z/OS and OS/390 Performance "HOT" Topics

Kathy Walsh
IBM

SHARE - Winter, 2002
EWCP Project
Session: 2500

Washington Systems Center

Agenda

- **Recent Performance Information**
  - Processor Announcements (z800)
  - z/Architecture Planning Information
  - Software Capacity Planning Information
  - Performance Topics
    - SADUMP Performance
    - IOP Information
    - IRD
    - Compat Mode
    - WLM, including Dynamic Alias management
    - CFCC Level 10
    - RMF
    - CICS and Logger
    - Enclaves
- **Addendum**
  - Older APARs or Performance Information

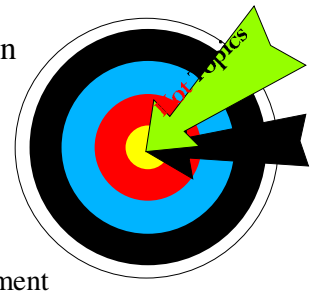

WSC Flashes

http://www.ibm.com/support/techdocs

Technical Support Technical Information Site

Note: for PC hardware, parts, accessories and thin client support go here

This site provides access to the Technical Support Technical Information Database. It gives you access to the most current installation planning and technical support information available from IBM pre-sales support, and is updated on a daily basis. You can view and search these databases by date, document number, product, platform or keywords.

View All by Number
View All By Date
View All by Product
All Documents Y2K
View

Flashes
Hints and Tips
Frequently Asked Questions
White Papers
Presentations & Tools
Search All Documents

New IBM @server zSeries 800

- **Introduction of new zSeries z800 Processor (2066)**
  - 8 general purpose models, z/Architecture capable
  - CF and Linux models also offered

Processor	CPs	MSU	MIPs*	CP MIPs*
2066-0A1	1	13	80	80
2066-0B1	1	20	115	115
2066-0C1	1	25	145	145
2066-0A2	2	44	260	130
2066-001	1	32	192	192
2066-002	2	60	250	175
2066-003	3	84	500	167
2066-004	4	108	636	159
9672-RX5	10	69	444	44
9672-X37	3	80	464	155

*** MIPs values are based on IBM's LSPR Default Mix workload, setting the 2064-1C1 equal to 250 MIPs. MIPs estimates are provided for comparison purposes only and should not be used as the basis for processor capacity planning.**

WSC White Paper WP100258

- **Performance considerations of moving to Fewer, Faster CPs**
  - Current processor design continues to deliver more CPU capacity with the same or fewer central processors (CPs)
- **WSC ran measurements to understand the issue**
  - Planning issues to consider are discussed in paper
    - CPU Capacity
    - Effects of LPAR on CP Capacity
    - LPAR Weight
    - Number of Logical Processors Assigned
  - Tuning options for Basic and LPAR modes
    - SRM parameters
    - LPAR Definitions
  - Impacts On:
    - CICS Throughput
    - CICS Response Time
    - LPAR Overhead

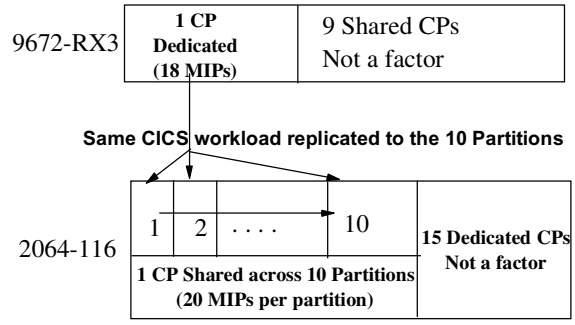

WSC FLASH W9634A

- **CPENABLE - trades off I/O responsiveness for throughput**
  - Re-issued 2/2002 for z800 recommendations
  - Re-issued 12/2000 for z900 recommendations
  - Re-issued 5/2000 for G5 / G6 recommendations

Processor Family	BASIC Mode	LPAR Dedicated	LPAR Shared
z800	(10,30)	(10,30)	**(0,0)**
z900	(10,30)	(10,30)	**(0,0)**
9672 G5, G6	(10,30)	(10,30)	**(0,0)**
9672 G1, G2, G3, G4	(10,30)	(10,30)	**(10,30)**
MP3000(1)	(10,30)	(10,30)	(0,0)
MP2000(2), ASP3000(3)	(10,30)	(10,30)	(10,30)
ES/3090, ES/9021	(10,30)	(10,30)	(0,0)

