



LSU Danmark

# **z/OS 1.8, Sysplex, Availability, Scalability and Performance Update**

**Nordic**

**Large Systems Update Seminar**

**October 2006**

*Henrik Thorsen*

**IBM Service Delivery Denmark (SDDK)**

**e-Mail: [henrikthor@dk.ibm.com](mailto:henrikthor@dk.ibm.com)**

# Trademarks

The following are trademarks of the IBM Corporation in the United States and/or other countries.

APPN*	HiperSockets	OS/390*	VM/ESA*
CICS*	HyperSwap	Parallel Sysplex*	VSE/ESA
DB2*	IBM*	PR/SM	VTAM*
DB2 Connect	IBM eServer	Processor Resource/Systems Manager	WebSphere*
DirMaint	IBM e(logo)server*	RACF*	z/Architecture
e-business logo*	IBM logo*	Resource Link	z/OS*
ECKD	IMS	RMF	z/VM*
Enterprise Storage Server*	Language Environment*	S/390*	z/VSE
ESCON*	MQSeries*	Sysplex Timer*	zSeries*
FICON*	Multiprise*	System z	z9
GDPS*	NetView*	System z9	
Geographically Dispersed Parallel Sysplex	On demand business logo	TotalStorage*	
		Virtualization Engine	

\* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

\* All other products may be trademarks or registered trademarks of their respective companies.

**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.

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.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.



**ENABLING BUSINESS.  
A THROUGH Z.**

## Notice & Disclaimers

Copyright © 2006 by International Business Machines Corporation.

No part of this document may be reproduced or transmitted in any form without written permission from IBM Corporation.

Product information and data has been reviewed for accuracy as of the date of initial publication. Product information and data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or programs(s) described herein at any time without notice.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Consult your local IBM representative or IBM Business Partner for information about the product and services available in your area.

Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead. It is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR INFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein.

## Notice & Disclaimers

The performance data contained herein was obtained in a controlled, isolated environment. Actual results that may be obtained in other operating environments may vary significantly. While IBM has reviewed each item for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere.

The responsibility for use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's or user's ability to evaluate and integrate them into their operating environment. Customers or users attempting to adapt these techniques to their own environments do so at their own risk. IN NO EVENT SHALL IBM BE LIABLE FOR ANY DAMAGE ARISING FROM THE USE OF THIS INFORMATION, INCLUDING BUT NOT LIMITED TO, LOSS OF DATA, BUSINESS INTERRUPTION, LOSS OF PROFIT OR LOSS OF OPPORTUNITY.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not necessarily tested those products in connection with this publication 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.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

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

## Agenda

- z/OS strategy / values
- z/OS 1.6 & 1.7 quick revisit
- z/OS 1.8
  - Scalability and Performance
  - Availability
  - Optimization and Management
  - Security
  - Enterprise-wide roles
  - Networking
  - New applications and standards
  - Usability and Skills
  - Bit bucket
    - ▶ Including sysplex and other updates not related to a specific z/OS release
    - ▶ Including IBM statements of direction
    - ▶ An early z/OS 1.9 preview
- Note: The presentation handout does not include all charts presented, nor do it 100% represent the slides presented





LSU Danmark



Danmark LSU October 2006


*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation

## A Different Era for Business

- People, organizations and economies are connected and highly interdependent
- National security / personal privacy
- New Regulatory Environment
- Pressure for revenue and profit growth



# The problem: How can I achieve more with less?



- More responsive to business needs
- More Apps
- More Data
- More security
- More availability





If you were starting from Scratch . . .


*What to require in an on-demand, 24x7, always-on world?*

- Nearly unlimited scalability
- Economies of scale
- Highly automated
- Comprehensive security





## It's Time to Re-THINK the Role of the Mainframe

- Leadership in a new era of business and technology
  - Workload Management
  - Business Resilience
  - Security
  - Data Hub & SOA
  
- New markets
  - Russia, India, China
  
- New uses
  - Enabled by new SW

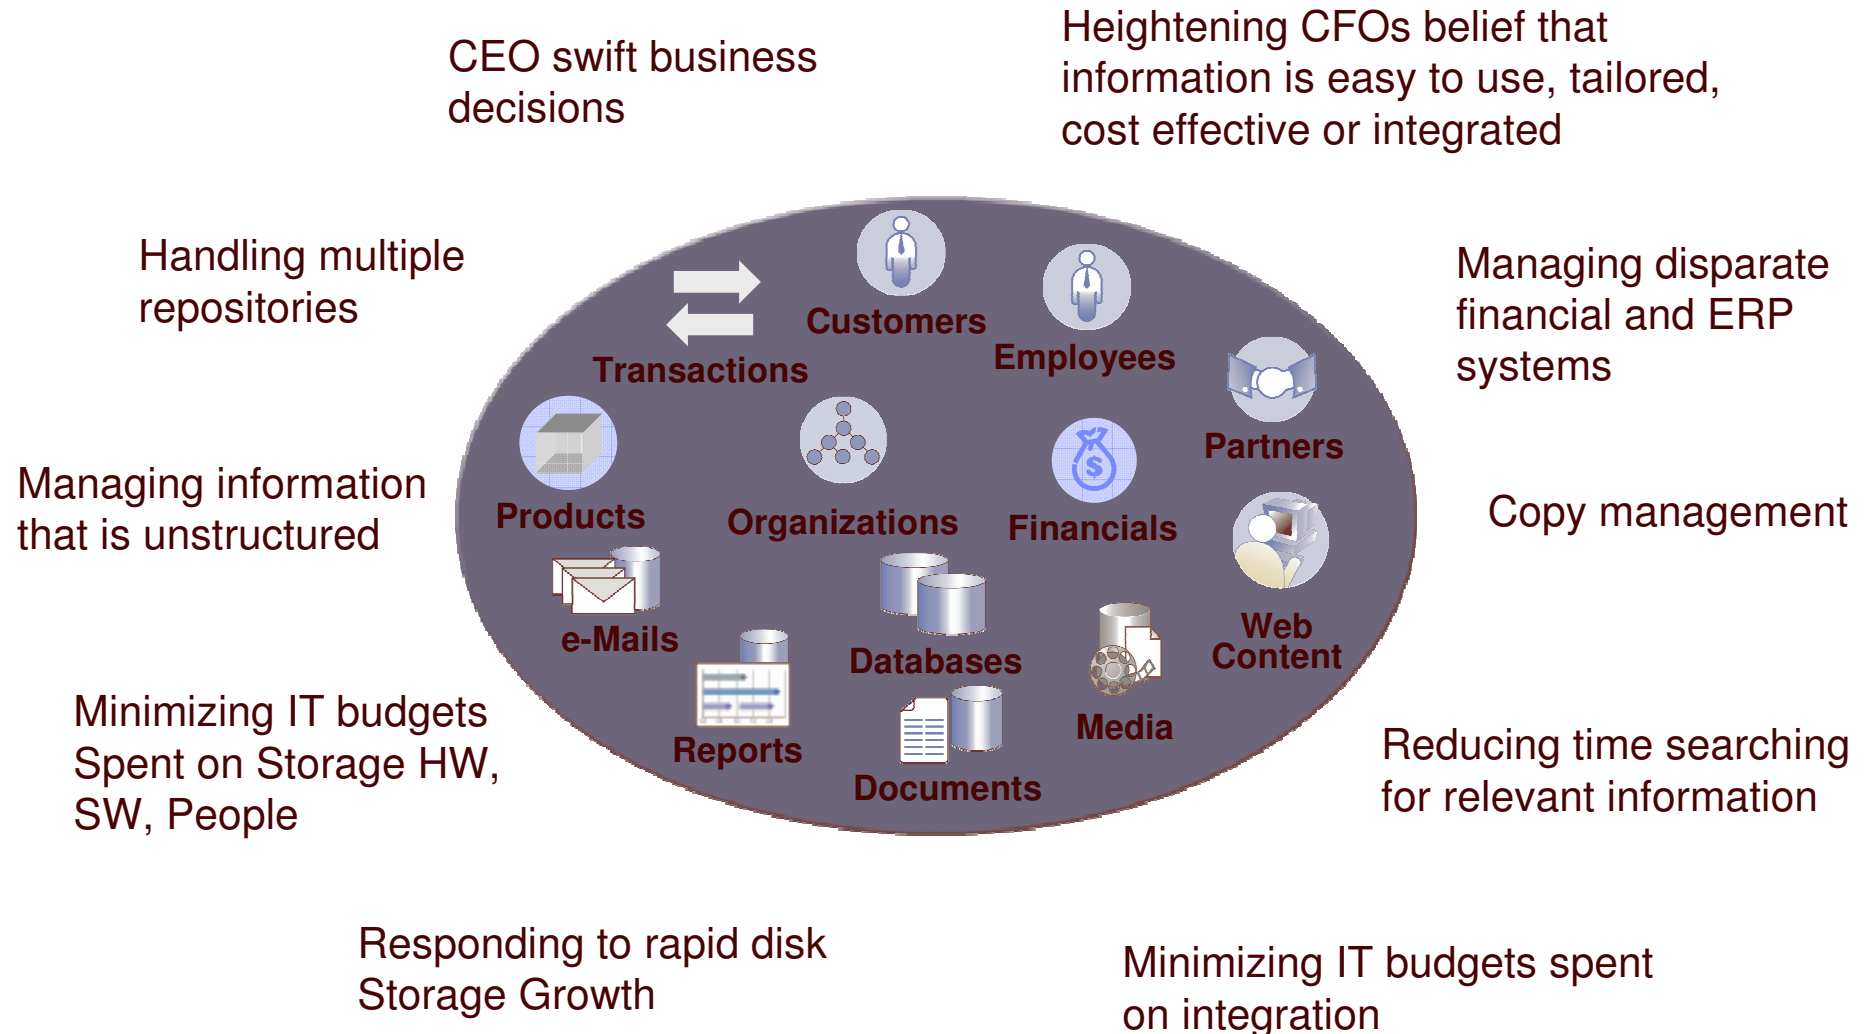


## 4 Enterprise-wide Roles of the Mainframe *extending qualities of service to your **enterprise***

- **Enterprise** business resilience manager
- **Enterprise** security manager
- **Enterprise** workload manager
- **Enterprise** hub for data & SOA



# Data Management Value



## Gartner DBMS survey March 2006

A surprising result was nearly **30 percent** of respondents who intend to transition some of their DBMS workload onto z/OS....


DBMS is one area in which organizations can leverage the power and throughput of the mainframe through the consolidation of systems....

This is further illustrated by the low number of respondents (**7 percent**) who indicated that they intend to migrate some of their workload off of the z/OS platform.

Finally, a percentage of those moving to Linux are moving to a zLinux LPAR on the z9 Series for the same reasons of resource reduction. Many of these shops are moving Oracle from Unix to zLinux, often consolidating hundreds of Unix servers to a single mainframe."

## What's driving the growth?

- Integration of Java-based workloads through zAAPs
- Linux
- Database / application integration
- SOA

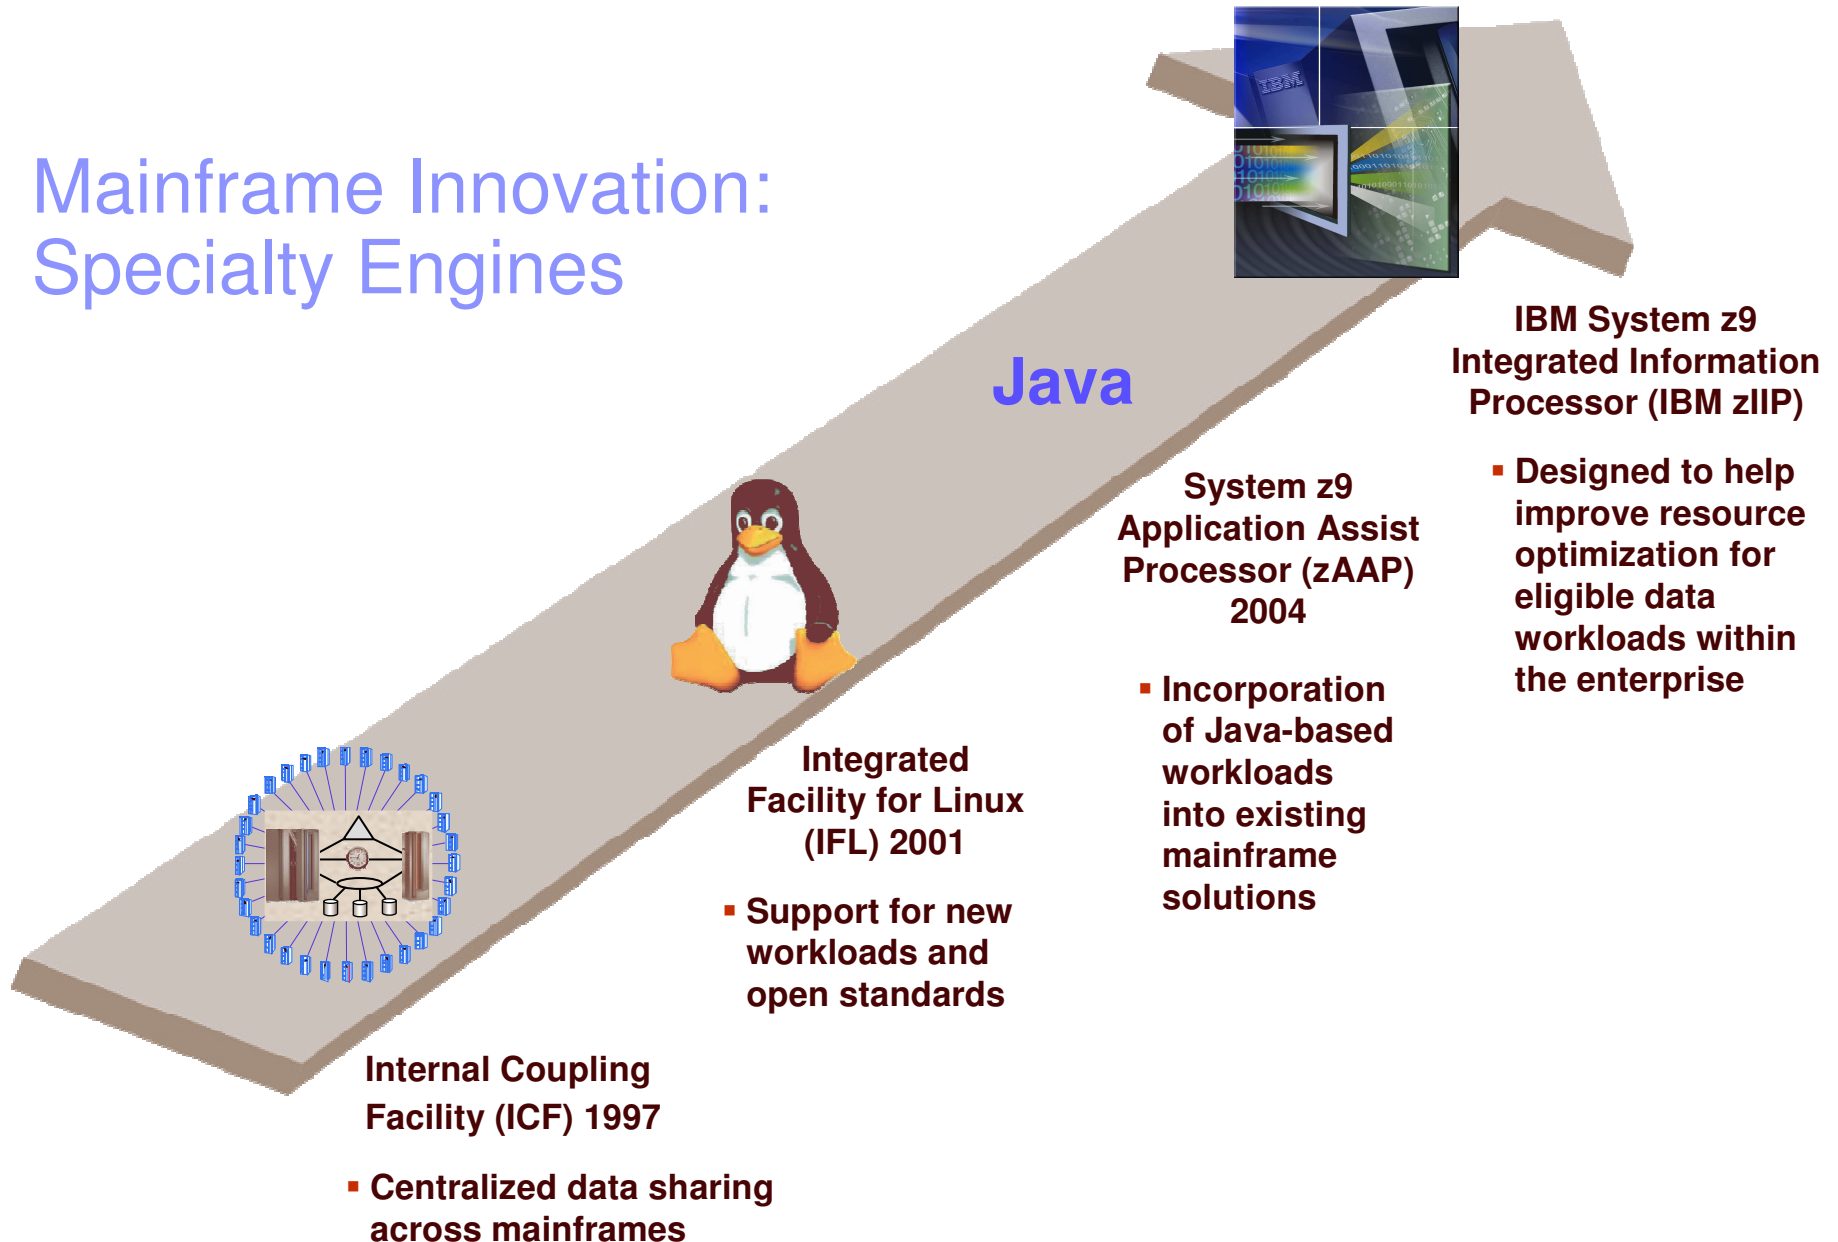


# System z Data Serving Roadmap



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

# Mainframe Innovation: Specialty Engines





## Mainframe Community Ecosystem

- **Customer Councils**
- **IBM Academic Initiative**
  - 20,000 new skills by 2010
  - 250 colleges and universities
  - 300 professors
- **Partners**
  - 1,500 mainframe partners
- **zNextGen**




**Mainframe Student Content Winners  
Mohonk Mountain House, New Paltz**

# WW Servers >\$250K

## Factory revenue 4 Quarter rolling average

Source: IDC Worldwide Quarterly Server Tracker; May, 2006



(1) System z plus other IBM

(2) HP is reported as combined entity of HP and Compaq



## LSU Danmark



z/OS 1.6 (3Q 2004) news  
*SMF,GRS, DFSMS,RRS.....*



Old news

- SMF
  - Buffer constraint relief
- GRS
  - Potential outage reductions
- WLM
  - 64-bit VS support for WAS
  - Sysplex Distributor support for WL balancing TCP/IP servers
  - DB2 stored procedures can prioritize dependent requests
  - DFSMS improved placement of critical DSs on ESS/PAV
- RRS
  - Restart anywhere, anytime...

z/OS 1.6 (3Q 2004) news  
*USS enhancements at a glance*



Old news

- Automove system list enhancements – wild card support
- Increase to 128K per process file descriptor limit
- Automount enhancements
- HFS to zFS automount
- ISHELL enhancements
- LFS support of sysplex zFS – PFS termination in Shared HFS
- Distributed BRLM enhancements
- Fork() accounting for FTP client with SSL
- Superkill command

## z/OS 1.6 (3Q 2004) news


*DFSMS news...*

- DFSMSdfp
  - ✓ SMS volume selection based upon Parallel Access Volume
  - ✓ PDSE restartable address space
  - ✓ MLS SECLABEL in ACS routines
- DFSMSdss
  - ✓ DFSMSdss Replace Unconditional
- DFSMShsm
  - ✓ DFSMShsm Secondary Space Management Multitasking Enhancements
- DFSMSrmm
  - ✓ DFSMSrmm as server for multiple systems
  - ✓ DFSMSrmm ISPF usability

z/OS 1.6 (3Q 2004) news  
*Security Server enhancements*



- 64-bit enhancements
- Enhancements to CDT and Router Table updates
  - ✓ Update the CDT dynamic.
  - ✓ Router Table should only be updated for exceptions
- Class Name restriction relaxed
- SETROPTS LIST enhancements



z/OS 1.6 (3Q 2004) news  
z890/z990 GA3



Old news

- zAAP
  - zAAPs are designed for Java code execution
  - z/OS JVMs assists with the execution of code from standard processors to zAAPs. JVM executes Java code on zAAPs
  - Calls (JNIs) to DB subsystems (e.g., DB2) switch back to standard processors
  - DB returns (JNIs) to Java switch back to zAAPs
- z/OS support for up to 24 CPs per LPAR
- Dynamic add of LPAR



## z/OS 1.6 (3Q 2004) news

*Sysplex news*


- SETXCF reallocate
  - Putting your CF structures where there belong – once and for all!
- New SETXCF START,REALLOCATE command:
  - ✓ Clears all CFRM "policy change pending" conditions
  - ✓ Moves all simplex structures into their "most preferred" CF location
  - ✓ Moves all duplexed structure instances into their two "most preferred" CF locations, in the correct order (automatically corrects "reversal" of primary and secondary structure locations)
  - ✓ Serial, one-structure-at-a-time processing to minimize any disruption caused by reallocation actions
  - ✓ Issues a message describing the evaluation process for each allocated structure. Issues a summary message upon completion of all structures, summarizing actions taken
  - ✓ **Can be used to simplify CF structure movement during disruptive CF maintenance/upgrade scenarios**

z/OS 1.6 (3Q 2004) news

*RMF enhancements - IFA and ESS*



- RMF support for IFA processors
  - ✓ z/OS 1.6 support zAAP(IFA) processors
  - ✓ RMF (PP and MIII) reports enhanced to show IFA processor utilization.
- RMF support ESS Link performance statistics
  - ✓ New ESS Activity report.





