶ No IXS on Standard to Enterprise
- Must load V5R2 (Feb 2003 level) and PTFs before upgrade

## Existing Models still available

- Earlier models supported by V5R2 includes:
  - ▶ 6xx, Sxx, 7xx, 170, 150, 250, 270, 820, 830, 840
- Earlier models which can still be ordered from IBM as "new boxes"
  - ▶ 250, 270, 820, 830, 840
- Earlier models which can be upgraded to newest models
  - ▶ 7xx, 270, 820, 830, 840 (7xx upgrades available until October 8, 2003)



# The New iSeries eServer Family

870



890



825



*iSeries 800*



810



# Notes: The New iSeries eServer Family

IBM iSeries eServers are built to handle the processing demands of core business applications and e-business. From the aggressively priced iSeries 800, to the superior performance of the i890, iSeries can give the customers business the tools it needs for success. The iSeries 800 and i810 bring outstanding price performance solutions for small and mid-sized business. The IBM Power4 technology used in the i825, i870 and i890 are now available in simplified editions with an outstanding value at attractive prices. The whole new family of more powerful iSeries servers bring also the possibility for enhanced web enablement. More details will be covered later in this presentation.

In order to help satisfying and balancing the processing needs a business may require, there are now enhancements available for Capacity Upgrade on Demand also. Temporary activation of processor by just switching them on and off as needed, increases the flexibility of the new servers for the i825, i870 and i890.

Other enhancements such as faster I/O processors and adapters, PCI-X, HSL-2 and 15K rpm drives are covered also during this presentation.

## **For your information:**

iSeries Models 800, 810, 825, 870, and 890 require OS/400 V5R2, plus February software updates. Refer to Information APAR II13365 for details at: [HTTP://WWW.IBM.COM/ESERVER/ISERIES/SUPPORT](http://www.ibm.com/eserver/iseries/support)

New features require OS/400 V5R2, plus February software updates. Refer to Informational APAR II13440 at: [HTTP://WWW.IBM.COM/ESERVER/ISERIES/SUPPORT](http://www.ibm.com/eserver/iseries/support).

The following is an excerpt from APAR II13440:

This informational APAR addresses the software requirements needed to support the following iSeries I/O features.

Feature # 2757, PCI-X Ultra RAID Disk Controller	Feature # 2782, PCI-X RAID Disk Unit Controller
Feature # 2844, PCI IOP	Feature # 5702, PCI-X Ultra Tape Controller
Feature # 5705, PCI-X Tape/DASD Controller	Feature # 4531/4631, DVD-ROM
Feature # 4326/4327 15K RPM disk drives	Feature # 5094 PCI-X Expansion Tower
Feature # 5294 1.8M I/O Tower	Feature # 5088 PCI-X Expansion Unit
Feature # 0588 PCI-X Expansion Unit in Rack	Feature # 5095 PCI-X Expansion Tower
Feature # 0595 PCI-X Tower Unit in Rack	Feature # 8094 Opt Base 1.8M I/O Rack

If you are planning to upgrade your existing server with any of these new features, you must have the following on your system: OS/400 V5R2 PTF cumulative package C3021520 or later for V5R2

If your server is partitioned, this level of code must be on each partition using the new feature(s).

To verify that PTF cum package C3021520 or later is installed, enter the following command from a V5R2 command line:

```
DSPPTF LICPGM(5722SS1) SELECT(TC03021)
```

If PTF TC03021 is found, cum package C3021520 is on this OS/400 partition.

In addition, you need installation media for the Licensed internal code labeled with a RSx where x is B or greater (label: LICENSED INTERNAL CODE - RSx). This does not need to be reinstalled, but is needed for potential system recovery purposes. If you need to get this media, order 5722SS1 feature code 2691.

Remember also that CUoD will require an additional set of PTFs.

# New and Changed Notions

- Server Feature Code
  - ▶ Determines total number and type of processors
- Processor Feature Code
  - ▶ Determines type of processor
- Edition Feature Code
  - ▶ Server Personality for i810, i825, i870 and i890
    - Standard Edition
    - Enterprise Edition
  - ▶ Value, Standard and Advanced Editions for the iSeries 800
- Capacity Upgrade on Demand (temporary)
  - ▶ Activation Feature Code
  - ▶ Enablement Feature Code
  - ▶ Billing Feature Code
- Capacity Upgrade on Demand (permanent)

# Notes: New and Changed Notions

The server feature code is unique to a combination of features and products included in the order. The server feature identifies the total number and the type of processors available in the server ordered. Specific ordering rules for each server feature are implemented in the configurator tool. The processor feature only identifies the type of processor(s) used. The processor type has a certain amount of commercial processing workload capabilities associated with it. For iSeries Models with multiple processors, the processor feature has a minimum and a maximum commercial processing workload capability associated with it. The available CPW capacity depends on the number of activated processors for the models that have CUoD capability.

The new iSeries servers come in different Editions. The full details of the available editions for each model will be described later in this presentation. What follows here is a brief general description.

The **Standard Edition** features a wide variety of e-business and client server solutions based on integrated middleware such as DB2, WebSphere Application Server and IBM HTTP Server (Powered by Apache). The Standard Edition supports multiple operating environments, such as OS/400 and Linux using dynamic LPAR. Capacity Upgrade on Demand (temporary and permanent) is also supported on i825, i870 and i890 servers. The iSeries 800 and the i810 servers do not support CUoD.

The **Enterprise Edition** builds on the functions available for the Standard Edition, and features a range of e-business, client server and 5250 OLTP solutions. In addition, it includes maximum interactive capacity along and a broad range of enterprise class tools for optimizing and managing e-business, client server and 5250 CPW. It incorporates an Integrated xSeries Server with the i825, i870 and i890 servers. A standby processor is also activated for Linux workload on the i870 and i890 server. IBM education and services vouchers are also included with the Enterprise Edition to accelerate deployment of new workloads such as WebSphere, Linux and Windows Integration. Look for details of the i810 Enterprise Edition later in this presentation.

The iSeries 800 offers three editions tailored specifically for small businesses, with a growth path within the model so you can start small and grow as your needs change. It incorporates a uni-processor processor with either 300 CPW or 950 CPW (25 or 50 - 5250 CPW), up to 8 GB of memory and up to 4 TB of disk. iSeries 800 is ideal for growing small enterprises that need an affordable, easy-to-use server with the flexibility to grow quickly to new demands. The iSeries 800 enables simple management of core business applications and your Intel-based servers all from the same system. The iSeries 800 **Value Edition** is the iSeries' entry-level server for customers who run core business applications (5250 CPW) and who want the flexibility to add features as necessary. The iSeries 800 **Standard Edition** offers a more robust configuration for customers who want run their core business applications, manage their Windows server workloads from the iSeries, and begin the move to e-business through optionally adding WebSphere Application Server – Express for iSeries. The iSeries 800 **Advanced Edition**, with its 950 CPW and 50 5250-CPW, is ideal for the customer seeking a multi-platform operating system with room to expand as they take on the increasing demands of e-business.


We will also cover full details of the different new options for On Demand processing later in this presentation. iSeries has exceptional flexibility to adjust dynamically to your business priorities. With temporary Capacity Upgrade on Demand, you can turn processors on when you need them – turn them off when you don't. You pay only for what you activate. iSeries servers (i825, i870 and i890) now offer Capacity Upgrade on Demand with both permanent and temporary activations for , which are equipped with standby processors when shipped.

## Hardware Support with V5R2 (Feb 2003 Level)

- Model 150, 6xx, S10, S20, S30, S40 and SB1
  - ▶ Not supported on the OS/400 **release following V5R2**
- ATM (Asynchronous Transfer Mode) technology
  - ▶ Not supported on the OS/400 **release following V5R2**
- Hardware **not supported at V5R2**
  - ▶ Model 4xx/5xx
  - ▶ 13GB QIC tape drives, Diskette drives, ISDN adapters
  - ▶ Upgrades from 6xx, Sxx boxes into the new models.
- Hardware not supported by new models
  - ▶ Migration towers are not supported on 800, i810, i825, i870 and the i890
  - ▶ Older SPD attached PCI towers are not supported (can convert)
  - ▶ Internal QIC tape drives: 1.2GB, 2.5GB and 13GB
  - ▶ SPD IOPs/IOAs are not supported
- Withdrawn hardware list at: [www.ibm.com](http://www.ibm.com)

# Agenda

- Server considerations
- **Main Storage**
- Expansion and IO
- High Speed Loop
- Capacity Upgrade on Demand
- Upgrade scenarios
- LPAR
- Linux
- Integrated Windows Environment
- SAN and External devices
- Software considerations

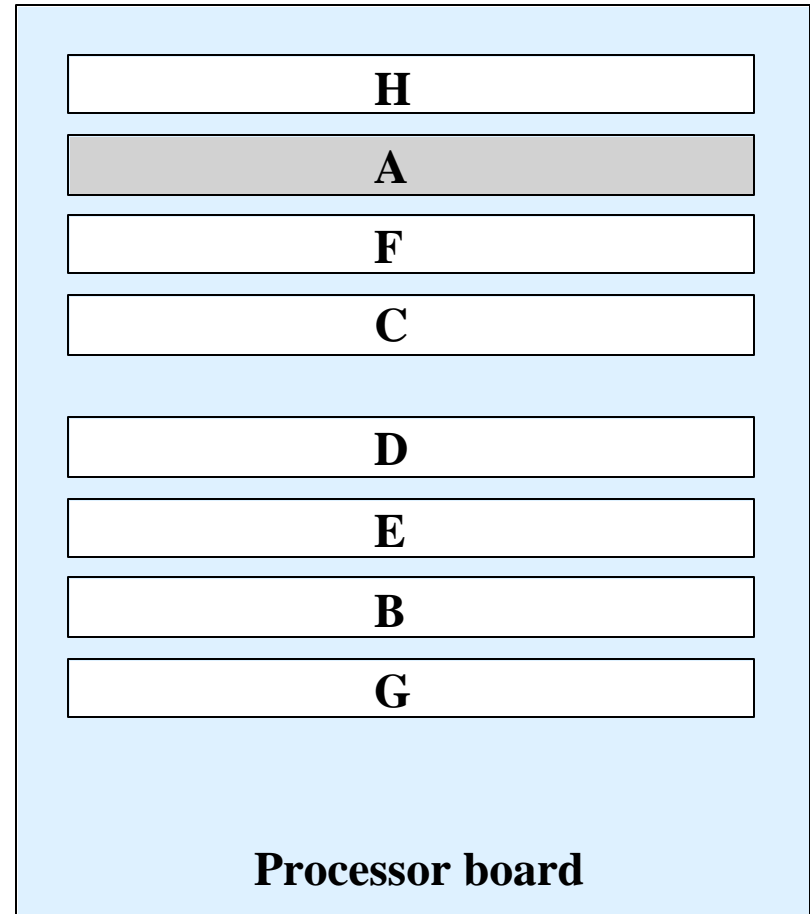


# Main Storage for iSeries 800 and i810 (#2466 and #2467)

- Memory on Processor board
  - ▶ Features, size and technology
    - For Server Feature 0863
      - Paired except for configurations with a single DIMM
      - Minimum 256 MB, maximum 8 GB
    - For Server Feature 0864 and 0865
      - Minimum 512 MB, maximum 8 GB
    - For Server Features 0866 and 0867
      - Minimum 512 MB, maximum 16 GB

Feature	Capacity (MB)	Technology
#3092	256	128 Mb DDR
#3093	512	256 Mb DDR
#3094	1024	256 Mb DDR*
#3096	2048	512 Mb DDR*

\* Stacked



# Notes: iSeries 800 and i810 (#2466 and #2467)

For iSeries Model 800 and i810 Server Features 0863,0864,0865, 0866 and 0867, the following main storage features are available

- #3092 - 256MB Main Storage
- #3093 - 512MB Main Storage
- #3094 - 1GB Main Storage
- #3096 - 2GB Main Storage

A single main storage DIMM feature is allowed on these servers as long as the minimum main storage requirements are met. When the total number of main storage DIMMs is increased greater than one, the single existing DIMM must be paired up (with a DIMM of the same capacity) and any additional DIMMs must be added in pairs of the same capacity. There are 8 memory DIMM slots available in the base system for main storage DIMMs which plug directly onto the backplane.

Allowable increments in MB for the iSeries Model 800 (server features 0863,0864 and 0865) main storage are :

256 (not allowed for server feature 0865) 512 1024 1536 2048 2560 3072 3584 4096 4608 5120 5632 6144 6656 7168 7680 and 8192

Allowable increments in MB for the i810 (server features 0866 and 0867) main storage are :

512 1024 1536 2048 2560 3072 3584 4096 4608 5120 5632 6144 6656 7168 7680 8192 8704 9216 9728 10240 10752 11264 12288 12880 13312 14336 and 16384

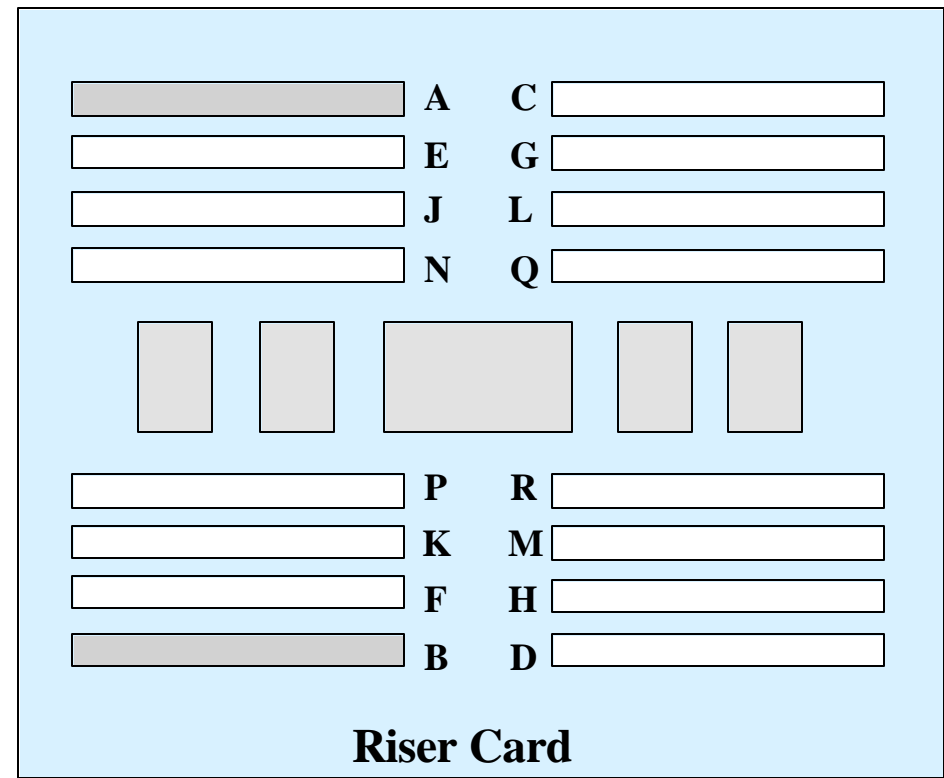
Stacked memory means that the memory DIMMs have memory chips mounted on top of each other. DDR stands for double data rate. These DIMMs are SDRAM DDR (Double Data Rate) technology cards. SDRAMs are based on standard dynamic RAM chips, but have sophisticated features that make them considerably faster. First, SDRAM chips are fast enough to be synchronized with the CPU's clock, which eliminates wait states. Second, the SDRAM chip is divided into two cell blocks, and data is interleaved between the two so that while a bit in one block is being accessed, the bit in the other is being prepared for access. This allows SDRAM to burst the second and subsequent, contiguous characters at a rate of 0ns, compared to 60ns for the first character. SDRAM provides up to 1.1 GBps data transfer, and DDR SDRAM (Double Data Rate SDRAM) doubles the rate to 2.2 GBps. DDR is also known as DSDRAM (Double-Speed DRAM), DDR DRAM (Double Data Rate DRAM) and SDRAM-II



# Main Storage for i810 (#2469)

- Memory on Riser Card
  - ▶ Quads of the same feature
    - Except for configurations with a single pair of DIMMs
  - ▶ Features, size and technology
    - Minimum 512 MB, maximum 16 GB

Feature	Capacity (MB)	Technology
#3022	128	64 Mb
#3024	256	128 Mb
#3025	512	128 Mb
#3026	512	256 Mb
#3027	1024	256 Mb
#3029	128	128 Mb



# Notes: Main Storage for i810 (#2469)

The Main Storage for Server Feature 0869 consists of the following features:

- #3022 - 128MB Main Storage (64Mb technology) is only orderable up to the minimum number of DIMMs required to meet a pairing or quad system memory requirement.
- #3024 - 256MB Main Storage (128Mb technology unstacked)
- #3025 - 512MB Main Storage (128Mb technology stacked)
  - #3025 can not be mixed with #3026 for pairing or quading.
- #3026 - 512MB Main Storage (256Mb technology unstacked)
  - #3026 can not be mixed with #3025 for pairing or quading.
- #3027 - 1GB Main Storage (256Mb technology stacked)
- #3029 - 128MB Main Storage (128Mb technology unstacked)
  - #3029 can not be mixed with #3022 for pairing or quading.

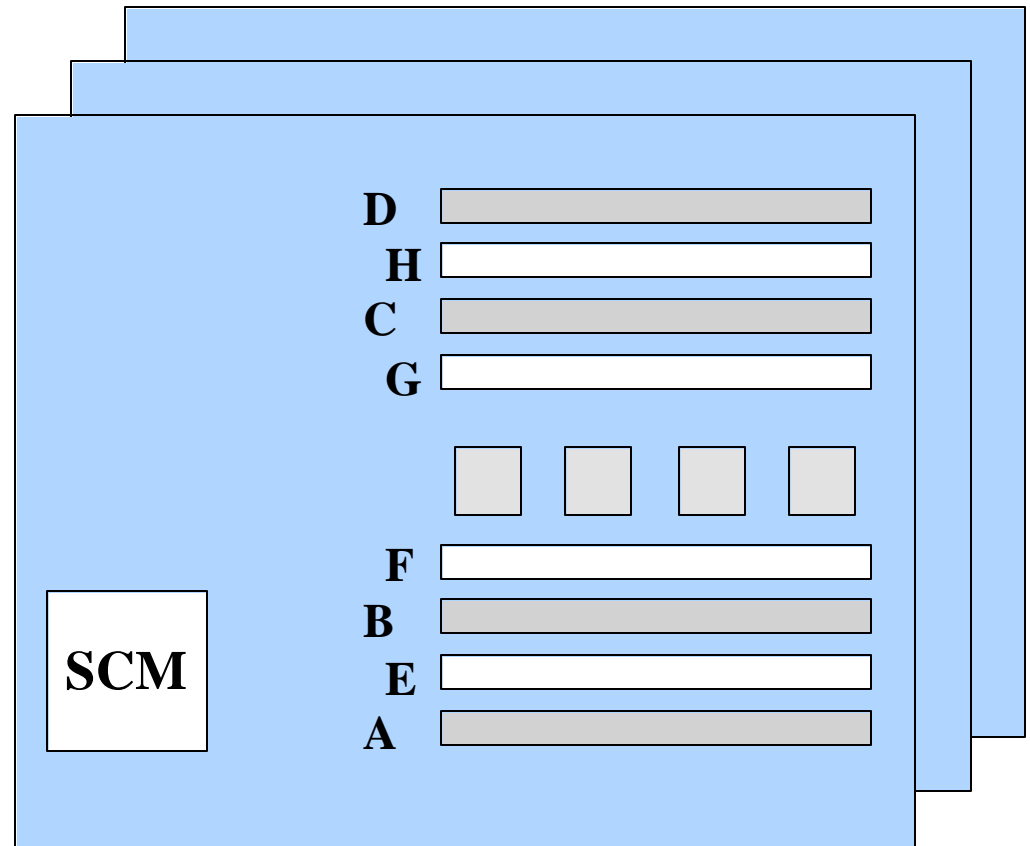
Maximum of 8 of these DIMM features allowed on a system. A minimum of two DIMMs (same feature code) must be selected. If more than two DIMMs are required, all memory features MUST be in sets of 4 (quads)

# Main Storage for i825

- Memory on Processor Cards
  - ▶ Each processor card has 8 memory slots
  - ▶ Minimum 2 GB, maximum 48 GB
    - 2 GB not recommended
  - ▶ Placement rules
  - ▶ Features, size and technology

Feature	Capacity (MB)	Technology
#3042	256	128 Mb DDR
#3043	512	256 Mb DDR
#3044	1024	256 Mb DDR*
#3046	2048	512 Mb DDR*

\* Stacked



## Notes: Main Storage for i825

For the i825, the Main Storage DIMMs are installed directly onto the processor cards. Each processor card has eight DIMM slots. These eight slots are arranged in two sets of four. The DIMMs must be installed in sets of 4 (quads). Each set of 4 DIMMs must be the same memory capacity and technology. Each system order must have at least one set of 3 main storage DIMM quads (12 DIMMs total) installed (exception is 2GB memory capacity point, where 2 quads (8 DIMMs total) are allowed).

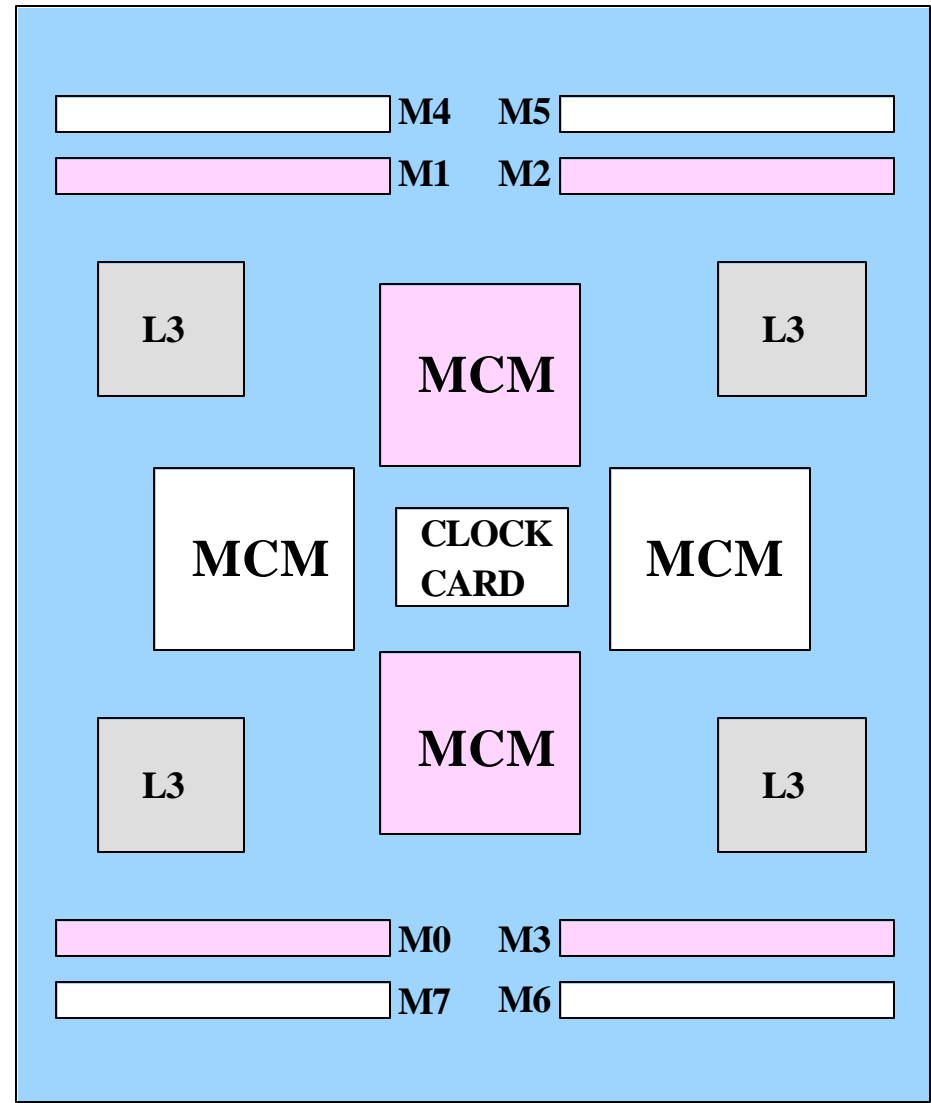
The i825 has three processor cards with a single chip module (SCM), each SCM has two processors. Within a given Main Storage configuration, a processor card can't contain more than the double of main storage than any other processor card in the server. As an example, a certain processor card can have four 256 MB DIMMs installed being 1024 MB in total, in this case the other processor cards can have a maximum of 2048 MB installed, either by quads of 512 MB or octals of 256 MB DIMMs.