(1) Model Type = 7060, (2) Model Type = 2003, (3) Model Type = 3000

New Information on IOP utilization

- **WSC Flash 9646 documented how to review IOP / SAP utilization to ensure sufficient SAP capacity**
  - XCF CTC traffic can place high demand on SAP capacity
  - No direct method to determine IOP / SAP capacity
    - High IOQ AVG LNGTH values could be attributed to either SAP utilization or device contention rate
    - IOP Activity Rate reflected only SSCH/RSCH, not I/O Interrupts
    - SAP capacity estimates dependent on redrive rate of the configuration
      - DASD and Tape assumed 2 redrives per I/O
      - CTC assumed 1 redrive per I/O
- **IOP / SAP utilization statistics now reported**
  - z/OS V1R2 on a z900 or z800
  - RMF APAR OW49806 to produce new RMF IO Queuing Activity report
  - SMF 78.3 records updated with new and changed information

IOP: I/O Processor
SAP: System Assist Processor

RMF IOQ Activity Report
New Information for IOP / SAP

- Each IOP supports different STIs - review the balance across IOPs
- As AVG Q Lngth approaches 1.0, I/O delays should be reviewed
  - ⇔ If IOP Busy is low, check contention rates (Retries)

IOP	- INITIATIVE QUEUE - ACTIVITY RATE	AVG Q LNGTH	------- IOP UTILIZATION ------- % IOP BUSY	I/O START RATE	INTERRUPT RATE
00	41.587	0.00	0.30	41.593	43.100
01	2.077	0.00	0.01	2.077	1.477
02	16.823	0.00	0.07	16.830	16.497
SYS	60.487	0.00	0.13	60.500	61.074

-- % I/O REQUESTS RETRIED -- ALL	CP BUSY	DP BUSY	CU BUSY	DV BUSY	-------- RETRIES / SSCH --------- ALL	CP BUSY	DP BUSY	CU BUSY	DV BUSY
8.8	4.8	0.3	3.7	0.0	0.10	0.05	0.00	0.04	0.00
0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
42.9	21.4	0.1	21.4	0.0	0.75	0.37	0.00	0.37	0.00
21.6	11.0	0.2	10.3	0.0	0.28	0.14	0.00	0.13	0.00

- High redrives, review portion of the path causing contention

WSC Flash 10086

- **z/Architecture Capacity Planning**
  - Describes how to evaluate capacity planning impacts of a migration to 64-bit mode
  - Techdocs FAQ item with common 64-bit mode Q&A's
  - New software capacity planning tool introduced
  - MSO issues in a large real environment
- **Capacity Planning factors of 31-bit mode vs 64-bit mode**
  - Storage contention
    - Demand pages moved to / from expanded storage
    - High Average UIC
  - Storage management intensity
    - Indicator of how much real storage management, (RSM), activity is being done by a workload.
    - Activity includes operating system services such as getmain, freemain, page fault, page fix, and page free.
    - Activity is not controlled by an installation, rather it is a byproduct of the type of workload the installation runs

Software Capacity Planning

- **SOFTCAP V2.2**
  - Completely rewritten tool, available to customers, and business partners via the web
  - Covers OS/390 and z/OS migrations
  - Covers IMS and CICS software migrations
  - Capacity Planning factors of 31-bit mode vs 64-bit mode

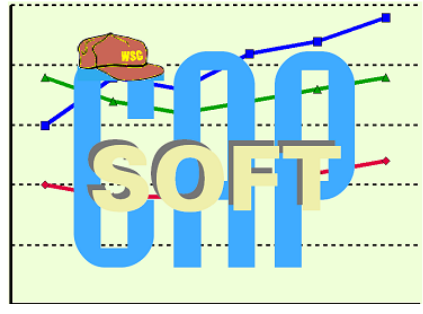

Customers: http://www.ibm.com/support/techdocs
Business Partners: http://www-1.ibm.com/partnerworld/sales/systems

WSC White Paper xxxxx
Managing an MVS system in an all "Real" Environment

- **Discusses changes in MVS performance metrics and impact on reviewing RMF reports**
  - UIC
  - Available Frame Queue
  - Effective Logical Swap
  - MSO
  - Transition Swaps pre z/OS 1.2
    - V=R
    - REAL=
    - RSU=
  - Fixed Frames below 16M line
- **Recommendations for Auxiliary Paging Subsystem**
  - Contiguous slots algorithm
  - PAVs
  - Allocation and performance

Virtual Storage

- **OW50225 - SQA Shortage - RAS Improvement**
  - OS/390 2.10 and up
  - Improves detection mechanism for a critical SQA Shortage condition so recovery routines can receive control earlier
  - May see increased IRA100E messages because of threshold change
    - Encourages tuning SQA and/or CSA allocations below the line

Message	Reason	Before OW50225	After OW50225
IRA100E	SQA Shortage	8 4K pages (32K)	128 4K pages (512K)
IRA101E	Critical SQA Shortage	4 4K pages (16K)	64 4K pages (128K)