## LSU Danmark



## z/OS 1.7 (3Q 2005) news

**Console Restructure Stage 1B**Old news

- The continuation of work started in Console Restructure Stage 1 (delivered in z/OS V1R4.2).
- EMCS Console removal support
  - No way to remove unused EMCS consoles which might lead to long IPL time, and console data refresh time in a sysplex
  - Information of all EMCS Consoles sent across systems
  - New Program IEARECLE now available. Can be executed to discard EMCS consoles
- 1-byte Console ID Elimination
  - z/OS 1.4.2 Track all 1-byte Console ID usage
  - z/OS 1.7 Remove externals (macros, commands) that support 1-byte Console IDs and migration IDs
  - **This was the last release to support 1-byte Console IDs**

## z/OS 1.7 (3Q 2005) news

**JES2**

- NJE over TCPIP
  - Will support established TCP/NJE protocol
  - Overall better RAS characteristics
  - Stronger authentication will also be available
    - Uses SAF/RACF APPCLU class
  - New **NETSERV** Adresspace
  - Major change to JES2 exits
- Availability enhancements & Migration considerations
  - JES2 checkpoint data corruption recovery
    - Designed to detect and correct certain additional kinds of checkpoint control block corruption when JES2 is restarted
  - Spool volume ENQ
    - A sysplex level ENQ is obtained when starting a spool volume to ensure only one MAS allocates this spool volume in the sysplex.
  - Can warm start z2 mode checkpoints only
    - R4 mode is no longer supported.



## z/OS 1.7 Health Checker



Old news

- Health Checker for z/OS integrated into 1.7
  - Back bone shipped in its own FMID HZS7720
  - Current and future checks shipped with individual components.
  - Additional checks can be shipped in the service stream
  - You can develop your own checks
  - Checks may be provided by vendors/3<sup>rd</sup> parties
  
- SDSF support
  - Support to modify checks and view output
  
- IBM health Checker for z/OS can coexist with the web deliverable prototype that has been available for a few years

## z/OS 1.7 Unix System Services



Old news

- Dynamic Service activation
  - **Not** intended to be used to as a complete replacement for regular preventative maintenance application
  - **Only** those PTFs with **++HOLD REASON(DYNACT)** data will be capable of dynamic activation
    - **New** SERV\_LPALIB and SERV\_LINKLIB Parmlib Parameters
    - **New** CMD: F OMVS, ACTIVATE / DEACTIVATE=SERVICE
- Latch contention analysis
  - Additional display capabilities are needed to better identify the reason for which the mount latch is being held and similar information is needed for outstanding cross system messages.
- Mounting from the console
  - SET OMVS=xx console command
- zFS preferred filesystem, HFS will be discouraged
  - Migration tool available to assist in migration from HFS to zFS

## z/OS 1.7 Security Server / RMF

Old news

- Security Server.
  - Mixed Case Password support
    - Switch to turn on Mixed case password
- RMF
  - zFS Support
    - Monitor III zFS reports provides data on
      - zFS response time / wait times
      - zFS cache activity
      - zFS activity / capacity by aggregate
      - zFS activity / capacity by filesystem
  - Disk space monitoring
    - New Monitor III report provides data on
      - Storage space (storage group based)
      - Disk space (volume based)





## LSU Danmark



Danmark LSU October 2006

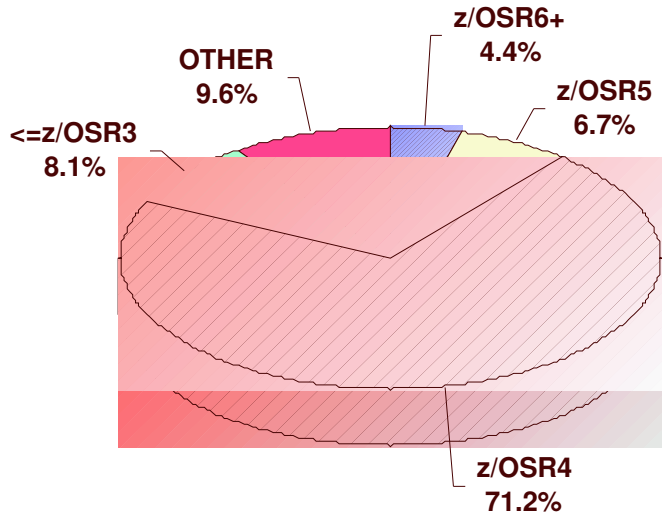
*Henrik Thorsen*

11/5/2007

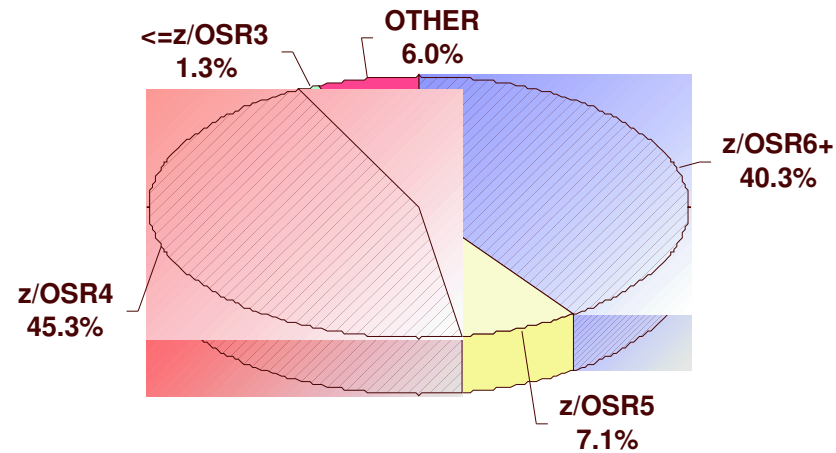
© 2006 IBM Corporation

# z/OS Release Distribution 2Q 2006

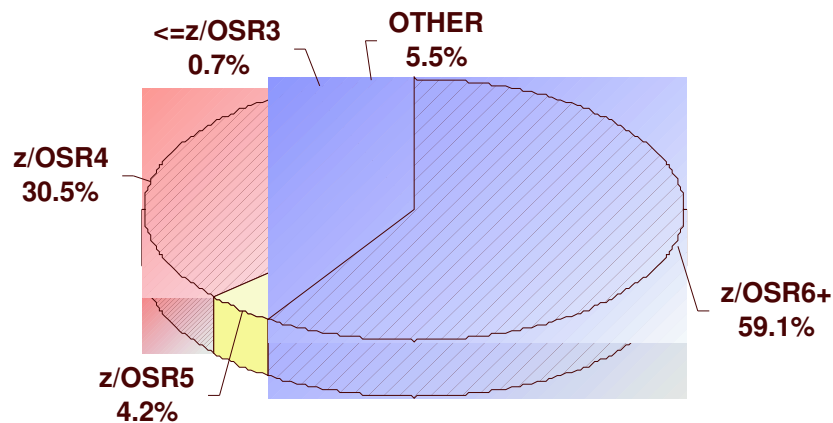
*Over 93% on supported releases*



**YE 2004**

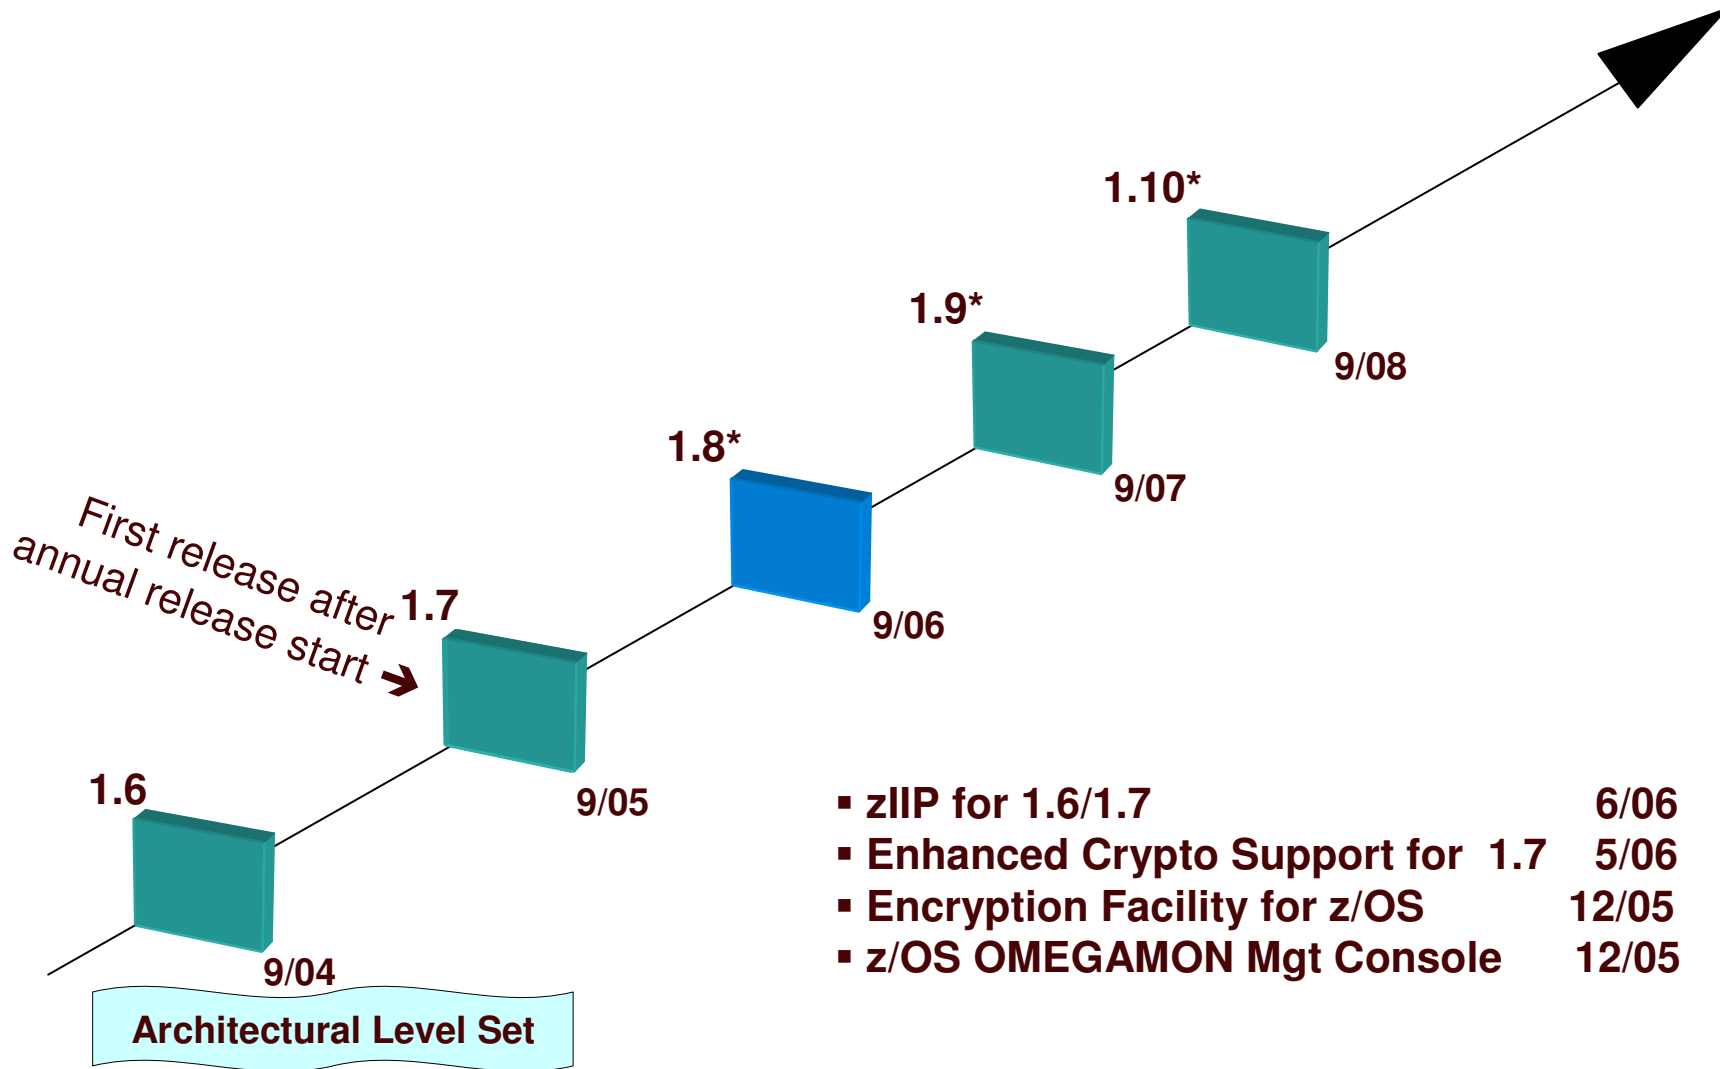


**YE 2005**









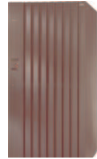
**% of z9 EC and z9xx  
WW System Inventory**

# z/OS Release Roadmap



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

## z/OS Support Summary Dates

									End of Service	Coexists with z/OS...	Planned Ship Date
		G5/G6 Multiprise® 3000	z800	z890	z900	z990	z9-109	DS8000 DS6000			
z/OS	1.1	X	X		X				3/04	1.4	
	1.2	X	X	X <sup>c</sup>	X	X <sup>c</sup>			10/04	1.5	
	1.3	X	X	X <sup>c</sup>	X	X <sup>c</sup>			3/05	1.6	
	1.4	X	X	X <sup>1</sup>	X	X <sup>1</sup>	X <sup>1</sup>	X	3/07	1.7	
	1.5	X	X	X	X	X	X	X	3/07	1.8	
	1.6		X	X	X	X	X	X	9/07	1.8	
	1.7		X	X	X	X	X	X	9/08*	1.9*	9/05*
	1.8*		X	X	X	X	X	X	9/09*	1.10*	9/06*
	1.9*		X	X	X	X	X	X	9/10*	1.11*	9/07*

- z/OS 1.5, z/OS 1.6 and z/OS 1.7 coexistence with z/OS 1.8
- z/OS 1.7 is planned to coexist with z/OS 1.9

x<sup>c</sup> - Compatibility support only

x<sup>1</sup> – z990 compatibility or exploitation feature required

There is no IBM Bimodal Accommodation Offering available for z/OS 1.5 or higher releases.

## z/OS.e Support Summary Dates



		z800	z890	End of Service	Coexistence Migration Policy	Planned Ship Date
z/OS.e	1.3	X	X <sup>c</sup>	3/05	1.6	
	1.4	X	X <sup>e</sup>	3/07	1.7	
	1.5	X	X	3/07	1.8	
	1.6	X	X	9/07	1.8	
	<u>1.7</u>	X	X	9/08*	1.9*	
	1.8*	X	X	9/09*	1.10*	9/06*
	1.9*	X	X	9/10*	1.11*	9/07*

Bimodal Accommodation offering is not available for z/OS.e

X<sup>c</sup> - Compatibility support – does not exploit z890 features: 30 LPARs, 2 LCSS

X<sup>e</sup> - Requires exploitation support feature

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

## z/OS Release Coexistence and Migration/Fallback

Release	General Availability (GA) of Release Identified in Column 1	Service Support of Release Identified in Column 1 is Available Through	Releases Coexistence and Migration/Fallback supported with Release Identified in Column 1
R10	Sept 2000	Sept 2004	R10, R9, R8, R7, R6 (1)
z/OS R1	March 2001	March 2004	{z/OS R1/R10}, R9, R8, R7, R6 (1,2)
z/OS R2	October 2001	Oct 2004	z/OS R2, {z/OS R1/R10}, R9, R8 (2)
z/OS R3	March 2002	March 2005	z/OS R3, z/OS R2, {z/OS R1/R10}, R9 (2)
z/OS R4	Sept 2002	March 2007	z/OS R4, z/OS R3, z/OS R2, {z/OS R1/R10} (2)
z/OS R5	March 2004	March 2007(4)	z/OS R5, z/OS R4, z/OS R3, z/OS R2
z/OS R6	Sept 2004	Sept 2007 (4)	z/OS R6, z/OS R5, z/OS R4, z/OS R3 (5)
z/OS R7	Sept 2005 (3)	Sept 2008 (4)	z/OS R7, z/OS R6, z/OS R5, z/OS R4
z/OS R8	Sept 2006 (3)	Sept 2009 (4)	z/OS R8, z/OS R7, z/OS R6, z/OS R5 (6)
z/OS R9	Sept 2007 (3)	Sept 2010 (4)	z/OS R9, z/OS R8, z/OS R7

Legend:

Where no product name is indicated in the table, OS/390 should be assumed

z/OS R7 - z/OS R9 used for illustrative purposes, release numbering not a guarantee of actual release number

(1) OS/390 R6 -> OS/390 R10 provided as an exception

(2) OS/390 R10 & z/OS R1 treated as one coexistence level

(3) Represents projected GA date

(4) Represents projected end of service support date (GA+3 years)

(5) z/OS R2 excluded since service for z/OS R2 ends within 1 mo. of z/OS R6 GA

(6) z/OS R4 is not planned to be coexistence, migration, or fallback supported with z/OS R8

## z/OS migrating to 1.8

For z/OS V1R8, the z/OS Migration book (**GA22-7499**) is expanded to cover three migration scenarios:

1. z/OS V1R7 to z/OS V1R8
2. z/OS V1R6 to z/OS V1R8
3. z/OS V1R5 to z/OS V1R8 (which is the entire book)



## z/OS migrating to 1.8...

- z/OS 1.8 *Migration* Book available now on z/OS 1.8 bookshelf:
  - [http://publibz.boulder.ibm.com/cgi-bin/bookmgr\\_OS390/Shelves/EZ2ZO10H](http://publibz.boulder.ibm.com/cgi-bin/bookmgr_OS390/Shelves/EZ2ZO10H)
- Remember--z/OS 1.8 requires an IBM System z server:
  - IBM System z9 EC (z9 EC)
  - IBM eServer zSeries 990 (z990)
  - zSeries 900 (z900)
  - System z9 BC
  - zSeries 890 (z890)
  - zSeries 800 (z800)
- Simplify migration, ordering, and maintenance:
  - ServerPac electronic delivery available since January 2005
  - SMP/E Internet Service Retrieval available since September 2005
  - SystemPac electronic delivery available since October 2005
  - ShopzSeries enhancements made in January 2006



Check out [www.ibm.com/zseries/zos/migration](http://www.ibm.com/zseries/zos/migration) for more details.



