International Technical Support Organization

# IBM System z9 - Availability Enhancements

ibm.com/redbooks

Redbooks Workshop

IBM ITSO - International Technical Support Organization






## Notices

This information was developed for products and services offered in the U.S.A.

Note to U.S. Government Users Restricted Rights — Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to: IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

# Agenda

- Positioning
- Technical Details
- I/O and Network Enhancements
- I/O growth solutions

- Availability Enhancements
- HMC/SE changes
- LPAR Management
- Cryptography





# Availability Enhancements

- Enhanced Book Availability
- Redundant I/O Interconnect
- Dynamic STI Management
- Enhanced Driver Maintenance

- MBA Fan-Out
- Dynamic Oscillator Switch
- On Demand Support
- Flexible Memory Option

# z9-109 Dynamic I/O Bus

- **STI Multipath Module (STI-MP)**
  - A Mutiplexer that resides between the MBAs STI link and the various I/O cards on the STI daughter card
  - The key attribute or change for STI-MP is the introduction of a second port. This allows for the alternate I/O path to the processor.
- **Key Usage**
  - Memory Upgrade
  - Dynamic MBA error recovery
  - Book Repair (EBA)
  - STI cable repair
  - MBA card repair

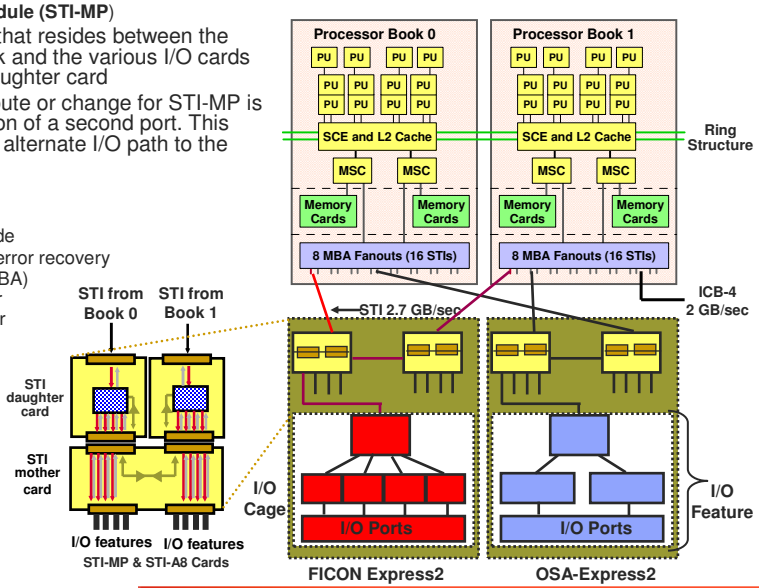






# Redundant I/O Interconnect

- All Adapters / Channels support swap except ICB's
- All other attachments will SWAP with NO Architectural visibility
- New SE Panel provided to Display Multiplexor / Path Status
- New R/V Service Scenarios to exploit Function
- I/O Performance is potentially degraded during the Service Window
  - It may be necessary to advise customers to perform Service "Non-Prime Time".

# Packaging Changes

- **New Passive Mother Card (STI-A8)**
  - For all STI domains except the 7th
- **New Passive Mother Card (STI-A4)**
  - Specially for the 7th domain, compact with PSCN-24 Pinout
  - Only a Single I/O Domain, but 2 STI-MPs plugged
- **New STI-MP Daughter Card**
- **New Full Book Carrier for the PSCN-24 Card**
  - Allows plugging anywhere
  - Plugging rules are restricted to slot 29 (used to be 28)
  - All other capability unaltered.





# Alternate Path I/O Adapter Attach

- **Function Provided**:
  - Support dynamic book maintenance while maintaining system connectivity to the attached I/O
  - Support MBA fanout card and/or STI cable maintenance while maintaining system connectivity to the attached I/O
  - Manually initiated for configuration and maintenance actions
  - System initiated for recovery support
- **Benefits:**
  - No loss of connectivity to I/O from any node .
  - No planning required or dependency on alternate I/O attachment interfaces (e.g. FICON, ESCON,..)
  - Automatic traffic re-routing and return to original configuration upon maintenance completion
  - Book removal does not cause loss of I/O in use by other Books
- **Traditional "Alternate Channel Path Support" was considered as a cheaper solution but rejected.**
  - Not All Devices are "Alternate Path Capable"
  - Cannot Guarantee Customer Configuration
  - CHPID de-configures painful for Customer

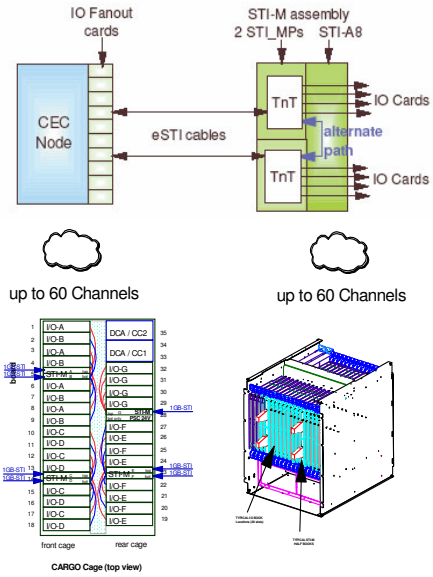

# Cargo I/O cage configuration

Danu *Cargo* Board

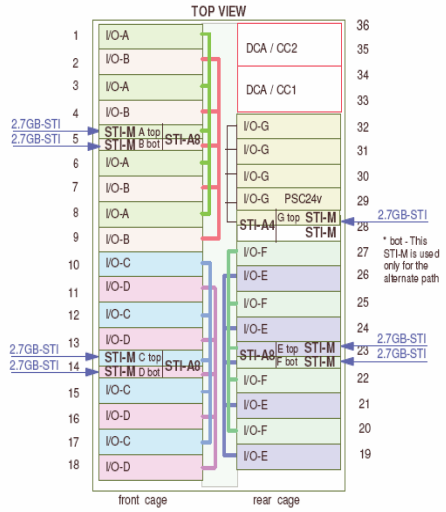


## Domains 1 - 6

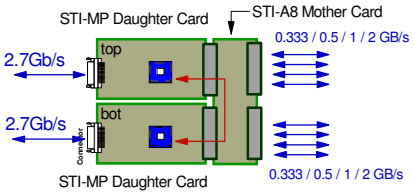


## Domain 7

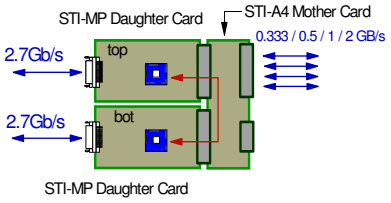






# MBA structure - z990 & z9-109

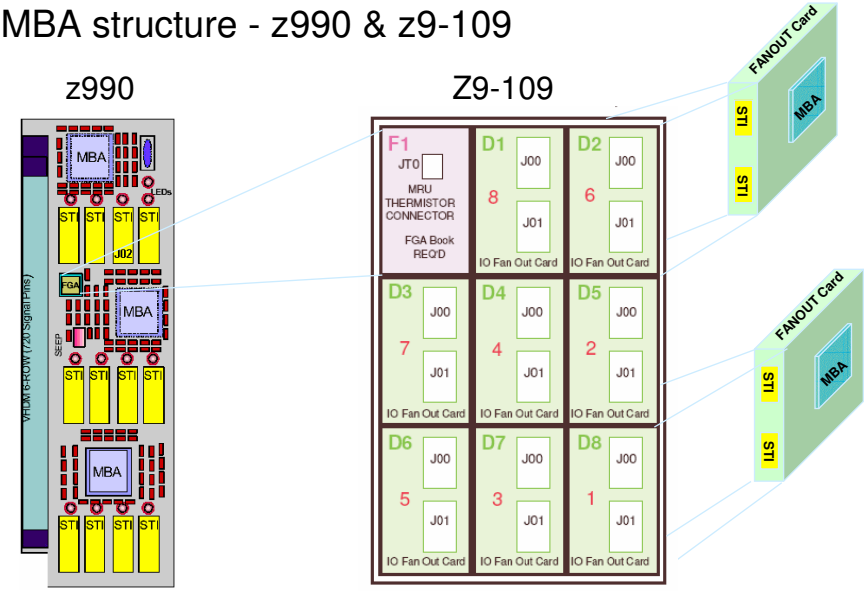

# Dual STI - Terminology

- **Primary Link: The Link Attached from a STI-MP to an MBA**
  - It is the primary, or default path for system traffic for the attached STI-MP, it is also the alternate path for the chained STI-MP
- **Secondary Link: Channel card port link on a STI-MP**
  - There are 4 channel port links per STI-MP card
- **Chained Link: The Link that provides the connection, via the mother-card, between STI-MP pairs**
  - This Link is part of the Alternate path for either STI-MP
- **Primary / Default Path**
  - The path that connects the Primary Link to the Secondary link on the same STI-MP card
- **Alternate Path**
  - The path that uses the Chained Link
- **Active Path: The STI-MP north port (Primary or Chained) that is the path for System Traffic. Normally, this is the primary, but may be the Alternate Path (Chained Link)**
- **Link States**
  - A **Primary** or **Chained** Link can be:
    - Operational
    - Standby (unusual)
    - Fenced or Check stopped
  - A **Chained** Link may be
    - Operational
    - Standby (normal)
    - Fenced or Check stopped