- **II13162 - INITSQA**
  - WAIT040 ABEND878 during IPL with ARCHLEVEL=2
  - In 64-bit mode the layout of storage has changed due to the elimination of expanded storage
  - Amount of ESQA previously specified may need to be increased when running in 64-bit mode
    - May have previously been piggybacking on the ESQA set aside during the IPL for the expanded storage tables
    - With this set aside eliminated insufficient ESQA may exist

Software Capacity Planning (Variable Workload License Charge)

- **Availability**
  - 100% WLC is available with z/OS V1R1
  - VWLC available 3Q01
  - IBM offers a transition program to ease migration to WLC
- **Planning**
  - WSC Planning for VWLC Seminar (course code: VWLC1)
  - Covers planning steps for VWLC
  - 2 1/2 day web-conferences
    - Schedules for East Coast and West Coast
- **Planning Tool for VWLC Products (WLCTOOL)**
  - Used to get information necessary for software capacity planning
    - Provides CEC and LPAR view, tool output is in CSV format
    - Provide utilization in MSU, LPAR weights, CEC Utilization
    - **http://www-1.ibm.com/servers/eserver/zseries/swprice/**
- **Planning Tool for Value Unit Priced Products (VUCTOOL)**
  - **http://www-1.ibm.com/servers/eserver/zseries/swprice/**
- **VWLC Reporting Tool (SCRT)**
  - Used to send capacity information to IBM to enable VWLC
    - **http://www-1.ibm.com/servers/eserver/zseries/swprice/**

WSC Flash 10099 (VWLC)

- **WLM support for Variable Workload License Charges**
  - Any z/OS partition with a defined capacity set in the LPAR panels will be managed by WLM to the defined capacity
    - WLM calculates the rolling four hour average and enforces PR/SM softcap
  - Support **is active** whether or not VWLC is being used
    - Eligibility requires all active MVS images to run z/OS but CEC's with both OS/390 and z/OS images active have capping enabled for the z/OS images
    - RMF PP can be used to see the potential impacts of WLM soft capping on system capacity (CPU Activity via PDR report)
  - **Recommend:** Don't specify defined capacity for any z/OS image unless the installation wants to enable the defined capacity support

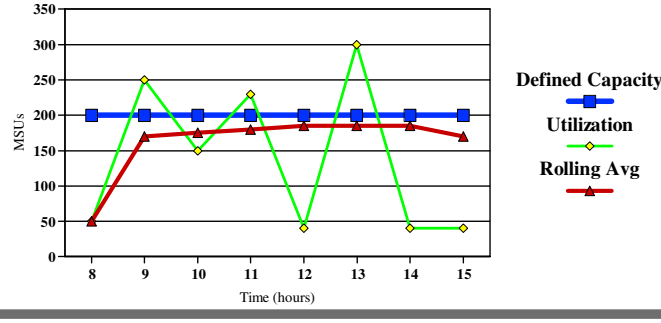

VWLC and License Manager
(SHARE Sessions)

★ More Information:

- **Session 2828: IBM License Manager Update**
  - Speaker: Guy Harrison (IBM Corporation)
  - Mon 4:30 PM Level 2, Presidential Ballroom,Presidential D
- **Session 1873: SAM in the zSeries - From Theory to Practice**
  - Speaker: David Vardi (Isogon Corporation)
  - Tue 3:00 PM, Level 2, Governor's B
- **Session 2526: WSC Software Capacity Planning**
  - Speaker: Kathy Walsh (IBM Corporation)
  - Wed 3:00 PM, Level 2, Presidential Ballroom, Presidential A
- **Session 1936: The Technology Behind IBM's WLC Licensing Model**
  - Speaker: Ron Higgin (BMC Software)
  - Wednesday 4:30 PM, Level 2, Governor's Chamber A

WSC Flash 10043
(SADUMP)

- **SADUMP Performance Benchmarks**
  - Produced by zSeries Performance team
  - Measurement variables were:
    - Placement of the dump output data sets on an IBM Enterprise Storage Subsystem (ESS) vs. on a non-ESS DASD (IBM 3390 mod 3)
    - Use of 2, 4, or 8 spanned ("striped") volumes
    - The Real Storage size
  - Measurement environments included:
    - Real Storage Dump
    - Virtual Storage Dump
    - IPCS Initialization

Intelligent Resource Director (SHARE Sessions)

★ **Components are:**

- **Dynamic Channel Path Management**
  - Basic mode, Single LPAR (not a cluster, no EMIF allowed), or LPAR cluster
- **LPAR CPU Management**
  - LPAR Weight Management
  - Vary CP Management (Requires LPAR Weight Management)
- **Channel Subsystem Priority Queuing**