The memory is a customer installable feature. Advise to carefully to read the instructions and documentation before installing main storage features since rules and restrictions apply.

# Main Storage for i870 and i890

- Memory cards on Processor board
  - ▶ i870 uses M0 through M3
    - Equal quads, except for 8GB minimum
      - Minimum 8 GB, maximum 128 GB
  - ▶ i890 16 to 24-way uses M0 through M5
    - Placement rules apply
      - Minimum 8 GB, maximum 192 GB
  - ▶ i890 24 to 32-way uses M0 through M7
    - Placement rules apply
      - Minimum 16 GB, maximum 256 GB

Feature M0-M3 / M4-M7	GB	Technology
#3020 / #3021	4	PDIMM / DDR
#3015 / #3016	8	PDIMM / DDR
#3035 / #3036	16	PDIMM / DDR
#3017 / #3018	32	PDIMM / DDR



# Notes: Main Storage for i870 and i890

The i870 uses only the inner memory slots, (colored slots M0 thru M3) due to the processor to memory connectivity wiring. The following memory features are available on the i870

- #3020 - 4096MB Main Storage Card (Inside) : Plugs into memory card slots 0, 1, 2 and 3
- #3015 - 8192MB Main Storage Card (Inside) : Plugs into memory card slots 0, 1, 2 and 3
- #3035 - 16384MB Main Storage Card (Inside) : Plugs into memory card slots 0, 1, 2 and 3
- #3017 - 32768MB Main Storage Card (Inside): Plugs into memory card slots 0, 1, 2 and 3

For plugging the above main storage cards in the i870 backplane, the following rules apply:

- Main Storage Cards must be installed in pairs of equal capacity
- Pairs of Main Storage Cards must plug into Memory Card slots under the same MCM (slots 0 and 1 or slots 2 and 3)
- Mixed pairs of Main Storage Cards must be of the adjacent capacity (Example : 4GB and 8GB is OK, 4GB and 16GB is not OK)

The i890 with processor feature #2497 (server feasture 0897 16 to 24-way) is using memory slots M0 thru M5. The i890 with processor feature #2498 (server feasture 0898 24 to 32-way) is using all available memory slots (M0 thru M7). The following memory features are available on the new i890:

- #3020 - 4096MB Main Storage Card (Inside) : Plugs into memory card slots 0, 1, 2 and 3
- #3021 - 4096MB Main Storage Card (Outside) : Plugs into memory card slots 4, 5, 6 and 7
- #3015 - 8192MB Main Storage Card (Inside) : Plugs into memory card slots 0, 1, 2 and 3
- #3016 - 8192MB Main Storage Card (Outside) : Plugs into memory card slots 4, 5, 6 and 7
- #3035 - 16384MB Main Storage Card (Inside) : Plugs into memory card slots 0, 1, 2 and 3
- #3036 - 16384MB Main Storage Card (Outside) : Plugs into memory card slots 4, 5, 6 and 7
- #3017 - 32768MB Main Storage Card (Inside) : Plugs into memory card slots 0, 1, 2 and 3
- #3018 - 32768MB Main Storage Card (Outside) : Plugs into memory card slots 4, 5, 6 and 7

For plugging the above main storage cards in the i890 backplane, the following rules apply:


- Main Storage Cards in installed in pairs of equal capacity
- Pairs of Main Storage Cards must plug into Memory Card slots under the same MCM (slots 0 and 1 must be pairs, slots 2 and 3 must be pairs, slots 4 and 5 must be pairs and slots 6 and 7 must be pairs)
- Different pairs of Main Storage Cards must be of the adjacent capacity (4GB and 8GB is OK, 4GB and 16GB is not OK)
- No mix Main Storage Cards of more than 2 capacities (4GB and 8GB is OK, 4GB and 8GB and 16GB is not OK) is allowed

# Agenda

- Server considerations
- Main Storage
- **Expansion and IO**
  - High Speed Loop
  - Capacity Upgrade on Demand
  - Upgrade scenarios
  - LPAR
  - Linux
  - Integrated Windows Environment
  - SAN and External devices
  - Software considerations

## PCI-X Towers and Expansions

- I/O Tower feature 5094
- I/O Tower feature 5294
- I/O Tower feature 0595 / 5095
  - ▶ 0595 is rack mounted
- I/O Tower feature 0588 / 5088
  - ▶ 0588 is rack mounted
- Base I/O Tower feature 9094
  - ▶ Base I/O tower for i870 / i890
- Optional Base I/O Tower feature 8094
  - ▶ Optional Base I/O tower for i870 / i890





# Notes: PCI-X Towers and Expansions

The #5094 has 15 disk unit slots, with an additional 30 slots available with feature #5108 (if #5115 not present) or #5118 (if #5115 is present). The #5094 also has 2 removable media slots and 14 PCI-X IOP/IOA card slots.

A #9844 Base PCI IOP is included as "base".

The #5294 is equivalent to two #5094 PCI Expansion Towers, with covers and casters removed and positioned in 1.8 meter rack. It includes two bus adapters to provide the HSL-2 interface to the system.

The #5095/#0595 (rack mountable - 5 EIA units) has 7 PCI-X IOP/IOA slots and supports up to 12 disk units. A #9844 Base PCI IOP is included as "base". It includes a bus adapter to provide the HSL-2 interface to the system.

The #5088/#0588 (rack mountable - 8 EIA units) has 14 PCI-X slots for installation of PCI IOPs and IOAs. It includes a bus adapter to provide the HSL interface to the system.

Bus adapters to provide the HSL interface to the system:

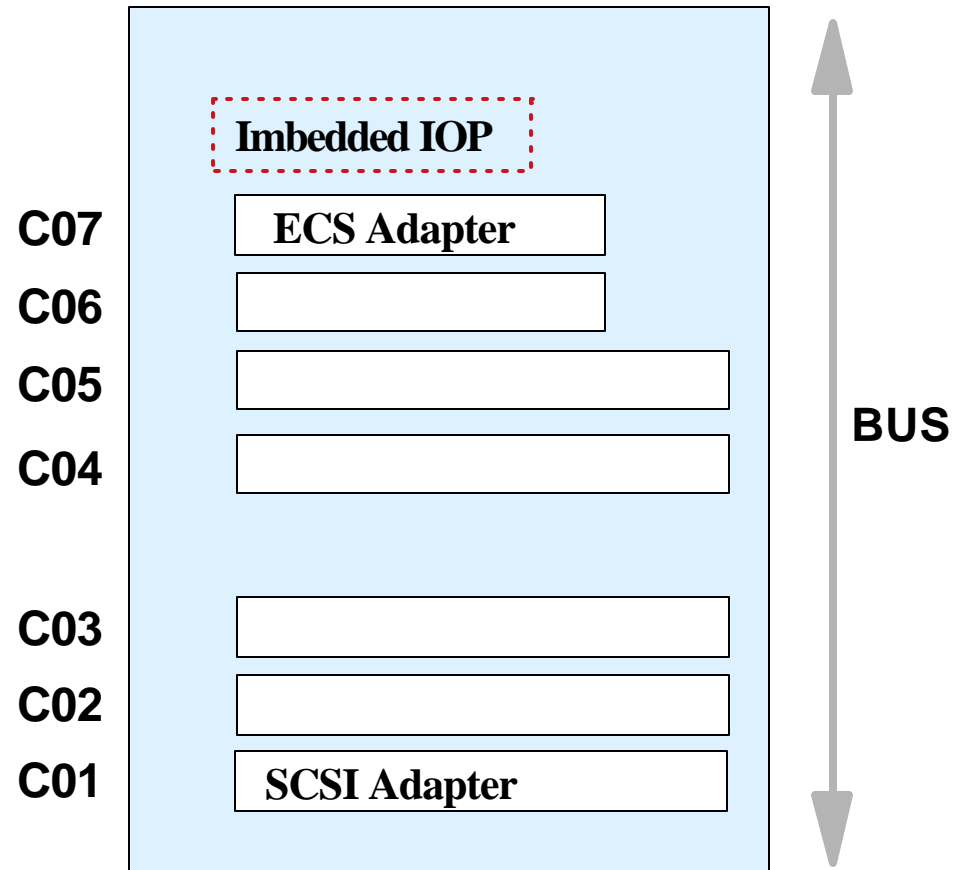
- #9886 Base Optical Bus Adapter is used when attaching to optical HSL ports but is not supported when attaching to models 810 and 820 (these models do not have optical HSL ports)
- #9887 Base HSL-2 Bus Adapter is used when attaching to copper HSL or HSL-2 ports

# PCI and PCI-X Tower, IOA and IOP Placement (V5R2)

- PCI-X Towers / Expansions
  - ▶ #0595 / #5095 supported on all iSeries
  - ▶ #5094 supported on all iSeries except Model 270
  - ▶ #5294 supported on all iSeries except Model 270 and iSeries 800
  - ▶ #0588 / #5088 supported on all iSeries except Model 270 and iSeries 800
  
- PCI Towers / Expansions
  - ▶ #5075 supported on Models 270 and 820, on i810 and i825
  - ▶ #5074 / #5079 supported on all iSeries except Model 270 and iSeries 800
  - ▶ #0578 / #5078 supported on all iSeries except Model 270 and iSeries 800

# Base PCI-X slots iSeries 800 and i810

- iSeries 800 and i810 base PCI slots
  - ▶ Imbedded IOP
  - ▶ Single system bus backplane
  - ▶ One Integrated xSeries Server allowed in slot C04
  - ▶ Two short PCI-X slots

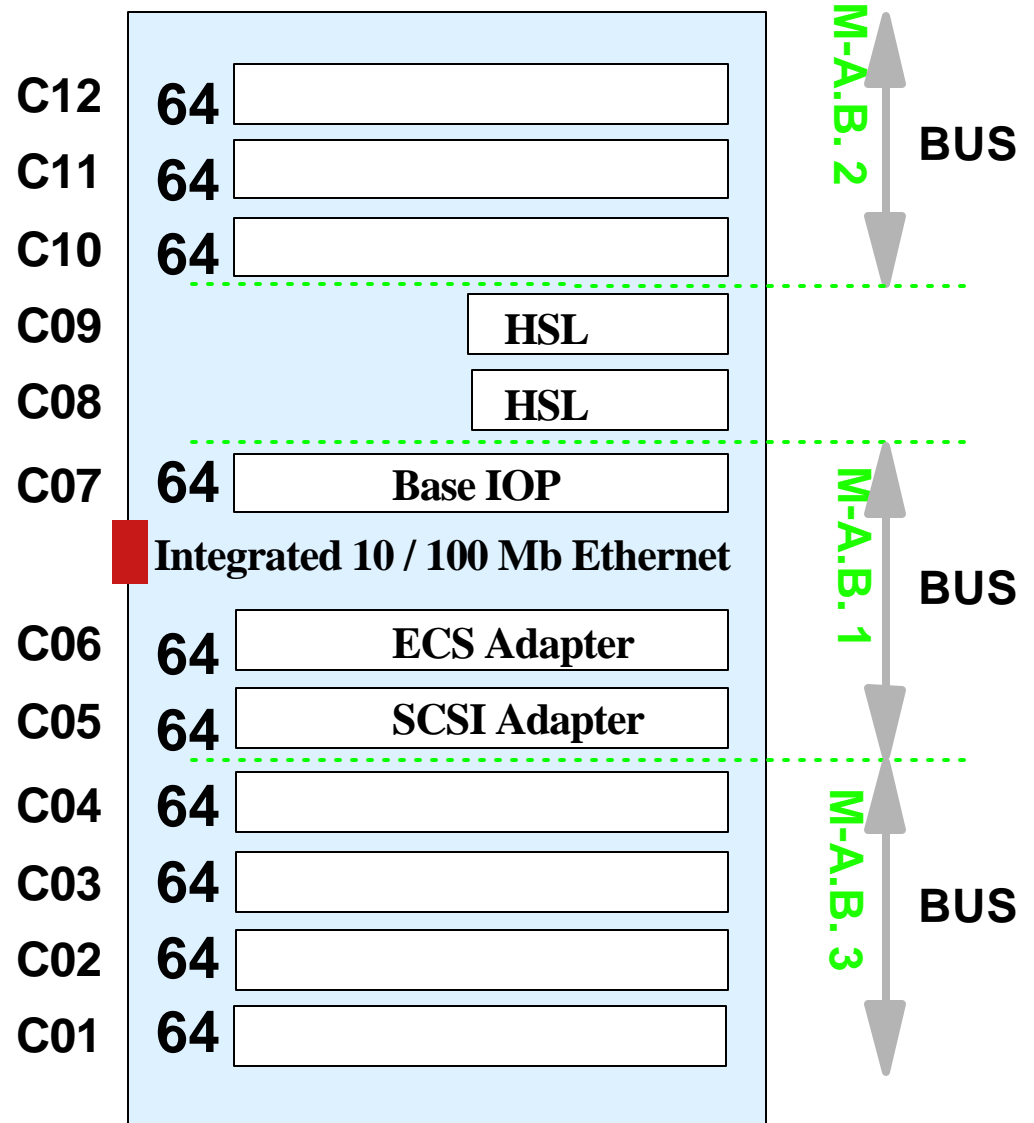


# Notes: Base PCI-X slots iSeries 800 and i810

iSeries 800 and i810 server CEC have seven PCI-X slots on a backplane with imbedded IOP, two slots are restricted for short PCI or PCI-X cards. The backplane is a single system bus unit.

# PCI-X slots i825

- i825 base PCI-X slots
  - ▶ Includes integrated 10/100 Mb Ethernet Adapter
  - ▶ Three system busses
  - ▶ One Integrated xSeries Server allowed in slots C04/C03
  - ▶ Support for "long" PCI cards in all positions




## Notes: PCI-X slots i825

The active backplane of the i825 contains three Multi Adapter Bridges. The first two Multi Adapter Bridges providing three PCI-X slots each and the third Multi Adapter Bridge is containing four PCI-X slots for a total of ten PCI-X slots in the i825 CEC. The i825 has an imbedded 10/100 Mb Ethernet adapter on the active backplane, this is the reason why there is no support for a LAN console over Token Ring. (no specify #5546 support for the i825)

# PCI-X Slots in 5094, 0588, 5088, 5294, 8094, 9094

- PCI-X cage contains 3 busses
  - ▶ Older PCI Tower cages contain 2 system busses
  
- All slots are 3 Volt
  - ▶ Older PCI Towers have four 5 Volt slots
  
- HSL adapter in slot C10
  - ▶ Copper # 9887 (#2887)
  - ▶ Optical # 9886 (#2886)
  
- Towers attach to iSeries Models 820,830 and 840.
  - ▶ Use HSL to HSL-2 cables



# Notes: PCI-X Slots in 5094, 0588, 5088, 5294, 8094, 9094

It is easy to differentiate a PCI from a PCI-X tower. Just by looking at the HSL adapter position you will know.

- The HSL adapter for a PCI-X tower is in C10
- The HSL adapter for a PCI tower is in C08

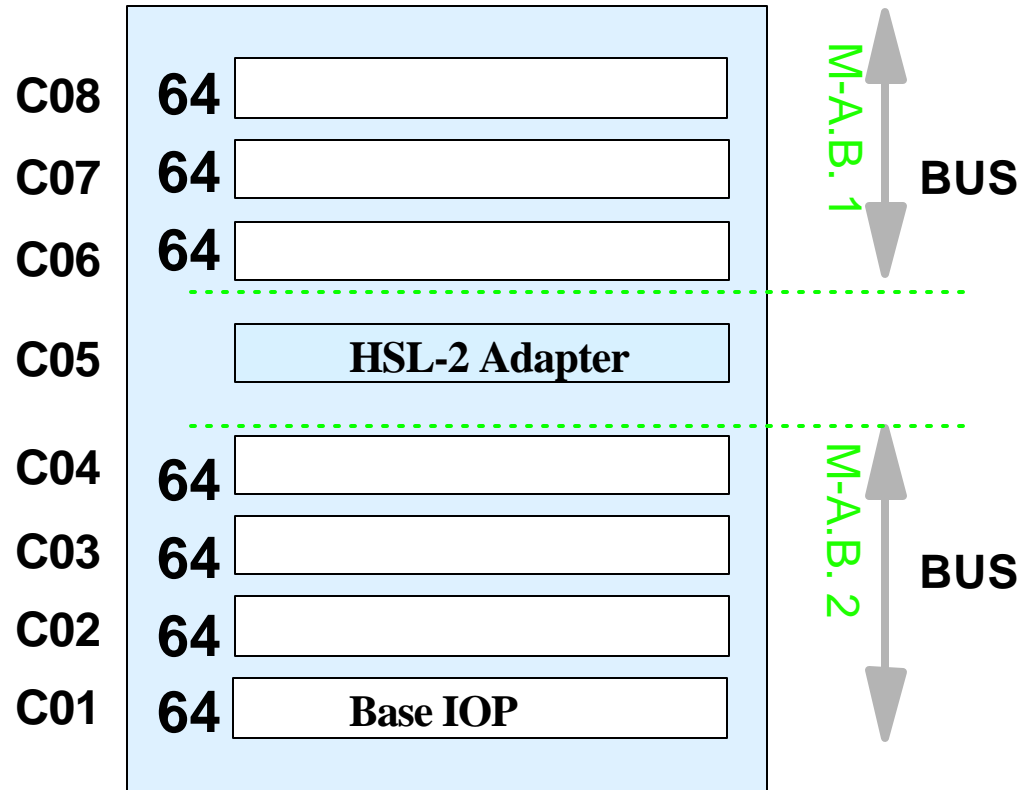
The PCI-X towers have a PCI card cage that can contain an optical or a copper HSL adapter. The initial order feature code for a HSL adapter is a #9887 (copper) or a #9886 (optical). When a configurator change is needed after install and you want to switch from copper to optical HSL or the other way around, you can order the #2887 (copper) or #2886 (optical) as required.

The new PCI-X towers can also attach to the iSeries model 820, 830 and 840 using the HSL to HSL-2 cables.



# PCI-X Slots and DASD options for Towers 5095 and 0595

- PCI-X cage specifics
  - ▶ Seven available 64-bit PCI-X slots
  - ▶ Two system busses
  - ▶ No card length restrictions
- Disk Unit Slots
  - ▶ 12 Ultra SCSI slots available
  - ▶ Up to two DASD IOAs
- Includes Base IOP
  - ▶ Input Output Processor #9844
- Redundant Power
  - ▶ Available with #5138 Power Supply



## Notes: 5095 and 0595 PCI-X Tower

The 5095 and 0595 PCI-X Tower contains two system busses. One with three the other with four PCI-X slots. The two 6-disk unit cages can be driven by separate disk adapters if required (for LPAR you can have two IOPs and each IOP a disk controller with a load source disk unit for a partition). The #5095 and the #0595 contains a base IOP (#9844) in slot C01. The initial order feature code for a HSL adapter is a #9887 (copper) or a #9886 (optical). When a configurator change is needed after install and you want to switch from copper to optical HSL or the other way around, you can order the #2887 (copper) or #2886 (optical) as required.

A 5095 is the standalone PCI-X tower, while a 0595 is the same unit, but mounted in a #0551 rack.

# Disk Unit Expansions

- 30 Disk Unit Expansion #5108
  - ▶ Orderable for Expansion Towers #5094 and #9094
  - ▶ Expansion Towers #5294 and #8094 contain two #5108
- 5 Disk Unit Expansion Feature 7124
  - ▶ For i825 system unit
- Optional Base 6 Disk Unit Expansion # 7137 with concurrent maintenance enabled for iSeries Model 800 processor feature 2463
- Expansion feature 7116 for iSeries 800 and i810
  - ▶ Base expansion supports 6 disk units and 6 optional with #7136 DASD expansion unit
- All New Disk Unit Cage slots support 7.2K, 10K and 15K RPM drives
  - ▶ 7.2K rpm : #4308 (4GB), #4314(8.5GB) and #4324(17.5GB) only

# Notes: Disk Unit Expansions

The new disk unit expansions all support the new announced 15K rpm drives. Important for upgrade and migration scenarios is the fact that all these expansions as well as the base disk slots of the new servers also support the drives that fit in any PCI tower

# Rack Offering Examples



# Notes: Rack Offering Examples

The pictures on this foil show a couple of possible combinations for mounting iSeries hardware in a rack. Rack placement can be useful for saving floorspace. During the planning consider the weight of a populated rack.

From left to right:

- i825 with two #0588 in a rack
- #5294
- two i825s in a rack
- iSeries Model 800 and i810 in a rack

The standard rack #0551 feature, when ordered by itself, provides an empty 1.8 meter rack which contains 36 EIA units of total space. The following specifies, when ordered in conjunction with (or for) a #0551, indicate the means of populating a #0551 rack.

- #0121 - 270 Lower Unit Specify (plant install)
- #0122 - 270 Upper Unit Specify (plant install)
- #0123 - #5074 Lower Unit in Rack (plant install)
- #0125 - #9079 Lower Unit in Rack (plant install)
- #0127 - 270 Field Install in Rack
- #0133 - iSeries 800, i810 Field Install in Rack
- #0134 - i825 Field Install in rack
- #0137 - 270 Field Install in Rack
- #0578 - PCI Expansion Unit in Rack
- #0588 - PCI Expansion Unit in Rack
- #0595 - PCI Tower Unit in Rack

1 to 4 Power Distribution Units may be specified with the #0551. The PDUs may be ordered on initial orders, model upgrades or on simple MES orders. Each PDU has six power sockets. The following PDUs are available:

- #5160 Power Distribution Unit 1 Phase NEMA
- #5161 Power Distribution Unit 1 Phase IEC
- #5162 Power Distribution Unit 2 of 3 Phase

Each PDU type has specific power cords associated with it. Use the iSeries System Builder for detailed information. See also power considerations in this presentation for guidelines on PDU ordering.

# Power Considerations

- Voltage
  - ▶ iSeries Model 800, i810 and expansions #7116, #0595 and #5095 operate at 110-127 Volt or 220-240 Volt
  - ▶ i825, and expansions #0588, #5088, #5094, #5294, #8094, and #9094 operate at 220-240 Volt only
  - ▶ i870 and i890 CEC require 200-480 Volt three phase
- Dual Line Cords
  - ▶ Standard for i825, i870 and i890 CEC
  - ▶ Standard or mandatory for expansion #0588, #5088 and #9094
  - ▶ Optional dual line cord enablement features for
    - Expansion tower 5094: #5115
    - Expansion tower 5294: #5116
    - Expansion units 0595 and 5095 require redundant power and cooling feature 5138 for dual line cord enablement
  - ▶ Special considerations for using dual line cords in a #0551 rack (PDUs !)
- Consider using a Uninterruptable Power Supply

# Notes: Power Considerations

The iSeries Model 800, i810 and expansions #7116, #0595 and #5095 are equipped with an autotransformer and automatically adapt to 110-127 Volt or 220-240 Volt input voltages. For the i825, and expansions #0588, #5088, #5094, #5294, #8094, and #9094 input voltages in the 220-240 Volt range are required.