IBM: z/OS V1R8.0 migration and installation - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Address <http://www-03.ibm.com/servers/eserver/zseries/zos/installation/>

Country/region [select] Terms of use

Home Products Services & industry solutions Support & downloads My account

Servers > Mainframe servers > z/OS >

## z/OS V1R8.0 migration and installation

You can find the following installation information topics on this Web page:

- [z/OS V1R8.0 installation planning](#)
- [Ordering z/OS and related products](#)
- [z/OS installation-related publications](#)  
[V1R8](#) | [V1R7](#) | [V1R6](#) | [V1R5](#) | [V1R4](#) | [V1R3](#) | [V1R2](#) | [V1R1](#)
- [Other useful resources](#)

**z/OS V1R8.0 installation planning**

[z/OS V1R8.0 and z/OS.e V1R8.0 Installation Planning Wizard](#)  
 Installing z/OS using ServerPac or CBPDO? The z/OS V1R8.0 Installation Planning Wizard can guide you through all the z/OS installation planning tasks.

**IBM fee offerings**

[Worldwide CustomPac Offerings](#)  
 Find out about customized software packages to install z/OS and

**z/OS migration & installation resources**

- [z/OS migration & installation Web pages](#)  
[V1.8](#) | [V1.7](#) | [V1.6](#)  
[V1.5](#) | [V1.4](#) | [V1.3](#)  
[V1.2](#) | [V1.1](#)

**SystemPac price reduced**

- [U.S. SystemPac now more affordable and easier than ever](#)

**z/OS migration manuals**

- [z/OS migration manuals available for z/OS V1.4 through z/OS V1.8](#)


**Internet delivery of ServerPac available for z/OS**

Slide 18 of 306 Blue Pearl Deluxe English (U.S.)

Start Inf... 5:3... 6 L... Rec... 2 W... Win... 1-2... 7 M... EN 97% 3:25 PM



## LSU Danmark



Danmark LSU October 2006

*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation

# z/OS 1.8

## Helping meet challenges of on demand business

### Improving Usability and Skills

More Health Checks and framework improvements, HCM, ISPF, InfoPrint Central, & RRS improvements, Extensions to the IBM Configuration Assistant for z/OS Comm Server, EE usability improvements, ...

### Integrating new Applications and Supporting Industry and Open Standards

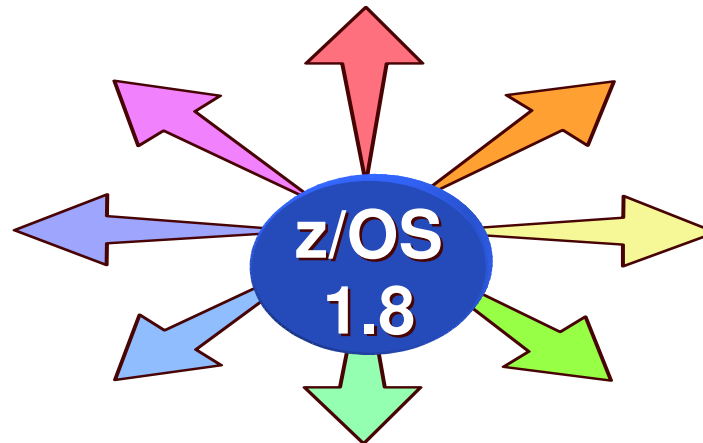
Unicode improvements; inittab support; LE improvements, new LDAP server, z/OS XML System Services, BPXBATCH, CIM, XPLINK, Networking APIs, XL C/C++ enhancements, ...

### Extending the Network

Sysplex partitioning support, Dynamic DNS registration, JES3 NJE via TCP/IP, TN3270 and telnet improvements, IP filtering, IKE, IPsec, WTS, Samba, improved sysplex failure recovery, REXX FTP interface, ...

### Scalability & Performance

Support for up to 4 TB memory, New zIIP processor, CFRM performance improvements, GRS ENQ limits, Parallel VARY ONLINE, z/OS UNIX file descriptors, Device Groups, DADSM/CVAF Rapid Index Rebuild, z/OS USS asynch socket fast path, 64-bit PDSE support ...



### Enterprise-Wide Roles

WLM/EWLM Service Class Correlation, zAAP Reporting for eWLM, DFSMSrmm CIM agent, RMFeServer OS Monitoring Stage II, WLM delay services mapped to ARM ...

### Improving Availability

Master console elimination, Logstream rename, test and production logstream separation, Recoverable BRLM, GDPS enhanced recovery, Fast replication, HyperSwap trigger, z/OS UNIX latch contention detection...

### Optimization and Management

WLM, zAAP Stage 2, Sysplex Distributor, Tape I/O priority, JES2 balancing of WLM-managed initiators, GRS CNS processing, Improved OAM BLOB support, SCLM improvements, Group Capacity Limit ...

### Enhancing Security

PKI extensions, Distributed Identity Support, RACF password phrase support, Tape data set protection, 128-bit AES for IPsec, SAF identity tokens, RACF virtual key rings, IDS configuration, ...

## z/OS 1.8 enhancements at a glance

### *the flagship operating system*

- Scalability and performance improvements
- Availability improvements
- Optimization and management
- Enhanced security
- Enterprise-wide roles
- Extending the network
- Integrating new Applications and Supporting Industry and Open Standards
- Improving Usability and Skills
- Bit bucket
  - **Statement of directions & functions withdrawn**

# z/OS 1.8

## Helping meet challenges of on demand business

### Improving Usability and Skills

More Health Checks and framework improvements, HCM, ISPF, InfoPrint Central, & RRS improvements, Extensions to the IBM Configuration Assistant for z/OS Comm Server, EE usability improvements, ...


### Integrating new Applications and Supporting Industry and Open Standards

Unicode improvements; inittab support; LE improvements, new LDAP server, z/OS XML System Services, BPXBATCH, CIM, XPLINK, Networking APIs, XL C/C++ enhancements, ...

### Extending the Network

Sysplex partitioning support, Dynamic DNS registration, JES3 NJE via TCP/IP, TN3270 and telnet improvements, IP filtering, IKE, IPsec, WTS, Samba, improved sysplex failure recovery, REXX FTP interface, ...

**Scalability & Performance**  
 Support for up to 4 TB memory, New zIIP processor, CFRM performance improvements, GRS ENQ limits, Parallel VARY ONLINE, z/OS UNIX file descriptors, Device Groups, DADSM/CVAF Rapid Index Rebuild, z/OS USS asynch socket fast path, 64-bit PDSE support ...



### Enterprise-Wide Roles

WLM/EWLM Service Class Correlation, zAAP Reporting for eWLM, DFSMSrmm CIM agent, RMFeServer OS Monitoring Stage II, WLM delay services mapped to ARM ...

### Improving Availability

Master console elimination, Logstream rename, test and production logstream separation, Recoverable BRLM, GDPS enhanced recovery, Fast replication, HyperSwap trigger, z/OS UNIX latch contention detection...

### Optimization and Management

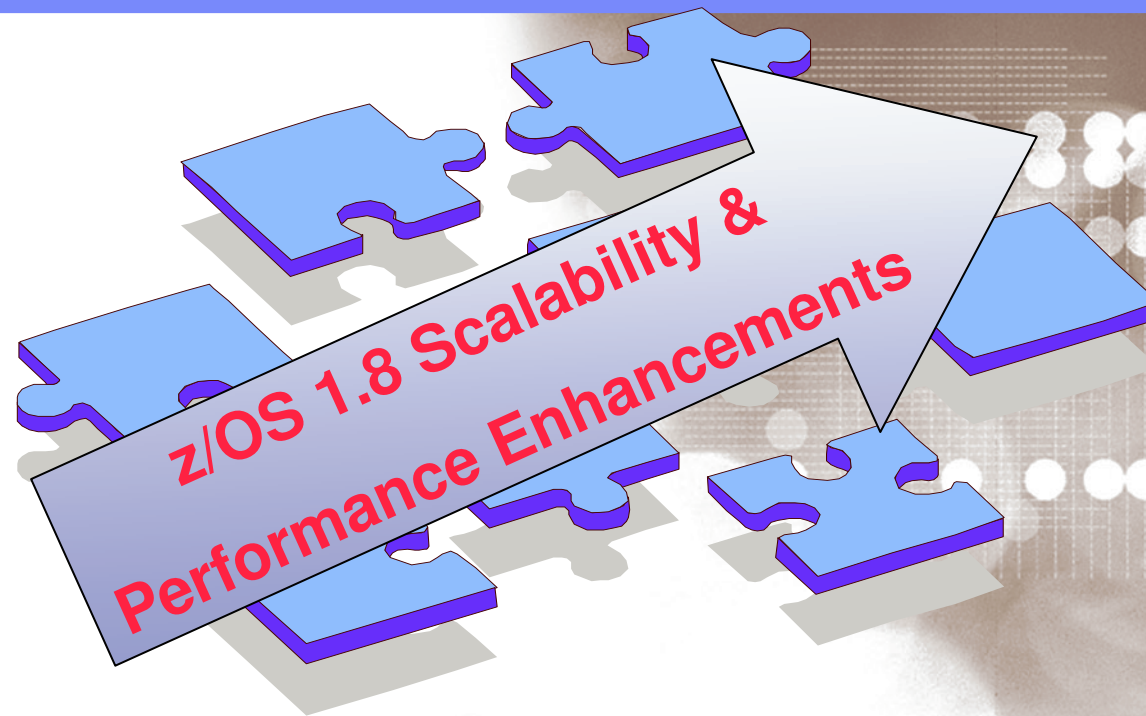
WLM, zAAP Stage 2, Sysplex Distributor, Tape I/O priority, JES2 balancing of WLM-managed initiators, GRS CNS processing, Improved OAM BLOB support, SCLM improvements, Group Capacity Limit ...

### Enhancing Security

PKI extensions, Distributed Identity Support, RACF password phrase support, Tape data set protection, 128-bit AES for IPsec, SAF identity tokens, RACF virtual key rings, IDS configuration, ...



LSU Danmark



Danmark LSU October 2006

*Henrik Thorsen*

11/5/2007


© 2006 IBM Corporation

# IBM System z: Balanced System Design

*z/OS designed to support faster CPUs, more memory and bandwidth*

System I/O Bandwidth

**Balanced System  
CPU, nWay, Memory,  
I/O Bandwidth\***



\*z9 EC exploits a subset of its designed I/O capability

## z/OS Scalability

*more: memory, GRS ENQs, GRS VSCR, faster UNIX*



- Support for more real memory
  - **New z/OS limit will be 4 TB; old limit was 128 GB**
  - Up to 512 GB supported on z9 EC, up to 256 GB on z990
- GRS support for more concurrent ENQs
  - Default limits unchanged: Unauthorized ENQ default limit remains 4096; authorized default ENQ limit is 250,000; can specify higher maximums than the defaults
  - **New maximum is 2G**
  - New ISGADMIN authorized interface, T GRS command, and GRSCNFxx parmlib support for setting higher maximums
- GRS 31-bit constraint relief
  - **Move star-mode global QCBs and QCBS control blocks above the bar**
  - **Better compaction for remaining QCBs**
- z/OS UNIX asynch socket read/write now designed to use fast-path processing
  - **Improve performance for applications that use asynchronous I/O or that run in SRB mode.**



## z/OS Scalability

*more performance: CFRM, VARY processing, LE*

- CFRM performance enhancements Stage 2
  - Designed to allow more systems, structures, and connectors to be added without availability impacts
- Parallel VARY ONLINE processing:
  - Reduce duration of VARY commands for large numbers of devices, reduced serialization contention
  - Complements parallel VARY OFFLINE processing, in z/OS 1.7
- Language Environment support for sequential data sets larger than 64K tracks and for VSAM extended addressability data sets
  - QSAM support in Language Environment for C/C++ programs (using noseek)
  - Support for ESDS, KSDS, RRS extended format data sets with extended addressability



# CFRM Performance Enhancement

## *Background*

- Coupling Facility Resource Management (CFRM) manages Coupling Facility resources and is critical to the operation of a high-availability parallel sysplex environment.
- **The CFRM couple data set (CDS) is the centralized “control point” used by CFRM to manage Coupling Facilities (CF) and CF Structures.**
- High availability in a Parallel Sysplex is achieved by enabling surviving systems to recover rapidly for a failed system within the sysplex and allowing CF structures that fail to be quickly rebuilt or recovered via duplexing failover.
- **Poor CFRM performance can adversely impact the availability of CF data and thus the availability of products and subsystem using the structures.**
- Surviving systems participate in cleanup of the failed connection which requires **“one system at a time”** to access the CFRM CDS. Recovery time is applicable to structure rebuild/duplexing failover times, and to removal of a failed system from a sysplex.

# CFRM Performance Enhancement

## Stage 0 and 1 - z/OS 1.4



Old news

- Recovery actions (CEC failure, CF failure, Structure rebuild) may result in large records requiring large I/O bandwidth
  - originally: 64 structures with 32 connectors
  - currently: 1024 structures with 255 connectors
    - >120X increase
- Up to 30% CFRM CDS I/O performance improvement
- **Reduced possibility of I/O bottleneck on CFRM CDS**
- Improved Parallel Sysplex availability
- CFRM performance enhancements Stage 0 delivered in the service stream added support for **optimized sysplex failure cleanup processing** for CF Structures.

# CFRM Performance Enhancement

## Stage 2 – z/OS 1.8

- **Improve Parallel Sysplex Availability by reducing the elapsed time for key Sysplex recovery events such as:**
  - CF structure rebuilds
  - CF Duplexing failover and establishing/re-establishing CF Duplexing
  - Connecting to/disconnecting from a CF structure
  - Sysplex partitioning cleanup for CF structures
- **Focus on significantly reducing CDS I/O**
  - Processing to manage the coordination of CF Structure rebuild/duplexing failover and system recovery use a XCF **message-based protocol** to reduce serialized access to the CFRM CDS.
  - The message-based protocol defines a single system as manager and all other systems as participants.
  - The manager system is responsible for coordinating the processes involved with the participating systems and updating the CFRM CDS when needed.
  - The manager system does the I/O to the CFRM CDS which greatly reduces the number and frequency of serialized access.

## CFRM Performance Enhancement

### *Installation and Migration*

- A new version of CFRM CDS required for systems to enable message-based processing.
- A CDS formatted to enable message-based processing also enables all functions supported by lower level versions of the CDS.
- All systems in sysplex must be z/OS 1.8 or higher.
- Systems that do not support message-based processing will not be able to join a sysplex that is using a CDS that was formatted for message-based processing.
- A CDS formatted for message-based processing cannot be brought into use by a sysplex while there are 1 or more down-level systems (z/OS 1.7 or lower) in the sysplex.
- [SA22-7625 – z/OS V1R8 MVS Setting up a Sysplex](#)


## z/OS Scalability

*more: UNIX file descriptors, device groups, PDSE VTOC rebuild performance*

- z/OS UNIX limit on file descriptors per process increased from **128K** to **512K**
  - MAXFILEPROC limit now 512K in BPXPRMxx and SET OMVS command
  - Can restrict individual users by setting FILEPROC MAX in the OMVS segment
  - Mostly a TN3270 CICS issue, raised from 64K in 1.6, **better performance**
- Support for more than **32K** device groups
  - DGs consumed by tape devices (1 per dev) and DASD esoteric names (1 per dev. range per esoteric name).
  - You can reach this limit with fewer than the 64K maximum number of devices
  - APAR OA02983 increased limit to **64K** devices in 03
  - The new limit will be **4G**
  - Note: check your system by looking at the number of entries in the Group Pointer Table using the IPCS LISTEDT command.
- DADSM/CVAF Rapid VTOC Index Rebuild
  - Designed to help speed VTOC conversions from non-indexed to indexed
- 64-bit exploitation by PDSE
  - SMSPDSE and SMSPDSE1 will use storage above the bar for control blocks
  - New SMS initialization parameter to specify amount of storage to use above bar
  - Relieves the prior limitation of about a million open PDSE members
  - Option to retain directory and member cache for closed data sets and SMF14/15 cache