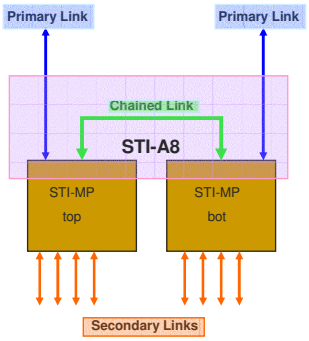






# STI Status Control Panel

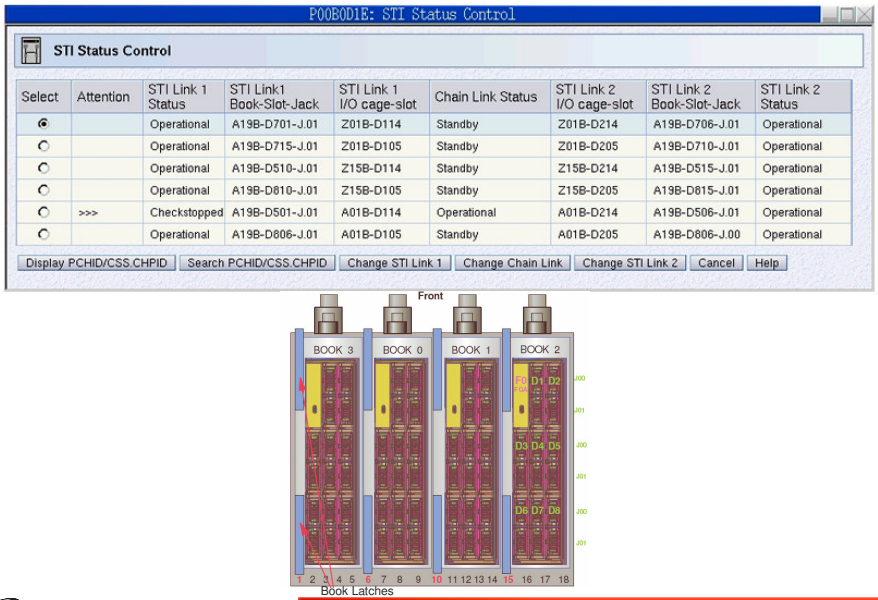

## z990 A08 concurrent upgrade to B16 without FC #2400
## - Concurrent STI reconfiguration not possible

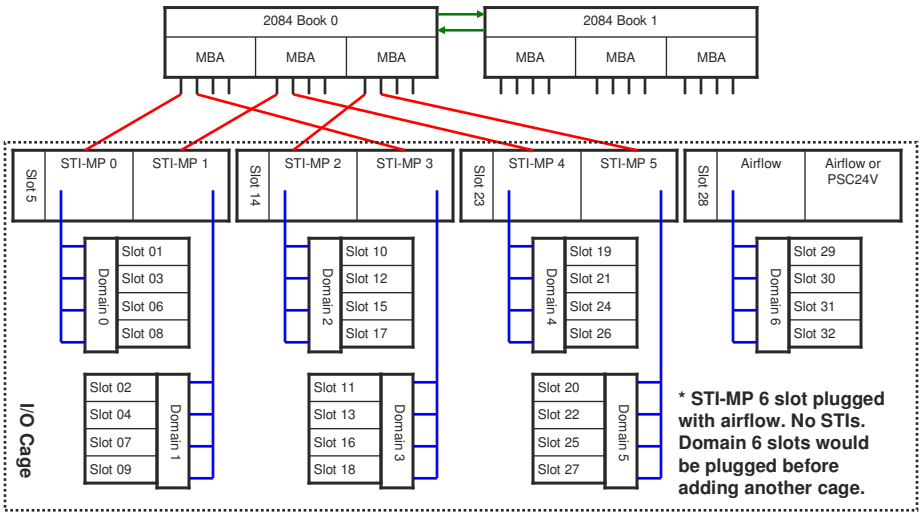






## z990 A08 upgrade to B16
## - Disruptive STI reconfiguration with FC #2400 (optional)

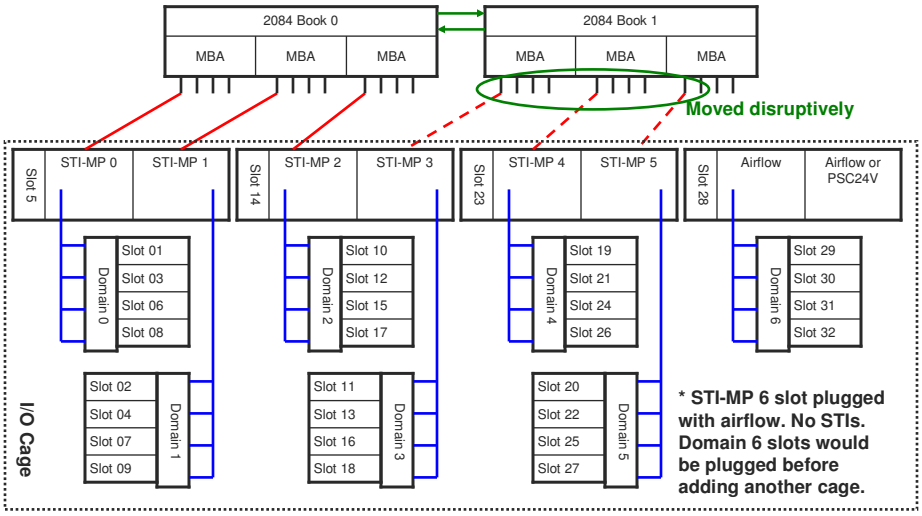

## z9-109 S08 concurrent upgrade to z9-109 S18
## - Concurrent move of MBAs supporting I/O domains

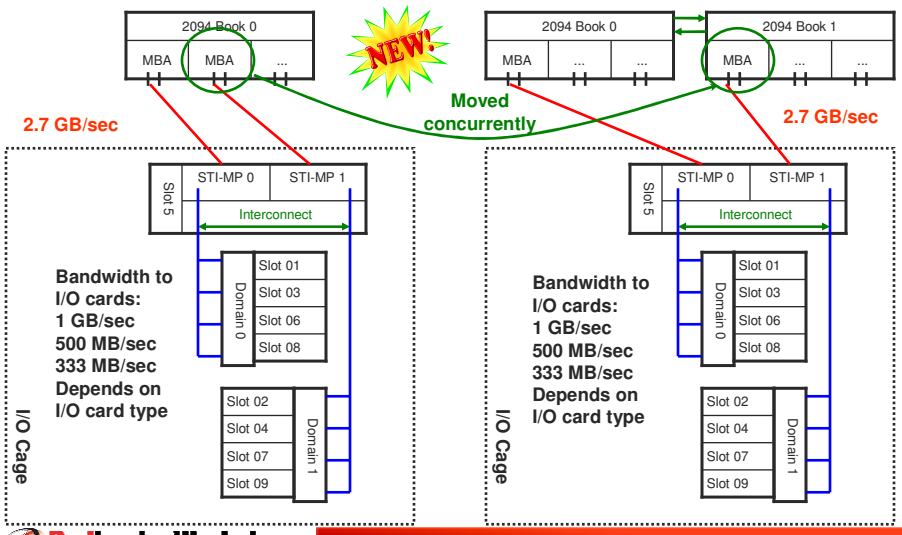






## z9-109 S08 concurrent upgrade to z9-109 S18 (with ICBs)
## - No Concurrent move of MBAs supporting ICB connectivity

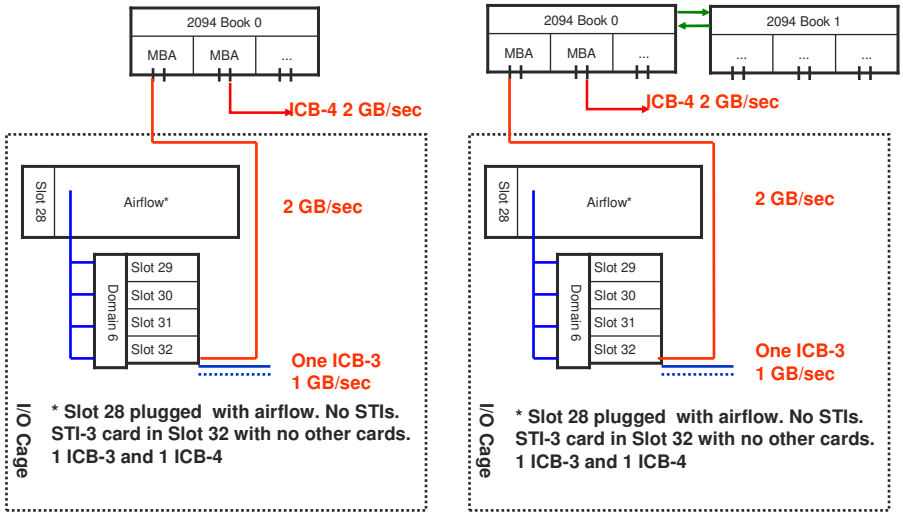

## z9-109 S08 upgrade to S18 (with ICBs)
## - Disruptive move of MBAs with FC #2400 (optional)

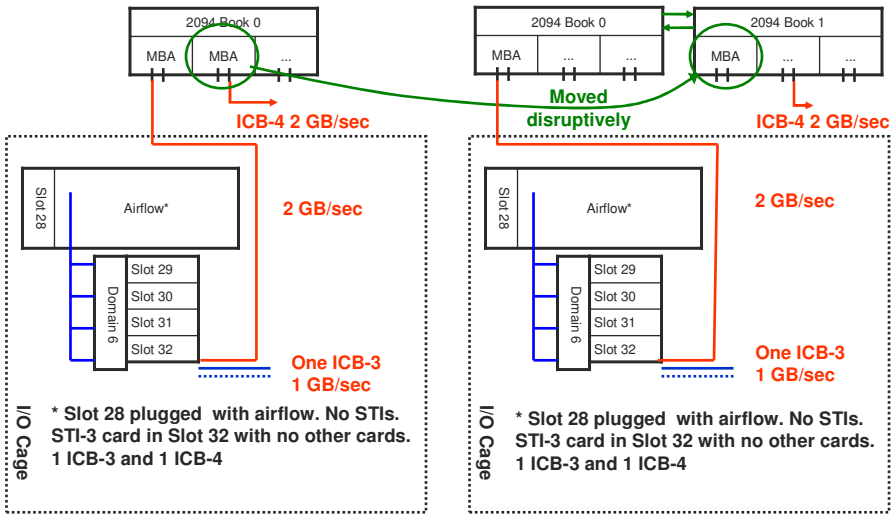






## z990 A08 with 6 domains plugged, one I/O cage

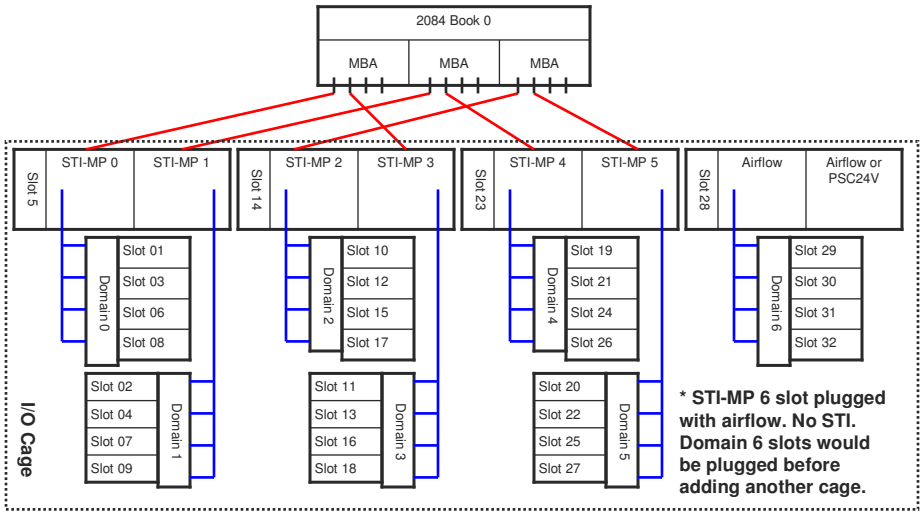

## z990 A08 concurrent upgrade to B16 without FC #2400 - Concurrent STI reconfiguration not possible

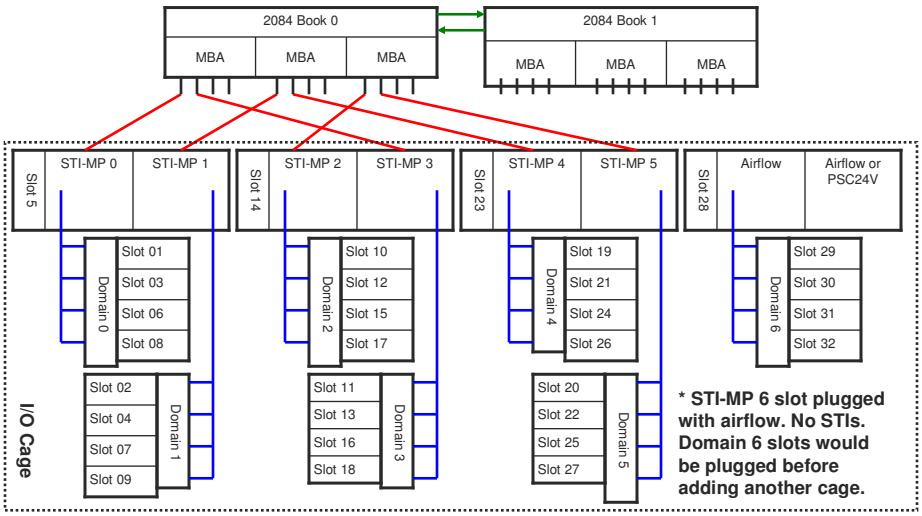


*** STI-MP 6 slot plugged with airflow. No STIs. Domain 6 slots would be plugged before adding another cage.**





## z9-109 S08 with 6 domains plugged, one I/O cage

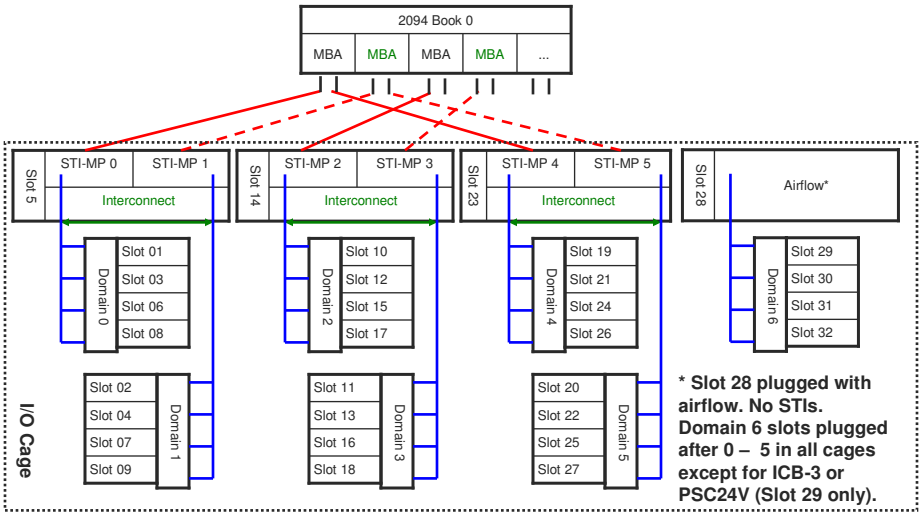