The i870 and i890 both function normally with a nominal input voltage in the range of 200-480 V, AC, three phase. Order the power cord feature that applies to the power source used.

Rated current for the i870 and the i890 per phase:

- For 200-240 Vac : 45 Amps
- For 380-415 Vac : 25 Amps
- For 480 Vac : 20 Amps

Dual Line cords can be used for protection against a single phase failures. Some iSeries hardware has dual line cords by default, other hardware requires the dual line cord feature to be ordered. When using dual line cord features to protect the configuration against a certain type of failure, make sure that all components have this feature the order and the power distribution is wired according to your needs. Marking the receptacles in a self-explainable way can be very useful.

When using a #0551 standard rack with a configuration using dual line cords as a protective vehicle, make sure that the rack itself is wired accordingly. Make sure that (depending of the number of units to be installed in the rack ) 2 or 4 Power Distribution Units are ordered. The Power Distribution Units can then be connected to the power receptacles just as for other expansions or units with the dual line cord feature. The dual line cord units mounted in the rack can then be connected to the Power Distribution Units in such a way that their protection is likewise assured.

To protect your critical equipment against various types of power failures, it is strongly advised to use a UPS. Most common failures such as power loss can be caused by a number of events such as a lightning strike, over demands, accidents and natural disasters. Protection against other but also very common malfunctions such as power sags and surges, noise, frequency variations and harmonic distortion can be assured by using a UPS too. Selecting the UPS that is right for your environment is a critical part of the planning cycle.



# New PCI-X Processor and Adapters - 1

- Input Output Processor #2844
  - ▶ Replaces existing IOP #2843
  - ▶ Imbedded PowerPC Processor
  - ▶ 64 bit PCI bandwidth and increased processor bus speed
  - ▶ 4 Gb per second instantaneous data transfer
  - ▶ Faster internal memory (DDR)
  - ▶ Increased internal L2 cache
- Performance considerations
  - ▶ Select the #2844 for obtaining the best performance for:
    - New RAID5 controllers
    - 2 Gb connections over optical adapters
    - Other high bandwidth or high throughput adapters such as Gb Ethernet adapters


# Notes: New PCI-X Processor and Adapters - 1

The new PCI-X IOP #2844 / #9844 replaces the existing PCI IOP #2843 / #9943. The new IOP has a 64 bit Power PC processor and DDR memory on board and is working with its own level two cache. The new technology provides a much higher bandwidth on the PCI or PCI-X bus. The following IOAs are supported by the #2843/#2844 PCI IOP:

#2742 - PCI Two-Line WAN IOA	#2743/#2760 - PCI 1Gbps Ethernet IOA
#2744 - PCI 100/16/4Mbps Token Ring	#2749 - PCI Ultra Mag Media Ctrlr
#2757 - PCI RAID Disk Unit Ctrl	#2763 - PCI RAID Disk Unit Ctrl
#2765 - PCI Fibre Channel Tape Ctrlr	#2766 - PCI Fibre Channel Disk Ctrlr
#2768 - PCI Magnetic Media Ctrlr	#2772/#2773 PCI Dual WAN/Modem
#2782 - PCI RAID Disk Unit Ctrlr	#2793/#9793 - PCI 2-Line WAN w/Modem
#2794/#9794 - PCI 2-Line WAN w/Modem	#2805 - PCI Quad Modem IOA
#2806 - PCI Quad Modem IOA (CIM)	288E CCIN - Imbedded 100/10 Ethernet IOA (model 825 only)
#2817 - PCI 155MBPS MMF ATM	#2849 - PCI 100/10Mbps Ethernet IOA
#4723 - PCI 10Mbps Ethernet IOA	#4745 - PCI WAN IOA
#4746 - PCI Twinaxial IOA	#4748/#9748 - PCI RAID Disk Unit Ctrlr
#4778/#9778 - PCI RAID Disk Unit Ctrlr	#4801 - PCI Crypto Coprocessor
#4805 - PCI Crypto Accelerator	#4815 - PCI ATM 155MBPS UTP OC3
#4816 - PCI ATM 155MBPS MMF	#4818 - PCI ATM 155MBPS SMF OC3
#4838 - PCI 100/10Mbps Ethernet IOA	#5700 - PCI 1Gbps Ethernet IOA
#5701 - PCI 1Gbps Ethernet UTP IOA	#5702 - PCI Ultra Controller
#9767 - Base PCI Disk Unit Ctrlr	#9771 - Base PCI 2-Line WAN w/Modem
#9793/#9794 - Base PCI 2-Line WAN w/Modem	

# Ultra RAID Disk Controller

- Ultra RAID Disk Controller
  - ▶ #2757
    - 235 MB write cache
      - Uncompressed: compression ratio depends on data structure
      - Compressed: up to 757 MB on average
    - New RAID-5 implementation and options for optimization
    - Four Ultra SCSI buses
    - Support for up to 18 disk units and two internal optical / tape devices



# Notes: Ultra RAID Disk Controller

The DASD controller plays an important role in determining the I/O busy time for a given process. Each time an I/O request has to be staged to online storage, the DASD controller is one of the prime vehicles to provide the components and the code to provide persistence of the data between main storage and online storage. The key elements that reduce the time spent for this operation are:

**Bandwidth** : the new 2757 controller has a up to 8 times larger bandwidth on the PCI bus compared to 2778/4778 controller, meaning that transfer of large blocks or transfer of multiple small blocks will improve its speed considerably.

**Processing capacity** : the use of a Power PC processor, running at 500 MHz and 256K of L2 cache allows faster execution (6.25 times) of all storage transfer commands, including the parity checking and data scrubbing.


**Cache memory** : the performance of I/O commands can greatly benefit from staging data into temporary RAM, especially if the memory has enough capacity to store large segments. Using data compression, the store write buffers in cache memory can support up to 757 MB. In comparison, the 2778/4778 controller supported up to 104 MB of compressed write cache.

**Number of SCSI buses** : the 2757 supports 4 SCSI buses, which allows to better distribute the disk units over the controller, thus optimizing the data flow.

**Renewed RAID5 Technology** : See next notes page for details on new rules and RAID5 technology details later in this presentation

## Ultra RAID Disk Controller -2

- Ultra RAID Disk Controller
  - ▶ #2782
    - 40 MB write cache
      - Uncompressed; compression ratio depends on data structure
    - New RAID-5 rules
    - Two Ultra SCSI buses
    - Support for up to 12 disk units and two internal optical / tape devices



# Notes: Ultra RAID Disk Controller -2

## **RAID 5 Rules** for the new disk unit adapters #2757 and #2782

A minimum of three disk units of the same capacity are needed for a valid RAID-5 configuration. A maximum of six arrays are allowed per controller, with a maximum of 18 disk units allowed per array. All disk units in an array must be of the same capacity. Parity is spread across either 2, 4, 8 or 16 disk units in an array. If an array of 3 disk units is started, parity is spread across 2 disk units. If an array of 4-7 disk units is started, parity is spread across 4 disk units. If an array of 8-15 disk units is started, parity is spread across 8 disk units. If an array of 16-18 disk units is started, parity is spread across 16 disk units. The number of arrays and size of each array can be influenced by specifying an optimization of either "Balance", "Performance" or "Capacity" in Operations Navigator when starting arrays. An optimization of "Balance" will be used by default when starting arrays from the green screens. If disk units are included into an existing array, parity may be spread across less than the preferred number of disk units. In this case, the RAID function must be stopped and then started in order to redistribute the parity.

## **RAID-5 Options** for the new disk unit adapters #2757 and #2782

Starting RAID-5 from a green screen will result in the creation of a "balanced" type of parity spreading across the disk units.

Optimization options selectable from iSeries Navigator when starting RAID-5 on disk units under a #2757 or #2782:

**Capacity** : A parity set optimized for capacity stores the most data possible. The I/O adapter may generate fewer parity sets with more disk units in each parity set. For example, if an I/O adapter has 15 disk units and is optimized for capacity, the result could be one parity set containing 15 disk units.

**Balanced** : A balanced parity set compromises between the ability to store large amounts of data and also provide fast access to data. For example, if an I/O adapter has 15 disk units and you choose balanced parity optimization, the result could be two parity sets, one with nine disk units and one with six disk units.

**Performance** : Parity sets optimized for performance provide the fastest data access. The I/O adapter may generate more parity sets with fewer numbers of disk units. For example, if an I/O adapter had 15 disk units and is optimized for performance, the result could be three parity sets with five disk units each.

# Ultra SCSI Disk Controller / Tape Adapters

- Ultra SCSI Disk Controller / Tape Adapters
  - ▶ Ultra SCSI Tape Controller #5702
    - Two external Ultra SCSI ports
    - Attachment for selected external tape and optical devices
    - Support for LVD SCSI LTO Devices (358x-Lxx)
  - ▶ Ultra SCSI Disk Controller / Tape Adapter #5705
    - No write cache
    - Internal Ultra SCSI bus with support for up to 6 Disk Units (no RAID support), for the required DVD and optional removable media feature
    - External Ultra SCSI bus with support for selected Magnetic Media and Optical devices
    - Support for LVD SCSI LTO Devices (358x-Lxx)



# Notes: Ultra SCSI Disk Controller / Tape Adapters

The #5705 adapter has two Ultra SCSI ports, the first port is an internal port supporting up to 6 internal disk units and up to two internal removable media devices, the second port is external, please find supported devices and combinations below.

The #5702 external SCSI adapter has two external Ultra SCSI ports.

Devices supported on the external SCSI ports of the #5705 and the two ports of the #5702

PHYSICAL INTERFACE for both #5702 and external port of #5705 : VHDCI --- compare with #2749 : HD68

LTO drives with LVD SCSI interface, more in general the 358X-Lxx tape devices and tape libraries.

Model 7207-122 QIC-SLR Tape Bridge Box

Model 7208-345 60GB 8mm Tape Drive

Model 7329-308 1/4-inch Tape Autoloader

Model 7210-020 CD-ROM Bridge Box

Model 7210-025 DVD-RAM Drive

Possible combinations on a single port:

One Model 7207-122 QIC-SLR Tape Bridge Box

One 7208-345 60GB 8mm Tape Drive

One Model 7329-308 1/4-inch Tape Autoloader

One Model 7210-020 CD-ROM Bridge Box

One Model 7210-025 DVD-RAM Drive

One Model 7210-025 and one Model 7210-025 (the two devices are daisy-chained)

One Model 7210-025 and one Model 7210-020 (the two devices are daisy-chained with the Model 7210-025 physically connected first)

One Model 7207-122 and one Model 7210-020 (the two devices are daisy-chained with the Model 7207-122 physically connected first)

One Model 7207-122 and one Model 7210-025 (the two devices are daisy-chained with the Model 7207-122 physically connected first)

One Model 7208-345 and one Model 7210-020 (the two devices are daisy-chained with the Model 7208-345 physically connected first)

One Model 7208-345 and one Model 7210-025 (the two devices are daisy-chained with the Model 7208-345 physically connected first)



# RAID, what RAID ?

- Redundant Array of Independent Disks
  - ▶ RAID-0
    - Data striped accross all disks in the array
    - No error checking or redundancy
    - Not used on iSeries, used on other platforms
  - ▶ RAID-1
    - Also called disk mirroring within an array
    - On iSeries implemented with Storage Management functions on ASP level
    - Disk, IOP and bus level mirroring supported with OS/400
    - Full redundancy
  - ▶ RAID-1/0
    - Combination of RAID-0 and RAID-1,
  - ▶ RAID-3 and RAID-4
    - Parity data stored on a dedicated disk unit, slow and hardly used
  - ▶ RAID-5
    - Striping of data and parity across all or some disks in the array
    - Specific rules by adapter type, used on iSeries internal disk units with RAID-5 adapters

# Notes: RAID, what RAID ?

RAID-0 / iSeries does not use the data striping on disk unit arrays. iSeries storage management functions however are spreading the data over all available and configured drives in a auxiliary storage pool. No error checking or redundancy. The failure of one drive in the array results in the loss of all data stored in the array. When used with file based operating systems, where parts of the files not residing on a failed drive can be recovered , not all data is lost when a fatal drive failure occurs.

RAID-1 / On the iSeries, the implementation is called mirroring and is part of the operating system functions. Since iSeries storage management functions are also spreading the data over all available and configured drives in a auxiliary storage pool (actually over half of the available units because the mirrored image is on the others), some call this way of protection a semi-RAID-1/0 being (for the iSeries) a combination of distributing data over multiple units and mirror them at the same time. For the iSeries there are different levels of RAID-1 protection. There is disk level mirroring, IOP level mirroring and bus level mirroring. For description of these levels, look for descriptions on mirrored protection specify codes : #0040, #0042 and #0043.

RAID-3 / A dedicated drive keeps the parity information for the other data drives in the RAID set. These data drives make use of byte-level striping to spread large record information across the entire array. RAID-3 makes use of each sector's embedded ECC to detect and repair errors. The utilization of byte-level striping makes RAID-3 a suitable choice for single-user environments where large files are in use.

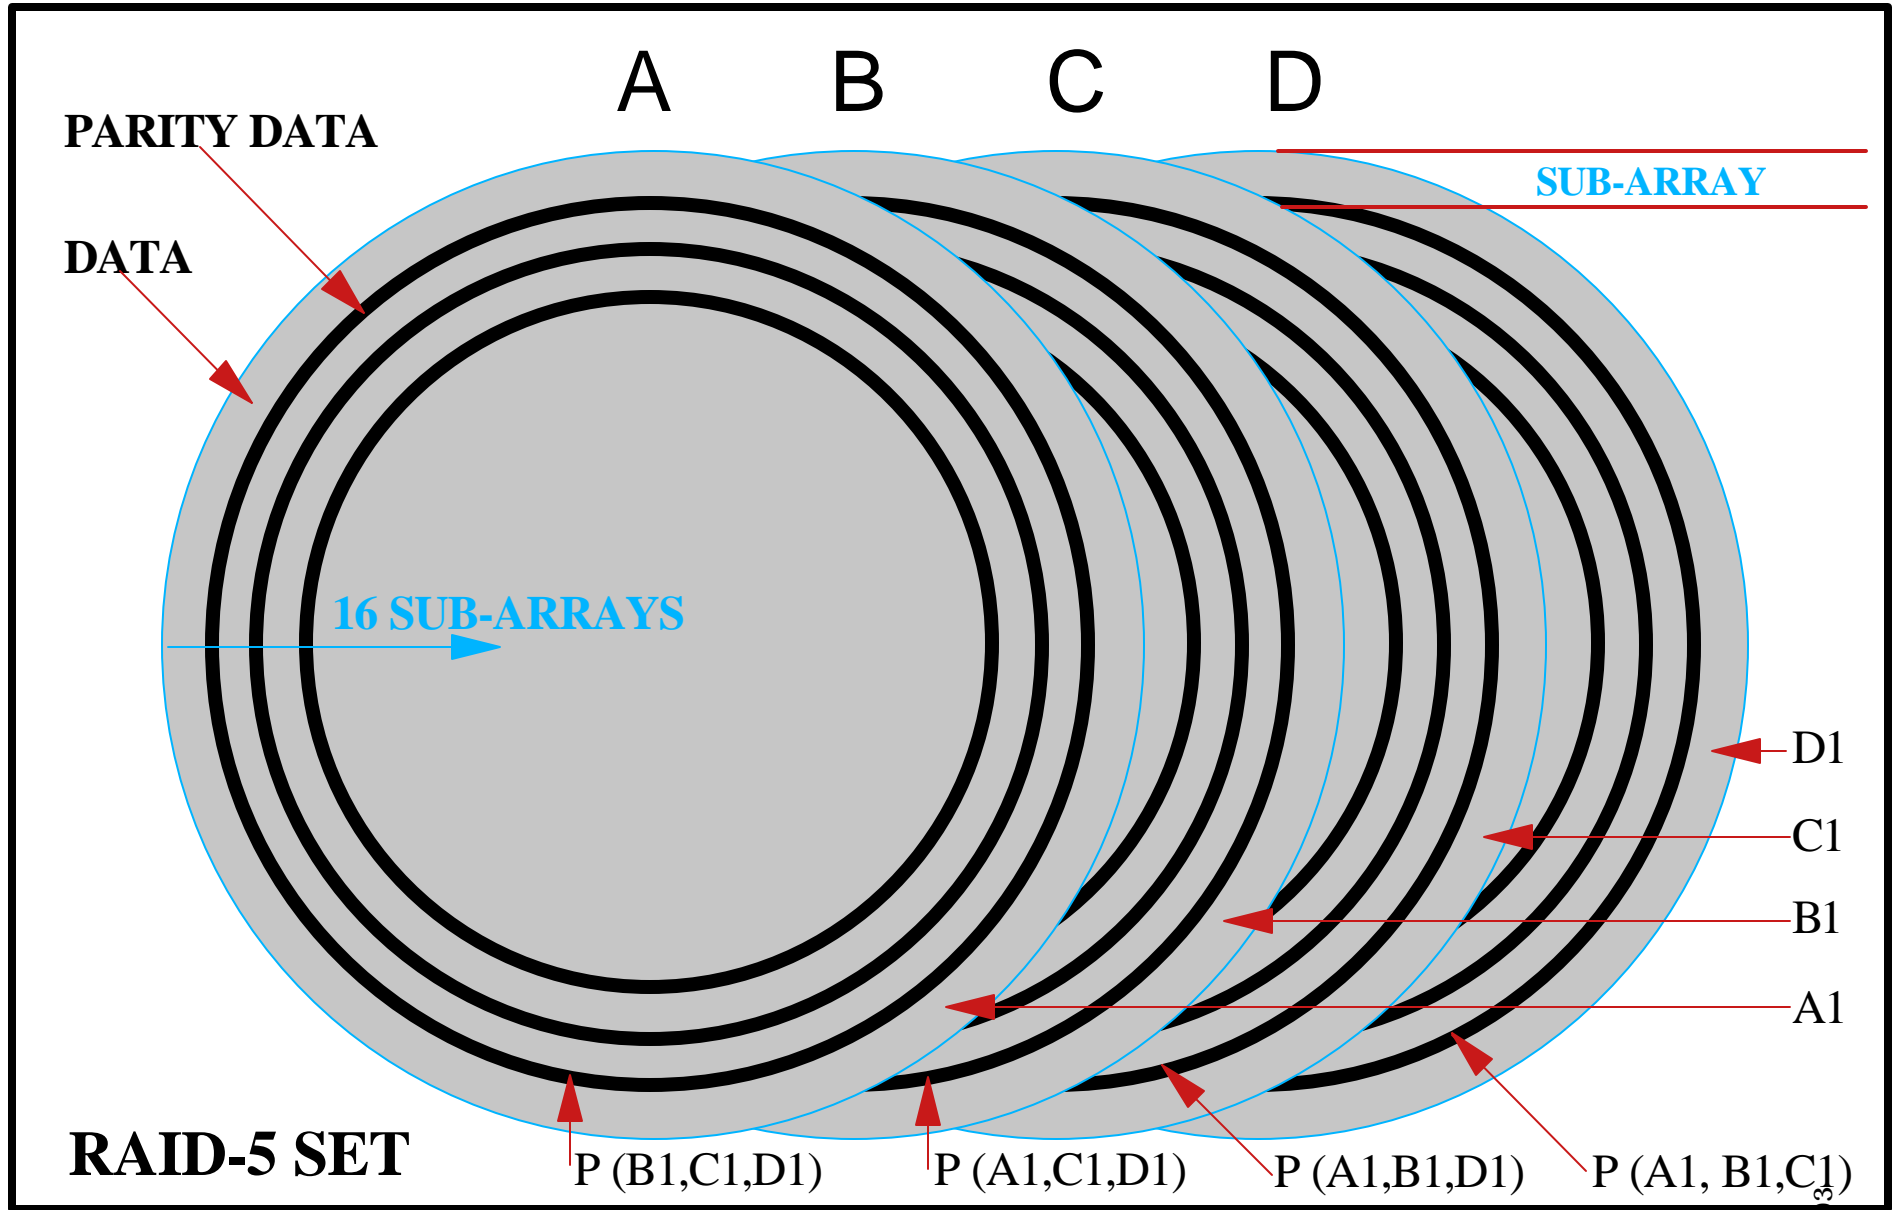
RAID-4 / Almost identical to RAID-3 protection, with a single dedicated parity drive and striped data drives. The difference is that unlike RAID-3, RAID-4 striping consists of block-level striping, allowing an entire record to be written to each stripe. This makes RAID-4 more suitable for systems with transactional environments where many smaller reads are needed at the same time.

RAID-5 / The optimal combination of redundancy, performance, and storage efficiency make RAID-5 one of the most widely favored data protection methods currently used. This controller implemented protection methodology spreads the parity locations across the entire disk set or a number of units in the set, installing a parity segment on each drive or a number of drives in the set . This arrangement is called distributed parity, and it greatly enhances the array's I/O performance, since the burden of parity write operations is shared throughout the array. RAID-5 is sometimes referred to as a Rotating Parity Array. While write efficiency is greatly improved over other RAID methods, parity calculations must still be performed for each write, and this still creates a slowdown. Using the new iSeries adapters with large, fast, intelligent AND protected write cache, physical write performance slow down can almost be ignored. Optimization of a RAID-5 array can be achieved by adjusting the stripe size to best deal with the environment's needs. Data protection is still optimal, with the full regenerative capabilities of parity and the fact that parity blocks are isolated from their associated data blocks. Cost effectiveness of RAID-5 is also good, comparable to other data protection methods. Hardware control is likewise a necessity. RAID-5 as implemented with the new iSeries adapters is also ideally suited to applications where large numbers of users are accessing information simultaneously.

# RAID-5 Redefined with the new RAID Adapters

- The new RAID-5 disk unit adapters use a improved striping mechanism
  - ▶ Parity data spread across drives
    - 16 subarrays
  - ▶ Short seeks
- Possible array disk unit members
  - ▶ Three unit array
    - Parity data across two units
  - ▶ Four to seven unit array
    - Parity data across four units
  - ▶ Eight to fifteen unit array
    - Parity data across eight units
  - ▶ Sixteen to eighteen unit array
    - Parity data across sixteen units

# Parity Data Striping Example with new RAID 5 Adapters



## Notes: Example with new RAID-5 Adapters

The previous foil is a graphical representation of the distribution of the data and the parity data on disk drives in a RAID-5 set driven by the latests iSeries RAID-5 disk controllers (#2757 or #2782). The example shows four drives(A,B,C and D) in a RAID-5 set, each drive represented by a single platter.


The RAID-5 striping of data and parity data now works with subarrays. There are 16 subarrays per drive. The size of such a subarray 6.25% of the disk unit capacity. The array includes data and parity data if the disk unit contains parity data. This existing of parity data in the subarray restricts the size of the subarray available for data with exactly the same percentage as if there were no subarrays whatsoever. Easy to understand that the total amount of disk space reserved for parity data stays the same as before. The difference with the new adapters is also the possible distribution of parity data over more or less units than before(see previous foils covering the new adapters).

Back to the details of the four disk unit example on the foil. The data in each subarray (A1,B1,C1,D1) has its parity data distributed over the units in the corresponding subarrays. The parity data of subarray 1 on disk A contains the parity bits for subarrays 1 of the other (B,C,D) disk units. The main advantage here is that when for example parity data has to be written on disk units with low occupancy, the disk unit actuator will only need to be moved over a very small distance on the platter before a write can start. This provides a serious advantage over the previous design, when data was written on one side of the platters and parity data on the opposite side. Seeks over more than 2/3 of the platter were happening all the time with the previously used method for parity data striping, especially when the used storage capacity of the disk(s) was rather low.