LSU Danmark



Danmark LSU October 2006

*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation

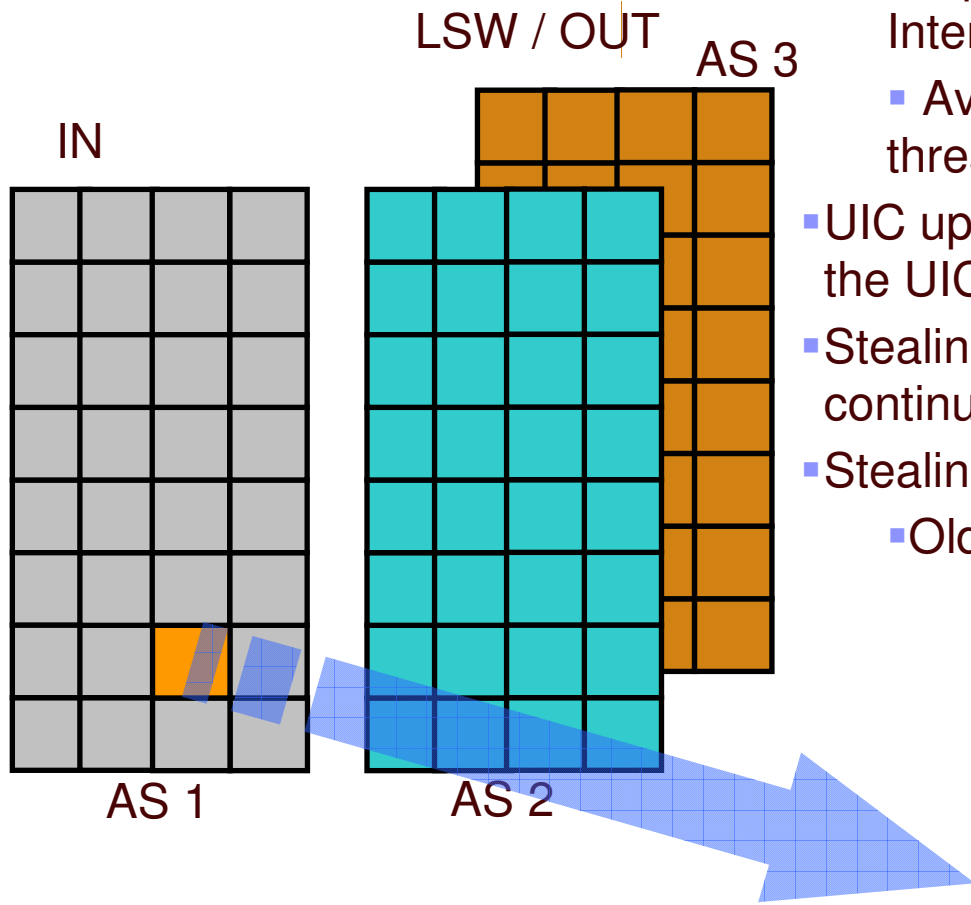
## z/OS 1.8 RSM/SRM Support for Large Real storage

- Support up to 4 Terabytes of real storage
- Page Replacement and UIC enhancements
  - The CPU cost of the current function continues to grow as we move to systems with more and more real storage while its effectiveness is diminishing. This algorithm was changed to perform more efficiently in large real systems.
- Physical Swap Processing
  - Physical Swaps to Auxiliary storage eliminated
- Pageable Storage Shortages
  - **Used to be resolved by physical swaps**
  - **Often RSM spin loops during swap-in**
  - They are now resolved by **frame exchanges**

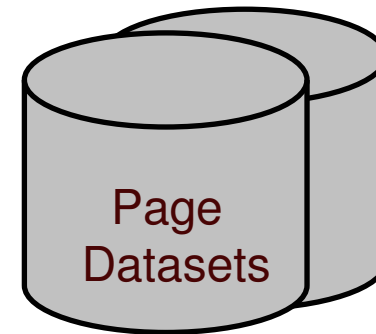


# Page Replacement

*prior to z/OS 1.8*




- Today's LRU algorithm
  - Implemented by keeping an Unreferenced Interval Count (UIC)
  - Available Frame Queue (AVQ) Low/OK thresholds
- UIC update process runs periodically to update the UIC of each in-use frame
- Stealing starts when we go below AVQ Low and continues until AVQ OK is reached
- Stealing happens on an AS basis
  - Oldest frames are stolen first




## Page Replacement

*z/OS 1.8 and beyond*

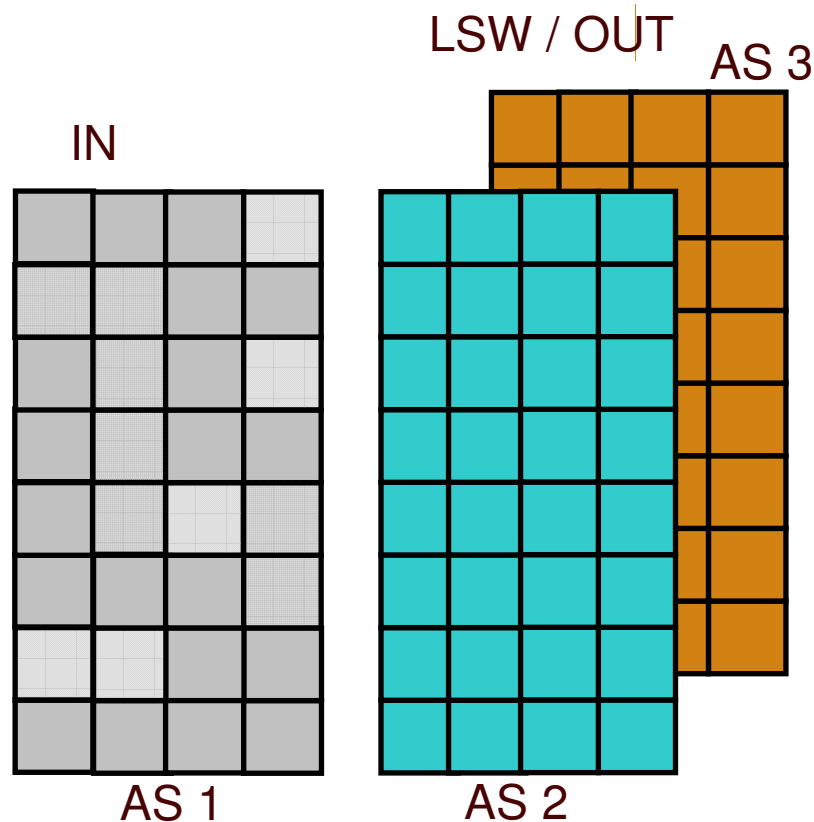


- Why Change?
  - Disruption of the UIC update process is intolerable as the amount of real storage allocated to ASs increases
- Enhanced page replacement algorithm:
  - **Runs when the AFQ needs to be replenished**
  - **Stealing happens on a global basis**



## UIC Calculation

*prior to z/OS 1.8*



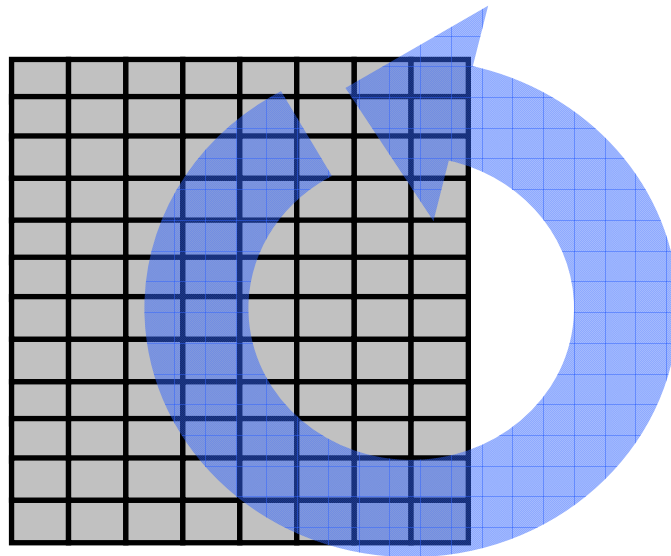
- UIC Update process scheduled every 10 secs.
- Process analyzes **every pageable frame** from each AS on the Swapped-IN queue
- The analyzed frames are counted, then distributed into age buckets.
- The oldest frame in an AS represents the **AS High UIC**
- The highest AS UIC represents the **Highest System UIC** (MCVSTCRI)

While stealing happens on a demand basis, UIC update still needs to run and consume **considerable CPU** while **holding AS locks**, even in environments where the system is not paging. While UIC update is running for a particular AS, RSM can not perform any other functions for that AS (such as service page faults) due to UIC is holding locks.

■

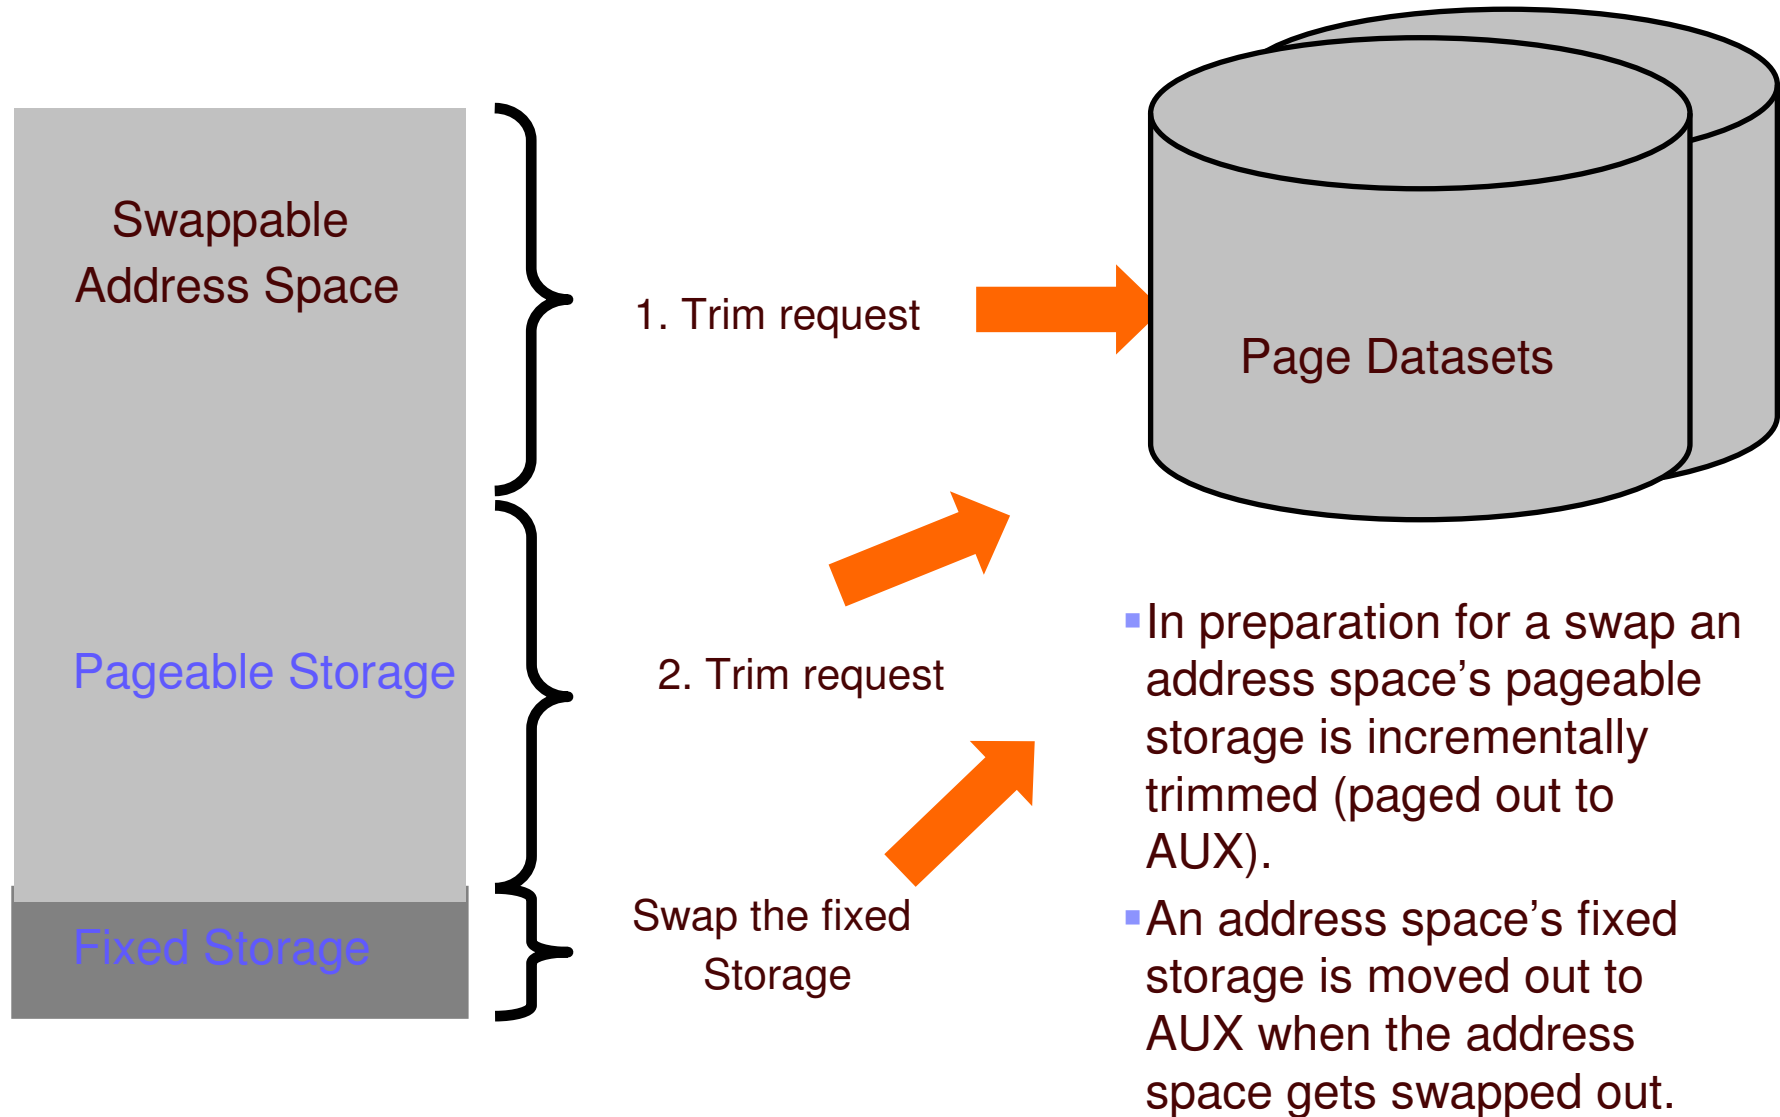
## UIC Calculation

### *z/OS 1.8 and beyond*



- Page replacement algorithm enhanced to efficiently process large amounts of RS.
  - z/OS 1.8 defines the UIC as a single walk through the whole storage in seconds.
  - As a result the UIC values you might see in an RMF report will vary from 0–65535 (18 hours), in the past the UIC value range was 0-2540.
    - The higher the UIC value the less contention for storage in the system
    - A very low UIC indicates that the system is storage constrained.
- 
- 3 UICs displayed by performance monitors:
    - Current UIC (calculated every 1 sec)
    - Minimum UIC (last walk through)
    - Maximum UIC (last walk through)

## Physical Swap - pre z/OS 1.8



## Physical Swap – z/OS 1.8 ...

- z/OS 1.8 will no longer physically swap AS to AUX
  - When fixed storage an AS owns increases, swapping an AS out no longer viable.
- Use RMF to analyze pre-z/OS 1.8 systems for the amount of physical swaps:
  - Next for each AS on the OUT Ready or OUT Wait queue you need to find out the number of fixed frames each AS owns.

C P U   A C T I V I T Y   R E P O R T

SYSTEM ADDRESS SPACE ANALYSIS                      SAMPLES = 3,600

NUMBER OF ADDRESS SPACES

	----- QUEUE TYPES -----				----- ADDRESS SPACE TYPES -----						
	IN READY	IN	OUT READY	OUT WAIT	LOGICAL OUT RDY	LOGICAL OUT WAIT	BATCH	STC	TSO	ASCH	OMVS
MIN	1	54	0	2	0	83	5	119	10	0	3
MAX	10	85	0	0	0	118	19	154	12	0	23
AVG	2.4	56.4	0.0	1.2	0.0	105.4	6.1	141.7	10.5	0.0	3.5

## Physical Swap – z/OS 1.8 ...

- Use RMF STORF Report to find number of frames each AS owns that are fixed.
  - This is storage that can never move out to AUX in z/OS 1.8, so in this case you may consider increasing the amount of real storage in your system.


Samples: 119    System: SYSE    Date: 03/15/06    Time: 16.29.00    Range: 120    Sec

Jobname	C	Class	Cr	-- Frame Occup. --			-- Active Frames --			AUX	PGIN	ES
			TOTAL	ACTV	IDLE	WSET	FIXED	DIV	SLOTS	RATE	RATE	
DFSZFST	S	SYSSTC	44189	44189	0	44189	388	0	0	0		
RMFGAT	S	SYSSTC	30053	30053	0	30053	199	0	0	0		
OMVS	S	SYSTEM	19717	19717	0	19717	394	0	0	0		
GRS	S	SYSTEM	15893	15893	0	15893	424	0	0	0		
SMSVSAM	S	SYSTEM	11158	11158	0	11158	473	0	0	0		
XCFAS	S	SYSTEM	9836	9836	0	9836	2885	0	0	0		
BENK6	O	OE	9003	9003	0	9003	119	0	0	0		
CANEM3	S	STCDEF	8731	8731	0	8731	245	0	0	0		
CONSOLE	S	SYSTEM	8442	8442	0	8442	145	0	0	0		
IXGLOGR	S	SYSTEM	7465	7465	0	7465	135	0	0	0		
RMF	S	SYSSTC	7388	7388	0	7388	110	0	0	0		
MVSNFSS	S	STCDEF	6093	6093	0	6093	181	0	0	0		
TCPIP	S	SYSSTC	4701	4701	0	4701	102	0	0	0		
*MASTER*	S	SYSTEM	3992	3992	0	3992	3231	0	0	0		
MVSNFSCS	S	SYSSTC	3870	3870	0	3870	111	0	0	0		
CANEM2	S	STCDEF	3714	3714	0	3714	119	0	0	0		

- Find out which AS were physically swapped
- Use STORF report to find out how many fixed frames each AS owns, when the address space is Swapped-IN.

## Pageable Storage Shortages – z/OS 1.8 ...

- Pre-z/OS 1.8 pageable storage shortages were resolved via physical swaps.
- In z/OS 1.8 an “In-Real Swap” occurs where frames in the shortage area are exchanged with other frames:
  - For a pageable storage shortage between 16M-2G frames will be exchanged with frames above 2G.
  - For pageable storage shortages below 16M frames will be exchanged with frames above 2G or above 16M.
- Message IRA404I lists the five largest users of fixed frames in the shortage area





## RMF 1.4 Support for UIC Overview

- RMF replaces the Highest System UIC (MCVSTCRI) by the Current System UIC (MCTCurSystemUIC) in several reports
- RMF Overview Conditions AVGHUIC and MAXHUIC are based on the new Current System UIC, too
- SMF record type 71 (paging activity) is extended
  - Minimum, Maximum and Current System UIC values observed during the reporting interval saved in paging data section
- SMF record type 79 subtype 3 (storage/processor data) and subtype 4 (paging activity data) is changed
  - Fields R793CRI and R794CRI no longer hold MCVSTCRI but MCTCurSystemUIC

## New Message to Indicate High AUX Usage

- **IRA205I 50% AUXILIARY STORAGE ALLOCATED**

The message is issued when more than 50% of the Auxiliary storage slots are in use.

- This allows an automation product to add additional page datasets before a critical shortage occurs
- The message is repeated every 2 hours as long as the Auxiliary slot usage is above 50%
- You should add more page data sets to your system already when utilization exceeds 30%,
  - the slot allocation algorithms become less efficient, and may degrade I/O performance.