*** Slot 28 plugged with airflow. No STIs. Domain 6 slots plugged after 0 – 5 in all cages except for ICB-3 or PSC24V (Slot 29 only).**

## z9-109 S08 to S18 concurrent upgrade, 6 I/O domains
## - Concurrent move of MBAs supporting I/O domains

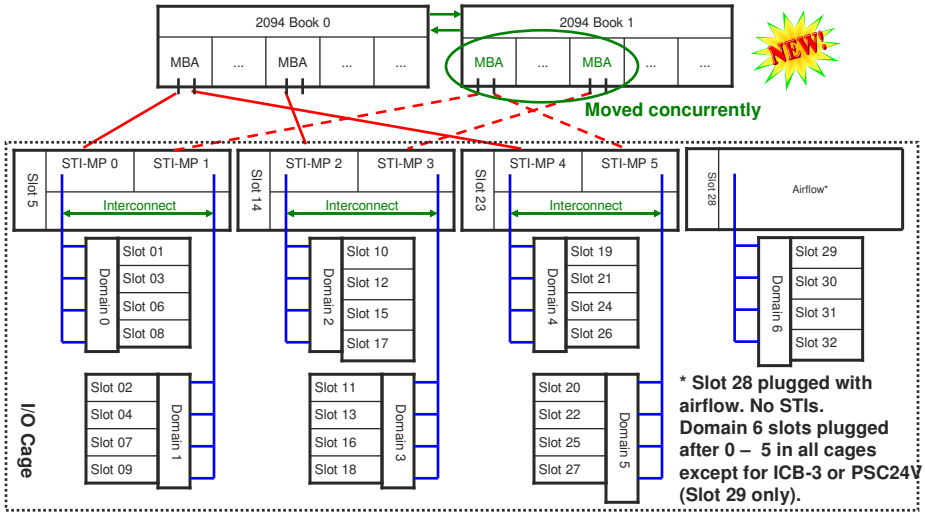






## Availability Enhancements

- Enhanced Book Availability
- Redundant I/O Interconnect
- Dynamic STI Management
- Enhanced Driver Maintenance


- MBA Fan-Out
- Dynamic Oscillator Switch
- On Demand Support
- Flexible Memory Option

## z9-109 Concurrent Memory Upgrades

- **LIC enable additional memory to the physical limit of the installed cards and memory configuration**
  - Designed to be possible and concurrent in many but not all configurations
- **Add a book with additional memory**
  - Designed to be possible except for Models S38 and S54
- **Exploit Enhanced Book Availability to change memory card configuration in existing books**
  - Not possible on Model S08
  - Exploits capability for concurrent book remove, upgrade and return
  - Designed to be possible with flexible memory and PU configurations
  - May be possible with standard memory and PU configurations depending on LPAR resources configuration

NEW!

Note: Concurrent memory upgrades above are designed not to require CEC activation (POR).
z/OS with "reserved memory" configured in the LPAR profile can add memory to a running partition. Otherwise adding memory to a partition requires deactivation, profile change and activation of the partition, which is designed to be disruptive to that partition only.