# New PCI-X Towers

## ■ New PCI-X architecture I/O Towers:


- ▶ 5094/9094
  - up to 45 Disk units (with # 5108 ), internal tape/DVD/CDROM
  - 14 PCI-X slots and up to 11 IOAs
  - 9094 alone on first loop
- ▶ 5294/8094 (double high rack)
  - 8094 same as 5294 but lower unit is system IO
  - 90 Disk units, internal tape/DVD/CDROM
  - Number of PCI-X slots twice as that of 5094
  - Separate HSL loop for each unit
- ▶ 5095/0595 (mini-tower)
  - 12 Disk units.
  - 7 PCI-X slots and up to 5 IOAs
  - 1 - 7 per 1.8m rack
- ▶ 5088/0588
  - 14 PCI-X slots and up to 11 IOAs
  - 1- 4 per 1.8m rack
  - Top hat on 5074/5094



**5094/9094  
Tower**



**8094/5294  
Tower**



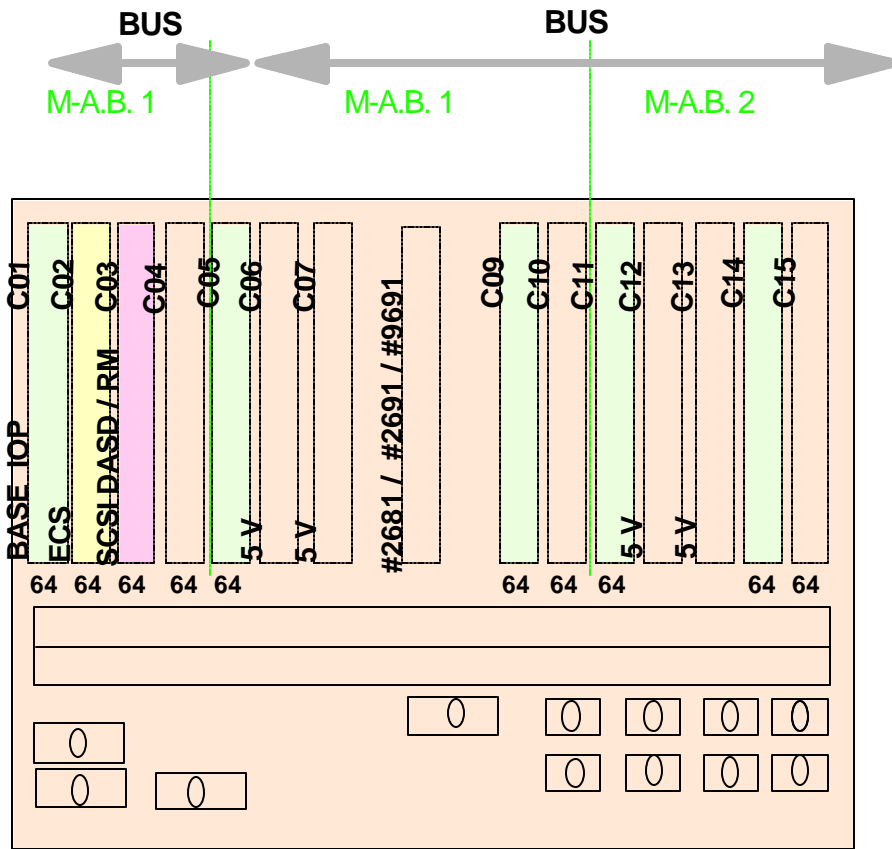
**5095  
Tower**

## Existing PCI I/O Towers

- ▶ New or existing expansion towers are supported on current 8xx model
  - New towers supported on 270 with restrictions
  - Must move to V5R2 (Feb 2003 Level)
- ▶ Consider separating on own loops for balance
- ▶ 5074/9074
  - 45 Disk units, internal Tape/CDROM
  - 14 PCI slots, up to 11 IOAs
- ▶ 5079/9079
  - 90 Disk units, internal tape/DVD/CDROM
  - Double PCI capacity of 5074
- ▶ 5075
  - 6 Disk units
  - 8 PCI slots, up to 7 IOAs.
- ▶ 5078/0578
  - 14 PCI slots, up to 11 IOAs
  - 1-4 per 1.8m rack
- ▶ Existing 840 System IO Unit becomes 5074 on upgrade

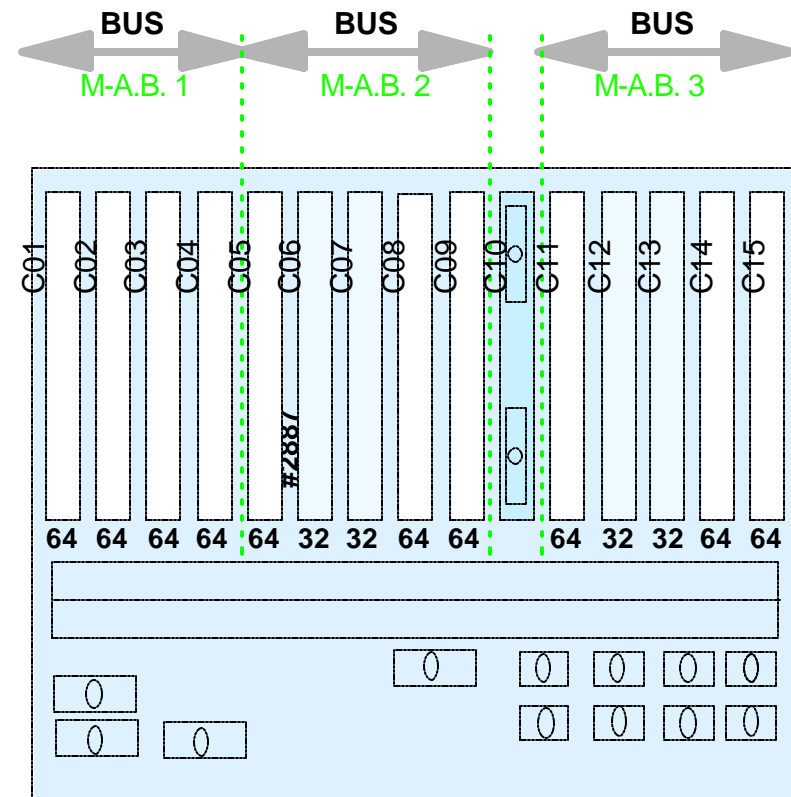
# PCI vs PCI-X Bus structure

Number of PCI buses increases from 2 to 3



**5074/5079/5078/0578**

**PCI cage**



**5094/5294/5088/0588/9094**


**PCI cage**



# Mini Towers 5075/5095



**5075  
Tower**



**0595 rack  
mount  
unit**

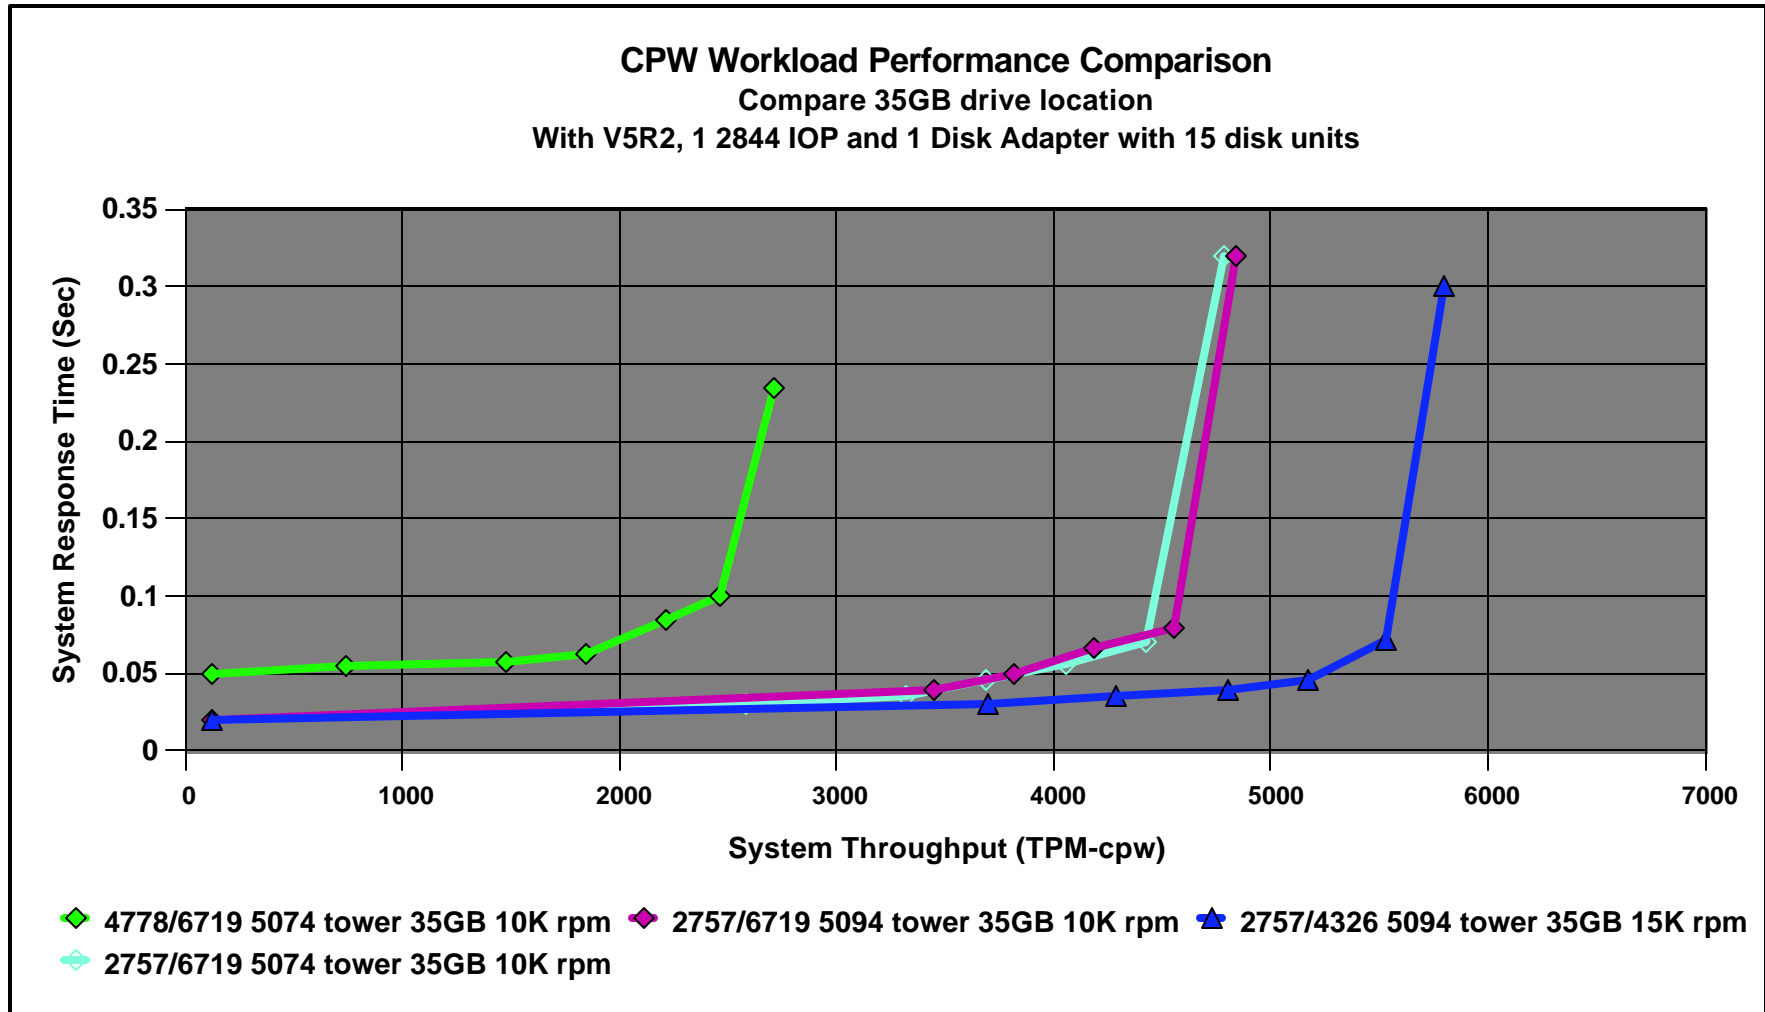


**5095  
Tower**


## Disk controller considerations

- Existing IOA can move to new systems/towers with no effect on RAID structure
- Can install new IOA on existing 270 and 820, 830, 840 models
  - ▶ V5R2 (Feb 2003 level) required
  - ▶ Installation will cause existing RAID structure to be removed and recreated
- Moving existing disks without IOA
  - ▶ On install the existing RAID structure detected, and removed
  - ▶ New RAID structure is created

# Performance Comparison: 35GB Drives and IOA



# Performance Comparison: 18GB with old and new IOAs



# Agenda

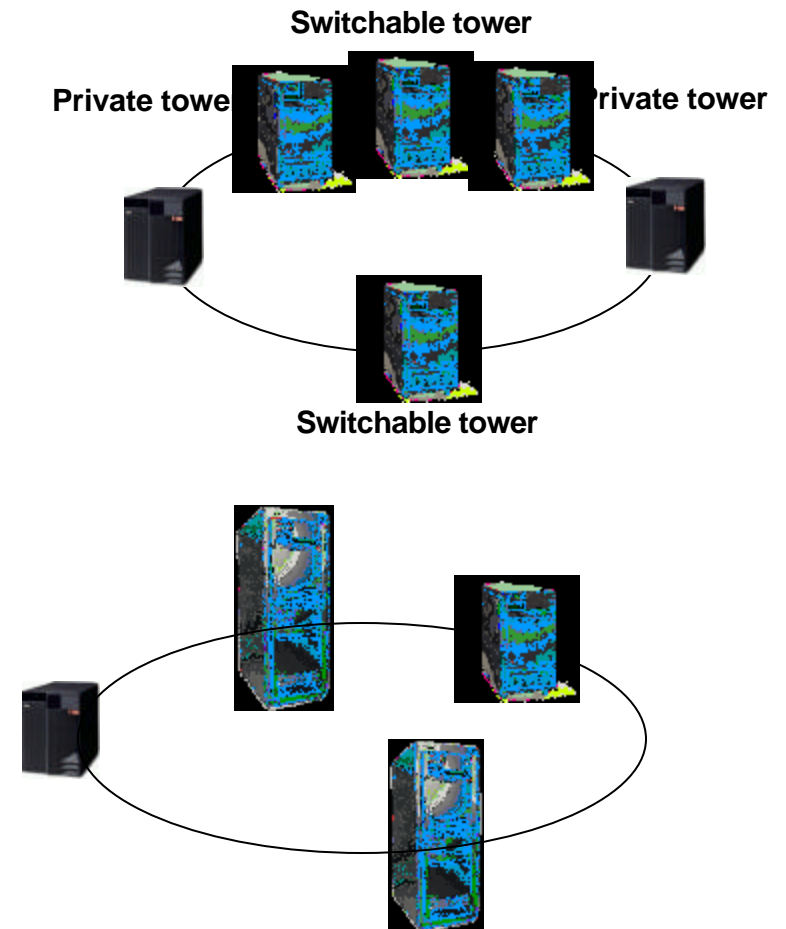
- Server considerations
- Main Storage
- Expansion and IO
- **High Speed Loop**
  - Capacity Upgrade on Demand
  - Upgrade scenarios
  - LPAR
  - Linux
  - Integrated Windows Environment
  - SAN and External devices
  - Software consideratons

## HSL on new Models at V5R2 (Feb 2003 level)

- Mixed HSL on new Models
- 800 and 810 are both HSL
- 825, 870 and 890 are all HSL-2
- HSL-2 with flexibility built in
- All new towers are HSL-2
- Copper HSL and Copper HSL-2 can be intermixed

# HSL Rules and where to find them

- New rules for new servers and IO
- e-Config has some base rules built in
- You must plan your upgraded HSL environment
- Aim to achieve 'balanced' loops
- Try to segregate types of high demand IO
  - ▶ Tape, Disk and Gbit Ethernet
- IXA a special case
  - ▶ Risk of loop split
  - ▶ Migration tower
- iSeries System Builder
- iSeries Server Migration Redbook
- <http://www-1.ibm.com/servers/eserver/series/ha/systemdesign.html>



# Number of HSL Loops by Model

Model	Max. Loops	Type	Max. I/O Towers per Loop	Total Max. I/O Towers	Max. IXA per Loop	Total Max. IXA	Combined Max. per Loop	Combined Max. Total
800	1	HSL	1	1	3	3	4	4
i810	1	HSL	4	4	7	7	8	8
i825	3	HSL-2	6*	(16)18*	5 / 8**	18	9	27
i870	8	HSL-2	6	32***	8	60	9	61***
i890	12 / 14	HSL-2	6	48***	8	60	9	61***

\* 5 PCI (6 PCI-X) towers max. on first and second loop ( i825 has TWO base HSL loops ! )

\*\* 8 on 3rd loop only

\*\*\* includes base I/O tower

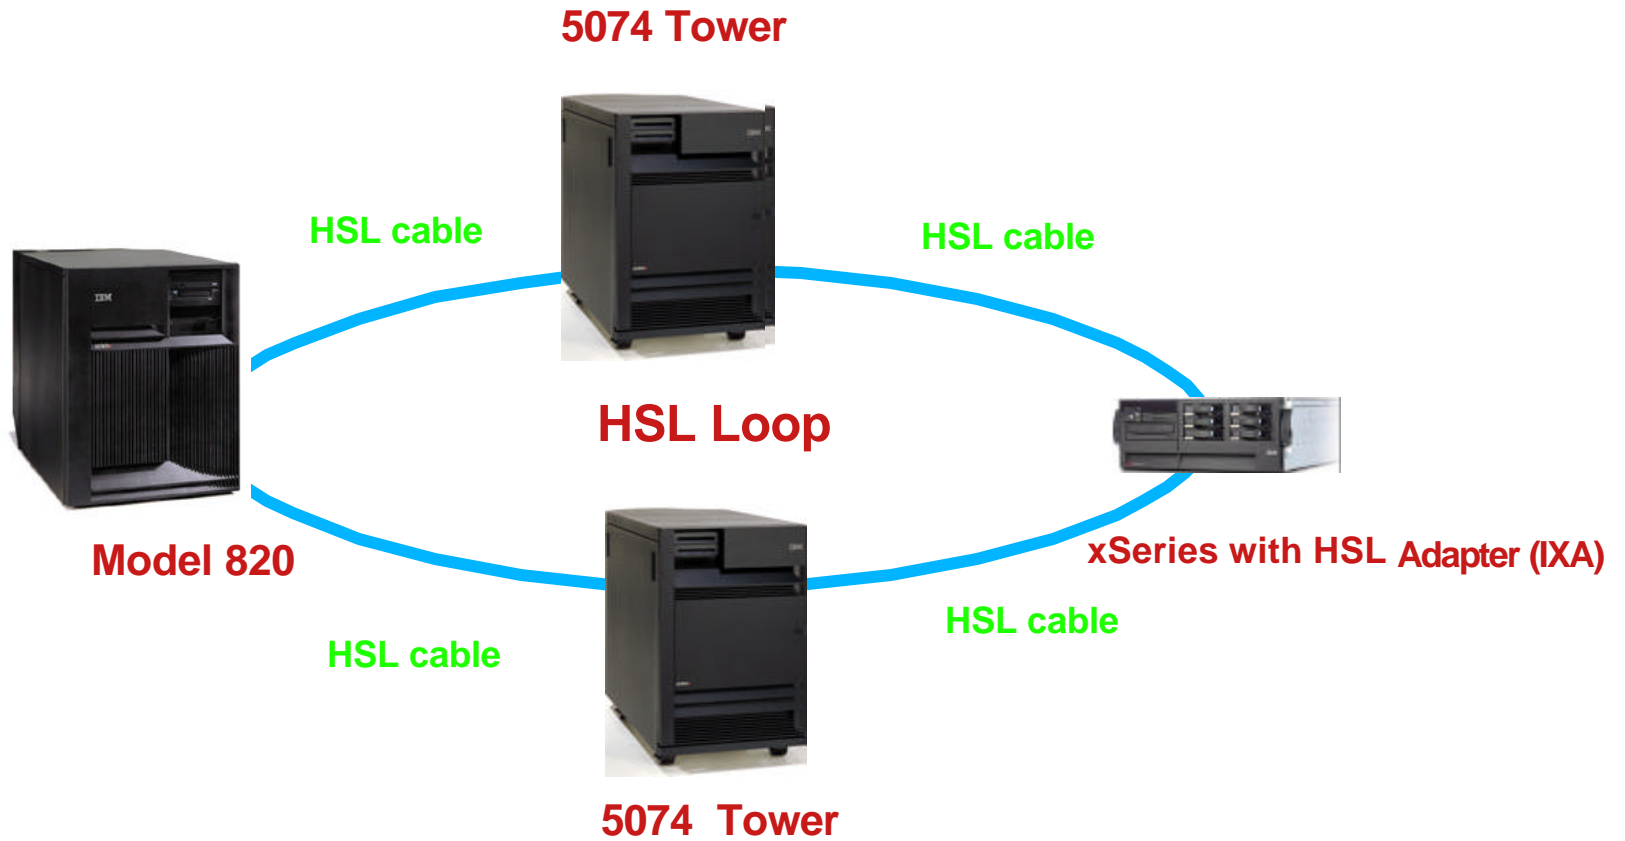
- Special attention required for migrated towers and upgrades
- No migration towers supported on any loop
- Tower placement rules and recommendations
  - Cabling rules and efficiency
  - Performance