# Usage & Invocation RMF

## Postprocessor

```

                P A G I N G   A C T I V I T Y

z/OS V1R8          SYSTEM ID SYSF          DATE 02/28/2006          INTERVAL 14.59.999
                   RPT VERSION V1R8 RMF    TIME 15.30.00          CYCLE 1.000 SECONDS

OPT = IEAOPT00    MODE = ESAME          CENTRAL STORAGE MOVEMENT RATES - IN PAGES PER SECOND
-----
HIGH UIC (AVG) = 60333    (MAX) = 65535    (MIN) = 50888

                WRITTEN TO          READ FROM          *--- CENTRAL STORAGE FRAME COUNTS ----*
                CENTRAL STOR        CENTRAL STOR        MIN           MAX           AVG

HIPERSPACE  RT          0.00          0.00          2             2             2
PAGES

VIO         RT          0.00          0.00          0             0             0
PAGES
    
```

- The HIGH UIC fields MIN, MAX and AVG do no longer format the old UIC fields SMF71LIC, SMF71HIC and SMF71ACA
- The new Current System UIC fields SMF71UAC, SMF71UHC and SMF71ULC
- ➔ Highest values possible change from 2540 to 65535

# Migration & Coexistence Considerations RSM/WLM

- Analyze your pre-z/OS 1.8 system and pay special attention to the amount of physical swaps that occur on your system. In most systems physical swaps never occur, so the fact that physical swaps were eliminated in z/OS 1.8 should have no impact.
- Use the RMF CPU Activity Report to find out how many physical swaps occur in your particular system configuration (OUT Ready and OUT Wait).
- If the RMF report shows many address spaces on the OUT Ready or OUT Wait queue, you should consider increasing the amount of real storage on your system.

CPU ACTIVITY REPORT

SYSTEM ADDRESS SPACE ANALYSIS                      SAMPLES = 3,600

NUMBER OF ADDRESS SPACES


	IN		QUEUE TYPES		LOGICAL		ADDRESS SPACE TYPES				
	READY	IN	OUT READY	OUT WAIT	OUT RDY	OUT WAIT	BATCH	STC	TSO	ASCH	OMVS
MIN	1	54	0	0	0	83	5	119	10	0	3
MAX	10	85	0	0	0	118	19	154	12	0	23
AVG	2.4	56.4	0.0	0.0	0.0	105.4	6.1	141.7	10.5	0.0	3.5

# OA14409 IEAOPT MCCAFC TH Enhancement

- MCCAFC TH (lowvalue, okvalue) Available on z/OS R6 and above
  - Specifies the **LOW** and **OK** threshold values for central storage.
    - The low value indicates the number of frames on the available frame queue when stealing begins.
    - The okvalue indicates the number of frames on the available frame queue when stealing ends.
    - SRM will automatically adjust the actual threshold values based on measurements of central storage usage
  - The **Initial** lowvalue and okvalue are calculated by the following formula:
    - $RCEAFCLO = \text{MAX}(\text{OPT Parameter (lowvalue)}, 400, 0.2\% \text{ of the pageable storage})$
    - $RCEAFCOK = \text{MAX}(\text{OPT Parameter (okvalue)}, 600, 0.4\% \text{ of the pageable storage})$
    - But capped to 35/70k, respectively
  - The formulas above will enforce the value, if the OPT parameter has specified a to smaller value
- RCEAFCLO will vary between Initial Value and 70K frames
- RCEAFCOK will vary between Initial Value and 140K frames



LSU Danmark



Danmark LSU October 2006  
*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation

# z/OS 1.8

## Helping meet challenges of on demand business

### Improving Usability and Skills

More Health Checks and framework improvements, HCM, ISPF, InfoPrint Central, & RRS improvements, Extensions to the IBM Configuration Assistant for z/OS Comm Server, EE usability improvements, ...

### Integrating new Applications and Supporting Industry and Open Standards

Unicode improvements; inittab support; LE improvements, new LDAP server, z/OS XML System Services, BPXBATCH, CIM, XPLINK, Networking APIs, XL C/C++ enhancements, ...

### Extending the Network

Sysplex partitioning support, Dynamic DNS registration, JES3 NJE via TCP/IP, TN3270 and telnet improvements, IP filtering, IKE, IPsec, WTS, Samba, improved sysplex failure recovery, REXX FTP interface, ...

### Scalability & Performance

Support for up to 4 TB memory, New zIIP processor, CFRM performance improvements, GRS ENQ limits, Parallel VARY ONLINE, z/OS UNIX file descriptors, Device Groups, DADSM/CVAF Rapid Index Rebuild, z/OS USS asynch socket fast path, 64-bit PDSE support ...

### Enterprise-Wide Roles

WLM/EWLM Service Class Correlation, zAAP Reporting for eWLM, DFSMSrmm CIM agent, RMFeServer OS Monitoring Stage II, WLM delay services mapped to ARM ...


**Improving Availability**  
**Master console elimination, Logstream rename, test and prod logstream separation, Recoverable BRLM, GDPS enhanced recovery, Fast replication, HyperSwap trigger, z/OS UNIX latch contention detection, ...**

### Optimization and Management

WLM, zAAP Stage 2, Sysplex Distributor, Tape I/O priority, JES2 balancing of WLM-managed initiators, GRS CNS processing, Improved OAM BLOB support, SCLM improvements, Group Capacity Limit ...

### Enhancing Security

PKI extensions, Distributed Identity Support, RACF password phrase support, Tape data set protection, 128-bit AES for IPsec, SAF identity tokens, RACF virtual key rings, IDS configuration, ...



## z/OS Some 1.8 Availability Enhancements at a Glance

- Master Console elimination
- Log Stream rename, test and production logstream separation
- z/OS UNIX latch contention detection
- GDPS enhanced recovery
- Fast replication
- HyperSwap trigger
- Recoverable BRLM







LSU Danmark



Danmark LSU October 2006

*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation

## Consoles - Problems to be solved



Old news

- (1) Message delivery:
  - Buffer shortages
  - Varying speeds and feeds
- (2) Synchronizing console state information:
  - Data replicated on every system in the sysplex
  - State changes, system join/leave cause flurries of activity
  - More systems => elongated startup, shutdown and recovery time
- (3) Limit of 99 MCS/SMCS/Subsystem consoles in a sysplex:  
Installation constraint

## Consoles Enhancement Staging Plan



Old news

- Enhancements to be staged:
- Stage 1a solves problem 1
  - Delivery in z/OS: 1.4.2 feature, 1.5 base
- Stage 1b delivers improved RAS, IPL/recovery time improvements
  - Available in z/OS 1.7
- Stage 2 solves problems 2 & 3
  - z/OS 1.8 and future release of z/OS (delay since announcement)

## Console Restructure - Summary

- Stage 1a (z/OS 1.4/1.5) solves message delivery
  - Targets WTO buffer shortages (accounting for ~45% of consoles multi-system outages)
  - Can re-IPL an image to “fall out” of stage 1
- Stage 1b (z/OS 1.7) provides:
  - Internal RAS / Problem determination enhancements
  - Ability to delete EMCS consoles
  - Operational changes in preparation for stage 2
- Stage 2 to solve the CONSOLE state data problem
  - Remove 99 console per sysplex constraint
  - Reduce serialization bottleneck on SYSZMCS global resources
  - Solves “the rest” of the consoles multi-system outages
  - Command activated – limited “reverse migration” all systems must be at required z/OS level
  - Future z/OS release; no “rollback” planned



Old news

## System z Console Restructure Road Map


### Stage 1 Overview: z/OS 1.5 – 1.8 Enhancements, more to come

- z/OS 1.8
  - Master Console Elimination
- z/OS 1.7
  - Console state information volume reduction, EMCS console removal, reduced IPL/join time
  - Serviceability Items
- z/OS 1.5
  - Message production/consumption imbalance eliminated
  - System Outage Avoidance, due to:
    - Buffer exhaustion avoidance
    - Workload slowdown and/or job abend reduction
    - Message throughput improvement
      - by reducing single task bottleneck



## z/OS 1.8 Console Enhancements

*Master Console Elimination, 1-byte consoles now history*




- Master Console elimination
  - **Removes a single point of failure for improved RAS**
    - Ability to define multiple master authority consoles **not** changed
    - Attributes unique to Master Console made available to any console
      - > Console ID zero “gone”
    - **“No Consoles Condition”** and **“No Master Console Condition”** no longer considered undesirable
  - Console Switch no longer needed to ensure Master Console function was never lost
    - Elimination of Master Console reduces importance of Console Switch
    - Complex task setting up “switch” definitions for all consoles eliminated
  - 1-byte console IDs no longer supported
  - Note: **Consoles Stage 2 delayed** (i.e. not in 2007)
    - Reduce configuration data passed around the sysplex under serialization
    - Increase max # of cons. from 99/sysplex to 250/system
      - Max 99 active/system in sysplex

## Master Console Elimination

### *Hardcopy Switch elimination details*

- Pre z/OS 1.8, if SYSLOG or OPERLOG was hardcopy medium and failed:
  - z/OS would switch to other log if log was defined and not already a hardcopy medium
  - If other log not available, hardcopy was suspended
- With z/OS 1.8, the switch has been eliminated and if hardcopy is essential, SYSLOG and OPERLOG should be defined as hardcopy mediums



# Master Console Elimination

*Changed externals: CNGRPXX changes and deleted CONSOLxx keywords*

## ■ CNGRPXX changes


- Vary MSTCONS
- Display CONSOLE MCONLY
- SWITCH CN
- Vary CN(xx),ALTGRP

## ■ Deleted CONSOLxx keywords

- ALTGRP
- NOCCGRP
- No console forced to receive routecode 1 and 2
- Master authority not forced to any MCS/SMCS console
- System Console will always have Master authority

## ■ New CONSOLxx keywords

- ▶ INTIDS (replaces sending messages to console ID zero)
- ▶ UNKNIDS (no longer supported 1-Byte console IDs)





## Master Console Elimination

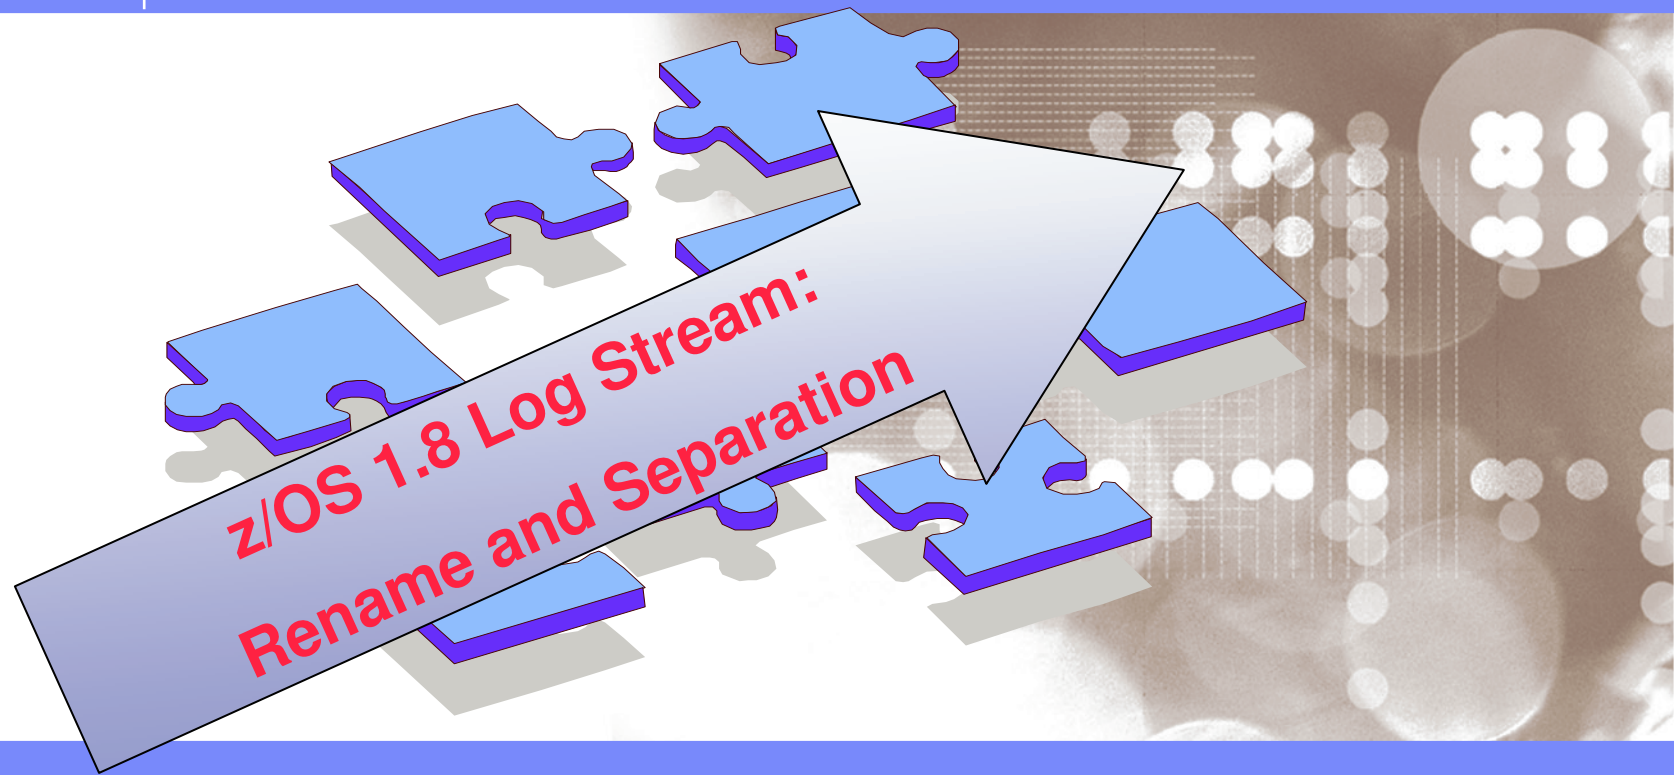
### *New Message Routing Attributes: "Internal and unknown IDS"*

- Since the master console no longer exists, predictable destinations for certain messages are needed
- New console routing attributes, **INTIDS** and **UNKNIDS**, ensure console destinations for messages which would previously have gone to the master console
- In a mixed sysplex, the toleration **APAR, OA10632**, has partially implemented these attributes on the lower levels of z/OS. It supports sending messages to consoles on 1.8, receiving INTIDS and UNKNIDS, from lower-level systems.





LSU Danmark



## z/OS 1.8 Logger **Rename** Logstream

- Issues with your log stream? "local fix" was often to delete and redefine the log stream in order to get applications up and running
  - **When log stream deleted all data associated with the log stream gone**
  - **No way to review data nor have application process the lost data**
- Logger utility and API to **rename a log stream** introduced in z/OS 1.8:
  - Update Log stream request indicate new name of the log stream:

```
UPDATE LOGSTREAM  NAME (xname) ...  
                NEWSTREAMNAME (xnewstreamname)
```



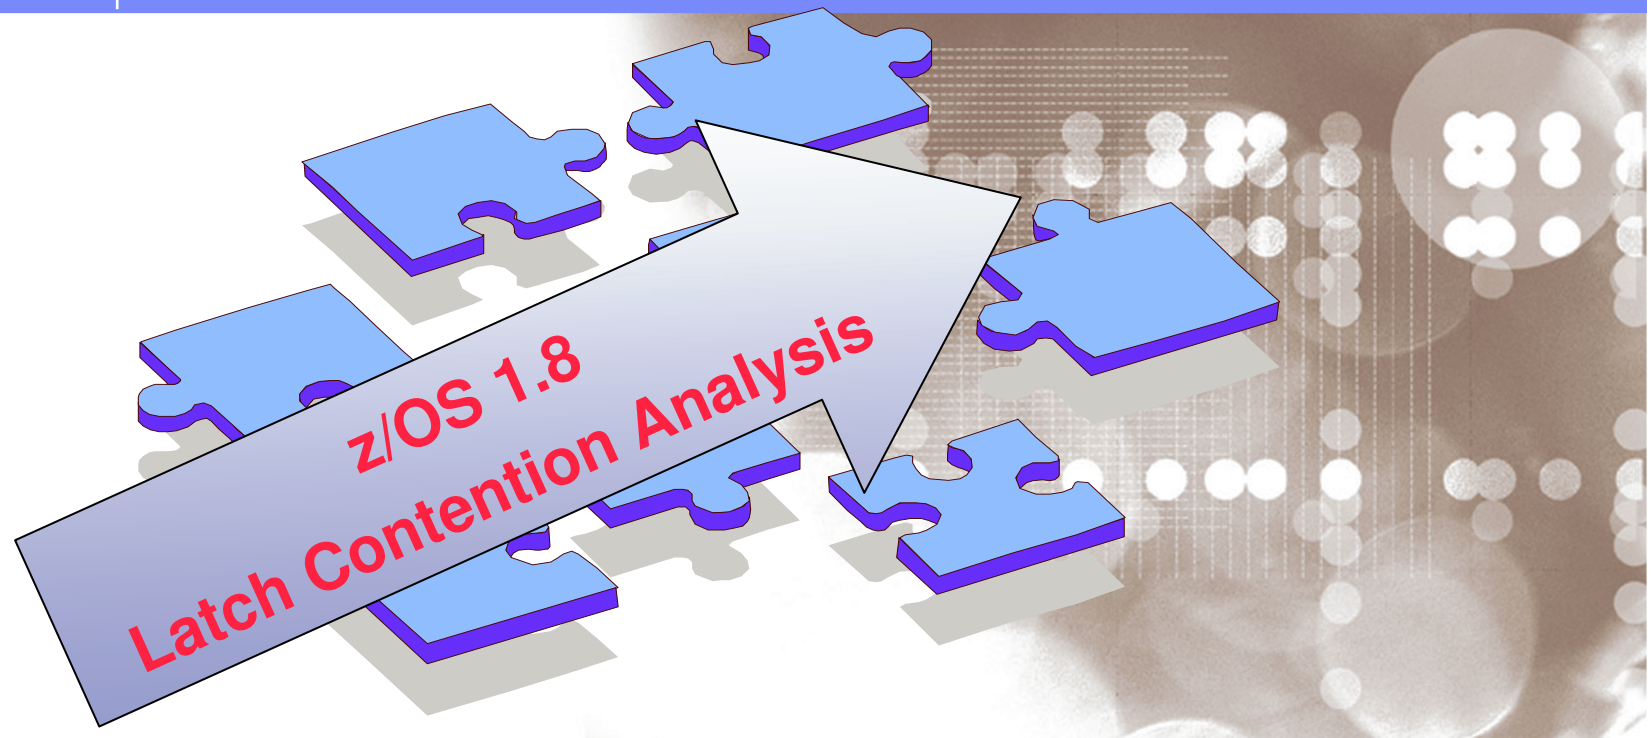
## z/OS 1.8 Log Stream Separation

- **Installations combine both production and test in the same sysplex**
  - Possible for work on test log streams to have an **adverse effect on production log streams.**
  - Problems encountered on a test log stream can lead to **outages**
  - Some Logger tasks, such as data set recall, and data set allocation are **single threaded**
    - Data set recall for the test log stream can hold off the data set recall for production log stream
- **This feature allows you to define both Test and Production logstreams in the same sysplex**
  - Reduces the effect of test log stream outages on production log streams
  - Benefit will be **less interference from test environment on production**
  - `DEFINE LOGSTREAM ...`
    - `STG_DUPLEX (NO) ,`
    - `GROUP (TEST)`
- **Coexistence APAR OA14009 required on pre z/OS 1.8 in sysplex**