★ **More Information:**

- **Session 2517: Introduction to LPAR CPU Management**
  - Speaker: Walt Caprice (IBM Corporation)
  - Tue 11:00 AM, Level 2, Presidential Ballroom, Presidential A
- **Session 2519: An IRD Case Study**
  - Speaker: Walt Caprice (IBM Corporation)
  - Tue 1:30 PM, Level 2, Presidential Ballroom, Presidential A
- **Session 2525: IRD: What LPAR Clusters Can Do for You**
  - Speaker: Alan Sherkow (I/S Management Strategies, Ltd.)
  - Wed 4:30 PM, Level 2, Presidential Ballroom, Presidential A

z/OS Compat Mode Support

- **On May 16, 2000 IBM announced withdrawal of Compat Mode support**
- **z/OS 1.3 available in 3/2002 will be first without compat mode support**
  - No WLM policy found, system IPLs with default policy
    - SYSTEM, SYSSTC, SYSOTHER, and TSO
  - Provides support to bypass WLM ISPF appl to install a service definition and activate a policy (batch interface)
    - IWMINSTL - sample JCL which runs REXX code to install a sample service definition
    - Sample service definition is not appropriate for production work
  - z/OS 1.2 shipped in 9/2001 will be the last with compat mode
    - General policy is to provide service to this product for 3 years
- **Goal Mode Migration Program**
  - Tool to convert IPS/ICS to Service Definition
  - WSC Migration Guide and Checklist V2
  - WSC Production Readiness Review
  - WSC WLM Sample Service Definition and Policy

Workload Manager
(SHARE Sessions)

★ More Information:

- **Session 2515: What's New in z/OS WLM**
  - Speaker: Juergen Holtz (IBM Corporation)
  - Tue 4:30 PM, Level 2, Governor's, Governor's A
- **Session 2518: IBM's Workload Manager Free-for-All**
  - Tue 6:00 PM, Level 2, Governor's, Governor's A
- **Session 2512: WLM, DB2, and Enclaves - The Inside Story**
  - Speaker: Juergen Holtz (IBM Corporation)
  - Wed 1:30 PM, Level 2, Governor's, Governor's A
- **Session 2513: CICS and the Workload Manager**
  - Speaker: Kathy Walsh (IBM Corporation)
  - Thur 1:30 PM, Level 2, Delta, Delta C
- **Session 2547: Using OS/390 WLM to Manage WebSphere Performance**
  - Speaker: Glenn Anderson (IBM Corporation)
  - Thur 9:30 AM, Level 2, Governor's, Governor's A
- **Session 2552: WLM: Revisiting Goals Over Time**
  - Speaker: Peter Enrico (Enterprise Performance Strategies Inc,)
  - Tuesday 8:00am, Level 2,Presidential Ballroom,Presidential A

WSC Flash 10101
(WLM)

- **WLM APAR OW47667 Regarding Disconnect Time**
  - Remove disconnect times from I/O using samples
    - Disconnect time is an indicator of poor performance, showing low cache hit rates, and possibly higher response times or lower throughput
    - WLM, prior to these PTF's, considered Disconnect time a Using Sample
    - Work with high disconnect time, a performance issue, was viewed by WLM as having a higher velocity due to increased I/O Using samples
  - If I/O Priority Management=YES then:
    - After application of this fix, WLM velocity goals will need to be reviewed
    - Amount of disconnect time will drive the amount of change needed
  - If I/O Priority Management = NO then:
    - After application of the fix may want to turn I/O Priority management on
    - Still need to review velocity goals due to possible impacts of Connect Time on I/O Using Samples

Parallel Access Volumes

- **New capability of Enterprise Storage Server (SHARK)**
- **Multiple concurrent data transfers to/from a DASD volume**
  - Writes to same Define Extent-Range are serialized
- **Multiple unit addresses (UCBs) per volume**
- **Base address**
  - Actual unit address of the volume
  - One base address per volume
  - Space associated with base
- **Alias address**
  - Maps to base address
  - Aliases are only visible to IOS
  - Up to 255 aliases assigned to a base

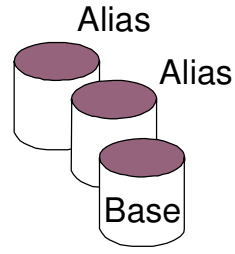

Dynamic PAVs

- **Managed by WLM**
  - Alias UCBs moved to minimize **IOS queue length**
    - Provides benefits for environment with significant IOS queue time
    - Required WLM definition

```
Dynamic alias management . . . . YES  (Yes or No)
```

  - Alias UCBs also moved to help meet service class period PI
    - Required WLM definition

```
I/O priority management  . . . . YES   (Yes or No)
Dynamic alias management . . . . YES  (Yes or No)
```