# Notes: Number of HSL Loops by Model

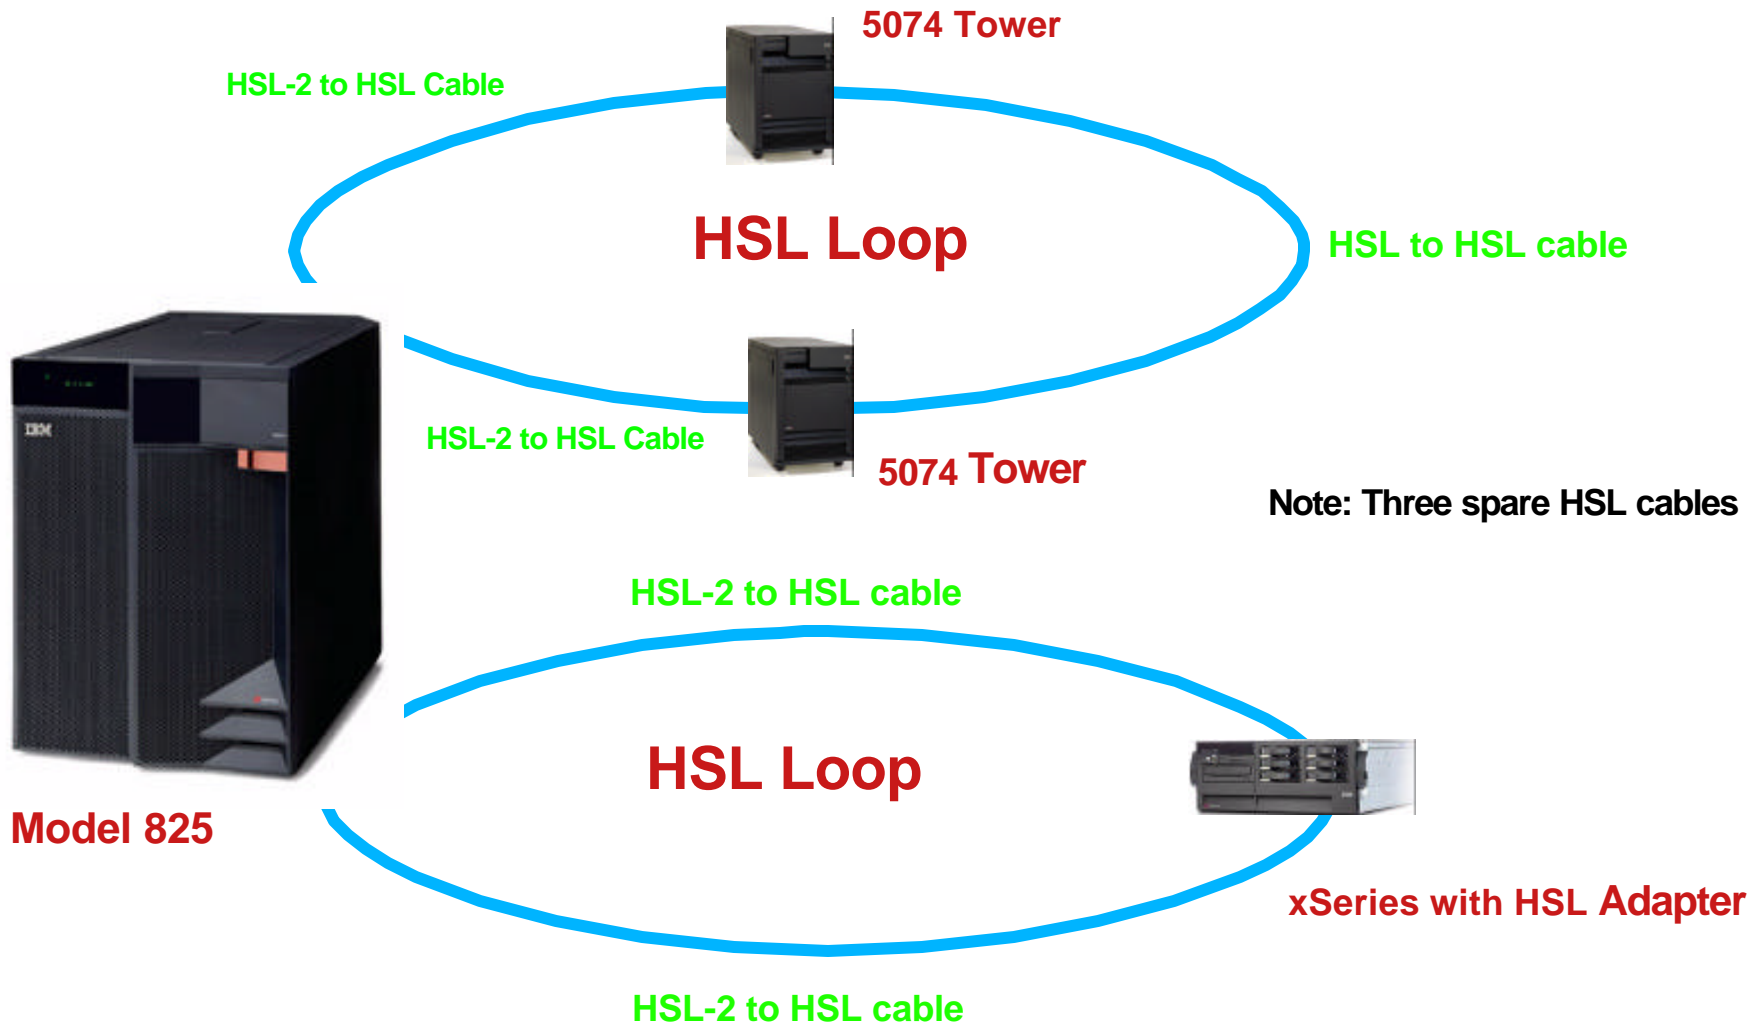
This table shows an overview of the iSeries servers and the associated HSL loops. You can also find the maximum number of towers and IXAs allowed on these servers. The following foils will cover more details on the HSL connectivity. There are no essential changes for HSL loop rules except for server maximums for PCI, PCI-X and IXA towers and their combinations allowed on the HSL loops.

# HSL Migration Example (source system)

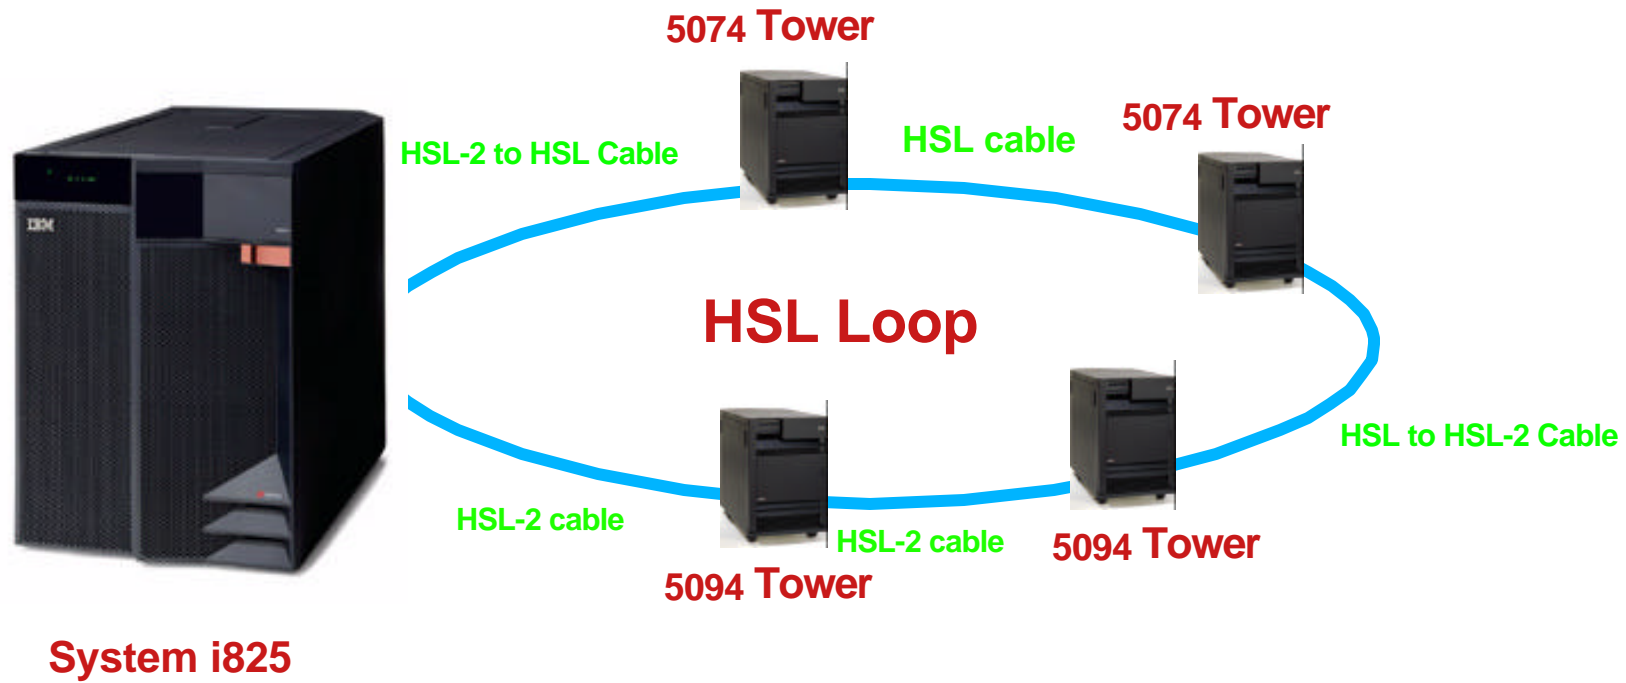


Note : If the bottom 5074 was a migration tower then the xSeries & the top 5074 positions must be reversed.

# HSL Migration Example (Target system)



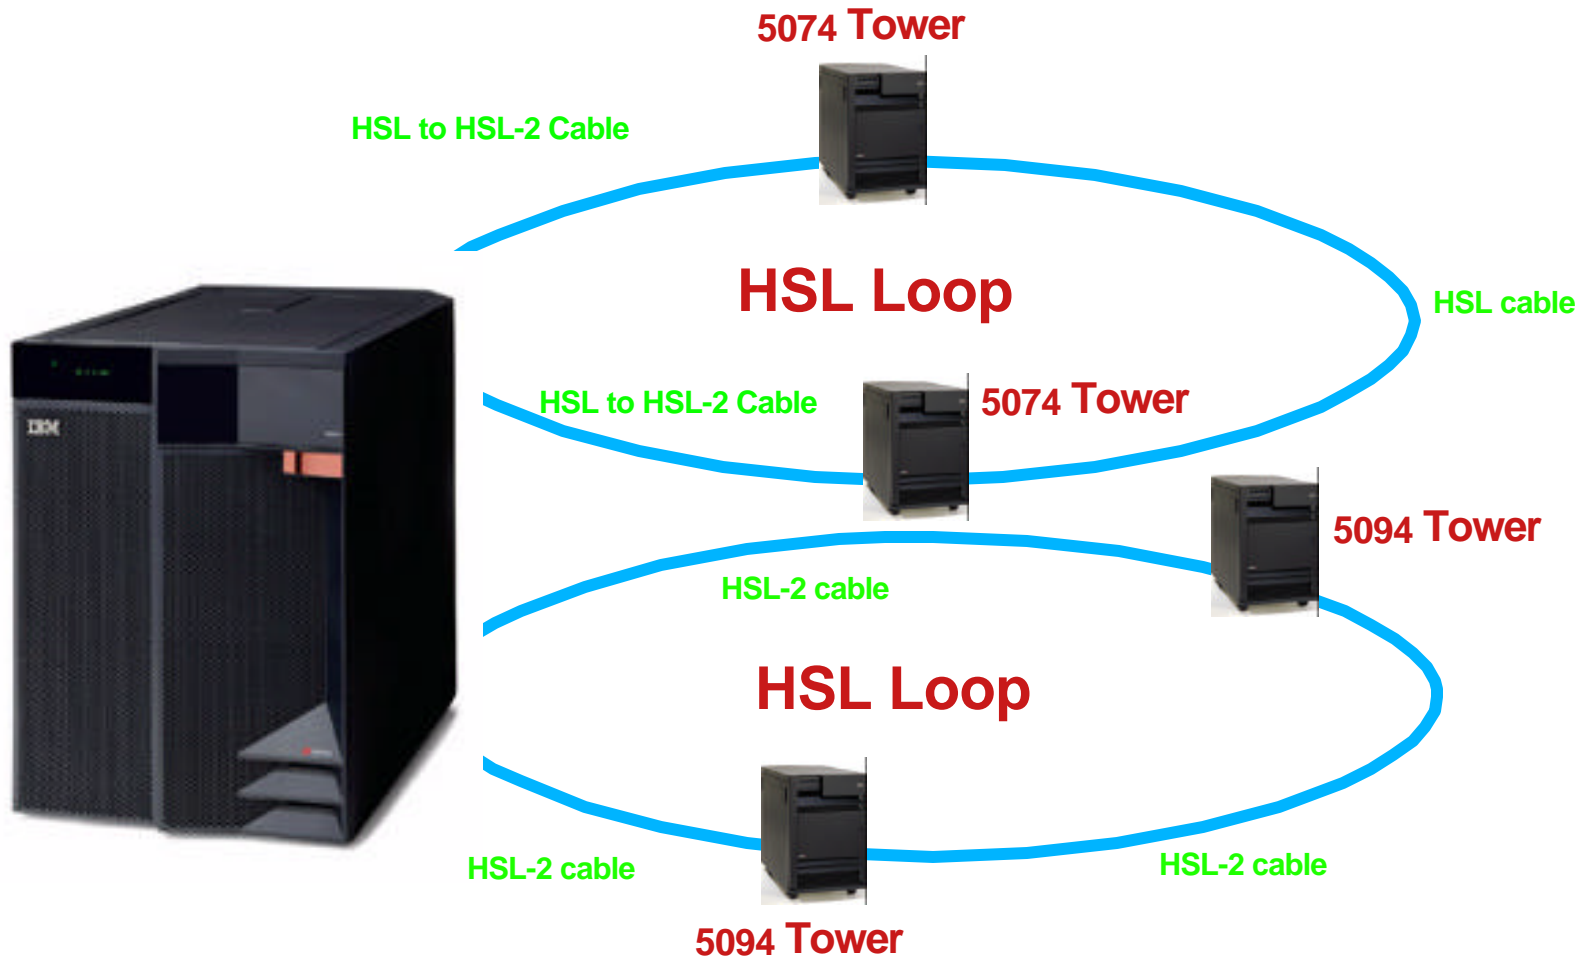
# Flexibility due to multiple HSL loops - 1:



## Scenario-1 : Single loop Scenario

**Note:** Mix of PCI and PCI-X towers in one loop

# Flexibility due to multiple HSL loops - 1:



## Scenario-2 : Multiple loops

Note: PCI and PCI-X towers are on separate loops

# HSL Adapter Overview

- Support for:
  - ▶ HSL-2 cables
  - ▶ HSL to HSL-2 cables
  - ▶ Optical HSL Cables
- All adapters are cluster enabled
  - ▶ Except base HSL adapter on i825
  - ▶ Three-server cluster loops
    - Selected loop on i825, Model 830, Model 840, i870 and i890
- #2776 (copper) and #2788 (optical)
  - ▶ 8 ports active
  - ▶ Only 4 ports are active
    - On i870
    - On i890 16-24 way in slot M39 see (\*) in table

iSeries	Adapter Feature	Active ports	Type
800	N/A	2	Cu / HSL
i810	N/A	2	Cu / HSL
i825	9787 (base)	2	Cu / HSL-2
i825	2785 / 9785	2	Cu / HSL-2
i825	2786 / 9786	2	Optical
i870 i890	9730 (base)	4	Cu / HSL-2
i870 i890	2776	4 8*	Cu / HSL-2
i870 i890	2788	4 8*	Optical


# Notes: HSL Adapter Overview

The new iSeries servers i825, i870 and i890 can have a variety of different HSL ports depending on the HSL adapters selected. It is important for planning the physical connections between the servers and the towers (can be PCI, PCI-X or IXA).

All HSL adapters for the new iSeries servers support HSL clustering. Only exception: the base HSL adapter (**HSL ports C0 and C1 !**) on the i825 does not support HSL clustering. The high end servers (i825, i870 and i890) can have dedicated three server loops. All participating servers on any HSL clustered loop must be at V5R2 when a new iSeries server is on the loop.

The i870 HSL adapters all have 4 ports enabled (2 loops per adapter), the i890 24 to 32-way adapters have all 8 ports enabled (4 HSL loops per adapter) and the i890 16 to 24-way adapters have all 8 ports enabled (4 HSL loops per adapter) except for the HSL adapter in position M39 which has only 4 ports (A0, A1, B0 and B1) enabled (2 loops per adapter).

# Single Loop configuration with IXA






# Notes: Single Loop configuration with IXA

Planning is important for configuring HSL loops, the associated cables and their type and length. Then there is the aspect of performance and availability. Since we will see more complicated combinations than before with the new type of towers (PCI-X), extra considerations must be taken into account. The HSL loop will be reduced to two HSL segments when the IXA is powered down from the xServer console. In the next foil we see an improved configuration.

# Dual Loop configuration with IXA




# Notes: Dual Loop configuration with IXA

On a server with more than one HSL loop, moving the IXA(s) to a separate loop only requires one additional HSL to HSL-2 cable. Chances to break the first HSL loop are reduced. This is a very simple example of improving the HSL configuration. The implementation of such improvements gets more complicated when adding more towers to the configuration since there are different type of towers available.

# Two Server Loop configuration

- Maximum 4 I/O Towers in a two server loop
- Maximum 3 I/O Towers in any segment of a two server loop



# Notes: Two Server Loop configuration

For loops that contain two servers the following rules apply:


- The maximum number of external towers per **loop** is 4.
- The maximum number of external towers per **segment** is 3.
- These limits apply to the sum of all types of external towers, including IXS towers.

Switchable towers in this type of setup must be placed according to other rules such as:

- A switchable tower must be immediately adjacent in the HSL loop to the alternate system, or to a tower (private switchable, or switched) owned by the alternate system.
- Two systems define the ends of an HSL loop segment on which a switchable tower resides. One of those systems must be the home system and the other must be the alternate system for that tower. (This rule is essentially a direct result of first rule.)
- IXS towers are not switchable towers

# Three Server Loop configuration

- No I/O Towers allowed in any segment of a three server loop

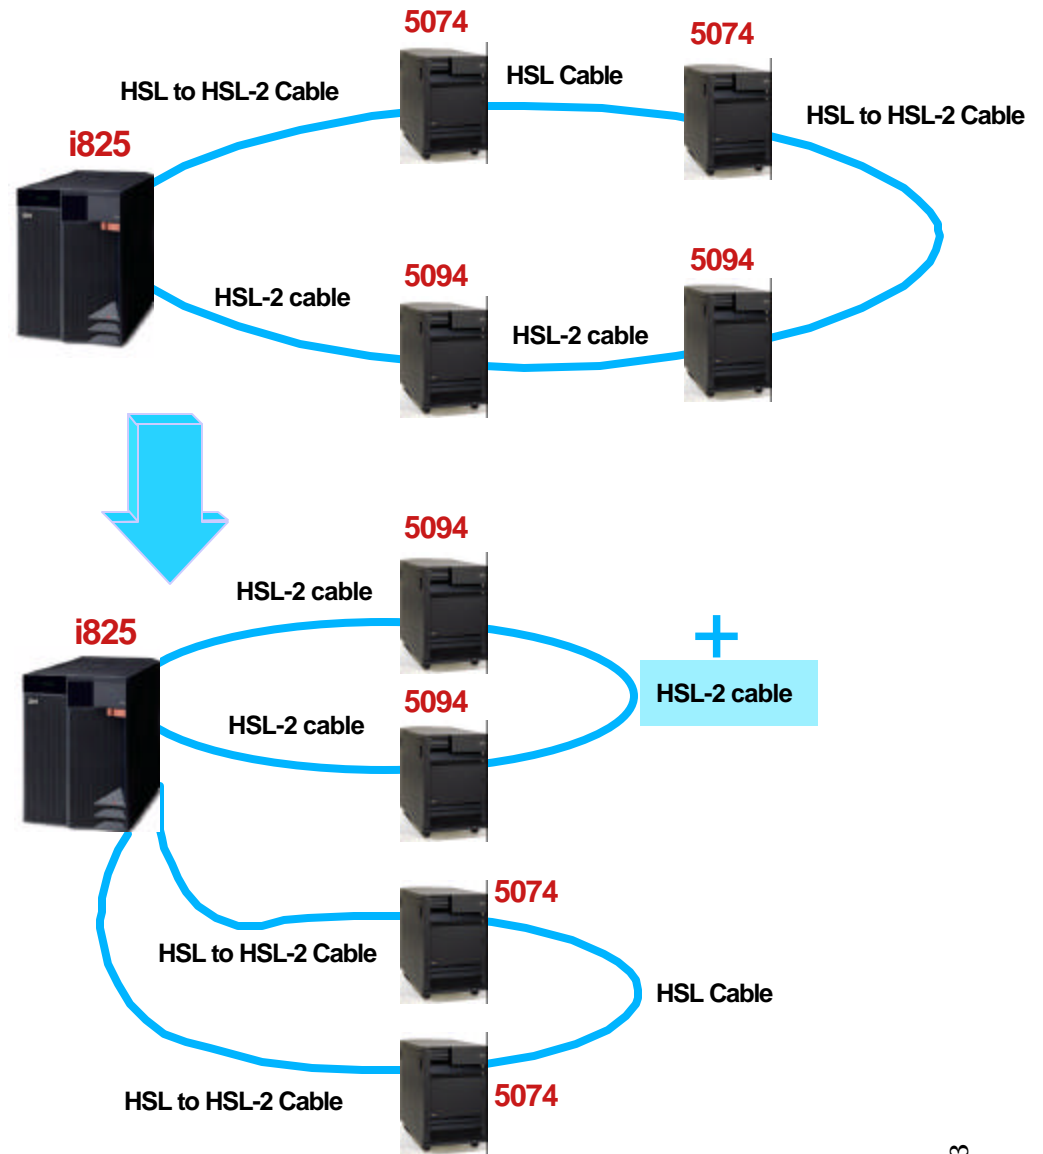


# Notes: Three Server Loop configuration

For the three server loop configuration there is one simple rule that applies: there can't be any towers on this loop. All segment cables need to run directly from one server to the other. All three participating servers must be at V5R2 !

# Cable Configurations and Tower Arrangements

- Tower arrangement for performance
  - ▶ Populate available loops
    - Balance number of towers
    - Find examples in System Builder
  - ▶ Depending on the number of towers and the arrangement:
    - More adapters required
    - More cables required
  - ▶ Group towers by type if possible
    - Order correct number and correct type of cables
      - HSL cables
      - HSL-2 cables
      - HSL to HSL-2 cables





# Notes: Cable Configurations and Tower Arrangements

In the example given on this foil, changing the I/O tower arrangement for better performance requires just a single additional cable. Since the i825 has two base HSL-2 adapters, it is much more efficient from a performance perspective to populate both loops with I/O towers. Depending on the requirements, one can decide to order additional HSL adapters with the correct type of HSL cables for a suitable tower cabling configuration. In the Tower chapter under High-speed link tower placement for iSeries in the iSeries System Builder SG24-2155-07, you can find a set of rules and recommendations for setting up the I/O towers on the different servers.

FIND DETAILS FOR CABLING IN **SUPPLEMENTAL INFORMATION SECTION**

# Agenda

- Server considerations
- Main Storage
- Expansion and IO
- High Speed Loop
- **Capacity Upgrade on Demand**
  - Upgrade scenarios
  - LPAR
  - Linux
  - Integrated Windows Environment
  - SAN and External devices
  - Software considerations

# Capacity Upgrade on Demand types


- Capacity Upgrade on Demand (Permanent capacity upgrade)
  - ▶ You buy the processors just before you need them
  
- On/Off CUoD (Temporary capacity upgrade)
  - ▶ Must sign a contract
  - ▶ Zero charge enablement feature that can be renewed
  - ▶ 14 days 'trial' available with restrictions
  - ▶ Pay as you use resources
  - ▶ Maximum of 192 days under single activation
  - ▶ Consider
    - memory
    - LPAR
  - ▶ Keep a balanced system

## Upgrading in CUoD

- **Simple upgrade process**
- Must activate all of the current model processors to upgrade to upgrade to new processor feature
- For example
  - ▶ 14 processors of 8/16 way 870
  - ▶ Permanently activate the other 2 processors to get upgrade to 16/24 way 890
  - ▶ Add any additional processor activation features for performance/capacity
- Remember to relocate processors, and memory to new partitions
- **Capacity planning**
- A 'business as usual' value-add of IBM's sales team.
- There is a comprehensive set of tools, resources and offerings available to aid IBMers, customers and IBM Business Partners with capacity planning assistance.
- Simply activating standby processors may not provide with the full potential of required capacity improvements.