## Availability Enhancements

Enhanced Book Availability

Redundant I/O Interconnect

Dynamic STI Management

Enhanced Driver Maintenance


MBA Fan-Out

Dynamic Oscillator Switch

On Demand Support

Flexible Memory Option

## MBA Fan-Out Card

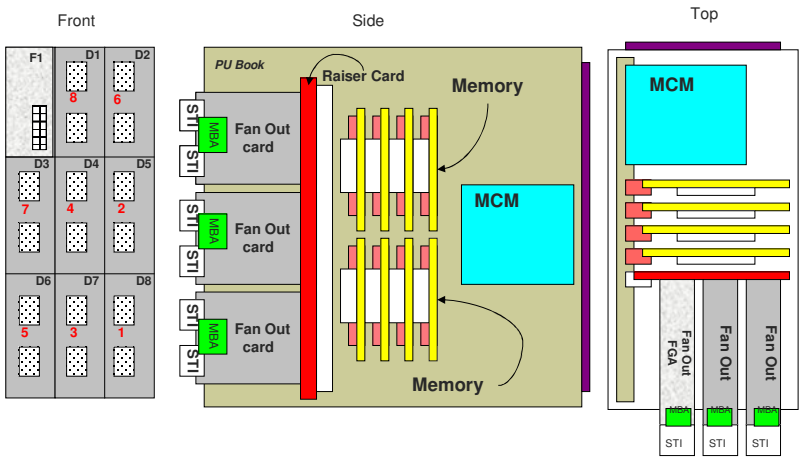






## MBA Fan-Out Card…

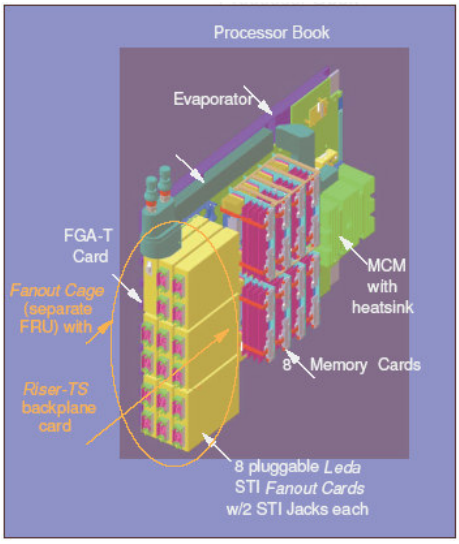


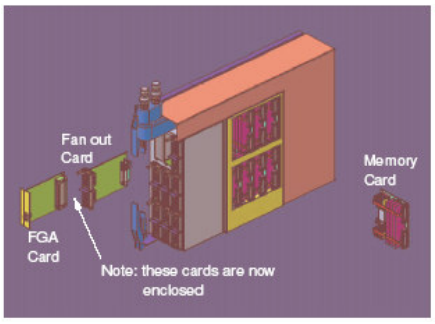

# Availability Enhancements

Enhanced Book Availability

Redundant I/O Interconnect

Dynamic STI Management

Enhanced Driver Maintenance

MBA Fan-Out

Dynamic Oscillator Switch

On Demand Support

Flexible Memory Option





# Concurrent Book Functions Supported in z9-109

- **Concurrent Book Add (z990 GA2)**
  - Concurrent Model Upgrade by the addition of a single new book adding physical processors, memory, and I/O Connections.
- **Continued Capacity with Fenced Book**
  - Make use of the LICCC defined resources of the fenced book to allocate physical resources on the operational books as possible.
- **“Cold” Concurrent Book Repair**
  - Concurrent replacement of a defective book when that book had been previously fenced from the system during the last IML.
- **Concurrent Book Replacement**
  - Concurrent removal and replacement of a book for either repair or upgrade (“Hot” Concurrent Book Repair).

# Concurrent Book Replacement Flow

- Prepare for Concurrent Book Replacement
  - Standalone SE Application
  - Customer Workload Reduction
  - Single Path I/O Connectivity Removal
- Perform Concurrent Book Replacement
  - Verification of Prepared System
  - Resource Evacuation
  - Fence Targeted Book
  - Physically Repair/Replace/Upgrade Targeted Book
  - Concurrent Book Add Operation
    - Rebalance Processor Resources
    - Restore I/O Connectivity
    - Optimize Memory Resources
  - Restore Customer Workload





# Preparing for Concurrent Book Replacement

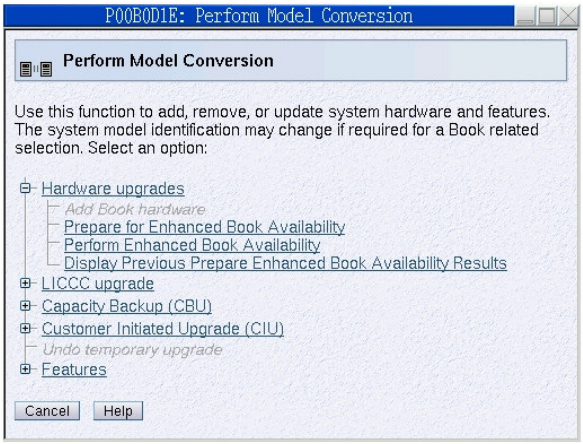


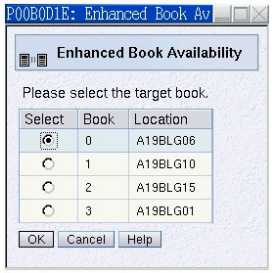

# z9-109 – Enhanced Availability Configurations

- **Concept: Configure enough physical memory and limit PUs configured so that all active purchased PUs and memory remain available with one book removed from the configuration with Enhanced Book Availability**
  - **Book removed concurrently for physical memory upgrade or repair**
  - **Restart with a fenced book following the rare event of a book failure**
- **How?**
  - **Select an S18, S28, S38, or S54 Model**
    - Configure no more than the following number of PUs
      - 8 active PUs on the S18
      - 18 active PUs on the S28
      - 28 active PUs on the S38
      - 40 active PUs on the S54
    - Requires no special feature codes for PU/model configuration.
  - **Select Flexible Memory configuration features**

| Configurable PUs (Standard) | 8 | 18 | 28 | 38 | 54 |
|---|---|---|---|---|---|
| Configurable PUs (Enhanced Availability) | NA | 8 | 18 | 28 | 40 |
| GB Memory (Standard) | 16 - 128 | 16 - 256 | 16 - 384 | 16 - 512 | 16 - 512 |
| GB Memory (Flexible) | NA | 32 - 128 | 32 - 256 | 32 - 384 | 32 - 384 |





# Prepare for Enhanced Availability Results

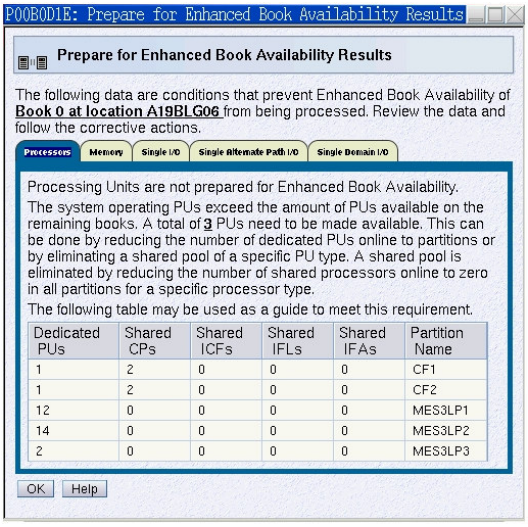

## Prepare for Concurrent Book Replacement

Stand alone SE application used by the customer/account team prior to the Service Action to help the customer prepare the system for the concurrent removal of the targeted book.

- **Memory Resources**
  - Total In Use Memory must be equal to or less than the physically installed memory on the remaining books.
  - Release Storage within a partition
  - Deactivate Partitions
- **Processor Resources**
  - Verify that sufficient spare processors or shared processors that can be downgraded exist on the remaining books. Customer may specify processor priority.
- **I/O Resources**
  - Customer must de-configure I/O with single path connectivity (i.e. ICB).

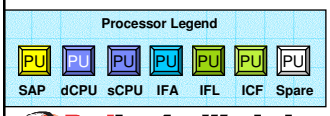


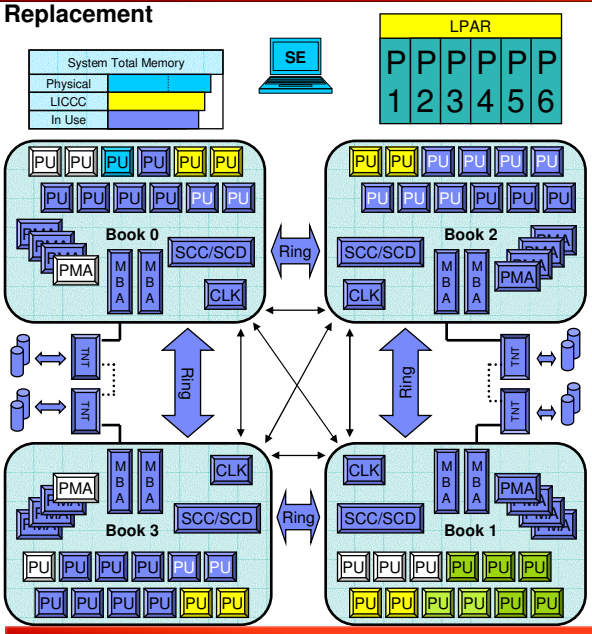






## Perform Concurrent Book Replacement

### Prepared System Verification

First step of the actual Concurrent Book Replacement Service Action.