LSU Danmark



Danmark LSU October 2006

*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation

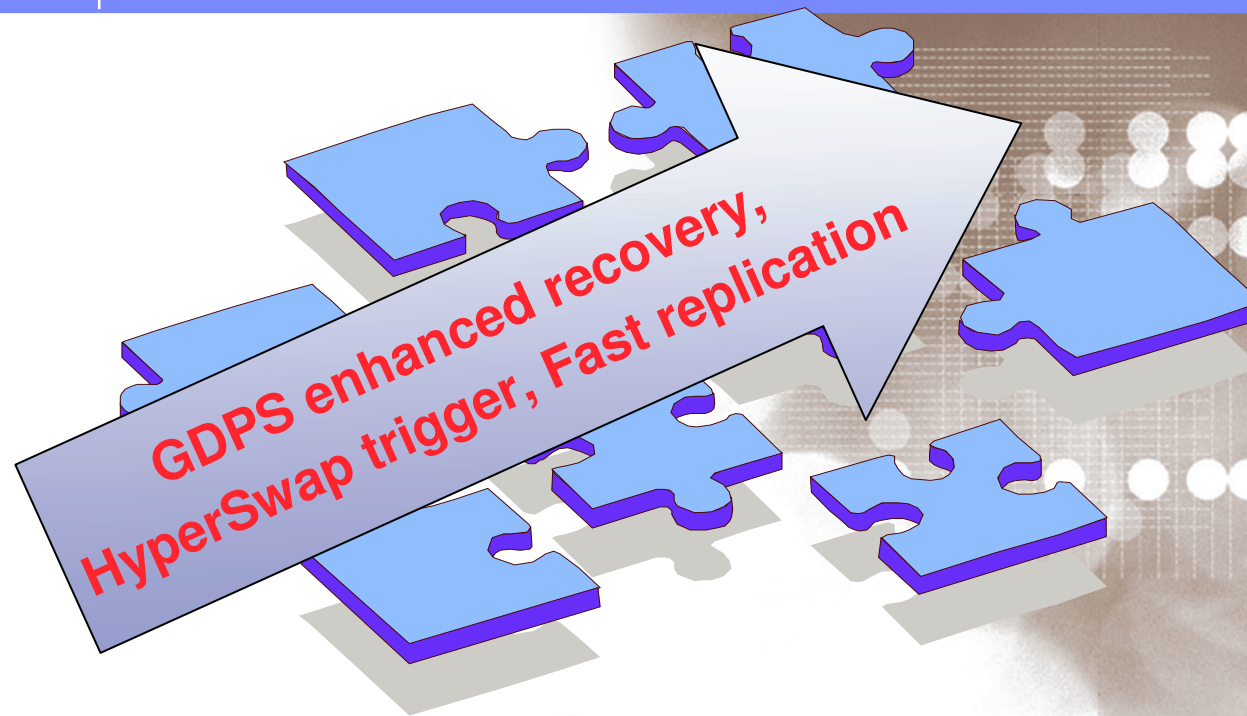
## z/OS 1.8 Latch Contention Analysis

### *Part 2: Follow-on to support introduced in z/OS 1.7*

- `DISPLAY OMVS, WAITERS` operator command extended to show **contention on File System Latches** (similar to what is shown for Mount Latch) and a table of all remaining **threads** that are waiting.
  - The information shown for a waiting thread enhanced to include:
    - Any File System Latch that it is holding,
    - And for operations in progress on a file
      - the file's name, device number and inode number
  - Benefit will be when a system gets hung, the display contributes to determining:
    - **What needs cancellation?**
    - **Actions needed to alleviate the system hang situation?**



## LSU Danmark



Danmark LSU October 2006

*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation

## GDPS Enhanced Recovery

### *CFRM Site Awareness improvement*

- In GDPS/sysplex XES can detect a connectivity failure **before** GDPS
  - XES may delete a structure inconsistently with GDPS policy
    - FREEZE=STOP: desired to keep structure at recovery site. Lossconn site failure might result in structure in recovery site being deleted instead of the copy in the primary site
    - GDPS does not allow duplexed structures to be used in DR operation, timeconsuming log-based methods needed for DR. **Application availability & performance may suffer**
- z/OS 1.8 provides CFRM *site awareness*:
  - Awareness of which CFs exist in each site. Duplexing failover decisions consistent with the recovery manager decisions (failover to a recovery site)
    - Duplexed structure available in recovery site, allowing GDPS to use it DR
    - Eliminates need for log-based recovery,
      - **and shortens recovery time**
  - Site awareness enabled by `SITE` keyword in CFRM policy.

```
CF NAME (FACIL01)  SITE (SITE1)
TYPE (002084) MFG (IBM) PLANT (EN) SEQUENCE (111111111111)
PARTITION (0) CPCID (00)
```



## z/OS 1.8 HyperSwap trigger enhancement

### *MIH – I/O Timeout*

- Prior to this HyperSwap function was invoked automatically for error indications returned for I/O operations:
  - Subsystem failures, boxed devices, or I/O errors.
- z/OS 1.8 support extended to act on I/O timeouts
  - Detected by Missing Interrupt Handler (MIH)
  - Allow transactions to **resume processing quickly on secondary volumes**

## Fast replication

### *Dataset recovery Overview*

- Using **FRRECOV** command a dataset can be recovered from a copy pool backup
- **FRRECOV** command enhanced, 1 or more fully or partially qualified data set specifiable:
  - Each data set must be cataloged and allocated on the same volume(s) that it resided on at the time the backup copy was created
- The backup version being recovered may reside on disk or tape.
  - If version resides on **both** disk and tape. **Default** will be to recover from **disk**.
  - If the backup version is recovered from disk, then the recovery can be performed using either fast replication or traditional copy methods.
  - Up to 64 concurrent data set recoveries are supported.

## Fast Replication

### *Tape support Overview*

Fast replication for tape support functions enhanced:

- **FRBACKUP**: Allows this command to initiate:
  - Dump of any **existing backup** copies or to **create a new copy** and **then dump** it.
  - In `VERSIONS=0 (NOCOPY)` environment, withdraw the relationships when the dump is finished.
  - Allow **resumption** of an **incomplete dump** version if any of the required dump classes **failed to complete**, only dumping the volumes not successfully completed by the previous attempt.

**FRRECOV**: Allows the **recovery** of a **dump copy** to a **single volume**.

**FRDELETE**: **Deletes** one or more **unneeded copy pool** dump versions.



LSU Danmark



Danmark LSU October 2006

*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation

## z/OS 1.8 Recoverable Byte Range Lock manager (BRLM)

### *BRLM History: OS/390 R9 – z/OS 1.8*

- z/OS 1.8 – Distributed BRLM with **recoverable** locks
- z/OS 1.6 – Distributed BRLM with **moveable** locks
  - Locks move when file system moves
  - Distributed BRLM now default in a R6 sysplex
- z/OS 1.4 – Distributed BRLM
  - One BRLM per system in a sysplex
  - Lock commands are routed to the file system owner
- OS/390 R9 - Central BRLM in a shared HFS sysplex
  - one BRLM is a **single point of failure**

## z/OS 1.8 Recoverable Byte Range Lock manager (BRLM)

### *Resolved BRLM USS Application locking remote file Issue*

- Problem when USS applications locks a remote file
  - **Lock is lost when the remote system is lost (I.e., fails).**
  - **Solution: backup remote locks locally, and recover them when system fails**
- Function will now be:
  1. Application issues lock/unlock command to a remote file
  2. USS forwards the lock/unlock command to local BRLM
  3. USS forwards the lock/unlock command to remote BRLM.. remote system fails
  4. USS recovers file system & declares a new owning system
  5. USS re-issues lock/unlock to new owner
- The MVS console command, **F BPXOINIT,FILESYS=DISPLAY,GLOBAL** identifies whether USS is using central BRLM or distributed BRLM in a sysplex.
  - z/OS sysplex => R6 – Distributed BRLM is the default
  - z/OS sysplex member < R6 – DISTBRLM(1) must be set in the USS CDS

# z/OS 1.8

## Helping meet challenges of on demand business

### Improving Usability and Skills

More Health Checks and framework improvements, HCM, ISPF, InfoPrint Central, & RRS improvements, Extensions to the IBM Configuration Assistant for z/OS Comm Server, EE usability improvements, ...

### Integrating new Applications and Supporting Industry and Open Standards

Unicode improvements; inittab support; LE improvements, new LDAP server, z/OS XML System Services, BPXBATCH, CIM, XPLINK, Networking APIs, XL C/C++ enhancements, ...

### Extending the Network

Sysplex partitioning support, Dynamic DNS registration, JES3 NJE via TCP/IP, TN3270 and telnet improvements, IP filtering, IKE, IPsec, WTS, Samba, improved sysplex failure recovery, REXX FTP interface, ...

### Scalability & Performance

Support for up to 4 TB memory, New zIIP processor, CFRM performance improvements, GRS ENQ limits, Parallel VARY ONLINE, z/OS UNIX file descriptors, Device Groups, DADSM/CVAF Rapid Index Rebuild, z/OS USS asynch socket fast path, 64-bit PDSE support ...

### Improving Availability

Master console elimination, Logstream rename, test and prod logstream separation, Recoverable BRLM, GDPS enhanced recovery, Fast replication, HyperSwap trigger, z/OS UNIX latch contention detection, ...

### Optimization and Management


WLM, zAAP Stage 2, Sysplex Distributor, Tape I/O priority, JES2 balancing of WLM-managed initiators, GRS CNS processing, Improved OAM BLOB support, SCLM improvements, Group Capacity Limit ...

### Enhancing Security

PKI extensions, Distributed Identity Support, RACF password phrase support, Tape data set protection, 128-bit AES for IPsec, SAF identity tokens, RACF virtual key rings, IDS configuration, ...

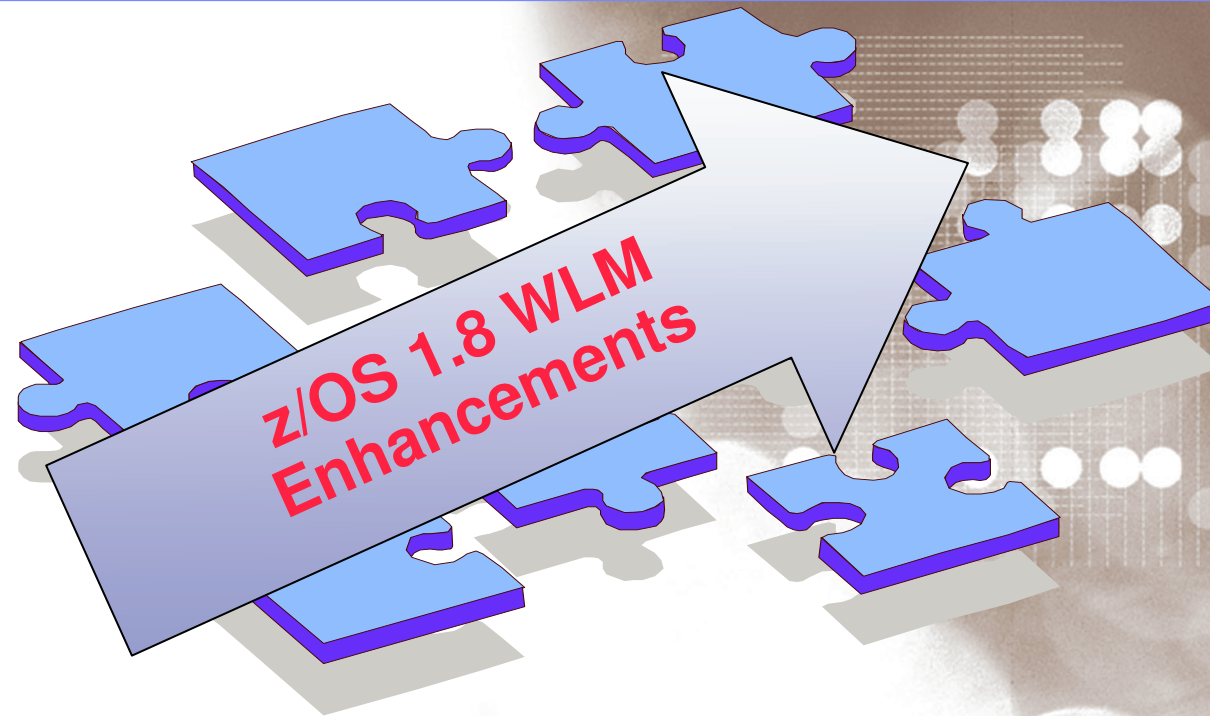
### Enterprise-Wide Roles

WLM/EWLM Service Class Correlation, zAAP Reporting for eWLM, DFSMSrmm CIM agent, RMFeServer OS Monitoring Stage II, WLM delay services mapped to ARM ...





LSU Danmark



Danmark LSU October 2006

*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation



## IBM Workload Manager enhancements for z/OS 1.7

*recent WLM enhancements at a glance*Old news

- WLM adjust to dynamic CP speed adjustment
  - Automatic policy reactivation for z890 without IPL
- WLM enhanced support for sub capacity pricing
  - Actual CPU time consumed without wait time for VM guests
- DB2 latch contention relief (retrofitted to z/OS 1.4)
  - Short time SRM/WLM ENQ promotion
- Enhanced computation of available system capacity
  - Used for WL balancing algorithms (WLM routing services)
- Support of multiple z9 Subchannel Sets
- Enhanced support for EWLM

## IBM Workload Manager enhancements for z/OS 1.8

### *WLM enhancements at a glance*

- Enhanced zAAP and zIIP Support
- RSM/SRM Support for Large Real storage > 128GB
- Preview: Group Capacity limit
  - Satisfies Danish WLM User Requirement!
- JES2 enhancements for WLM-managed batch (enhanced)
- New Resource Group types
- Routing enhancements
- User Friendly Interface for WLM Admin Application

# Overview of New WLM Functions

## *zIIP and zAAP*

Release Function	z/OS R8	z/OS R7 w/ JBB772S	z/OS R7	z/OS R6 w/ JBB77S9	z/OS R6
zIIP Reporting Support	☺	☺ +OA16005		☺ +OA16005	
zAAP Management Support	☺				
zAAP Processing Enhancements	☺	☺	OA14131 + OA13953	☺	OA14131 + OA13953
zAAP Reporting Support	☺				
z/OS R8 Coexistence Support	n/a	OA13837	OA13837	☺	OA13837

# Overview of New WLM Functions

*large RS, OPT, Subchs, LPAR cap, Routing, CP Speed change*


Release Function	z/OS R8	z/OS R7	z/OS R6	z/OS R5	z/OS R4
z/OS R8 Coexistence	n/a	OA13837			
RSM/SRM Support for >128GB real storage	☺				
IEAOPT MCCAFC TH Enhancement	☺	OA14409			
WLM Support of Multiple Subchannel Sets (z9-109)	☺	☺			
Enhancements in LPAR free capacity computations	☺	☺	OA10006		
Routing (Storm Drain) Enhancements	☺	OA14310 (planned)			
Sysplex Routing Enhancements	☺	☺			
Dynamically handle processor speed changes	☺	☺ OA12155	OA07510 OA12155	OA07510	

## Overview of New Functions ...

Enhancement	z/OS R8	z/OS R7	z/OS R6	z/OS R5	z/OS R4
Process Entitlement Type 2 and 3 Resource Groups	☺	z/OS R8 Coexistence Support (OA13837) on previous chart			
EWLM Service Class Correlation	☺	OA12784			
Lightweight EWLM Instrumentation for DDF	☺	OA12005			
EWLM Monitoring Support 1	☺		OA07196		
WLM-managed Initiator Start Considers System Affinities	☺	OA10814			
Controlling Initiator Dispatch Priority in Goal Mode	☺				OW55344

## zSeries Application Assist Processor (zAAP)

- **What is it?**
  - Powerful z/OS Java execution environment
  - Allows you to offload Java work from general purpose CPs to zAAPs
- **Advantages**
  - Significantly less expensive than general purpose CPs
  - Possible reduction of demand on general purpose CPs
  - Upgrade of addtl. zAAP processing power does not affect total MSU rating or machine model designation
  - No IBM SW charges on zAAP capacity



IBM eServer zSeries 990, 890, System z9-109 or later

## zAAP Exploitation

- **System z can only be deployed on z890/z990 or later with z/OS V1.6 or later and the IBM SDK for z/OS Java 2 Technology Edition, V1.4 with PTF.**
- **The following subsystems provide zAAP execution potential:**
  - **WAS 5.1**
  - IMS V7, V8, V9
  - CICS 2.3
  - DB2 V7, V8
  - WBI Brokers V5
  - **Java batch workload**
- **Additional information can be found on:**
  - zAAP home: <http://www.ibm.com/servers/eserver/zseries/zaap/>
  - FAQ: <http://www.ibm.com/servers/eserver/zseries/zaap/faq/>
  - Redbook: <http://www.redbooks.ibm.com/redbooks/pdfs/sg246386.pdf>
  - Whitepaper: <http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/FLASH10337>

## Initial Handling of zAAP-Related IEAOPTxx Parameters

IFACROSSOVER	IFAHONORPRIORITY	Meaning
<u>YES</u>	<u>YES</u>	<ul style="list-style-type: none"> <li>CP's help zAAP's honoring priorities until discretionary work is reached.</li> <li>CP's always help zAAP discretionary work after all CP work.                             <ul style="list-style-type: none"> <li>With soft-capping in effect: See next row</li> </ul> </li> </ul>
YES	NO	<ul style="list-style-type: none"> <li>CPs always help zAAP's after all CP work</li> </ul>
NO	YES	
NO	NO	<ul style="list-style-type: none"> <li>CPs never help zAAPs, except when no zAAPs are operational</li> </ul>

Changed




## OA14131/OA13953: IFAHONORPRIORITY Enhancements

IEAOPT Parameters		New Behavior
IFACROSSOVER	IFAHONORPRIORITY	R6 and R7 with OA14131+OA15297/OA13953
YES	YES	Standard processors can run zAAP eligible work in priority order <b>if the zAAPs become unable to process all queued work</b>  Notice: <b>regular CPs are asked for help therefore not all CPs may process zAAP work at the same time</b>
NO	YES	
YES	NO	unchanged
NO	NO	

- **Old Behavior**
  - IFAHONORPRIORITY became only effective if IFACROSSOVER was set to YES
- **New Behavior** (R6 and R7 with OA14131+OA15297/OA13953)
  - This change is intended to allow more zAAP eligible work to run on zAAP processors while still remaining responsive to the zAAP demand
  - **IFAHONORPRIORITY is independent from IFACROSSOVER**
  - **IFAHONORPRIORITY now has the meaning that the zAAPs may “need help”**

# 5272N1 Sample zAAP and CP Utilization ("before") - Chart

System Utilization  
I5272N1



## Sample zAAP and CP Utilization (“before”) - Explanations

- On the preceding chart we see the CP, zAAP and total (=average) processor utilizations under different workload conditions and settings of IEAOPT IFACROSSOVER and IFAHONORPRIORITY.
- The workload consists of a small number of batch processes (i.e. is “non-stochastic”)
- In the last part of the diagram (after 13:38) CrossOver is set to No and HonorPriority to Yes.
- The zAAP work is cancelled. The existing zAAPs can accommodate the zAAP-demand, though the utilization is still pretty high.
- We see that the **CPs do not pick up zAAP work, even when CP workload is cancelled and CP utilization decreases.**

## WLM Treatment of the zAAP Resource

- **A System z Application Assist Processor is a new WLM resource type:**
  - Contributes using and delay samples, and service times
  - zAAP utilization is reported by RMF (SPE OA05731)
- **Up to z/OS 1.7 zAAPs are managed by WLM as extension of CPs**
  - Java work executing on zAAPs **inherits the dispatch priority from regular CPs**
  - **Execution is accounted for in execution velocity and goal achievement (PI)**
- **zAAP management support**
  - **Beginning with z/OS 1.8 work on zAAP is managed independently**
- **zAAP service not:**
  - Included in defined capacity computations
  - Part of resource group management
  - Included in routing decisions
  - varied by IRD Vary CPU Management
- **On z9-109 specialty engines, such as zAAPs, are in separate pools.**
  - On z890 and z990: zAAPs inherit weight from CP pool but are part of the ICF/IFL pool

## New IBM System z9 Integrated Information Processor (IBM zIIP)

- **New specialty engine for System z9 EC/BC mainframes designed to help:**
  - Customers integrate data across the enterprise
  - Improve resource optimization
- **z/OS manages & directs work between general purpose CPs and zIIPs**
  - Number of zIIPs per z9-109 not to exceed number of standard processors
  - No changes anticipated to DB2 for z/OS V8 applications
- **DB2 for z/OS V8 will be first IBM exploiter of the zIIP with:**
  - System z9 EC/BC, z/OS 1.6 or later, DB2 for z/OS V8
- **More information available on the zIIP web site**
  - <http://www.ibm.com/systems/z/ziip/about.html>

## How does the zIIP work? .....

Work on z/OS may have all or a portion of its Enclave Service Request Block (Enclave SRB) work directed to the zIIP. The types of DB2 V8 work listed below are those executing in enclave SRBs, portions of which can be sent to the zIIP.