  - WLM is monitoring **IOS queue length**, not **IOS queue time**
    - IOS queue time may not go to 0 even if there are unassigned PAV's in the pool
  - Recommendation
    - Spread alias UCBs across the base addresses to expedite movemen

Integrity Implications

- **PAVs allow multiple, simultaneous I/Os.**
- **Control unit is responsible for Data Integrity**
  - Managed at the data set extent level
    - Multiple reads allowed
    - Only one write I/O allowed
    - Waiting for the write to complete can generate PEND time

DEV NUM	VOLUME SERIAL	PAV	SMF SYS ID	DEVICE ACTIVITY RATE	AVG RESP TIME	AVG IOSQ TIME	AVG DPB DLY	AVG CUB DLY	AVG DB DLY	AVG PEND TIME	AVG DISC TIME	AVG CONN TIME	% DEV CONN	% DEV UTIL
A000	PTSSPA		*ALL	365.016	9	7	0.0	0.0	0.1	0.4	0.1	0.8	30.26	32.68
		1	SYSC	183.491	9	7	0.0	0.0	0.1	0.4	0.1	0.8	15.20	16.43
		1	SYSD	181.525	9	7	0.0	0.0	0.1	0.4	0.1	0.8	15.06	16.24

DEV NUM	VOLUME SERIAL	PAV	SMF SYS ID	DEVICE ACTIVITY RATE	AVG RESP TIME	AVG IOSQ TIME	AVG DPB DLY	AVG CUB DLY	AVG DB DLY	AVG PEND TIME	AVG DISC TIME	AVG CONN TIME	% DEV CONN	% DEV UTIL
A004	PTSSPE		*ALL	404.286	7	3	0.0	0.0	2.6	2.8	0.2	1.0	9.94	12.05
		4	SYSC	199.885	7	3	0.0	0.0	2.5	2.8	0.2	1.0	4.92	5.95
		4	SYSD	204.401	8	4	0.0	0.0	2.6	2.9	0.2	1.0	5.02	6.10

PAV related Info

- **OW53205 - IOS - Open**
  - ABEND073 RC28 can occur when WLM dynamic alias tuning is used to move an alias from one base to another
- **OW47493 - IOS**
  - ABEND073 when alias devices transition from bound to unbound and vice versus due to serialization issues with UCB lock
  - Install PTFs before using dynamic alias management
- **OW52196 - IOS**
  - VARY ONLINE or VARY ONLINE,UNCOND causes WLM data to be lost impacting dynamic alias tuning for PAVs
  - OPEN but ++APAR is available
- **OW43946 - Media Manager**
  - Performance enhancement for d/t 2105 (Shark) to allow bypassing extent collision checking and validation checking of Define Extent and Locate Record commands
- **OW50243 - DFSMS**
  - Allow CVAF VTOC READ IO (EXCP) to exploit PAVs
  - Concurrent VTOC read access performance improved 6 times
- **OW51030 - IOS**
  - C.U.I.R. processing may incorrectly issue messages IOS283I and IOS290I due to incorrectly issuing VARY PATHs to PAV-Alias devices
  - Possible performance and storage problems

WLM PAV Algorithm

- **OW48647 - WLM**
  - Enhancement to dynamic PAV efficiency algorithm
  - WLM is not always moving to help devices with non-zero **IOS queue length**
    - Aliases are available on idle devices or a device has more aliases than it needs
  - WLM does not move an alias if the average IOS queue length falls below a threshold of 0.5
    - Considered too high when unused aliases exist in the same I/O subsystem
  - WLM will add a second pass to the efficiency algorithm to scan for base devices with an average IOS queue length of 0.05
    - If receiver devices are found and there are available aliases or aliases on low-usage devices WLM will move aliases to help the receiver device

PAV - WLM Assignment

- **OW51835 - (ESCON Environment)**
  - When using dynamic alias management an excessive number of aliases may be assigned to a base address
  - WLM keeps adding aliases even when there are control unit or channel constraints
    - Reserve contention, high DISConnect time, high PENDing time
  - As device response times slow, IOSQueue times start to build, attracting WLM's attention and WLM starts to move aliases to the device even though additional aliases cannot relieve source of the contention
    - Aliases may be used more productively elsewhere
  - WLM does not add an alias to help a device with IOS queuing when IOS queuing is:
    - Due to a pending reserve
    - Associated with high PEND time on the local system
    - Associated with high DISConnect time at the device level

WLM, XCF and PAVs

(Support Requires all Images in Sysplex to have fixes applied)