Verifies that the system is still prepared for the concurrent removal of the targeted book.

Total In Use Memory must be equal to or less than the physically installed memory on the remaining books.

Sufficient spare processors or shared processors that can be downgraded exist on the remaining books to meet the needs of the processor resource algorithm.

Only I/O with an alternate path to an MBA on another book is attached to MBAs on the targeted book.

If any of the above are not satisfied, Then the CBR Service action is terminated.

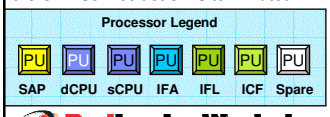


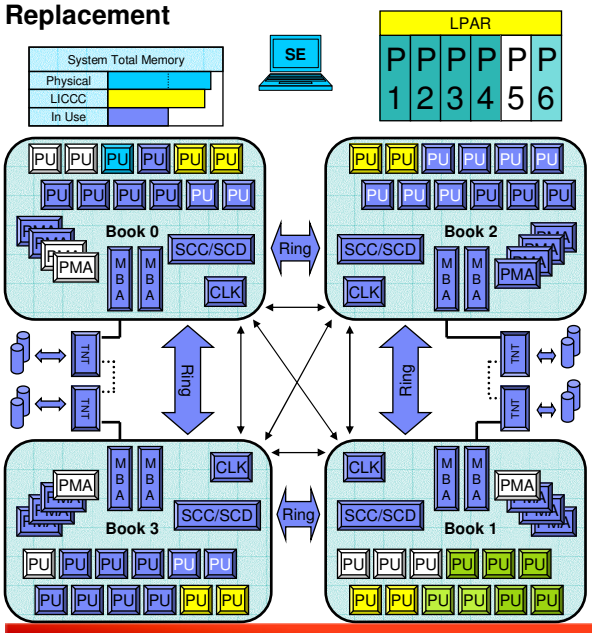

## Perform Concurrent Book Replacement

### Resource Evacuation

- **Resource Evacuation BFYCALL**
  - Initiates the Book Resource Evacuation
  - Targeted Book (2) and Degraded System Total Number of Processors by type (21 CPUs, 1 IFAs, 2 ICFs, 5 IFLs, 7 SAPs) passed from SE.
- **Verify Resources**
  - Fence (Lock) Target Book Physical Memory Resources
  - Memory Resource Requirements
  - Processor Resource Requirements
- **Respond to BFYCALL**
  - If all verification checks have passed then inform the SE that the Resource Evacuation has begun.
  - If any verification checks have failed, then terminate the service action.

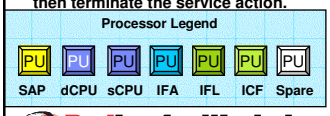


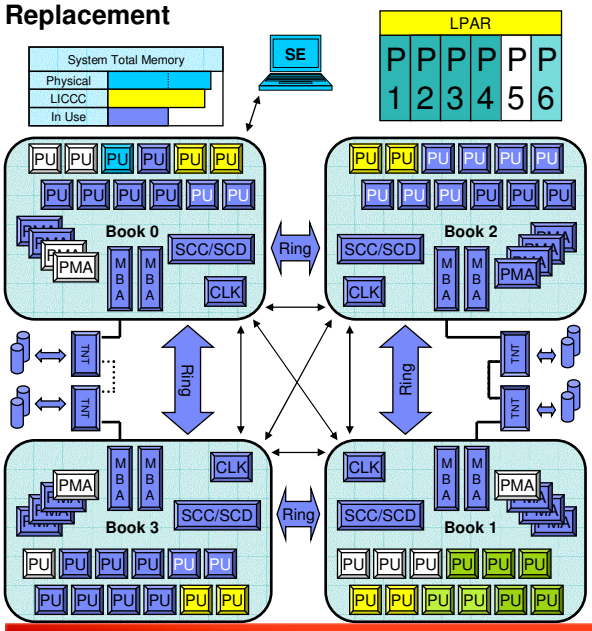






## Perform Concurrent Book Replacement

### Resource Evacuation (Continued)

- **Resource Evacuation BFYCALL**
- **Verify Resources**
- **Respond to BFYCALL**
- **Memory Evacuation**
  - Copy/Reassign
- **Processor Evacuation**
  - De-configure PUs
  - Downgrade Operation
  - Physical PU Reassignment
- **Completion Notification**
- **I/O Alternate Path Switch**

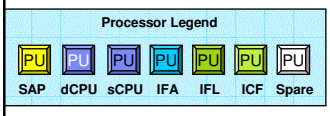


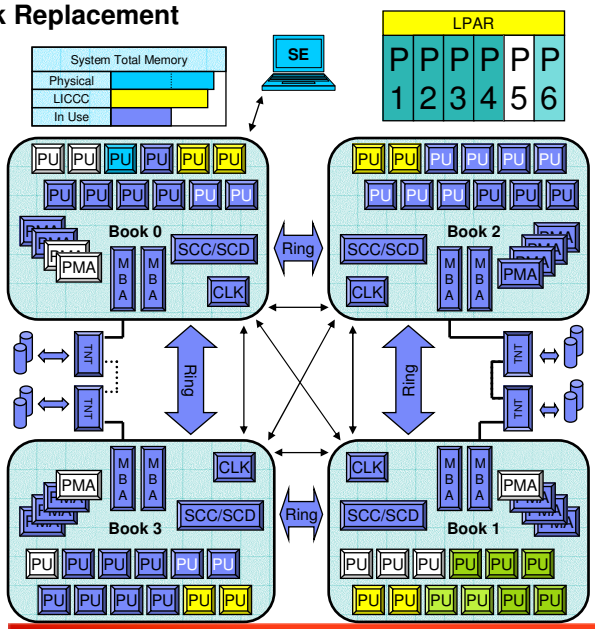

## Perform Concurrent Book Replacement

### Fence Targeted Book

- **Fence Book Service Word**
  Book Fencing operations in the CEC are initiated with an 'E8'x SW.
- **Book Fencing Operations**
  - Fence Book Physical Resources
  - Fence Ring Interfaces
  - Completion Service Word ('AC'x)
- **Clock to Clock Fencing**
- **Book Power Off**

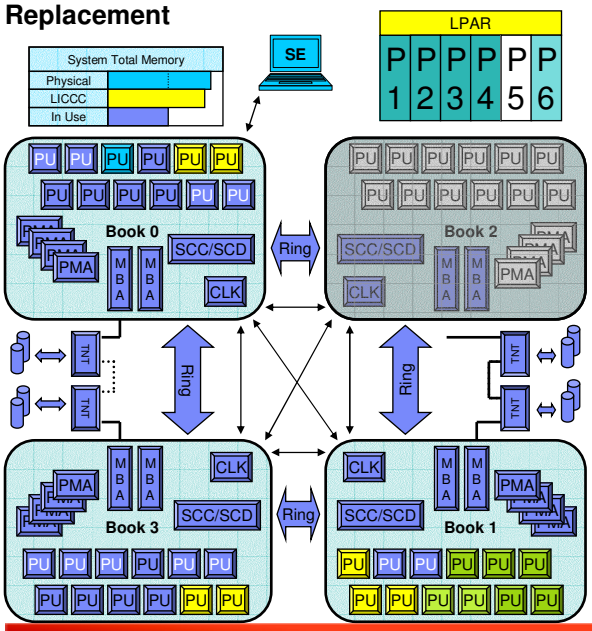


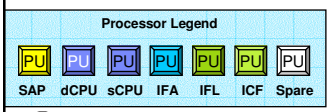






## Perform Concurrent Book Replacement

### Repair/Upgrade/Replace Targeted Book

- **Physical Book Removal**
  - UnCable
  - I/O Cagelet Removal
  - Physical Book Removal
  - Physical Book Upgrade or Replace