### Example 1 = Distributed SQL requests (DRDA)

Queries that access DB2 for z/OS V8 via DRDA over a TCP/IP connection are dispatched within z/OS in enclave SRBs. z/OS directs a portion of this work to the zIIP.

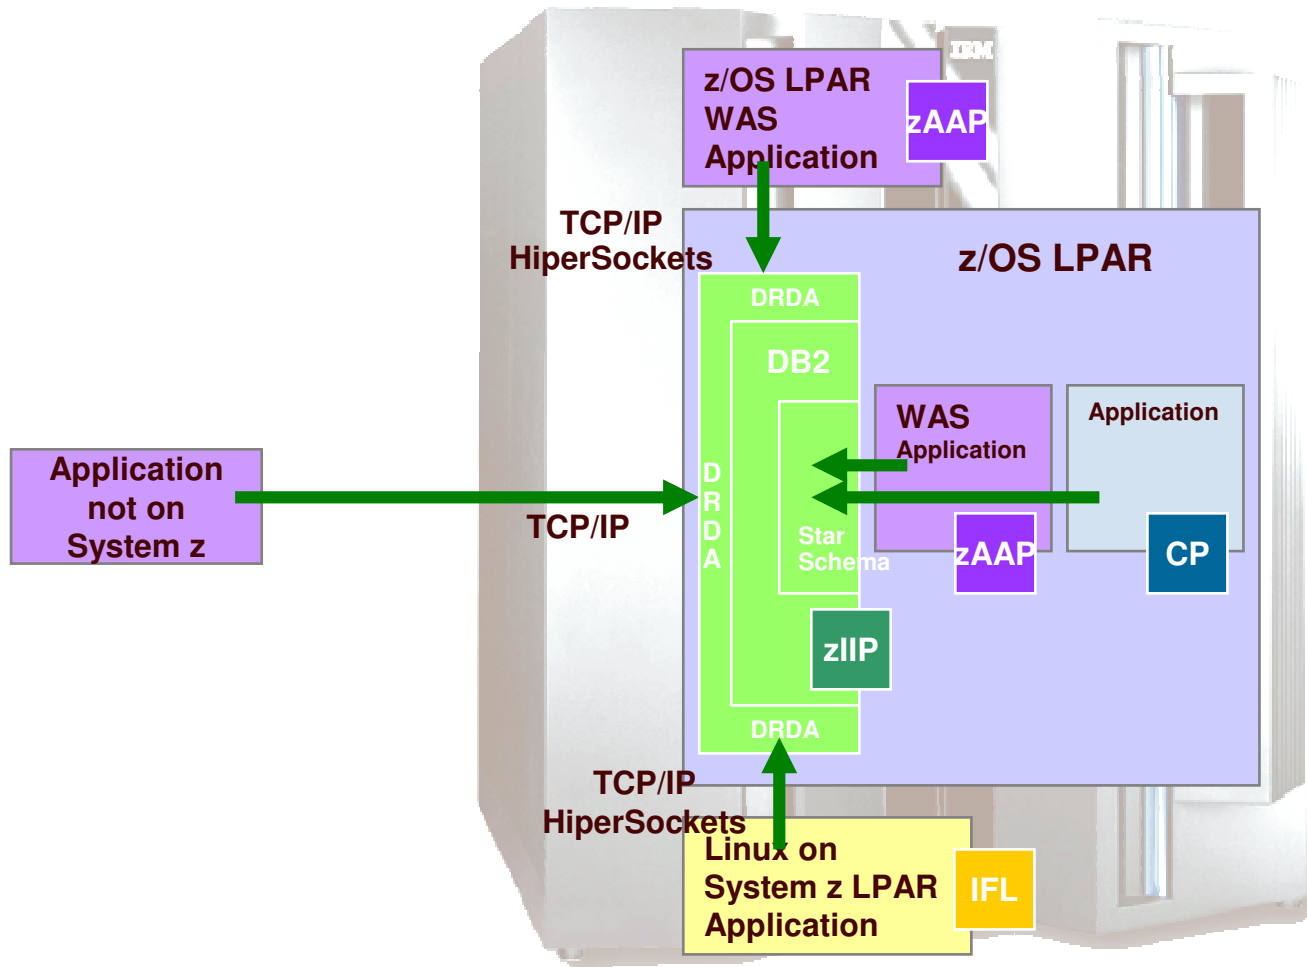
### Example 2 = Complex parallel query (BI)

Complex star schema parallel queries will now use enclave SRBs. z/OS directs a portion of this work to the zIIP.

### Example 3 = DB2 utilities for index maintenance

DB2 utilities LOAD, REORG, and REBUILD will now use enclave SRBs for the portion of the processing that is related to index maintenance. z/OS directs a portion of this work to the zIIP.

# Specialty Engines work Together!



The IBM System z9 specialty engines can operate on the same machine together

(shown at left, instances where zIIP can be employed)

## Handling of zIIP Work

- **In many respects WLM/SRM support of zIIPs is equivalent to the zAAP monitoring support with the zAAP processing enhancements**
  - But: support only available via distinct FMIDs
- **zIIP work managed as an extension of CP work**
- **zIIP work will flow over to general purpose CPs**
  - Like zAAP work with the zAAP processing enhancements (“needs help”)
  - But: No external controls
- **Not included in routing decisions**
- **zIIP service not included in defined capacity computations and resource group management**
- **zIIPs are not varied by IRD Vary CPU Management**
- **On System z9 specialty engines, such as zIIPs, are in separate processor pools.**



## zIIP Reporting via RMF

- WLM Services provide data on zIIP usage comparable to data on zAAP
- **RMF Monitor I**
  - CPU and Workload Activity reports
- **RMF Monitor III**
  - CPC, SYSINFO and ENCLAVE reports
- **zIIP Usage projection available without real zIIPs in the configuration**
  - Via IEAOPT parameter PROJECTCPU=YES
- **Detailed information available via <http://www.ibm.com/systems/z/ziip/>**
  - See "Introduction to System z9 Integrated Information Processor (zIIP)" at <http://publibz.boulder.ibm.com/zoslib/pdf/zosr6r7ziip.pdf> and the Program Directory <http://publibz.boulder.ibm.com/epubs/pdf/i1007730.pdf>

## Overview of New Functions – zIIP and zAAP

Release Function	z/OS R8	z/OS R7 w/ JBB772S	z/OS R7	z/OS R6 w/ JBB77S9	z/OS R6
<b>zIIP Reporting Support</b>	☺	☺ +OA16005		☺ +OA16005	
<b>zAAP Management Support</b>	☺				
<b>zAAP Processing Enhancements</b>	☺	☺	OA14131 + OA13953	☺	OA14131 + OA13953
<b>zAAP Reporting Support</b>	☺				
z/OS R8 Coexistence Support	n/a	OA13837	OA13837	☺	OA13837

## Handling of zIIP Work - Scenario

In this simple scenario we see that zIIP work can be taken over smoothly by CPs.


The upper figure shows the service consumed on CPs (by importance level). The lower figure shows the total service consumed by the SC that the zIIP work is using (by processor type).

Imp1 work is started on zIIPs such that work can be fully contained by 2 zIIP engines. In addition, Imp3 work is started on CPs.

When one of the zIIPs is taken offline to simulate a shortage of zIIP resources, some of the work is taken over by CPs (orange area).

When the zIIP becomes online again, the CPs no longer need to help the zIIPs.

The total service of the zIIP work's SC class is decreased during the period of reduced zIIP capacity, because the share of the zIIP work running on CPs needs to compete against higher important CP work.



## Preview: Group Capacity Limit on IBM System z9

*prompted by Danish user (IBM SO) requirement*

- IBM plans to enable LPAR group capacity limit on System z9 and later. This function allows you to specify a group of LPARs and apply a capacity limit. This allows you to manage the group such that the capacity used by the LPARs will not exceed the limit.
  
- Define a capacity group by entering the group name and group limit value for the LPARs which should belong to the same group:
  - Group Limit and Group Name are specified on the HMC/SE
  - Group Limit defs are independent from Defined Capacity defs. Both limits can be defined and both work together
  - Consists of multiple z/OS 1.8 and up LPARs on the same Machine
    - If not, group limit may be enforced incorrectly
    - It is possible to define multiple groups on a CPC and an LPAR can only belong to one group.
    - A capacity group is independent of a sysplex and an LPAR cluster (IRD)
  - WLM will only manage partitions with shared CPs
    - Dedicated partitions and partitions with wait completion equal to YES are ignored
    - > only LPARs with wait completion NO and shared processors are managed towards capacity limit

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

## Group Capacity Limit

### *WLM Management*

- Each LPAR managed independently from all other LPARs
- Group capacity based on defined capacity
  - 4 hour rolling avgs of group MSUs used as base for managing the group
  - Only general purpose CPs considered
- **Each partition sees the consumption of all other partitions on the machine**
  - If the LPAR belongs to a group it identifies the other partitions of the same group
  - Calculates its defined share of the capacity group
    - Based on the partition weight
      - This share is the target for the partition if all partitions of the group want to use as much CPU resources as possible
  - If one or more LPARs do not use their share, this donated capacity will be distributed over the LPARs which need additional capacity
- **Even when a LPAR receives capacity from another LPAR, it never violates its defined capacity limit (if one exists)**

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

# Group Capacity Limit

## Example

- Partitions A, B, and C are part of GROUP1
  - D and E do not belong to any group
- The limit for the group is 200 MSUs
  - Assuming the capacity is > 200
- The guaranteed minimum for each partition based on weights.

The target consumptions for each partition within the group are:

LP	Group Name	Group Limit [MSU]	Weight	Target MSU consumption based on Weight	LP Defined Capacity (Softcap) [MSU]
A	GROUP1	200	70	93	none
B	GROUP1		50	67	80
C	GROUP1		30	40	30
D	none	n/a	100	n/a	120
E	none	n/a	50	n/a	n/a

- Partition A: 93
  - Can use up to 200 MSU based on 4 hour rolling average. If all 3 partitions want to use as much as possible, A will get 93 MSU (if possible)
- Partition B: 67
  - Can use up to 80 MSU because an individual softcap is defined. If all 3 partitions want to use as much as possible, B will get 67 MSU (if possible)
- Partition C: 40
  - Can use up to 30 MSU because the defined capacity is smaller than the target based on partition weight

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

## Group Capacity Limit

### *Considerations*

- A 4 hour bonus exists after IPL for the group
  - Group wide capping starts when the 4 hour rolling group average reach group capacity limit
- When new LPARs join a group
  - If new system is IPLed and joins a group it does not have the history of the MSU consumption of the entire group
  - Therefore it can take up to 4 hours until all systems in the group have the same view
    - During that time period the group limit cannot be guaranteed
    - If that happens the other partitions will be reduced to their capacity based on their weight in the group
- A partition can dynamically be removed from a group and/or added to another group
  - The changed partition has no knowledge about unused capacity of the new group and it does not keep a history from previous activity
  - All systems must learn again about the new situation - therefore the group limit cannot be guaranteed
- Works together with IRD Weight Management and Vary CPU Management
  - IRD Weight Management may change the weight of partitions in a capacity group and thus changes the target share of the partition in the capacity group

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

## Group Capacity Limit

*Example 1: 3 Partitions with Different Load Pattern - Test Setup*

- Scenario setup:
  - 3 LPARs with different weights and no individual defined capacity limits
- Test scenario:
  1. Initially only IRD3 and IRD4 demand CPU
  2. All partitions have high CPU demands
  3. IRD5 demand drops to minimum (just system)
- **System load not constant - the test includes load changes on the systems**

Partition	Limit	Weight	Target MSU
IRD3	n/a	52	~8.5
IRD4	n/a	102	~16.7
IRD5	n/a	152	~24.8
<b>Group</b>	<b>50</b>	<b>306</b>	<b>50</b>

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

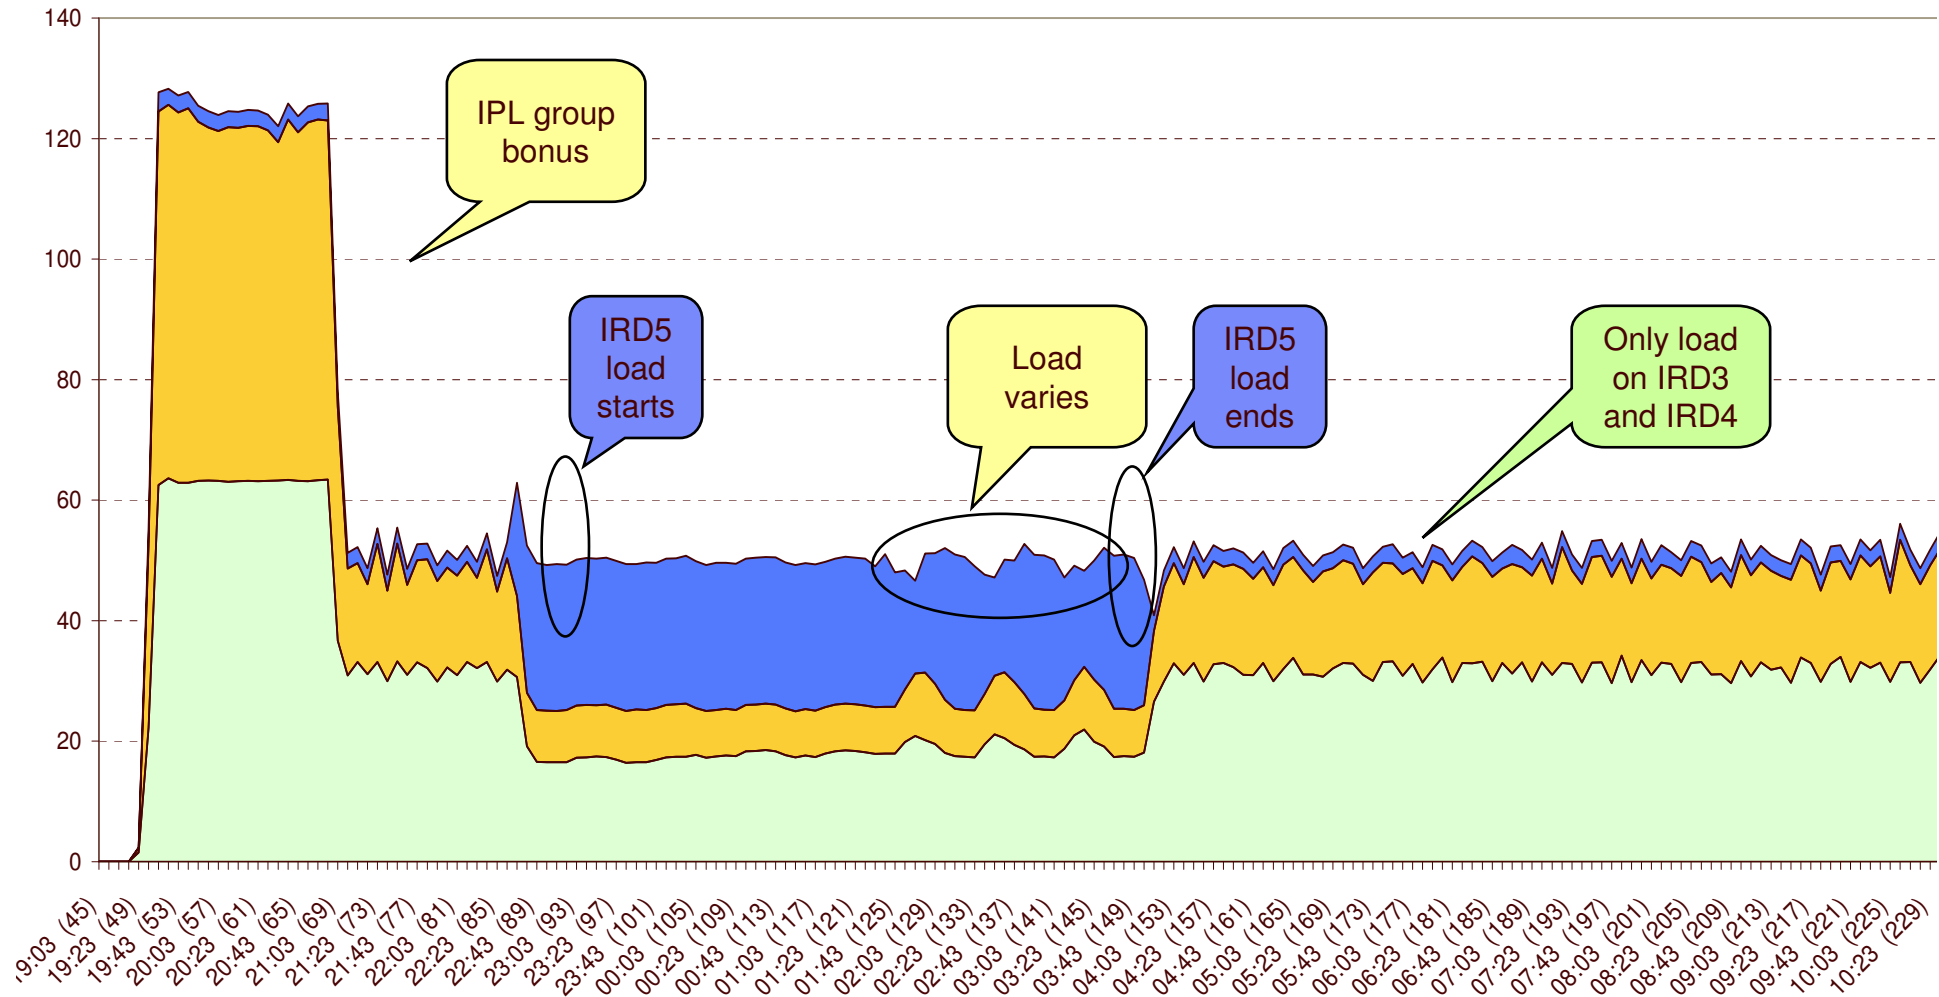


# Group Capacity Limit

## Example 1: 3 Partitions with Different Load Pattern - Test Execution

Actual MSU values


IR4 IR3 IR5



# Group Capacity Limit

## Example 2: New Partition Joining a Group


4 Hour Averages



## Group Capacity Limit

### Example 3: Partitions with Defined Capacities

Partition	Limit	Weight	Target MSU	Based on Limits
IRD3	n/a	51	~8.5	~40
IRD4	5	101	~16.7	5
IRD5	5	151	~24.8	5
Group	50	303	50	



## RMF Support of Group Capacity Limits:

- **RMF extends its existing CPU Activity reporting with information about capacity groups**
  - The Monitor I / Postprocessor CPU activity report provides a new report section for capacity groups with information about
    - Defined MSU limits
    - The share each partition in the group can take
    - The guaranteed share of the partitions
    - The overall MSU consumption within the group
  - SMF record type 70 subtype 1 is extended
  - Monitor III CPC report extended
  - An asterisk in the reports indicates that the partition is a member of the capacity group for less than four hours with regard to the RMF interval start.