# Permanent Capacity Upgrade on Demand

- Known growth plan
- may have used temporary CUoD to determine requirements
- Chart shows typical business cycle
  - ▶ Continuous planned growth
  - ▶ Known or predictable peak periods





# Capacity Upgrade on Demand (Temporary)

- Unknown growth plan
- Sudden unexpected growth
  - ▶ Acquisition
  - ▶ New application or version change
- Chart shows typical business cycle
  - ▶ Continuous growth with peak periods
  - ▶ Plus huge growth spike



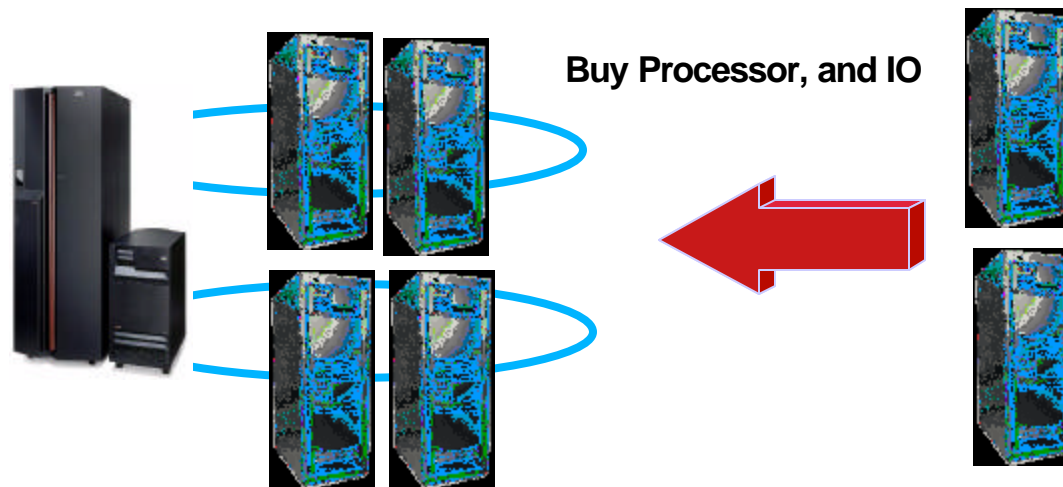
# Temporary upgrade on demand example 1

- Recovery Services
  - Buy a minimal number of processors
  - Activate processors temporarily
    - ▶ As needed in case of disaster
    - ▶ Charge straight back to customer
    - ▶ Turn off



## Temporary upgrade on demand example 2

- ASP provider
- Buy a minimal number of processors
- Buy temporary processors
  - ▶ As needed for new customers
  - ▶ Charge straight back to customer
  - ▶ Convert to Permanent after certain time





# Permanent Capacity Upgrade on Demand

- How does it work ?
  - ▶ Order activation feature
    - Can be included with initial order (no further user action required)
    - Separate as an upgrade on request
  - ▶ Receive key
    - By mail and at the same time
    - Via the Web :  
[HTTP://WWW.IBM.COM/ESERVER/SERIES/HARDWARE/ONDEMAND](http://www.ibm.com/eserver/series/hardware/ondemand)
  - ▶ Introduce activation key and start using the additional processing capacity (no IPL required)
- Changes for the new iSeries servers
  - ▶ Trial period replaced by Temporary Capacity Upgrade on Demand
    - More flexibility
    - Simplified ordering and activation
- Permanently activated processors can not be deactivated

# Work with System Capacity Interface

## System Service Tools (SST)

Select one of the following:

1. Start a service tool
2. Work with active service tools
3. Work with disk units
4. Work with diskette data recovery
5. Work with system partitions
6. Work with system capacity
7. Work with system security
8. Work with service tools

Selection

6

F3=Exit

F10=Com

# PLANNING

# !

## WORK WITH SYSTEM CAPACITY

SYSTEM: SYS1

ATTENTION: THIS SERVICE FUNCTION IS USED TO DISPLAY AND ALTER SYSTEM PROCESSING CAPACITY. SEE SERVICE DOCUMENTATION.

SELECT ONE OF THE FOLLOWING:

1. DISPLAY SYSTEM CAPACITY INFORMATION
2. ACTIVATE PERMANENT SYSTEM CAPACITY (PCOD)
3. ENABLE TEMPORARY SYSTEM CAPACITY (TCOD)

SELECTION

3

F3=EXIT

F10=DISPLAY HISTORY LOG

F12=CANCEL

# Notes: Work with System Capacity Interface

From this foil on we will show the flow for enablement and activation of Temporary Processing Capacity on an iSeries server. The use of Service Tools is required, using DST (the full paging environment using function 21) or SST can be used for minimal server operations interruption. The first screens shown are used to obtain the data needed for requesting the CUoD(temporary) enablement key. These steps may not be necessary for servers that have the Machine Serial number and/or the Capacity Card Serial number have not been altered outside of the IBM manufacturing or field maintenance process.

## Planning !

When you are requesting CUoD(temporary) enablement, you must also plan for making those processors active in your production environment without minor disruption. Whether you have a partitioned system or not, make sure that all the values in the partition configuration for minimum and maximum capacities are set in order to be able to adjust the active capacity to the needed value without requiring a IPL for a partition or the server. This issue does not only concerns the processor minimum and maximum values, but also the main storage minimum and maximum values.

Activating a number of processors for a certain number of days when you require them creates a condition that disables the activation of extra processors. This means for example that when you underestimate the number of processors you really need and want an additional processor, you must return your already activated processors first (release them from the partition(s) and stop CUoD(temporary)) before you will be able to activate the number of processors you really need.

Activations are done based on the number of processors and the number of days you need the extra capacity; the CUoD(temporary) counter provides statistics based on the multiplication of processors times the number of days. Before you can activate any processor, you need to order an enablement feature, orderable via the web, that will provide you a key that gives you the capability to use up to 192 units of the TCOd counter; when these have expired, a new key has to be ordered.

Once a number or processor days has been requested and these processors have been assigned to a partition, they will remain active until you remove them from the partition configuration and deactivate them. The Service Tools interface allows you to view and print the activation requests.

Let us proceed and walk through the most important steps and their matching screens.

# Enable Temporary System Capacity

## ENABLE TEMPORARY SYSTEM CAPACITY

SYSTEM: SYS1

TEMPORARY CAPACITY ON DEMAND (TCOD) IS NOT ENABLED ON THIS SYSTEM. TO LEARN MORE ABOUT CAPACITY UPGRADE ON DEMAND (CUOD) FOR THE ISERIES, REFER TO THE PLANNING GUIDE AVAILABLE AT:

[HTTP://WWW.IBM.COM/ESERVER/ISERIES/HARDWARE/ONDEMAND](http://www.ibm.com/eserver/series/hardware/ondemand)

IN ORDER TO ENABLE TCOOD ON YOUR SERVER, YOU MUST FIRST PURCHASE A TCOOD ENABLEMENT FEATURE. TO PURCHASE THIS FEATURE:

- 1) REQUEST YOUR SALES CHANNEL TO PLACE AN ORDER
- AND 2) COLLECT AND FORWARD MACHINE INFORMATION TO IBM (USE F10 TO DISPLAY MACHINE INFORMATION).

WHEN THE ORDER IS FULFILLED, YOU CAN RETRIEVE THE TCOOD ENABLEMENT CODE PROVIDED BY IBM AND ENTER IT ON THIS MENU. ONCE THE TCOOD ENABLEMENT CODE IS ENTERED, YOU WILL BE ABLE TO REQUEST AN AMOUNT OF TEMPORARY CAPACITY (PROCESSOR-DAYS) FOR YOUR SERVER.

COMPLETE BLANKS, PRESS ENTER.

TCOD ENABLEMENT CODE . . . . . \_\_\_\_\_

F3=EXIT    F6=PRINT    F10=DISPLAY MACHINE INFORMATION    F12=CANCEL

# Notes: Enable Temporary System Capacity

Since some vital product data needs to be obtained from the server that must be send to IBM, the user has to use function F10 Display Machine Information for retrieving the required data.

# Enable Temporary System Capacity - 2

## ENABLE TEMPORARY SYSTEM CAPACITY

SYSTEM: SYS1

COLLECT AND FORWARD THE FOLLOWING MACHINE INFORMATION TO IBM  
WHEN ORDERING A TCOB ENABLEMENT FEATURE.

SYSTEM TYPE . . . . .	:	9406
SYSTEM SERIAL NUMBER . . . . .	:	10-12345
CAPACITY CARD CCIN . . . . .	:	7421
CAPACITY CARD SERIAL NUMBER . . . . .	:	10-200010
CAPACITY CARD UNIQUE IDENTIFIER . . . . .	:	2035092805B86508
ACTIVATION FEATURE . . . . .	:	1776
ACTIVATED TCOB UNITS . . . . .	:	0000
SEQUENCE NUMBER . . . . .	:	0022
ENTRY CHECK . . . . .	:	5C

SYSTEM: SYS1

F3=EXIT   F6=PRINT   F10=DISPLAY ORDERING INFORMATION   F12=CANCEL



OF TEMPORARY CAPACITY (PROCESSOR-DAYS) FOR YOUR SERVER.

COMPLETE BLANKS, PRESS ENTER.

TCOB ENABLEMENT CODE . . . . . D11FB64C4BDA77EB17750000019200232E



F3=EXIT   F6=PRINT   F10=DISPLAY MACHINE INFORMATION   ~~F12=CANCEL~~

## Notes: Enable Temporary System Capacity -2

In order to receive the key for enabling CUoD(temporary), the data shown in the " Enable Temporary System Capacity " screen needs to be send to IBM together with the required contracts.

- IBM Customer Agreement, Attachment for Temporary Capacity On Demand
- IBM Supplement for Temporary Capacity On Demand
- And Optional: IBM Addendum for Temporary Capacity on Demand Alternative Reporting

After a number of necessary administration steps, IBM generates a TCOD Enablement Code and mails/posts it. The code can be obtained from a secure We site. Then the code needs to be entered in the "Enable Temporary Capacity" screen as shown in this foil.

# Confirm Enable Temporary System Capacity

## CONFIRM ENABLE TEMPORARY SYSTEM CAPACITY

SYSTEM: SYS1

**PRESS ENTER** TO CONFIRM YOUR CHOICE TO ENABLE TEMPORARY CAPACITY ON DEMAND.

PRESS F12 TO RETURN TO CHANGE YOUR CHOICE.

NUMBER OF SYSTEM PROCESSORS . . . . . : 16

**NEW NUMBER OF TCOD PROCESSOR-DAYS . . . . . : 192**

SYSTEM TYPE . . . . . : 9406

SYSTEM SERIAL NUMBER . . . . . : 10-12345

CAPACITY CARD CCI

CAPACITY CARD SER

CAPACITY CARD UNI

ACTIVATION FEATUR

ACTIVATED TCOD UN

F12=CANCEL

## WORK WITH SYSTEM CAPACITY

SYSTEM: SYS1

ATTENTION: THIS SERVICE FUNCTION IS USED TO DISPLAY AND ALTER SYSTEM PROCESSING CAPACITY. SEE SERVICE DOCUMENTATION.

SELECT ONE OF THE FOLLOWING:

- 1. DISPLAY SYSTEM CAPACITY INFORMATION
- 2. ACTIVATE PERMANENT SYSTEM CAPACITY (PCOD)
- 3. ENABLE TEMPORARY SYSTEM CAPACITY (TCOD)
- 4. WORK WITH TEMPORARY SYSTEM CAPACITY (TCOD)



SELECTION

4

F3=EXIT F10=DISPLAY HISTORY LOG F12=CANCEL

TCOD SYSTEM CAPACITY ENABLE WAS SUCCESSFUL.



# Notes: Confirm Enable Temporary System Capacity

With the correct code entered in the previous screen a confirmation screen appears and pressing the Enter button on the keyboard of your workstation will enable the server for CUoD (temporary). As can be noticed, a total of 192 processor days is now available for use. If this is the first billing period for a given iSeries serial number, the billing system will credit the first 14 processor days so that for the first period 178 processing days will be billed.

If the server is in bypass mode for the System Password, access to the Service Tools screens associated with the CUoD functionality may not be available. If this is the case, you should obtain a valid System Password by contacting your IBM business partner or IBM Sales representative. After entering a TCOD Enablement Code, the server must be run for 15 minutes to insure that the TCOD Enablement Activation Code is persistently stored (remembered) within the server. If you enter a TCOD Enablement Code and shut down the server before it has run for 15 minutes, on your next IPL, the TCOD Enablement Code may be lost by the server and has to be re-entered.

Using the new option 4 that now appears on the "Work with System Capacity" screen we reach the screen on the next foil.

# Work with Temporary System Capacity

```

                                WORK WITH TEMPORARY SYSTEM CAPACITY

                                SYSTEM:  SYS1
NUMBER OF SYSTEM PROCESSORS . . . . . : 16                01/15/03
STANDBY PROCESSORS CURRENTLY REQUESTED . . . . . : 0        10:52:39
STANDBY PROCESSORS AVAILABLE FOR REQUEST . . . . . : 8      ←
HISTORY OF REQUESTED PROCESSOR-DAYS . . . . . : 0
HISTORY OF UNRETURNED PROCESSOR-DAYS . . . . . : 0

TYPE OPTION, PRESS ENTER.
1=START

      REQUESTED   TCOD ENABLEMENT
OPT  PROCS  DAYS   COUNTER
-    -     -     192

F3=EXIT   F5=REFRESH   F11=DISPLAY TEMPORARY SYSTEM CAPACITY   F12=CANCEL

```

# Notes: Work with Temporary System Capacity

On the "Work with Temporary System Capacity" screen one can find the total processors for available for temporary activation. The option to start a number of processors for a certain amount of days together with the available amount of processor days is present on this screen. Before proceeding with the real activation let us first explore a number of interesting green screens that are relevant for CUoD. Using function F11 we go to the "Display Temporary System Capacity" screen.

# Display Temporary System Capacity

```

                                DISPLAY TEMPORARY SYSTEM CAPACITY

                                SYSTEM:  SYS1
NUMBER OF SYSTEM PROCESSORS . . . . . : 16          01/15/03
STANDBY PROCESSORS CURRENTLY REQUESTED . . . . . : 0          10:54:33
STANDBY PROCESSORS AVAILABLE FOR REQUEST . . . . . : 8
HISTORY OF REQUESTED PROCESSOR-DAYS . . . . . : 0
HISTORY OF UNRETURNED PROCESSOR-DAYS . . . . . : 0

SYSTEM TYPE AND SERIAL NUMBER . . . . . : 9406  10-12345
CAPACITY CARD CCIN . . . . . : 7421
CAPACITY CARD SERIAL NUMBER . . . . . : 10-2000010
CAPACITY CARD UNIQUE IDENTIFIER . . . . . : 2035092805B86508
ACTIVATION FEATURE . . . . . : 1775
SEQUENCE NUMBER . . . . . : 0023

REQUESTED  --PROCESSOR-DAYS--  CALENDAR DAYS  EXPIRATION  TCOD ENABLEMENT
PROCS  DAYS  EXPIRED  REMAINING  REMAINING  DATE          COUNTER
                                192

F3=EXIT  F5=REFRESH  F6=PRINT
F10=DISPLAY PERMANENT SYSTEM CAPACITY  F12=CANCEL

```



# Notes: Display Temporary System Capacity

The "Display Temporary System Capacity" screen shows the status of server for as far as CUoD (temporary) concerns. From this screen we can easily toggle from the "Display Temporary System Capacity" screen to the the "Display Permanent System Capacity" screen by using the function F10. This is ofcourse a step that is completely unnecessary for CUoD (temporary) but shown for their interest.

# Display Permanent System Capacity

```

                                DISPLAY PERMANENT SYSTEM CAPACITY
                                SYSTEM:      SYS1
NUMBER OF SYSTEM PROCESSORS . . . . . : 16
NUMBER OF PROCESSORS IN USE . . . . . : 8
NUMBER OF PROCESSORS AVAILABLE TO ORDER . . . . . : 8

USE THE FOLLOWING INFORMATION WHEN ORDERING A PERMANENT
(PCOD) ACTIVATION CODE.

SYSTEM TYPE . . . . . : 9406
SYSTEM SERIAL NUMBER . . . . . : 10-12345
CAPACITY CARD CCIN . . . . . : 7421
CAPACITY CARD SERIAL NUMBER . . . . . : 10-2000010
CAPACITY CARD UNIQUE IDENTIFIER . . . . . : 2035092805B86508
ACTIVATION FEATURE . . . . . : 1776
ACTIVATED PCOD UNITS . . . . . : 0000
SEQUENCE NUMBER . . . . . : 0020
ENTRY CHECK . . . . . : 46

F3=EXIT      F5=REFRESH      F6=PRINT
F10=DISPLAY TEMPORARY SYSTEM CAPACITY      F12=CANCEL
    
```

# Notes: Display Permanent System Capacity

As stated earlier this screen just to show that regardless the number of any enablement for CUoD(temporary) or any temporary activated processors , the number of processor eligible for CUoD(permanent) stay the same.

# Work with Temporary System Capacity -2

```

                                WORK WITH TEMPORARY SYSTEM CAPACITY

                                SYSTEM:  SYS1
NUMBER OF SYSTEM PROCESSORS . . . . . : 16          01/15/03
STANDBY PROCESSORS CURRENTLY REQUESTED . . . . . : 0          10:59:44
STANDBY PROCESSORS AVAILABLE FOR REQUEST . . . . . : 8
HISTORY OF REQUESTED PROCESSOR-DAYS . . . . . : 0
HISTORY OF UNRETURNED PROCESSOR-DAYS . . . . . : 0

TYPE OPTION, PRESS ENTER.
1=START

                                REQUESTED   TCOD ENABLEMENT
                                OPT  PROCES  DAYS      COUNTER
                                ---  ---  ---
1    3    4          192

F3=EXIT   F5=REFRESH   F11=DISPLAY TEMPORARY SYSTEM CAPACITY   F12=CANCEL

```



## Notes: Work with Temporary System Capacity -2

By returning to the "Work with Temporary System Capacity" screen we can now select to start a number of processors for a number of days. When starting these processors for a limited period of time one needs to know that once started, you can't start additional processors anymore before the already started processors are returned to the server. A condition of activated or unreturned processors will inhibit the start option to be functional.

# Confirm Temporary System Capacity Change

## CONFIRM TEMPORARY SYSTEM CAPACITY CHANGE

SYSTEM: SYS1

VERIFY INFORMATION, PRESS ENTER TO START THE REQUEST, OR  
PRESS F12 TO CHANGE YOUR CHOICE.

REQUESTED:

PROCESSORS . . . . .	:	3
DAYS . . . . .	:	4
PROCESSOR-DAYS . . . . .	:	12

YOU HAVE SPECIFIED REQUESTED UNITS (A QUANTITY OF PROCESSORS FOR A DURATION OF DAYS) OF TEMPORARY CAPACITY. AFTER YOU CONFIRM THIS REQUEST FOR SUCH TEMPORARY CAPACITY, BY PRESSING ENTER ON THIS MENU, YOU MUST ASSIGN THE NEWLY ACTIVATED PROCESSORS TO A PARTITION(S) OF YOUR CHOICE TO ACTUALLY BEGIN USING THE TEMPORARY CAPACITY. PLEASE NOTE, ACCOUNTING FOR BILLING PURPOSES BEGINS WHEN YOU CONFIRM THIS REQUEST. ACCORDINGLY, IT IS IMPORTANT THAT YOU ASSIGN THE NEWLY ACTIVATED PROCESSORS TO A PARTITION TO TAKE FULL ADVANTAGE OF YOUR REQUESTED UNITS. ONCE THE NUMBER OF DAYS YOU HAVE REQUESTED HAVE PASSED, YOU ARE REQUIRED TO MAKE THE TEMPORARY CAPACITY YOU HAVE REQUESTED 'AVAILABLE' (LPAR RESOURCE STATE) TO BE RECLAIMED BY THE SERVER. IN THE EVENT YOU FAIL TO MAKE SUCH TEMPORARY CAPACITY AVAILABLE TO BE RECLAIMED, THE SERVER WILL CONTINUE TO COUNT THE NUMBER OF UNRETURNED UNITS (DURATION OF DAYS MULTIPLIED BY THE QUANTITY OF PROCESSORS WHICH HAVE NOT BEEN RECLAIMED) AND YOU WILL BE BILLED FOR THE UNRETURNED UNITS, PURSUANT TO THE TERMS AND CONDITIONS OF THE AGREEMENT YOU HAVE SIGNED WITH IBM OR YOUR IBM BUSINESS PARTNER. DEFINED TERMS USED HEREIN ARE MORE FULLY DEFINED IN SUCH AGREEMENTS. FOR MORE INFORMATION ABOUT USING LPAR TO ASSIGN PROCESSORS TO/FROM PARTITIONS, SEE THE LOGICAL PARTITIONS TOPIC IN THE ISERIES INFORMATION CENTER.

Bottom

F12=CANCEL

# Notes: Confirm Temporary System Capacity Change

The text of the message on this screen is split over two screens on your workstation when you receive it, so you need to scroll down in order to read the message completely. The information makes you aware of the fact that even if you don't put the processors to work, you will be billed for the time they are activated. The way the activated processors must be assigned to a partition (whether or not the server is configured for LPAR) will be covered in the next foils.

# Work with Temporary System Capacity -3

```

                                WORK WITH TEMPORARY SYSTEM CAPACITY

                                SYSTEM:  SYS1
NUMBER OF SYSTEM PROCESSORS . . . . . : 16      01/15/03
STANDBY PROCESSORS CURRENTLY REQUESTED . . . . . : 3      11:02:52
STANDBY PROCESSORS AVAILABLE FOR REQUEST . . . . . : 5
HISTORY OF REQUESTED PROCESSOR-DAYS . . . . . : 0
HISTORY OF UNRETURNED PROCESSOR-DAYS . . . . . : 0

TYPE OPTION, PRESS ENTER.
2=STOP

      REQUESTED  --PROCESSOR-DAYS--  CALENDAR DAYS  TCOD ENABLEMENT
OPT  PROCS  DAYS  EXPIRED  REMAINING  REMAINING  COUNTER
_    3      4      0        12         4          180

F3=EXIT  F5=REFRESH  F11=DISPLAY TEMPORARY SYSTEM CAPACITY  F12=CANCEL
TCOD START REQUEST SUCCESSFUL.

```

## Notes: Work with Temporary System Capacity -3

When you enter on the confirmation screen the server returns you again the "Work with Temporary System Capacity Screen" but with a changed option (you can now only stop processors). The function F12 brings you back to the next screen.

# Work System Capacity

## WORK WITH SYSTEM CAPACITY

SYSTEM: SYS1

ATTENTION: THIS SERVICE FUNCTION IS USED TO DISPLAY AND ALTER SYSTEM PROCESSING CAPACITY. SEE SERVICE DOCUMENTATION.

SELECT ONE OF THE FOLLOWING:

1. DISPLAY SYSTEM CAPACITY INFORMATION
2. ACTIVATE PERMANENT SYSTEM CAPACITY (PCOD)
3. ENABLE TEMPORARY SYSTEM CAPACITY (TCOD)
4. WORK WITH TEMPORARY SYSTEM CAPACITY (TCOD)

SELECTION

—

F3=EXIT F5=REFRESH F10=DISPLAY HISTORY LOG F12=CANCEL  
TCOD PROCESSORS REQUESTED, LESS THAN 4 DAYS REMAIN.

# Notes: Work System Capacity

Back at the main menu for working with system capacity. Using the function F10 w will have a look at the system capacity history log.

# Display History Log (CUoD)

## DISPLAY HISTORY LOG

```

                                SYSTEM:  SYS1
SYSTEM TYPE . . . . . : 9406      01/15/03
SYSTEM SERIAL NUMBER . . . . . : 10-12345  11:04:49
CAPACITY CARD CCIN . . . . . : 7421
CAPACITY CARD SERIAL NUMBER . . . . . : 10-2000010
CAPACITY CARD UNIQUE IDENTIFIER . . . . . : 2035092805B86508

```

TYPE OPTION, PRESS ENTER.

5=DISPLAY DETAIL

OPT	CHANGE DESCRIPTION	DATE	TIME
<input type="checkbox"/> 5	TCOD START PROCESSOR-DAYS	01/15/03	11:02:52

F3=EXIT F5=REFRESH  F6=PRINT F11=HELP

## DISPLAY HISTORY LOG DETAIL

```

                                SYSTEM:  SYS1
CHANGE DESCRIPTION . . . . . : TCOD START PROCESSOR-DAYS
DATE . . . . . : 01/15/03
TIME . . . . . : 11:02:52
CHANGE:
  PROCESSORS . . . . . : 3
  DAYS . . . . . : 4
  PROCESSOR-DAYS . . . . . : 12

```

F3=EXIT F9=PREVIOUS DETAIL F10=NEXT DETAIL F12=CANCEL



# Notes: Display History Log (CUoD)

You can display details of the system capacity history log or print the full version of this log. Use option 5 for displaying details or function F6 for printing the system capacity history log. The system capacity history log may be used for keeping record of the days and number of processors already used or left available for future use.