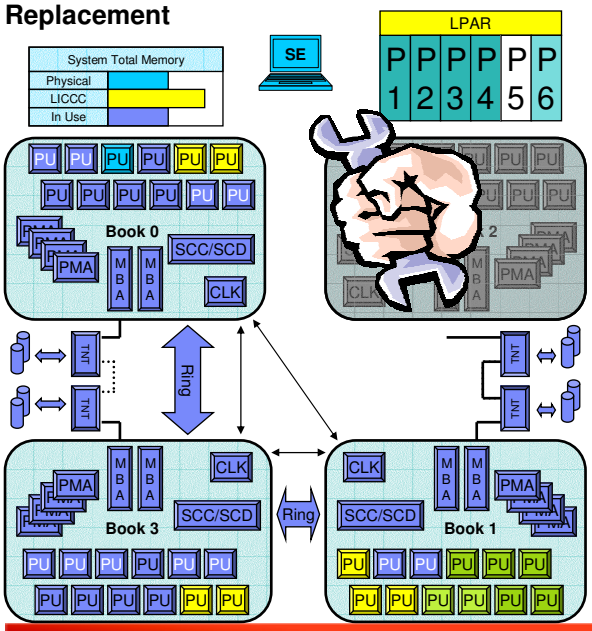


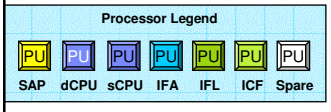

## Concurrent Book Add Operation

- Physical Hardware Installation
- Power Application and Selftest

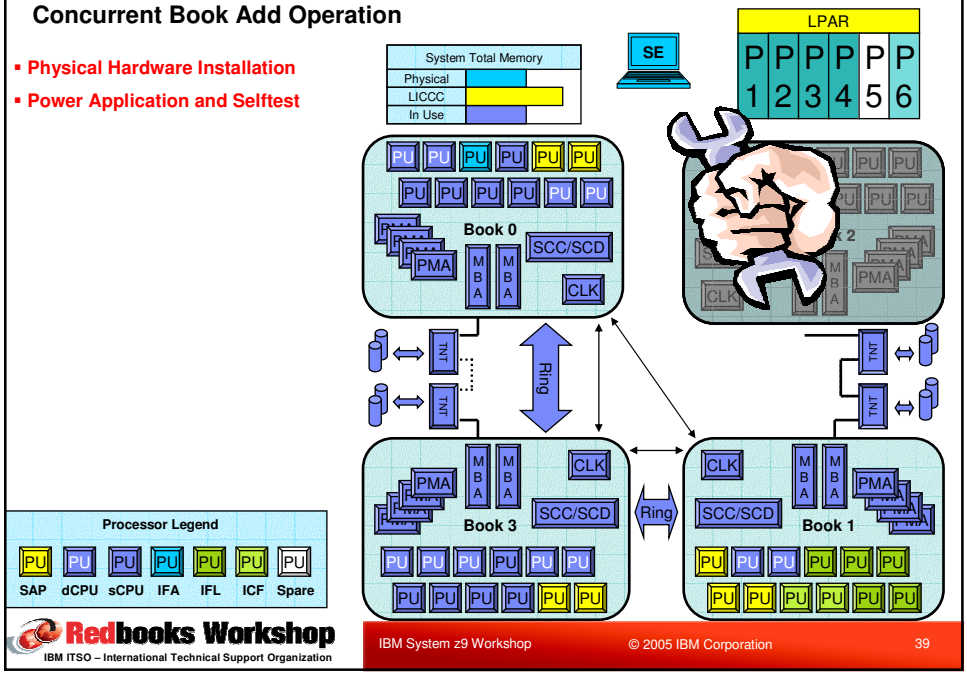






## Concurrent Book Add Operation

- Physical Hardware Installation
- Power Application and Selftest
- Book Basic Initialization
- Book Unfencing
- New Resource Allocations
- Processor Resource Optimization
- LPAR Notification
- Memory Resource Optimization
- Restore I/O Connections
- Restore Customer Workload

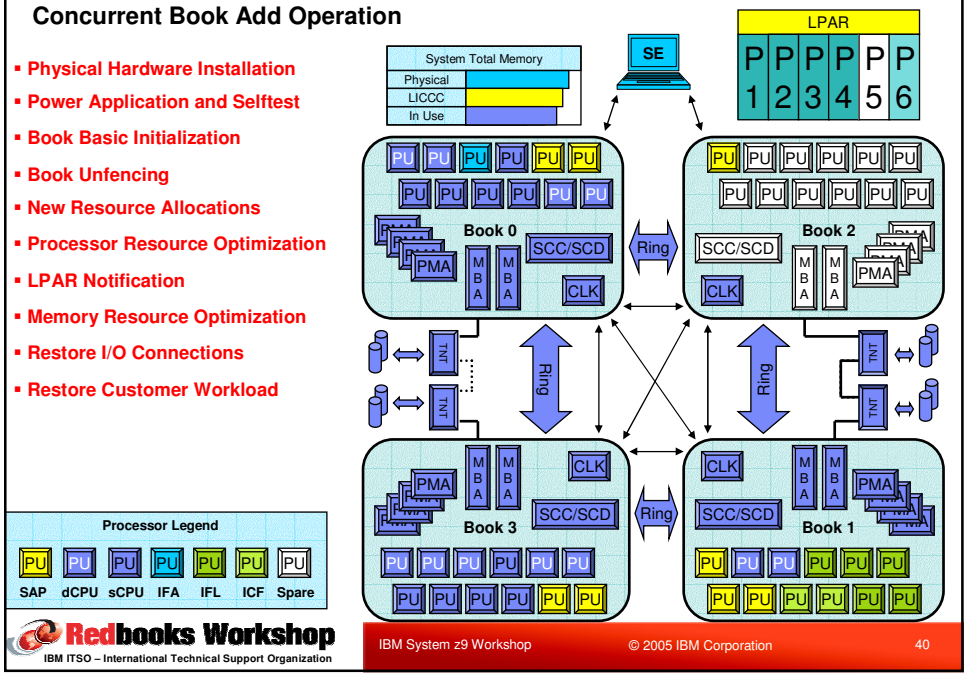

# Availability Enhancements

- Enhanced Book Availability
- Redundant I/O Interconnect
- Dynamic STI Management
- Enhanced Driver Maintenance
- MBA Fan-Out
- Dynamic Oscillator Switch
- On Demand Support
- Flexible Memory Option





# z9-109 Enhanced Driver Maintenance (EDM)

- **Ability to concurrently perform driver upgrades**
  - **Eliminates common planned outage**
  - **Window of opportunity within code maintenance stream**
    - Based on MCL sync points
    - MCL sync points may exist throughout the life of the current LIC level
      - Note: Once a sync point has passed, you will be required to wait until the next sync point supporting concurrent activation of a new LIC level.
    - EDM panels will display enablement / availability
  - **Like some concurrent patches today**
    - may need to vary off/on certain devices

# EDM Background

- Today
  - new GA Drivers/AROMs – disruptive via EC Upgrade (Alt SE Preload)
- T-Rex GA2
  - 1st experience of Concurrent new GA via MCLs
- Customer Feedbacks
  - Liked the concurrency aspect
  - Didn't like mix of fixes for previous GA with new GA
- EDM
  - addresses the customer feedback as well as provides concurrent release for all future GA1 + drivers (not GA1 drivers)





# EDM Design

- **New driver**
  - with totally new set of firmware ECs
- **Based on existing**
  - Alternate SE Preload
  - Concurrent Switch
  - Concurrent Patch
  - additional special enablement controls
- **EDM only goes concurrently upward**
  - GA downgrade (ie. GA2 to GA1) is always disruptive
- **EDM GA upgrades are mandatory sequential**
  - To go from GA1 to GA3:
    - If done in one step (GA1 to GA3), disruptive
    - If done in two steps (GA1 to GA2, GA2 to GA3), concurrent
    - (assuming all other requirements are met)
- **Just like patch**
  - Service Required State and open problems should not exist before starting EDM.
- **OOCoD (On Off Capacity on Demand) & CBU (Capacity BackUp) features**
  - can be active when doing EDM.

# Min/Max File Requirement

- **Min/Max control file**
  - Defines specified GAx From MCL level(s) and GAy To MCL level(s) required before EDM can be done.
  - The EDM requires the "from GA" to be within a given range (min/max requirement) of MCL levels before the upgrade can be performed
- **Min/Max control file defines which 'bundles' are supported.**
- **File will be patched into GA1 CDUMM EC stream**
- **MCLs Retrieved in CDUMM stream will automatically be Installed/Activated when panel or code System Information query.**
- **Not every EC stream required to have an entry in Min/Max file**
- **Not every patch MCL bundle released to field will support EDM**
- **Need to discuss with customers when they anticipate to move to next GA level in order to understand if they should apply MCLs above the Max EDM levels.**





# Min / Max Control file example

| EC Stream Number | Min MCL Applied Level | Max MCL Applied Level |
|---|---|---|
| J12964 | 001 | 001 |
| J12966 | 001 | 005 |
| J14382 | 004 | 006 |
| J14283 | 001 | 001 |
| J14384 | 012 | 012 |
| J14385 | 001 | 001 |
| J14386 | 003 | 005 |
| J14388 | 001 | 001 |
| J14390 | - - | - - |