- **OW50276 - Part 1 - Increase in 16K XCF messages**
  - XCF every 10 seconds sends state info on dynamic aliases
  - WLM collects and sends to all systems in the sysplex 64 bytes per device entry
    - Potentially large amount of data (256 device entries per LCU, and 8 LCU's possible in an 2105 ESS)
    - Spiky arrival rates
  - Spikes in the XCF message rates and increase pathout busy delays
  - Recommendation: Review robustness of XCF implementation when implementing dynamic PAV's
- **OW50276 - Part 2**
  - Changes the calculation for the sysplex view of average IOS queue length for a base device
  - Sysplex view of IOSQ length was an average across all images
    - In a 8-way sysplex, if one system is issuing I/O to a base device, IOSQ length seen by the system is divided by 8 for the sysplex average
    - Sysplex average would be below WLM's threshold for moving aliases until queuing is quite high.
  - Sysplex IOSQ length is calculated as the *sum* of the average IOSQ lengths on each image
    - High-order sysplexes will see PAV movement at lower levels of IOS queuing than before

FICON Connect Times

- **OW51126**
  - I/O priority management is triggered when a service class is missing it's goal (PI >1) and major reason for delay is I/O contention
    - ESCON attached DASD: Connect time is always considered productive I/O transfer time
    - FICON attached DASD: Connect time can represent productive data transfer time AND delay time due to the multiplexing of requests
  - If the multiplexing delay is significant, WLM may overstate a service class's velocity and understate its PI
    - I/O priorities may not be adjusted as quickly as they should be to help important work meet its goal
  - WLM now uses a constant connect time per FICON I/O request of 1ms rather than using the measured connect time when calculating velocity and performance index

WSC Flash 10118 (CFCC Level 10 Storage Increase)

- **CFCC Level 10 introduces a new function called System Managed Coupling Facility Structure Duplexing (SM CF Duplexing)**
  - May result in increased sizes for List and Cache structures
    - Directory only cache and Lock structures are not impacted
  - Increase the affected structure sizes by at least 768K bytes
  - Especially important for smaller structures where storage increase represents a proportionally larger change than it does for large structures
  - Impact is felt even if SM CF Duplexing is not used

Structure Type	Control Storage	Control Storage or non-Control Storage
List	List Entry Controls	Data Elements
Lock	ALL	
Directory Only Cache	ALL	
Cache	Cache Directory	Data Elements

CICS and Logger

- **PQ48163 (CICS TS) - FIN**
  - CICS batchs logger writes to reduce overhead of calling logger via parm LGDFINT system initialization parameter
  - When the log defer time happens to cross into the partition exit time the IXGWRITE is not issued until the partition exit time expires and CICS is re-dispatched
    - If a partition exit occurs it can mean a delay of anything up to the region exit interval (ICV) value
    - Default ICV value is 1000 ms
  - Impact is a noticeable increase in response time, especially in a less busy CICS region
    - Seen only with CF logstreams
  - Circumvention:
    - Set LGDFINT to 00 for CF logging
    - Set LGDFINT to 05 for DASDONLY logging
- **PQ57850**
  - Because of the pervasive nature of the problem this APAR was taken to describe the actions of moving the support back into CICS TS 1.3

RMF

- **OW51003**
  - New ICSF Support to Enable RMF to provide performance measurements on select ICSF services and functions
  - WSC (M. Allmond) presentations on crypto performance:
    - PRS333 - A Presentation of Crypto Concepts
    - PRS334 - Hardware Crypto Benefits
    - WP100245 - Understanding the Crypto Hardware Available for zSeries and S/390
    - TD100158 - Reference for Number of Crypto Coprocessor Available to a Server by Model
- **OW50084**
  - DASD response times and IOSQ time now reported in tenths of a millisecond
    - Monitor II and Post-processor
- **OW49536**
  - CPU Partition Report uses spare / reserved CPs
  - Impacts Effective Utilization Data
  - Due to introduction of new fields in the SMF 70 records

Enclaves

- **OW49571 - Supervisor**
  - ABEND322 may incorrectly occur in batch jobs using enclave TCBs
    - TCB timing is done whenever a TCB relinquishes control of the processor
    - When the TCB was an enclave TCB, time was compared against the the wrong field, (ASSBPHTM vs ASCBEJST), causing tests for time excession to fail incorrectly
- **PQ52813 - TCP/IP**
  - IPMAIN Lock contention results in ECSA / ESQA storage
  - TCP/IP issues SYSEVENT ENQHOLD to cause holder of the resource in contention to be promoted
  - Dispatchable Unit holding the lock was an enclave, and not the address space, so ENQ promotion did not cause promotion of the enclaves dispatch priority