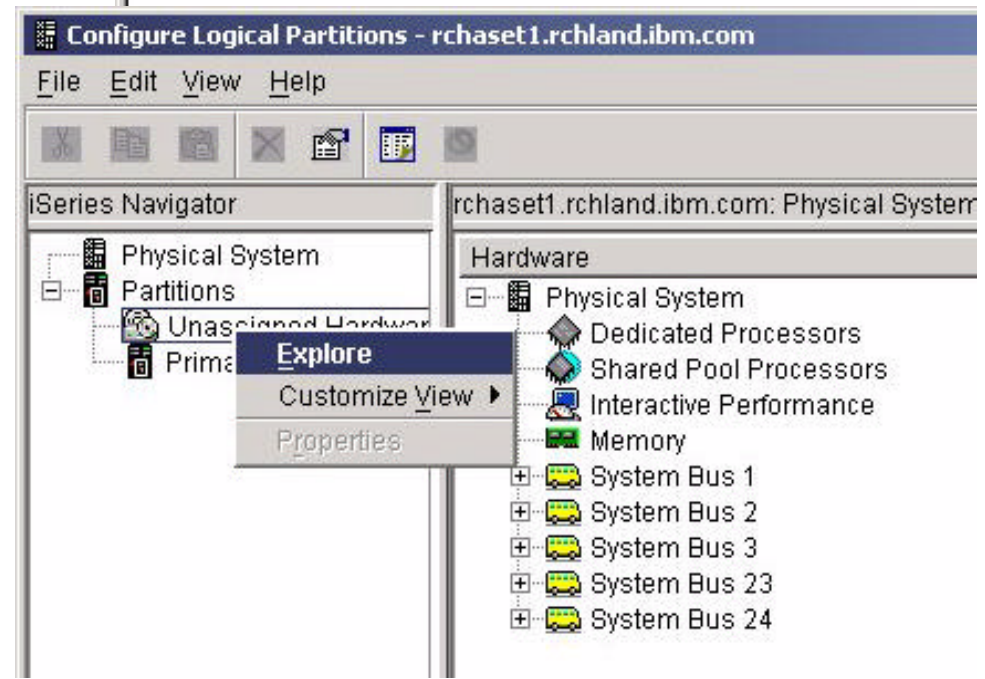
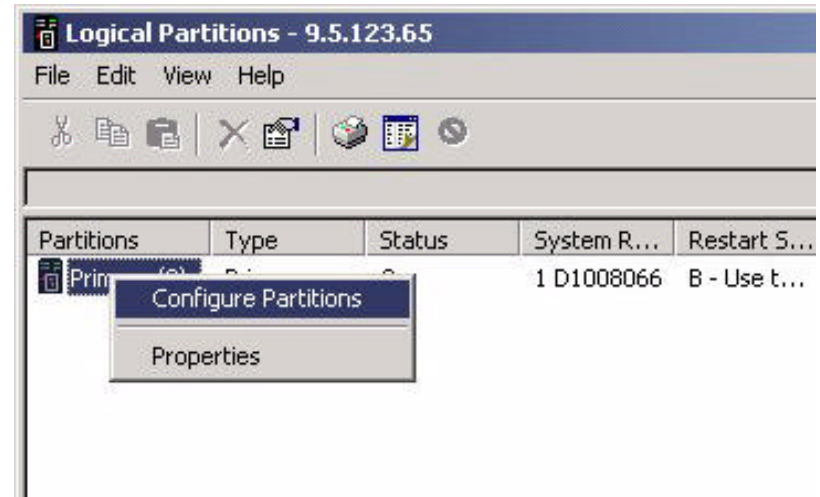
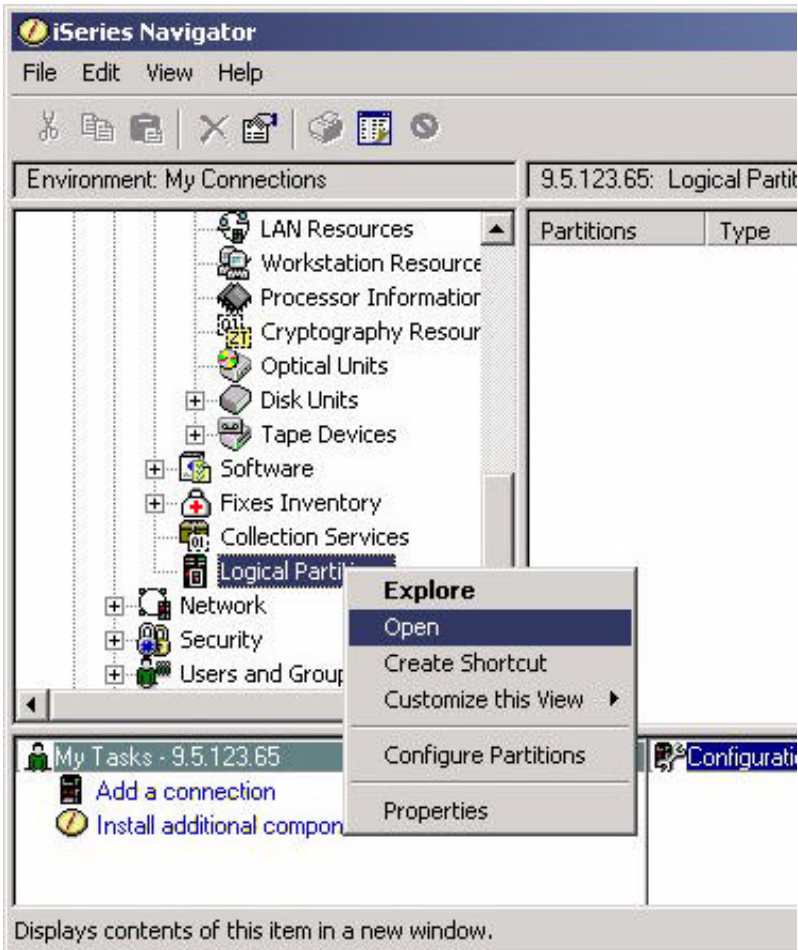
# New Structure: Simplification, all on one single table

MODEL N-WAY	SERVER FEATURE SW TIER	PROCESSOR FEATURE CPW	STANDARD EDITION FEATURE	800 EDITION FEATURE 5250 CPW	ENTERPRISE EDITION FEATURE	CUoD ACTIVATION FEATURE	TEMPORARY ENABLEMENT FEATURE	TEMPORARY BILLING FEATURE
i890 24 - 32	0898 P50	2498 29300 / 37400	7425	N/A	7427	1613	1778	1792
i890 16 - 24	0897 P50	2497 20000 / 29300	7422	N/A	7424	1612	1777	1789
i870 8 - 16	0886 P40	2486 11500 / 20000	7419	N/A	7421	1611	1776	1786
i825 3 - 6	0873 P30	2473 3600 / 6600	7416	N/A	7418	1609	1773	1783
i810 2	0869 P20	2469 2700	7428	N/A	7430	N/A	N/A	N/A
i810 1	0867 P10	2467 1470	7410	N/A	7412	N/A	N/A	N/A
i810 1	0866 P10	2466 1020	7407	N/A	7409	N/A	N/A	N/A
800 1	0865 P10	2464 950	N/A	7408 50	N/A	N/A	N/A	N/A
800 1	0864 P05	2463 300	N/A	7400 25	N/A	N/A	N/A	N/A
800 1	0863 P05	2463 300	N/A	7400 25	N/A	N/A	N/A	N/A

# Notes: Simplification, all on one single table

This single table shown here contains the complete overview of the new family of iSeries servers with the different feature codes and some of the most important specific properties of the different server features that come available with this announcement.

# Partition Management

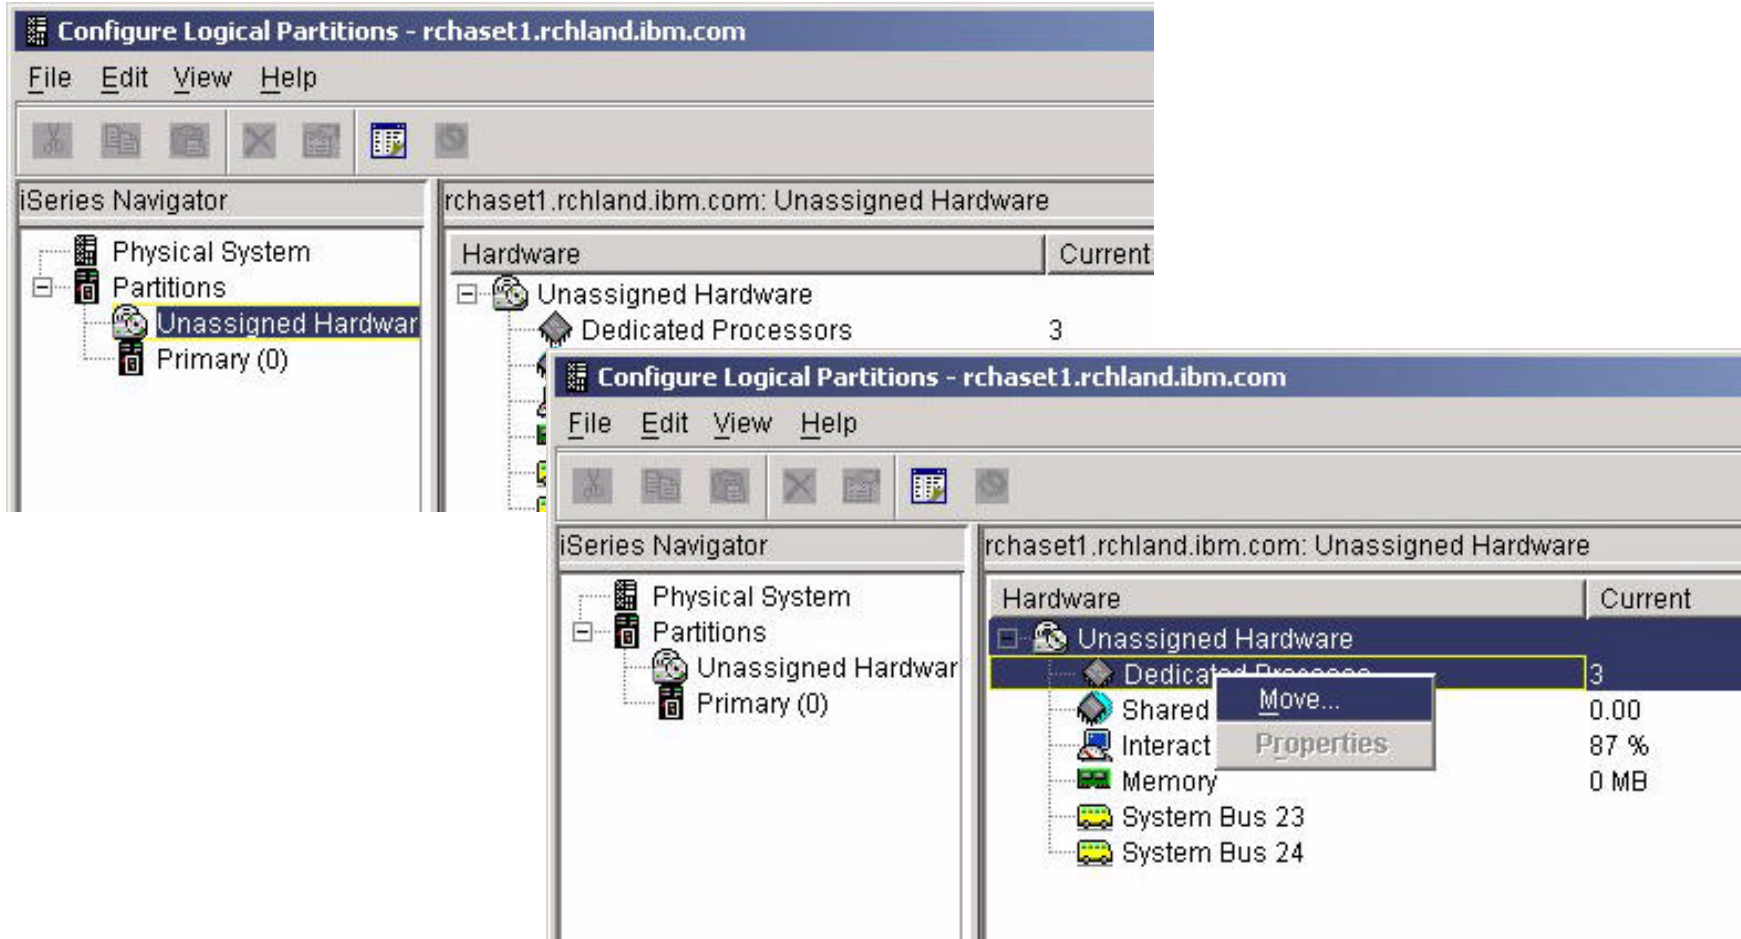


# Notes: Partition Management

You need a service tools user ID and password again to access the iSeries Navigator functions that follow. You may also select to use the green screen service tools functions and options for performing the same actions that will follow, but the iSeries Navigator is more self-explanatory and user friendly, that is why we use those for explaining the next steps.

Start with opening the Logical Partitions item under Configuration and Service of the server. Since we have only a single partition on the server in our example, we select the configure partitions option after a right hand mouse click on the partition. We can see the Unassigned hardware and the primary partition items under the system partition. If we now explore Unassigned hardware we can see the physical system configuration.

# Partition Management -2



## Notes: Partition Management -2

Since we have activated three additional processors for use, these processors appear in the current unassigned hardware. Unassigned hardware is eligible for move. With a right hand mouse click we can open the option list for the unassigned processors and then select the "move" function.

# Move Processing Power

**Move Processing Power**

Move from logical partition "Unassigned Dedicated"

	Current	After Move	Amount to Move
Dedicated processors:	3	0	3 processing units
Interactive performance:	87	87	0 percent

Move to

Logical partition: Primary (0)

	Current	After Move
Dedicated processors:	8	11 processing units
Interactive performance:		

**Configure Logical Partitions - rchaset1.rchland.ibm.com**

File Edit View Help

iSeries Navigator

- Physical System
  - Partitions
    - Unassigned Hardware
      - Primary (0)
        - Explore
        - Customize View
        - Properties

rchaset1.rchland.ibm.com: Primary (0)

Hardware	Current	Pending
Primary (0)	On	
Dedicated Processors	8	11
Interactive Performance	13 %	13 %
Memory	65536 MB	65536 MB
System Bus 1		Dedicated
System Bus 2		Dedicated
System Bus 3		Dedicated
System Bus 22		Shared



# Notes: Move Processing Power

The Move Processing Power window appears and we are now able to start the internal process for moving the activated processors from unassigned hardware to a selected partition. In a partitioned environment, it is also possible to move partial processor to partitions that use shared processors.

In order to deactivate the processors in a decent way, you should select the correct number of processors or 'processor units' for movement from a partition(s) to unassigned hardware so that the number of processor units in Unassigned hardware matches the number of activated CUoD(temporary) processors. The next step to take will then be to take the stop option in the "Work with Temporary System Capacity" screen of the Service Tools. After this step, display and print the system capacity history log.

To stop future use of temporary capacity, a TCOd Termination Code must be applied to the server. Once you obtained a termination key, you should use the same options as for introducing the enablement key in order to terminate the CUoD(temporary) option on your server.

# Agenda

- Server considerations
- Main Storage
- Expansion and IO
- High Speed Loop
- Capacity Upgrade on Demand
- **Upgrade scenarios**
  - LPAR
  - Linux
  - Integrated Windows Environment
  - SAN and External devices
  - Software considerations

# Upgrade Paths- Summary

	<i>800</i>	<i>810</i>	<i>825</i>	<i>870</i>	<i>890</i>
<i>270</i>		X			
<i>720</i>		X	X		
<i>730</i>		X	X	X	
<i>740</i>			X	X	X
<i>820</i>		X	X	X	
<i>830</i>			X	X	X
<i>840</i>				X	X
<i>800</i>	X				
<i>810</i>		X			
<i>825</i>			X	X	
<i>870</i>				X	X
<i>890</i>					X

## Model Upgrades to new models

The supported model upgrades to new models can be broadly classified as:

- 270 to 810
- 800 to 800
- 810 to 810
- 7xx to 810, 825, 870, 890
- 820,830,840, 825, 870, 890 to 825, 870, 890
- Upgrades to rack mounted servers

**Note:** For supported upgrade paths, refer Upgrade paths in Hardware overview.

The methods for achieving these model upgrades are:

- Load Source Migration
- Convert or relocate disks
- Unload/Reload
- Side-by-side

**Generally you should carry out a Solution Assurance for any upgrade**

## General upgrade preliminary tasks

- Ensure installed release at latest fix level
- Install V5R2 (Feb 2003 level) before full system saves
- Full system saves before and after upgrade
  - ▶ At least two saves
  - ▶ Full tape clean before and after each save
- Ensure there is space for relocated and new components
  - ▶ CD, disks, IOAs

# Upgrade/Migration Scenarios

## ■ Example-1

- ▶ Customer has a 270 (Processor feature #2434) system with 6 disks on a #9767 controller with mirrored protection. Existing server is running V5R2 (Feb 2003 Level).
- ▶ Customer would like to upgrade to an 810 (processor feature #2469)
- ▶ RAID parity protection in the new server, is a requirement. New DASD controller #2757 ordered.
- ▶ Additional Ethernet IOA ordered.

# Upgrade/Migration Scenarios

- Upgrade Tasks:
  - ▶ Preliminary tasks
  - ▶ Stop Mirrored protection on the source system
  - ▶ Processor and Backplane replaced in the new machine.
  - ▶ Install old memory modules (Same memory modules supported) + any new ones
  - ▶ New controller #2757 installed on the new backplane. (#9767 is not supported on 810.)
  - ▶ New Ethernet IOA installed.
  - ▶ Disks, CDROM, tape unit (not 13GB) and other IOPs may be migrated.
    - Migration depends on space/cost
  - ▶ At IPL RAID protection started on the new machine
  - ▶ Old Processor/Backplane returned to IBM.
  
- This completes the upgrade activity.

## Upgrade/Migration Scenarios:

- Example - 2:
  - ▶ Customer has a 270 (#2452 processor) and wants to replace with a 810(#2469 processor feature)
  - ▶ Customer has five disks on the source machine with RAID on a #4748 controller.
  - ▶ Target machine has a new controller #2757 IOA.



## Upgrade/Migration Scenarios:

- Migration tasks:
  - ▶ Preliminary tasks
  - ▶ Upgrade not possible. New machine is brought in.
  - ▶ Old memory cards not supported. New memory installed on the machine.
  - ▶ DASD units removed from the old machine and installed on #2757 on the 810.
  - ▶ CDRom/DVD/Tape units migrated from 270 to 810 provided there is space.
  - ▶ IOPs/IOAs migrated from 270 to 810 [provided there is space.
  - ▶ During first IPL, RAID is restructured for each existing RAID set under a new #2757.
    - NOTE: during this restructure activity, data is unprotected.
  - ▶ Old 270 remains with customer.
- This completes the upgrade activity. FULL SYSTEM SAVE

# Upgrade/Migration Scenarios

- Example - 3:
  - ▶ Customer has an AS/400 720 machine with #2748 controller with an 8GB load source. Customer upgrading to an 825 machine with 17.5GB load source. Customer has an expansion tower 5065 connected to the source machine.
  - ▶ Current server is at V5R2. This causes an additional step to install and test software APARs.
- Possible scenarios depending planned disk configuration changes
  - ▶ Many newer 10k RPM drives in base that will be reused
  - ▶ Mixed disk drives configuration in base (size and speed)
  - ▶ Few or no 10k RPM drives in base that will be reused
- Mixing disk sizes on target machines may impact performance

# Upgrade/Migration Scenarios

- Upgrade steps with **load source migration**:
  - ▶ Preliminary tasks
  - ▶ Install 5065 or 5066 and drives to meet planned needs
    - One drive to remain non-configured for Load Source Migration
  - ▶ Roll-in/Roll-out upgrade. All new memory DIMMs. Old memory not supported.
  - ▶ Move user data from drives to be removed and spread across ASP
    - The ASP includes the new drives in the 5065 tower
  - ▶ Copy load source to non-configured disk unit in 5065
  - ▶ Convert 5065 to 5074 (replace IOPs/IOAs as required)
  - ▶ Connect 5074 to new 825.
  - ▶ Convert feature CDROM/tape to #46xx before migrating to 825 CEC.
    - Migrate/convert supported PCI IOAs.
    - Migrate/convert supported DASD features.
    - Migrate load source with LSM utility
  - ▶ Return old 720 processor, frame to IBM.
- This completes the upgrade activity. FULL SYSTEM SAVE

# Upgrade/Migration Scenarios

- Upgrade steps **with disk relocation** (no load source change):
  - ▶ Perform preliminary tasks
  - ▶ Install 5065 or 5066 to meet planned needs
  - ▶ Roll-in/Roll-out upgrade. All new memory DIMMs. Old memory not supported.
  - ▶ Move user data from drives to be removed across ASP
    - This includes the new drives in 5065
  - ▶ Convert 5065 to 5074 (replace IOPs/IOAs as required)
  - ▶ Connect 5074 to new 825.
  - ▶ Convert feature CDRROM/tape to #46xx before migrating to 825 CEC.
    - Migrate/convert supported PCI IOAs.
    - Convert supported DASD features.
  - ▶ IPL from converted Load Source Unit
  - ▶ Return old 720 processor, frame to IBM
- This completes the activity. FULL SYSTEM SAVE

## 820, 830, 840 to 825, 870, 890 upgrade

- Very simple upgrades
- Carry out preliminary tasks
- Power down
- Unplug disks and IO from old CEC
- Plug disks and IO into new CEC
- Mixing disk sizes on target machines may impact performance

# Agenda

- Server considerations
- Main Storage
- Expansion and IO
- High Speed Loop
- Capacity Upgrade on Demand
- Upgrade scenarios
- **LPAR**
  - Linux
  - Integrated Windows Environment
  - SAN and External devices
  - Software consideratons

# LPAR

- Planning and information gathering is vital before LPAR upgrade
- Existing 820, 830, 840 to 820, 830, 840 follow existing V5 rules
- 270 to 810
- 820, 830, 840, 890 to 825, 870, 890
  - ▶ All partitions V5R2 (Feb 2003 level)
  - ▶ Simple plug and go **if** P0 owns all resources in system unit
    - After automatic bus renumbering, fix bus ownership
    - Add new resources to partitions as required
    - Adapt partition to new processing resources
- 7xx to 825, 870, 890
  - ▶ All partitions V5R2 (Feb 2003 level)
  - ▶ DMWA to move user data to new or existing disk units in PCI towers
  - ▶ Upgrade to new hardware
    - Simple load source migration
    - After automatic bus renumbering, fix bus ownership
    - Recreate LPAR and adapt partitions to new processing resources


# LPAR Upgrade considerations

- Can be simple - many simple tasks, overall complex process
- Ensure all partitions at most current release
- May be Roll-in/Roll-out
- Fix resource allocations and bus ownership
- Assign new features, memory, processors
- Moving to and from CUoD systems
  - ▶ Processor activation
  - ▶ Memory installation
  - ▶ May need IPL to make changes to P0 Min/Max
- Upgrade to Model 825, Ethernet LAN console only on P0




# 7xx to 8xx example

- Simple upgrade process for complex environment
- OS/400 to V5R2 (Feb 2003 announce)
- Two LPARs (Primary and Secondary)
- Migrate storage to PCI tower
- Copy load source
- Convert 5066 to 5079
- Copy load source to new server
- Bus renumbering
- Recreate LPAR config
- Reallocate resources



# 8xx to 8xx LPAR upgrade

- Even simpler upgrade .....
- Partitions to V5R2 (Feb 2003 Level)
- Add new tower to existing system
- Remove SPD hardware
- Unplug disk from Model 820
- Plug disk into Model 825
- Reallocate resources
- Go .....
- If new disk IOA you must remember RAID change



## Side-by-side

- Existing side-by-side P847156
  - ▶ RPQ with a fee
  - ▶ Up to 56 days
  - ▶ Limited availability for upgrades to pre-Feb 2003 Models 830 and 840
  
- New Side-by-side P847156
  - ▶ Information Only RPQ
  - ▶ Only available through a special bid approval process -requires a negotiated price for temporary software license, and services.
  - ▶ Up to 56 days

# Logical Partitioning

iSeries	Max # of Processors	Processor Sharing	OS/400 Partitions per Processor	Linux Partitions per Processor	Max # of OS/400 Partitions per Server	Max # of Linux Partitions per Server	Max # of Partitions per Server
800	1	Y	4	9	4	9	10
i810	2	Y	4	10	8	19	20
i825	6	Y	10	10	32	31	32
i870	16	Y	10	10	32	31	32
i890	32	Y	10	10	32	31	32

# Notes: Logical Partitioning

OS/400 V5R1 was a breakthrough release for iSeries logical partitioning (LPAR) with the introduction of dynamic movement of processor and other I/O resources, plus the ability to create partitions of less than one processor unit. Since the introduction of LPAR, over 44% of i840 customers have exploited this mainframe-class technology and since V5R1 partitioning usage has soared on i820 uni-processor servers. The exceptional granularity of being able to move 100th of a processing unit between partitions is another example of how the iSeries continues to exploit the concept of resource virtualization - for memory, disk storage, or processors.