# Checking MCL levels

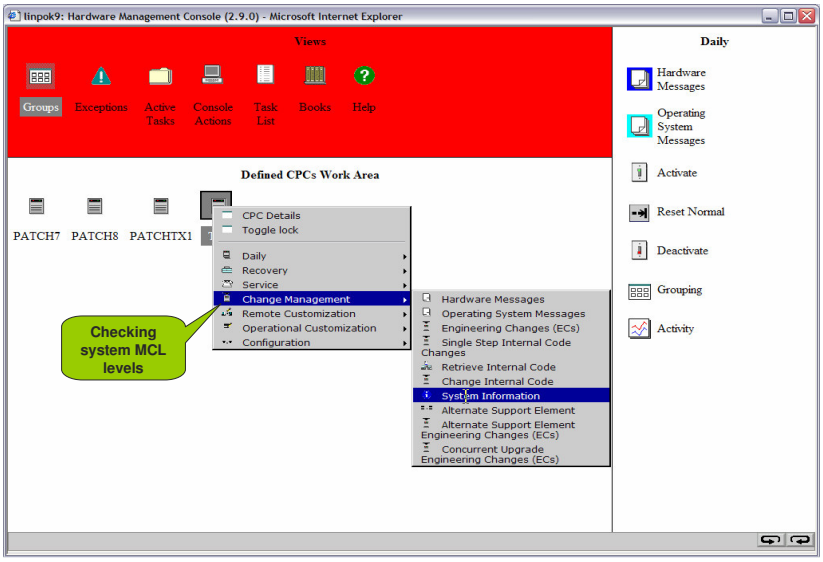






# Installed MCL levels – SE Panels

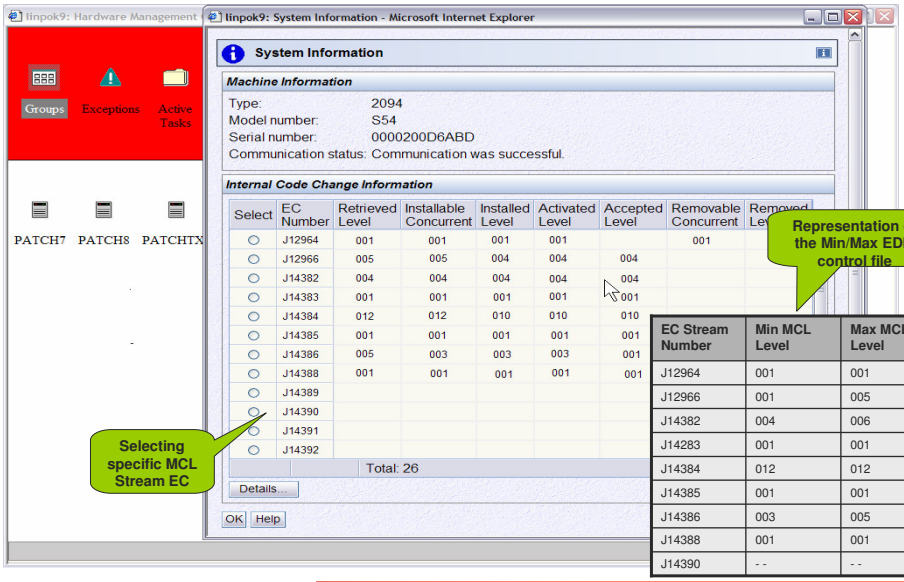


| EC Stream Number | Min MCL Level | Max MCL Level |
|---|---|---|
| J12964 | 001 | 001 |
| J12966 | 001 | 005 |
| J14382 | 004 | 006 |
| J14283 | 001 | 001 |
| J14384 | 012 | 012 |
| J14385 | 001 | 001 |
| J14386 | 003 | 005 |
| J14388 | 001 | 001 |
| J14390 | - - | - - |

# Individual MCL level information

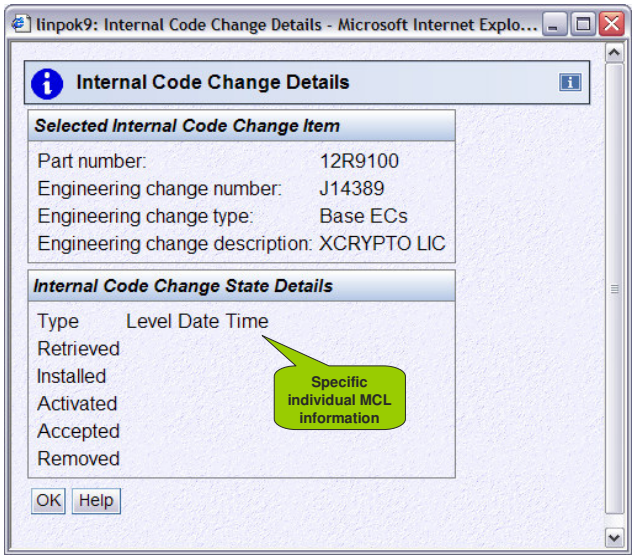






# EDM applicability chart

| MCL Bundle | Target Date | EDM Supported |
|---|---|---|
| 1 | GA | Yes |
| 2 | GA + 2 weeks | No |
| 3 | GA + 4 weeks | No |
| 4 | GA + 6 weeks | No |
| 5 | GA + 8 weeks | Yes |
| 6 | GA + 10 weeks | No |
| 7 | GA + 12 weeks | No |
| 8 | GA + 14 weeks | No |
| 9 | GA + 16 weeks | Yes |

# EDM SE Panel

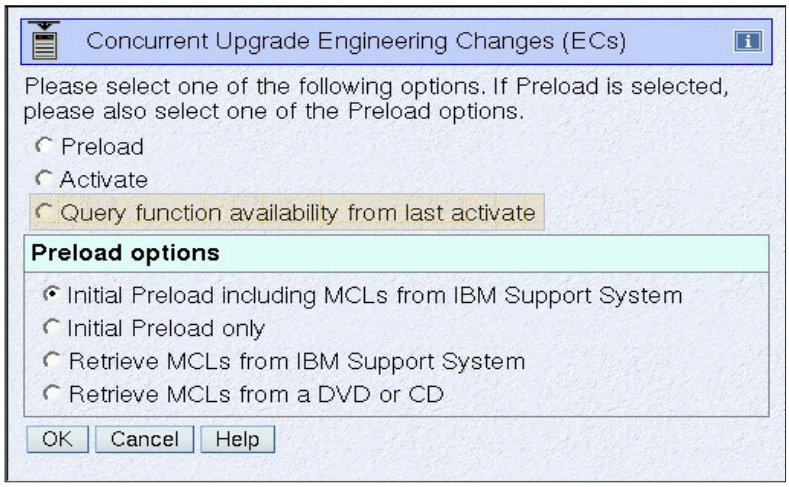






# Query Function Availability Panel

- Enabled
  - function enabled via EDM
- Not Enabled
  - EDM activate did not complete successfully
- Available
  - indicates function can be used by customer
- Not Available
  - Most likely next Power On Reset will make function available.
  - Maybe, you have to change the IOCDS, and then reIML.
  - It could vary by function.

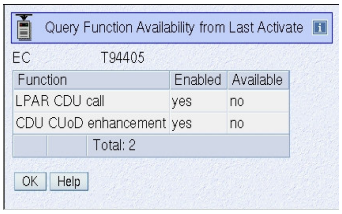


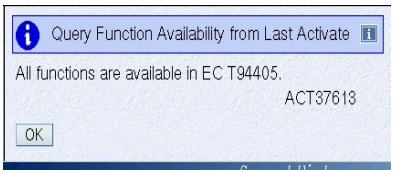

# Potential to temporarily Inhibit EDM

- Apply MCLs above EDM max levels
  - Based on Customer schedule for EDM
  - Change Internal Code will only show option if it truly exists
  - Single Step MCL will always show it.
  - Default is don't break EDM.
- During a EDM Activate, additional HSA (Hardware System Area) will most likely be allocated.
- If enough free memory isn't available, EDM won't be allowed (message will be given as to how much customer storage (in MB) must be freed to allow EDM.

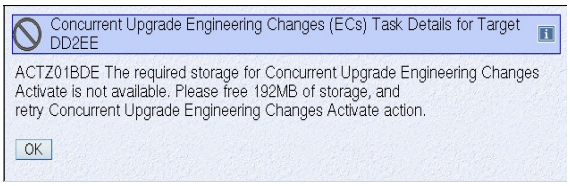






# Managed Driver Maintenance

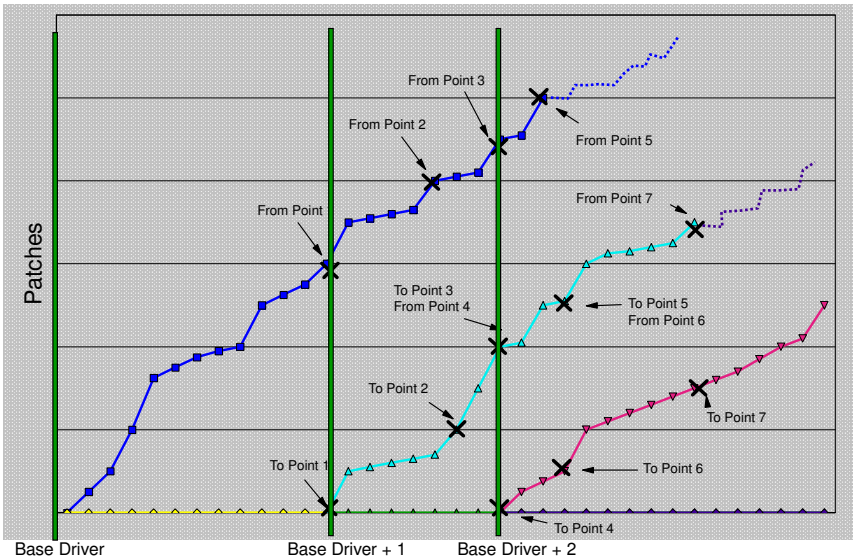

# Availability Enhancements

Enhanced Book Availability

Redundant I/O Interconnect

Dynamic STI Management

Enhanced Driver Maintenance

MBA Fan-Out

Dynamic Oscillator Switch

On Demand Support

Flexible Memory Option





## Dynamic Oscillator Switchover

- z9-109 has two oscillator cards
  - Primary and Backup
- In the event of a failure of the primary oscillator card, the backup is designed to detect the failure, switch over, and provide the clock signal to the server transparently.
- Previously, in the event of a failure of the active oscillator, a server outage would occur. IML was required to recover and switch oscillators.