SDSF Management of Enclaves

- **New with z/OS V1R2**
- **Entered via the enclave panel (ENC)**
  - SSType - Subsystem Type (DB2, CB (WAS), MQ, etc.)
  - Status - ACTIVE | INACTIVE (Best field to sort on)
  - SrvClass - will also support PGNs
  - CPU-Time - total CPU time
  - Enclave Owner Information - (system, jobname, ASID)
  - Enclave Attributes - (Scope (LOCAL or MultiSys), Type (IND or DEP))

```
Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------
SDSF ENCLAVE DISPLAY  SYSD     ALL                        LINE 1-10
PREFIX=TS*  DEST=(ALL)  OWNER=*  SORT=Status/A  SYSNAME=SYSD
ACTION=//-Block,=-Repeat,+-Extend,I-Info,M-Match
NP    TOKEN             SSType Status   SrvClass RptClass   CPU-Time
      240000142B        CB     ACTIVE   CBSLOW   RTSIVP2         2.65
      2800001441        CB     ACTIVE   CBSLOW   RTSIVP2         0.01
      2000001449        CB     INACTIVE CBSLOW   RTSIVP2         0.00
      2C00001446        CB     INACTIVE CBSLOW   RTSIVP2         0.00
      3000001448        CB     INACTIVE CBSLOW   RTSIVP2         0.00
      3400001443        CB     INACTIVE CBSLOW   RTSIVP2         0.00
      3800001447        CB     INACTIVE CBSLOW   RTSIVP2         0.00
```

Addendum

- **Older flashes which should still be understood, or make you go Hmmmm.**
- **APARs which are still causing issues, even though they are old.**

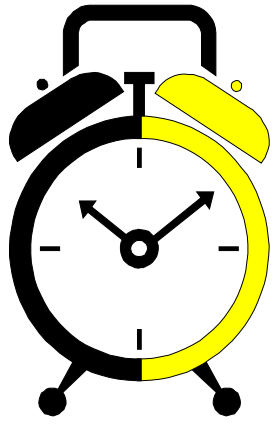

CF Storage Allocation
List Structures

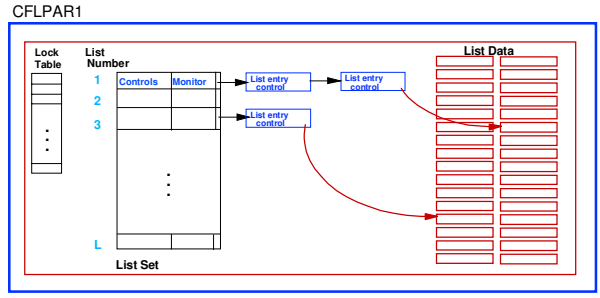

Lock Table is optional

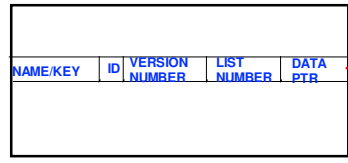

LIST ENTRY CONTROLS

Control Storage Resident

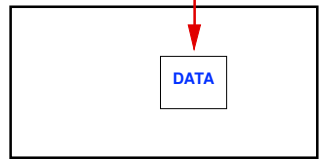

LIST DATA

non-Control Storage or Control Storage Resident

CF Storage Allocation
Lock Structures

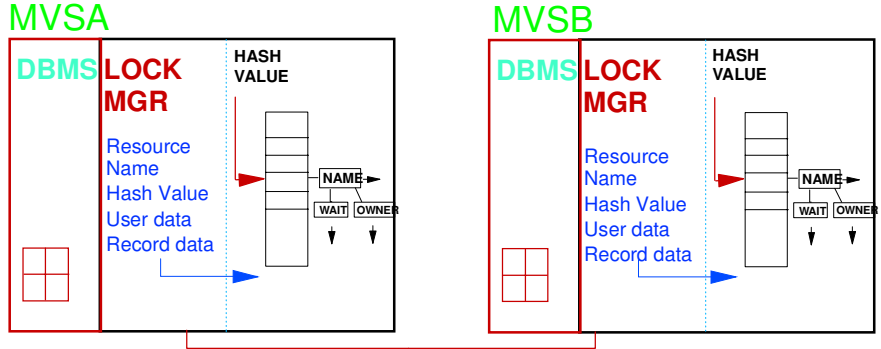

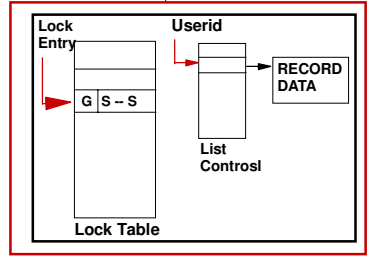

CF1LPAR

Control Storage Resident

CF Storage Allocation
Lock Structures

- **Directory only**
  - Contain no "data"
  - IMS uses these for OSAM and VSAM
- **Store Thru**
  - Contain hardend/ unchanged "data"
  - RACF uses this type of structure
- **Store In**
  - Contain unharded/ changed "data"
  - DB2 uses this type of structure