OS/400 V5R2 extends dynamic partitioning to Linux with the ability to move processing units between OS/400 and Linux partitions or between multiple Linux partitions. It also includes support to create and manage Linux partitions through iSeries Navigator.

V5R2 now supports up to 4 or 10 OS/400 partitions (see the next foil for more details) per processor with a total of up to 10 partitions per processor. That is, a single processors can support up to 10 Linux partitions. You can use combinations of OS/400 and Linux partitions per processors. 10 times the number of processors, minus the number of OS/400 partitions, equal the number of Linux partitions allowed. Maximum of 32 partitions on a system, with up to 31 Linux partitions.

OS/400 V5R2 also allows customers better manage their partition configuration data with new support to save partition configuration data into an HTML file for storage with other business continuity information.

A key enhancement to the flexibility of LPAR is the new sub-capacity pricing for WebSphere Commerce Suite, and WebSphere Application Server. Previously, WebSphere Application Server was priced per processor, so that on a 32-way machine, you would pay 32 times the charge for a uni-processor. Now, with sub-capacity pricing, you just pay for the number of processors that you are using in partitions running WebSphere (rounded up to the next full processor where you are using partial processor partitioning).

For example, if a 4-way i825 has WebSphere Advanced Application Server running in two partitions sized at 0.9 processor units and 0.8 processing units respectively, the total number of processing units allocated will be 1.7. Customers would require a license for 2 processors and will have the flexibility of moving resources between two WebSphere partitions as long as it does not exceed a total of 2 full processors.

This more flexible approach to pricing clearly benefits customers who want to deploy WebSphere workloads on a large consolidated iSeries system where only fewer number of processors are required to meet their web transaction serving requirements. A future release of WebSphere Application Server (after version 4) will monitor number of processors usage compared to license agreement and issue message.

# Upgrading HA environments

# Upgrading in HA

- Two types of HA environments:
  - ▶ Full replication environment
  - ▶ iASP switchable environment
  
- Planning is essential
  - ▶ You must test
  - ▶ Will business accept no protection during upgrade?
    - If no, third system or LPAR will be required to maintain protection
  - ▶ Operating system environment
    - Create switchable test environment
  - ▶ Application upgrade
    - Create switchable test environment - LPAR
    - Is application available at new release?
    - With a data migration program be required?

# HA Replication Upgrade

- Ideally need three servers for full protection
- Upgrade process (two servers)
  - ▶ Plan and test on development server or LPARs
  - ▶ end clustering on backup
  - ▶ upgrade backup system
  - ▶ Add backup system into cluster
  - ▶ resync
  - ▶ switch over (First production down time)
  - ▶ upgrade primary
  - ▶ add primary system back into cluster
  - ▶ resync
  - ▶ switch back to primary (Second production downtime)
  - ▶ Downtime at minimum (potentially 30 minutes total)





# Upgrading Independent ASPs

# Upgrade iASPs

- No change to support for iASP at V5R2 (Feb 2003 level)
- V5R1 to V5R2 upgrade
  - ▶ Simple upgrade, all conversion automatic at vary on
  - ▶ Cannot go back once V5R1 iASP varied on to V5R2 operating system
- Switchable iASPs
  - ▶ create test environment
  - ▶ switch to backup server
  - ▶ upgrade primary
  - ▶ switch back
  - ▶ upgrade backup server

# Switchable iASP upgrade

- Ideally need three servers for full protection
- Upgrade process (two servers)
  - ▶ Plan and test on development server or LPARs
  - ▶ remove backup from cluster
  - ▶ upgrade backup server
  - ▶ add backup into cluster
  - ▶ switch over iASP (first downtime)
  - ▶ remove primary from cluster
  - ▶ upgrade primary
  - ▶ add primary back into cluster
  - ▶ switch over iASP (second downtime)
  - ▶ Downtime at minimum (potentially 90 minutes)

# Agenda

- Server considerations
- Main Storage
- Expansion and IO
- High Speed Loop
- Capacity Upgrade on Demand
- Upgrade scenarios
- LPAR
- **Linux**
- Integrated Windows Environment
- SAN and External devices
- Software considerations

# Upgrading Linux Partitions to V5R2

- 2766 Fiber Channel Disk adapter support for external disk
  - ▶ The LP9002L PowerPC Linux device driver
    - available from Emulex Corporation
    - Distribution v4.20p support
      - Turbolinux Server 7.0 and SuSE Linux Enterprise Server 7 **Yes**
      - Red Hat Linux 7.1 **No**
  - ▶ No support for 2765 Tape adapter
- Up to 10 Linux partitions per iSeries processor
  - ▶ 1-way processor = 9 Linux partitions
  - ▶ 2-way processor = 19 Linux partitions
- 32bit to 64bit upgrades - V5R1 to V5R2
- New IOA support - Included in SLES8 or driver update to SLES7
- Linux version upgrades - no upgrade path, re-install Linux

# Linux Native Adapter Support

- Native Linux support for:
  - ▶ New DASD adapters
    - #0618 for the DASD adapter #2757
    - #0619 for the DASD adapter #2782
  - ▶ New external tape and optical device adapter
    - #0624 for the #5702 adapter
  - ▶ V5R2 (Jan. 2003) required
- i825, i870 and i890
  - ▶ POWER4 processors provide support for 64-bit Linux kernels only



# Notes: Linux Native Adapter Support

A number of new direct attachment features for Linux are available with V5R2. Cards being controlled by the Linux operating system do not use/require PCI IOPs. Linux direct attach PCI cards are supported only in a secondary LPAR partition. The existing direct attachments for Linux with V5R1 were limited to features #0607, #0601 and #0602 100/10Mbps Ethernet and 1Gb Ethernet adapters, #0603 100/10Mbps Token Ring adapter and the #0604, #0605 and #0606 Disk adapters. The addition of the direct attachment features mentioned in this foil largely extend the possibilities to directly address a much wider set of device attachments directly from the Linux environment on the iSeries server.

# Agenda

- Server considerations
- Main Storage
- Expansion and IO
- High Speed Loop
- Capacity Upgrade on Demand
- Upgrade scenarios
- LPAR
- Linux
- **Integrated Windows Environment**
- SAN and External devices
- Software considerations



# Upgrading integrated windows environments

- Coming from older internal PC servers
  - ▶ FSIOP, IPCS, INS, IXS to new IXS or IXA
- The latest 1.6GHz Integrated xSeries of iSeries (IXS)
  - ▶ Only supports Windows 2000 Server and Advanced Server
  - ▶ You must perform new Windows install to migrate
  - ▶ Windows PnP detects new hardware
- From existing Models 270, 8xx with IXA to new models of 8xx
  - ▶ HSL-2 supports IXA, but need HSL-2 to HSL conversion cable
  - ▶ No specific upgrade requirements
- When in HSL loop with migration tower, IXA position is critical
- Enterprise Edition includes a IXS

# Agenda

- Server considerations
- Main Storage
- Expansion and IO
- High Speed Loop
- Capacity Upgrade on Demand
- Upgrade scenarios
- LPAR
- Linux
- Integrated Windows Environment
- **SAN and External devices**
- Software consideratons

# SAN and upgrades

- Automatic switch from 1GB/sec (V5R1) to 2 GB/sec (V5R2)
- New #2757 IOP supports existing Fiber Channel adapters
  - ▶ No performance improvement
- Switch support
  - ▶ No change since 2002
- SAN unaware of server upgrade
  - ▶ Just add any new Fibre Channel IOAs
  - ▶ Connect to SAN device (tape or disk)
  - ▶ New LUNs appear as non-configured disk

# Ultrium Stage-2 and Second Generation LTO

- 200 GB per cartridge and 35 MB / sec data transfer rate (uncompressed)
- Ultrium Stage-2 implemented in
  - ▶ 3580 L23 and 3580 H23
  - ▶ 3584 with tape drive canister features
    - # 1474 (LVD SCSI)
    - # 1475 (HVD SCSI)
    - # 1476 (Fibre)
- LVD SCSI connect to
  - ▶ # 5702 and External port of # 5705
- HVD SCSI connects to # 2749
- Fibre channel connects to # 2765
- Ultrium stage 2 drives can **read and write** Ultrium stage 1 media with improved performance



# Notes: Ultrium-2 and Second Generation LTO

IBM is announcing the new Linear Tape-Open (LTO) IBM TOTALSTORAGE ULTRIUM-2 Tape Drives, which more than doubles tape drive performance over the LTO Ultrium Tape Drive (Ultrium 1) with up to 35 MB/sec native data transfer rate (70 MB/sec with 2:1 compression). In addition, with the use of the new IBM TotalStorage LTO Ultrium 200 GB Data Cartridge, the Ultrium 2 Tape Drive doubles the tape cartridge capacity up to 200 GB native capacity (400 GB with 2:1 compression).

**3580 MODELS H23 AND L23** offer high capacity, performance, and technology designed for the midrange open systems environment. The Model L23 comes with a SCSI Ultra160 LVD attachment, while the Model H23 comes with a UltraSCSI HVD attachment.

The IBM LTO **Tape Library 3584** now supports the IBM TOTALSTORAGE ULTRIUM-2 Tape Drives. These IBM Ultrium 2 Tape Drives come with SCSI Ultra160 LVD, Ultra SCSI HVD, or switched fabric 2 Gbps Fibre Channel attachments for connection to open system servers including the iSeries.

The 3580 Model H23 and the UltraScalable Tape Library 3584, with the LTO Ultrium 2 HVDDrive Canister (#1475), can be attached to iSeries and AS/400 models that support Version 5 Release 1 (V5R1), or later (depending on adapter type), and the following adapters:

- Feature 2729 - PCI Magnetic Media Controller
- Feature 2749 - PCI Ultra Magnetic Media Controller
- Feature 6534, Magnetic Media Controller

The 3580 Model L23 and the UltraScalable Tape Library 3584, with the LTO Ultrium 2 LVD Drive Canister (#1474), can be attached to iSeries and AS/400 models that support Version 5 Release 2 or later (depending on adapter type), and the following adapters :

- Feature 5702 - PCI-X Ultra Tape Controller
- Feature 5705 - PCI-X Ultra Tape Controller (external Ultra SCSI port)

The UltraScalable Tape Library 3584, with the LTO Ultrium 2 Fibre Drive Canister (#1476), can be attached to iSeries models that support Version 5 Release 1 (V5R1), or later (depending on adapter type), and the following adapter:

- PCI Fibre Channel Tape Controller (#2765)

**PRODUCT PREVIEW** : IBM intends to further introduce and integrate the IBM TotalStorage Ultrium 2 Tape Drive technology in additional or alternative IBM LTO Ultrium 358x products. Previews provide insight into IBM plans and direction. Specific availability dates, ordering information, and terms and conditions will be provided when the product is announced.


**The IBM Ultrium 2 Tape Drives can read and write original LTO Ultrium Data Cartridges at original Ultrium-1 capacities and with improved performance up to 20 MB/sec native data transfer rate (40 MB/sec with 2:1 compression).**

# Tape

- New IOAs
- SAN support for tape
  - ▶ Great support
    - distance, up to 10km for Tape Vaulting
    - flexibility, zoning allows multi-system attach to single subsystem
    - speed, 1GB/sec or 2GB/sec
  - ▶ No change to #2765 support from 2002

## External Disk

- No change to #2766 support from 2002
- Multiple disk targets 32 devices per #2766
- #2766 supports 1Gbit (V5R1) or 2 Gbit (V5R2)
- New IBM TotalStorage Enterprise Storage Server F800
  - ▶ up to 55.9 TB
  - ▶ up to 64 GB cache
  - ▶ 2 Gbit SAN support
  - ▶ RAID - 5 or RAID -10
  - ▶ 15 K rpm drives
    - 18.2 GB, 36.4 GB and 72.8 GB
  - ▶ Standard or Turbo Processor feature
    - 2 x 4-way SMP for Standard feature
    - 2 x 6-way SMP for Turbo feature



# Notes: Enterprise Storage Server

The ESS is an intelligent disk system designed to provide high-performance high-availability, and high-capacity disk storage. The Model 800 is the third generation of the ESS and maintains the functionality, stability and reliability. Flexible capacity and performance are further enhanced with support for multiple disk drive capacities (18.2, 36.4, 72.8, and 145.6 GB) at 10,000 and (18.2, 36.4GB and 72.8GB) 15,000 rpm. A storage subsystem's ability to scale upward, is in large measure gated by the performance needed to take advantage of the sheer capacity it offers. Matching performance to scalability is thus a key component in delivering world class storage consolidation solutions. More details can be found at the following web site: <http://www.storage.ibm.com/hardsoft/products/ess/ess.htm>

Connectivity with iSeries over 2Gb Fibre Channel is enabled with #2766 under V5R2 for higher performance, while supporting current customer investments by maintaining compatibility with existing 1Gb solutions. All components in the Fibre Channel connection need to be 2Gb enabled to make use of this ability.



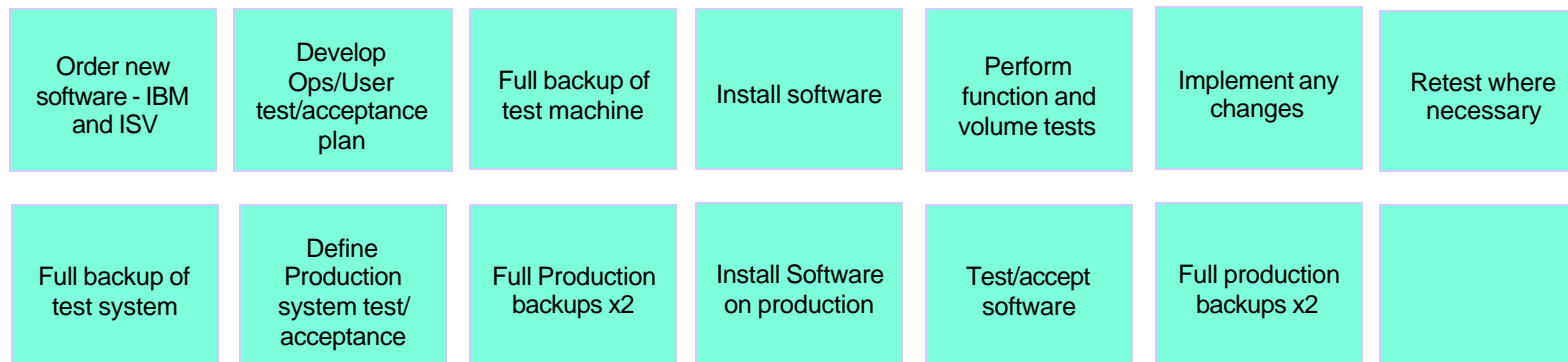
# Agenda

- Server considerations
- Main Storage
- Expansion and IO
- High Speed Loop
- Capacity Upgrade on Demand
- Upgrade scenarios
- LPAR
- Linux
- Integrated Windows Environment
- SAN and External devices

-  **Software consideratons**

# How to lessen the impact of upgrade

- Plan ahead
  - ▶ Services available to assist with planning and implementation
- Move to new software now
- Define test period
  - ▶ Test your applications
- Hardware upgrade should be much simpler



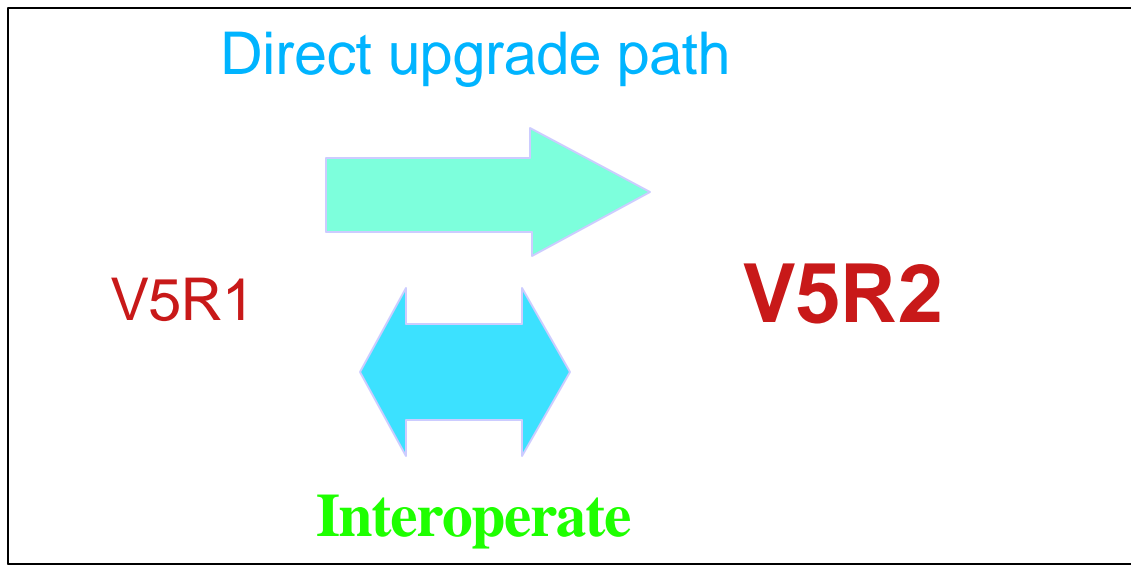

## V5R2 (Feb 2003 level) prereq for new hardware

- New hardware requires February 2003 level of V5R2
  - ▶ Still V5R2, with device-driver-like changes grouped together to support the new hardware
- Info APAR II13365
  - ▶ Software required for new models 800, 810, 825, 870, and 890 including new I/O
  - ▶ Information on new OS/400 software licensing
- Info APAR II13440
  - ▶ Software required to support new I/O and towers on previous models
- Info APARs available 01/20/03 with details on ordering and installing software
- Customers can start preparing now and be ready for the hardware upgrade on 2/21/03
  - ▶ APARs planned to be available by February 3, 2003

# Software documentation availability

- Memo to Users
  - ▶ Available at announce from Information Center home page
  - ▶ <http://www.ibm.com/eserver/series/infocenter>
  - ▶ This will allow you to access national language versions
- iSeries Information Center
  - ▶ Software Installation manual available from the following path
    - Installation, Upgrades, and Migration -> Install the OS/400 release and related software
  - ▶ Upgrade Information available from the following path
    - Installation, Upgrades, and Migration -> Upgrades

# V5R2 Interoperability



## OS/400 Migrations - Earlier Releases

OS400 supported Two-step upgrade process to V5R2:



No OS/400 supported upgrade paths to V5R2:

V3R2

V4R1

V4R2

V4R3

And any earlier OS/400 releases

V5R2 does not support 4xx/5xx models



# Software Maintenance for iSeries

# Software Maintenance Agreement (SWMA) for iSeries:

- Advantages:
  - ▶ Support solution which is simple, consistent and inclusive
  - ▶ Should improve Customer Satisfaction
  - ▶ Access to latest version and release
- How Does it work?
  - ▶ All new purchases include one year SWMA (3 year with delta.)
  - ▶ SWMA is inclusive of
    - Upgrade entitlement (Software Subscription)
    - Remote Technical support (Supportline)
  - ▶ Price of SWMA comparable to combination of current offerings.
- What about existing Software Subscription (SS) and upgrade?
  - ▶ Existing SS contract honored until renewal
  - ▶ Move to SWMA after renewal



# IBM iSeries Service from IGS

# IBM Services for Upgrade and Migration

- Installation and Implementation Services
  - ▶ AS/400 BRMS Installation Services
  - ▶ AS/400 System Transition Services
  - ▶ IBM Implementation Services for iSeries clustering and independent auxiliary storage pools
  - ▶ IBM Implementation Services - for pSeries and iSeries LPAR
  - ▶ Implementation Services - for iSeries security
  - ▶ ServicePac for Installation Services
- Maintenance
  - ▶ Hardware Maintenance Services
- Migration
  - ▶ iSeries Upgrade and Migration Services V5R2
  - ▶ Migration Services for AS/400 - Data Migration
  - ▶ AS/400 Server Consolidation Services
- <http://www.ibm.com/services/its/us/inteapps.html>

## IBM Services for Upgrades and Migrations 2

- High Availability
  - ▶ High Availability Services for AS/400®
  - ▶ High Availability Server Assessment for Small and Medium Business
  
- Operational Support
  - ▶ Account Advocate
  - ▶ Advanced Support
  - ▶ Consult Line
  - ▶ Onsite Technical Advocate
  - ▶ Onsite Technical Assistant
  - ▶ PM/400e
  
- <http://www.ibm.com/services/its/us/inteapps.html>

# iSeries upgrades at a glance

## Your flexible friend - iSeries upgrades

- 820, 830, 840 to 810, 825, 870, 890 are plug and go upgrades -
  - ▶ Roll-in new server
  - ▶ Physically move disk or migrate load source
  - ▶ Plug cables and IPL
  - ▶ Could use new tower (507X or 509X) as migration tower
  
- 7xx to new 810, 825, 870, 890
  - ▶ No SPD or Migration towers
  - ▶ Could use PCI tower + conversion (5065>5074) as migration tower
  - ▶ Use load source migration to save a reload
  
- Must do prep work
  - ▶ Load new release and test
  - ▶ Plan, Plan..... HSL, IOPs, IOAs, etc..
  - ▶ Must maintain RAID sets if moving disks

# Overall upgrade considerations

- Hardware prereqs
  - ▶ Make sure all migrated hardware will be accommodated
- Software prereqs
  - ▶ APAR installation for new models and IO
- Editions
  - ▶ New terms and conditions for new models
- Disk Migration
  - ▶ Remember RAID to RAID
  - ▶ Number of existing disks into new server
- SPD
  - ▶ No SPD hardware or towers on new models
- LPAR
  - ▶ 2 to 3 buses reallocate resources
  - ▶ recreate LPAR config on 7xx LPAR upgrade