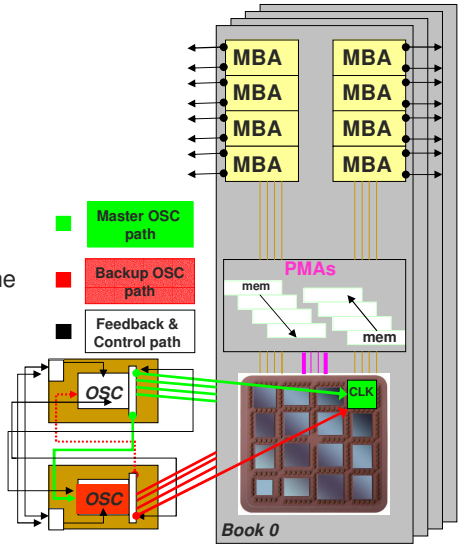

# Availability Enhancements

Enhanced Book Availability

Redundant I/O Interconnect

Dynamic STI Management

Enhanced Driver Maintenance

MBA Fan-Out

Dynamic Oscillator Switch

On Demand Support

Flexible Memory Option





## Concurrent Upgrade - Customer Controlled

- **CBU – Capacity Backup - Temporary emergency capacity upgrade**
  - Non-disruptive temporary addition of CPs, zAAPs, IFLs and ICFs in an emergency situation
  - CBU contract required to order CBU features and CBU LIC CC
  - Customer (or IBM) activates upgrade for test or temporary emergency
  - Non-disruptive downgrade required after test or recovery completed
- **CIU – Customer Initiated Upgrade - Express - Permanent upgrade**
  - Customer capability to order and install permanent upgrade
  - CUoD capabilities NOT included:
    - Upgrades requiring parts (e.g. S08 to S18 upgrade)
    - Channel upgrades by LIC enable of existing ports
  - CIU feature - ordered to initiate contract and administrative setup
  - Customer orders and installs upgrade via Resource Link™ and IBM RSF
- **On/Off Capacity on Demand - Temporary upgrade**
  - Nondisruptive temporary addition of CPs, zAAPs, IFLs, and ICFs in any situation
    - Upgrades requiring parts (e.g. S08 to S18 upgrade) not supported
  - "Right to use" feature - ordered to initiate contract and administrative setup
  - Customer orders and installs upgrade via Resource Link and IBM RSF
  - Nondisruptive removal when capacity is no longer wanted

NEW!

# z9-109 Concurrent Upgrade – CUoD

- CUoD -Capacity Upgrade on Demand- Standard machine function
  - Nondisruptive addition of CPs (software model upgrade), ICFs, IFLs, zAAPs, and memory
  - LIC enable additional 16 GB memory increments
  - Concurrent model upgrade (capacity setting) to add active PUs, memory, and STI busses
  - Downgrades are nondisruptive
    - Note: I/O feature adds and removes are also nondisruptive but not really "CUoD"







# z9-109 CBU Features


- **CBU features are ordered using eConfig**
  - Features are: CBU CP, CBU zAAP, CBU IFL and CBU ICF
  - There is NO limitation on the mix of CBU features except CBU zAAPs
  - Limitations:
    - **Active PUs plus CBU PUs cannot exceed configurable PUs** (Unassiged CP capacity and Unassigned IFL features do not limit CBU PUs)
    - **Adding active PU features may require removal of CBU PU features**
- **Conversion among CBU PU types is concurrent**
  - Order using eConfig

| From \ To | CBU CP | CBU zAAP | CBU IFL | CBU ICF |
|---|---|---|---|---|
| CBU CP | x | Yes | Yes | Yes |
| CBU zAAP | Yes | x | Yes | Yes |
| CBU IFL | Yes | Yes | x | Yes |
| CBU ICF | Yes | Yes | Yes | x |

# zSeries CIU and On/Off CoD

- **Order CIU and CoD "right to use" features**
  - Qualification, contracting, and pricing
  - Resource Link ID Authorization
- **Customer CIU or On/Off CoD order or On/Off CoD test order (up to 24 hours)**
  - Configure upgrade on Resource Link
  - Secondary Approval (Option)
  - Resource Link communicates with Remote Support Facility (RSF) to stage order and prepare download
  - Customer notified order ready
- **Access Support Element (SE) using Hardware Management Console (HMC)**
  - "Perform Model Upgrade"
  - Code obtained using RSF and installed on target machine

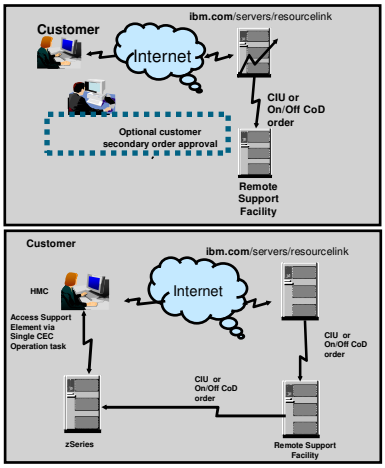






# z9-109 On/Off Capacity on Demand

- **Prerequisite for use:**
  - **Customer Initiated Upgrade** (FC #9898) and **On/Off CoD** (FC #9896) "right-to-use feature"
  - Signed CIU contract with specific Ts & Cs governing temporary capacity
- **Order temporary capacity – Resource Link**
  - **Can at most add capacity equal to active permanent capacity of the same type**
    For example – Go from 2 CPs to 4, 1 IFL to 2, or do both in the same order
    (Note: CIU upgrades and CBU do NOT have the this restriction)
  - PUs that have never been characterized can be activated as CPs, zAAPs, IFLs or ICFs
  - Unassigned IFLs can be activated only as IFLs – **Cost advantage on z9-109**
  - Unassigned CP capacity can be activated only as CPs – **Cost advantage on z9-109**
- **Order is manufactured: A LIC record is established and staged to RETAIN**
  - Multiple orders to meet different customer requirements can be staged
  - Orders remain on RETAIN for an extended period until:
    - Downloaded and activated (Initiates billing except for the 24-hour test)
    - Customer cancels order
    - Machine is no longer under warranty or IBM Maintenance Service Agreement
    - Change to Permanent PU and/or memory configurations invalidates order
  - A record, once activated, has no expiration date
- **On/Off CoD activation and CBU can coexist**
  - Must deactivate one function to activate the other one.
  - CBU PUs configured do not reduce On/Off CoD temporary capacity orderable

# z9-109 Capacity Upgrade on Demand

- CUoD is concurrent addition of processors and/or memory or concurrent type conversion among CPs, IFLs, and ICFs without disruption to workloads running on the machine - no power-off, power-on. Includes:
  - Addition of CP, ICF, IFL and zAAP
    - includes turning on (assigning) Unassigned" IFL features
  - LIC enabling additional 16 GB memory increments
  - Concurrent z9-109 model upgrade (book add) to add active PUs, memory, and MBA fanout cards with STIs
- All CUoD capabilities can be exploited by IBM ordered/installed MES upgrade
- Some CUoD capabilities can be exploited by customer controlled upgrades:
  - **Capacity Backup (CBU)** – temporary emergency upgrades
  - **Customer Initiated Upgrade (CIU)** – permanent upgrades
  - **On/Off Capacity on Demand (On/Off CoD)** – temporary on-demand upgrades
- **Notes:**
  1. CUoD is built on a base of concurrent "hot-plug" maintenance
  2. I/O feature adds and removes are also nondisruptive but not really "CUoD"





# z9-109 Concurrent PU Feature Conversions

- **Flexibility to meet changing business environments**
- **Decreasing the number of CP or IFL features is designed to be concurrent.**
  - Can be ordered by MES or by CIU. (No RPQ needed.)
  - Like z990 and z890, z9-109 unassigned IFL capacity is recorded by **Unassigned IFL** features
  - Unlike z990, z9-109 does NOT have **Unassigned CP** features
  - Like z890, z9-109 unassigned CP capacity is recorded by a **Capacity Marker** feature


- **PU type conversions shown below with Yes are designed to be concurrent**
  - Can be ordered by MES or CIU (No RPQ needed)
  - Example: From z9-109 S08 with eight CPs, convert one CP to an IFL

| From \ To | CP | IFL | Unassigned IFL | ICF |
|---|---|---|---|---|
| CP | x | Yes | No | Yes |
| IFL | Yes | X | Yes | Yes |
| Unassigned IFL | No | Yes | x | No |
| ICF | Yes | Yes | No | x |

**Exceptions: Disruptive if ALL current PUs are converted to different types**
**May require individual LPAR disruption if dedicated PUs are converted.**

# IBM System z9 - RAS Comparisons

| | z900 | z990 | Z9-109 |
|---|---|---|---|
| Microcode Driver Updates | **6 Hr Scheduled outage** | **2 Hr Scheduled outage** | **Concurrent** |
| Book Replacement | **Not Applicable (1 Book System)** | **Scheduled Outage** | **Concurrent** |
| Memory Replacement | **Scheduled Outage** | **Scheduled Outage** | **Concurrent (Node Offline)** |
| ECC on Memory Control Circuitry (EX: SMI) | **Unscheduled Outage** | **Unscheduled Outage** | **Transparent** |
| Memory Bus Adapter (MBA) Replacement | **Scheduled Outage** | **Scheduled Outage** | **Concurrent** |
| Oscillator Failure | **Unscheduled Outage** | **Unscheduled Outage** | **Transparent** |
| Processor Upgrades | **Concurrent** | **Concurrent** | **Concurrent** |
| Physical Memory Upgrades | **Scheduled Outage** | **Scheduled Outage** | **Concurrent (Node Offline)** |
| IO Upgrades | **Concurrent** | **Concurrent** | **Concurrent** |
| Spare PUs | **1 System** | **2 / Book** | **2 / System** |