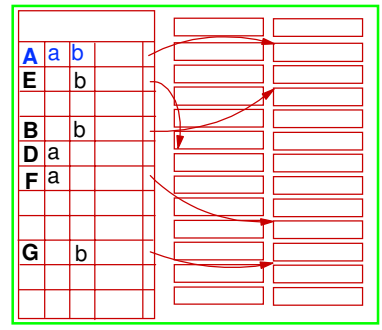

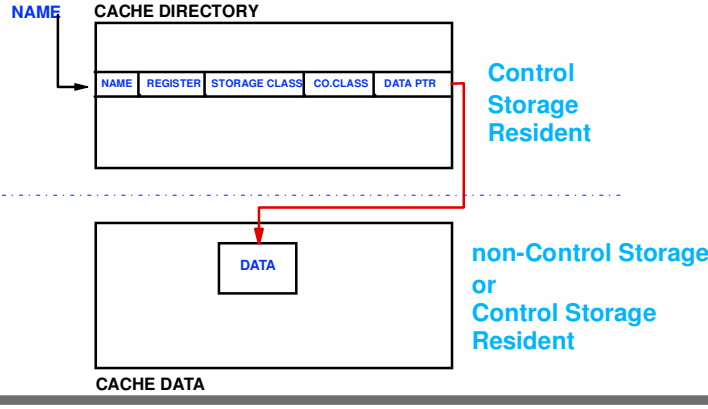

CF Link Utilization

- **INFO APAR II05151**
  - Indicates some CF Link utilizations are recorded in RMF
  - Data is coming from the STCPS instruction and not the CPMF, so the data is not available on all processors or for all CF link types
  - Available data is questionable
- **How do I know when more CF links are needed?**
  - Review the set of subchannels as a pool of capacity to the CF
  - If CF links are shared, sum subchannel activity across all images
  - New ISC-3, and ICB-3 links, (peer mode links), have 7 subchannels per link, while other links have 2 subchannels per link

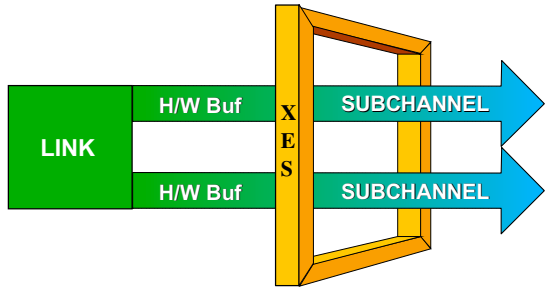

CF Subchannel Utilization

- **Use an RMF report for an appropriate peak interval**
  - SYSRPTS(CF)
  - Non-peer mode links keep SCH Utilization below 30%
  - Peer mode links keep SCH Utilization below 20%

Total SCH Busy Time per Sec

```
=SYNC ops/sec * SYNC serv time + ASYNC ops/sec * ASYNC serv time
=14176K/3600 * .0000337      + (871929+99787)/3600 * .0003706
= .233 secs per second
```

AVG SCH Utilization

= 100 * SCH Busy Time per Sec / SCH USE

= 100 * .233 / 4

= 6%

```
                                    INTERVAL 001.00.00
                                    SUBCHANNEL  ACTIVITY
   SYSTEM
   NAME   -- CONFIG --    ----------- REQUESTS
  WSC1    SCH GEN     4   -------
          SCH USE     4              #    -SERVICE TIME
          SCH MAX     4              REQ      AVG
          PTH         2    SYNC      14176K    33.7
                           ASYNC    871929    370.6
                           CHANGED   99787   INCLUDED
```

Disclaimer and Trademarks

IBM @server zSeries

DFSMS/MVS	IMS/ESA	RAMAC	RMF
Language Environment	S/390	Multiprise	S/390 Parallel
DFSORT	MQSeries	Enterprise Server	Enterprise System/3090
MVS/DFP	MVS/ESA	ES/9000	Sysplex Timer
C/370	ES/3090	System/390	CICS*
ESA/390	OpenEdition	ESCON	OS/390
CICSPlex	Parallel Sysplex*	VTAM	WebSphere*
DB2*	PR/SM	IBM*,	IBM logo*
z/OS	z/VM	zSeries	RACF

Notes:

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

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

IBM considers a product "Year 2000 ready" if the product, when used in accordance with its associated documentation, is capable of correctly processing, providing and/or receiving date data within and between the 20th and 21st centuries, provided that all products (for example, hardware, software and firmware) used with the product properly exchange accurate date data with it. Any statements concerning the Year 2000 readiness of any IBM products contained in this presentation are Year 2000 Readiness Disclosures, subject to the Year 2000 Information and Readiness Disclosure Act of 1998.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the 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.