PARTITION DATA REPORT

z/OS V1R8                      SYSTEM ID SYS1                      DATE 30/04/2006                      INTERVAL 15.00.999  
 RPT VERSION V1R8 RMF                      TIME 13.30.00                      CYCLE 1.000 SECONDS

MVS PARTITION NAME	A	NUMBER OF PHYSICAL PROCESSORS	16	GROUP NAME	GROUP1
IMAGE CAPACITY	120	CP	8	LIMIT	200 *
NUMBER OF CONFIGURED PARTITIONS	6	IFA	2		
WAIT COMPLETION	NO	IFL	5		
DISPATCH INTERVAL	DYNAMIC	ICF	1		

## JES2 Improved Batch Initiator Balancing



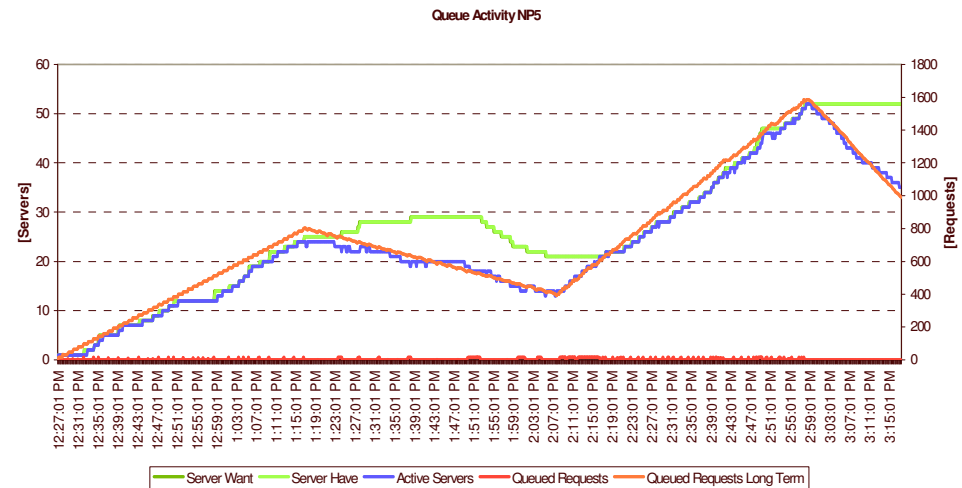
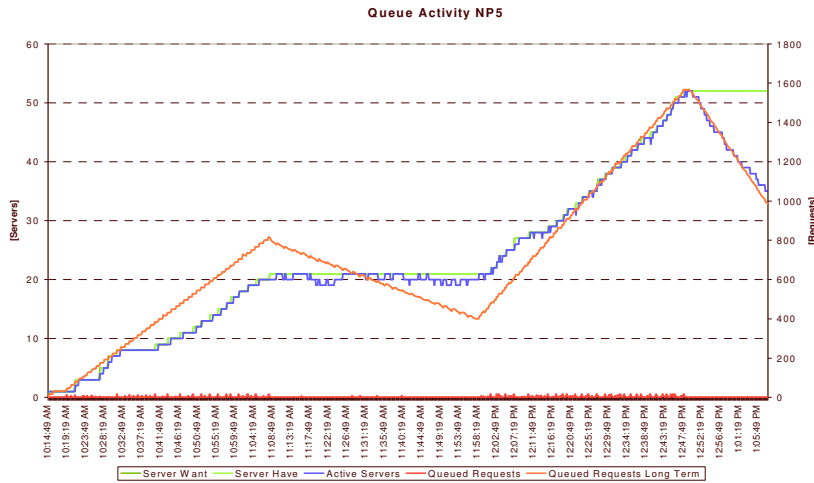
### *Problem prior to z/OS 1.8*

- In a WLM managed initiator environment JES2 will **preferably start jobs on the *submitting system***
- Enhancement in R8 attempts to use *approximately the same percentage of active WLM-managed initiators in each Service Class on each system.*
  - May help balance batch workloads across a JES2 MAS configuration within a sysplex
  - In the scenario on the next chart we submit one job per minute (for a total of 50 min) on SYS1.
    - With R7, the blue lines (Active servers = initiators) differ on the two systems. E.g. SYS1 peaks at 52, SYS2 at 45.
    - With R8, the lines are very similar and peak at 50.

# JES2 Improved Batch Initiator Balancing - Comparison

z/OS R7

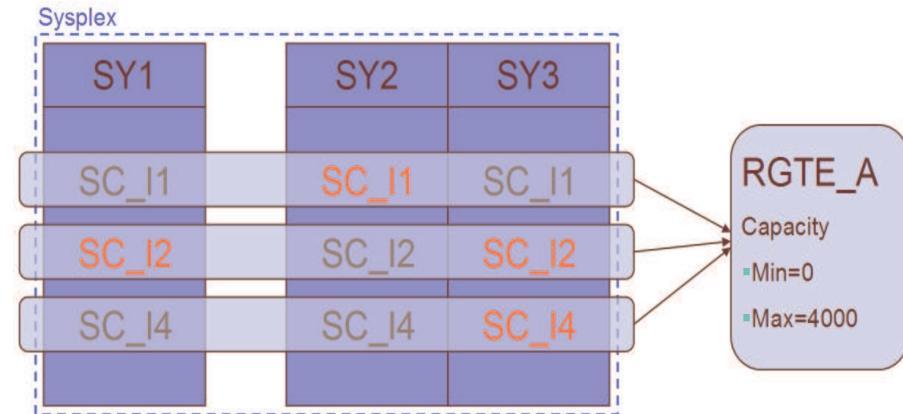
z/OS R8



## Existing WLM Resource Group Concept

*Now referred to as a type 1 Resource Group*

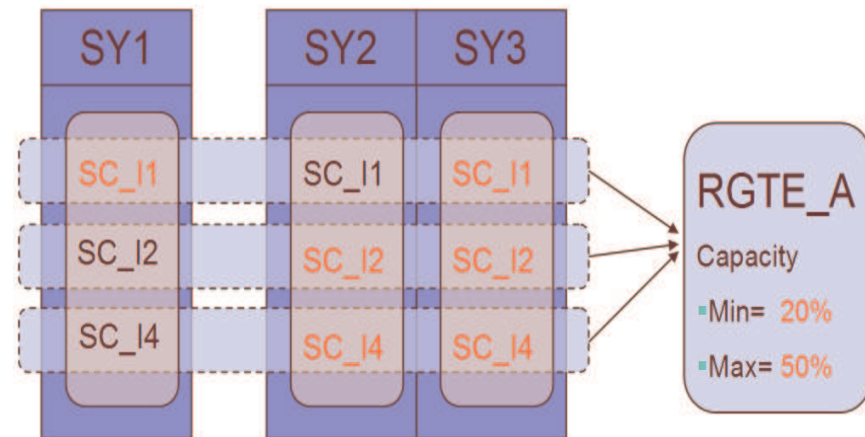
- **Sysplex-wide defined in unweighted service units per second**
- **Sysplex-wide managed**
- **General Considerations**
  - Multiple service classes may be assigned to a resource group
    - With different utilizations on different systems
- **Systems may have different capacities, resulting in:**
  - **Not easy to understand how much is consumed on which system**
  - **Consumption depends highly on the capacity of the systems!**
  - **Resource Group definitions need to be revisited every time**
    - **Systems are upgraded**
    - **Workload utilization changes**
- **Resource Groups are sometimes hard to understand**



## New Types of WLM Resource Groups

*Introduced in z/OS 1.8: Type 2 & Type 3 RGs*

- **Sysplex-wide defined, but definition applies to each system**
- **Managed by each system**
- **General Considerations**
  - Multiple service classes can be assigned to a resource group but this has no sysplex-wide effect anymore
  - Definition is based on one of two possible units:
    - **LPAR capacity: based on system weight (+...)** ⇒ **Type 2 Resource Group**
    - **LCP capacity** ⇒ **Type 3 Resource Group**
- **Results:**
  - **New resource groups are managed by, system thus they must be evaluated on a per system basis**
  - **Resource groups grow automatically if systems are upgraded**
- **Easier to understand**





## Type 2 WLM Resource Group Specification

- **Capacity is specified as a percentage of the effective LPAR share**
  - Range 0..99
  - Effective LPAR share is the minimum of
    - LPAR share:  
CPC shared pool capacity · (LPAR weight / sum of weights)
    - Logical processor capacity
    - Defined capacity limit (only when soft capping is in effect)
  - Use WLM Planning: Appendix B - CPU Capacity Table to locate the CPC capacity that would be equivalent to your shared pool configuration

Numbers should be meaningful for configuration, e.g. sum of all minimum LPAR shares  $\leq 99$

$$RG_{[SU_i]} = \begin{cases} \text{Softcap}_{[SU_i]} \cdot RG_{Limit(\%)} & \text{if softcap} < \text{ShareCapacity or LCPCCapacity} \\ \text{LCPCCapacity}_{[SU_i]} \cdot RG_{Limit(\%)} & \text{if LCPCCapacity} < \text{ShareCapacity or softcap} \\ \text{ShareCapacity}_{[SU_i]} \cdot RG_{Limit(\%)} & \text{if ShareCapacity} < \text{LCPCCapacity or softcap} \end{cases}$$

$$\text{ShareCapacity}_{[SU_i]} = \text{CECCapacity}_{[SU_i]} \cdot \text{LPARShare}$$

CECCapacity<sub>[SU<sub>i</sub>]</sub> = Capacity based on shared physical processors for the CEC

$$\text{LPARShare} = \frac{\text{Weight}(\text{Current Partition})}{\sum_i \text{Weight}(i)}$$

LCPCapacity<sub>[SU<sub>i</sub>]</sub> = Capacity based on shared processors available to the LPAR

## Type 3 WLM Resource Group Specification

- Capacity is expressed as an equivalent number of general purpose processors
  - Scaled by 100, i.e. 100 equivalent to 1 CP  
Range: 0...999999
  - Use *WLM Planning: Appendix B - CPU Capacity Table* to locate the model that would be equivalent to your virtual processor configuration.

Numbers should be “meaningful” for configuration.

## Comparison of WLM Resource Group Types

*...all three may coexist*

Definition	Scope	Advantages	Possible Disadvantages
<b>Service Units</b>  <b>Type 1</b>	Sysplex	<ul style="list-style-type: none"> <li>Allows balancing of resources across members in a sysplex</li> </ul>	<ul style="list-style-type: none"> <li>Requires adjustments for Migration</li> <li>Difficult to monitor</li> <li>Not applicable to constrained work on a single system in a sysplex environment</li> </ul>
<b>% of LPAR</b>  <b>Type 2</b>	System	<ul style="list-style-type: none"> <li>Allows to control work on single members in a sysplex</li> <li><input checked="" type="checkbox"/> <u>Stable for migrations</u></li> </ul>	<ul style="list-style-type: none"> <li>Perhaps difficult to understand what LPAR capacity really is</li> </ul>
<b>% of LCPs</b>  <b>Type 3</b>	System	<ul style="list-style-type: none"> <li>Allows to control work on single members in a sysplex</li> <li><input checked="" type="checkbox"/> <u>Easy and straight forward definition</u></li> </ul>	<ul style="list-style-type: none"> <li>Requires adjustments for migration</li> </ul>

## Resource Groups: Administrative Application

```

Resource-Group  Xref  Notes  Options  Help
-----
Command ===> _____ Modify a Resource Group

Enter or change the following information:

Resource Group Name . . . . . : DIVB
Description . . . . . : Division B service


Define Capacity:
2  1. In Service Units (Sysplex Scope)
   2. As Percentage of the LPAR capacity (System Scope)
   3. As a Number of CPs times 100 (System Scope)
Minimum Capacity . . . . . : 10
Maximum Capacity . . . . . : 20

```

- **Resource group definition panel has been extended and does now allow for definitions of the three resource group types:**
  1. In service units with a sysplex-wide scope (existing)
  2. As a percentage of the LPAR capacity (new concept)
    - Based on LPAR weight
    - System-wide managed
  3. As a percentage of the LCP capacity
    - System-wide managed

# Resource Groups: Example

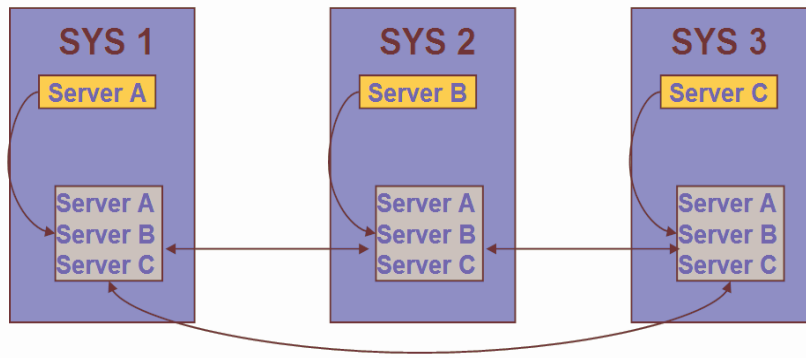
**Resource Group Overview**  
**System: IRD1, Resource Group: ELPMAX , Type: 0-60% of LPAR**



## Resource Groups: Migration Considerations

- **New system service classes and new resource group concept introduce incompatible changes to the WLM service definition**
  - Introduce new WLM policy level to LEVEL019
  - APAR OA13837 is required on pre-R8 systems
    - listed in GA22-7499 z/OS V1R8.0 Migration: z/OS V1Rx coexistence and fallback PTFs
  
- **The compatibility APAR OA13837**
  - Ensures that type 2 and 3 resource groups are not managed on pre-1.8 systems
  - As long as policies are only installed / activated from back level systems the compatibility code is not required
    - Compatibility PTFs can be installed via rolling IPL
      - no Sysplex-wide IPL necessary

# Routing Services: Before z/OS 1.7



▪ Step 1

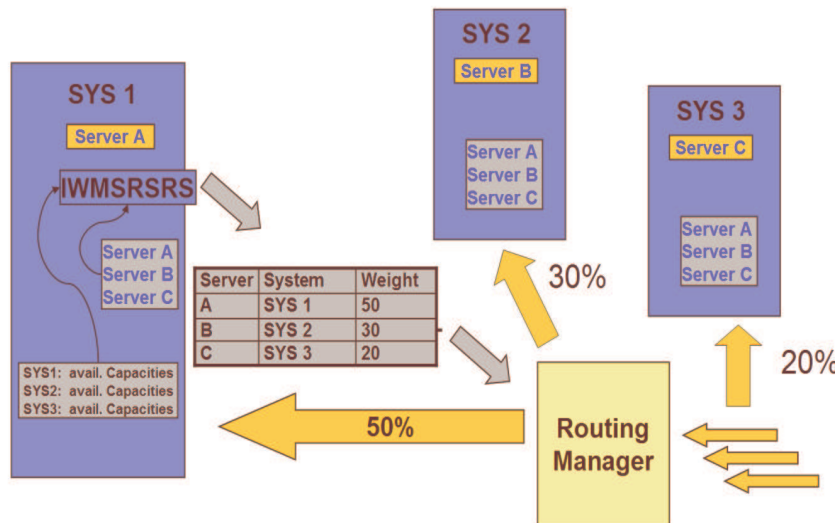
– Registration of servers by IWMSRSRG services

- Each server registers on the local system as a routing candidate
- WLM communicates registered servers and capacity data between systems

▪ Step 2

– Router asks WLM for recommendations by using IWMSRSRS service

- WLM calculates a weight for each registered server based on capacity
  - Servers on systems with serious constraints are excluded (weight 0) unless all systems have serious constraints
- The returned weight reflects the available or replaceable system capacity



## Sysplex Routing: Observations Before z/OS 1.7



Old news

- **Current recommendation is purely capacity based**
  - Weight reflects system capacity and is not server specific
  - Some observed situations:
    - If the LPAR is loaded with low important work, but the server has high importance  
→ **weight too low**
    - If the LPAR is lowly utilized, but the server has a bad performance index or is waiting for resources  
→ **weight too high**
- **TCP/IP stack registration does not reflect how work is processed on the system**
  - TCP/IP address space is usually running in a high important service class
  - The registered address space does not consume the work, it just passes the work to the real consumer on the system
- **Abnormal conditions of the work receiver or consumer are not factored into the routing recommendations**
  - This can lead to *storm drain scenarios* where a system seems to be eligible for processing many requests while in fact the work receiver or consumer is ill
  - In many cases this requires that the work receiver consumer provides additional information to WLM
- **LPAR Wait Time includes wall-clock time**
  - In addition *soft capping and LPAR weights are not correctly reflected*
  - This can result in an over estimation of the available capacity of a logical partition
  - *As a result to many requests may be routed to certain systems (partitions) in a sysplex while other systems receive to few requests*



# Sysplex Routing: Enhancements

## *z/OS 1.4 and z/OS 1.7 enhancements*



Old news

- **Current recommendation is purely capacity based**
  - New option on routing services will allow a router to receive weights which consider the **goal achievement of the work**
  - **z/OS 1.7 (IWMSRSRS FUNCTION=SPECIFIC)** May be exploited by DB2 in the future
- **TCP/IP only registers the stack and doesn't know how the work is being processed**
  - Introduce new routing service e.g. for TCP/IP which allows to provide **routing recommendations for TCP/IP backend applications**
  - **z/OS 1.7 (IWM4SRSC)**
- **Abnormal conditions of the work receiver or consumer are not factored into the routing recommendations**
  - In a first step the **queue time of the work is accounted for DDF regions** and its proportion to the execution time is factored into the routing recommendations
  - **z/OS 1.7 (incorporated in IWMSRSRS FUNCTION=SPECIFIC)**
- Current activities to provide broader solutions
  - Include abnormal terminations which are recorded by **CICS**
  - New interface/option which allows a server to **provide its state to WLM**
- **Measured Wait time put in perspective with real available system and partition capacity**
  - Over estimation is reduced
  - This part is also available as **OA10006 on z/OS 1.4 and above**

# Routing Services: Summary

Functions in grey/non-bold font are considered for a future release. All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only..

Service	Function	Remark	Description
IWMSRSRS	Existing function SELECT		Returns capacity for the system relative to other systems in the sysplex (accounts for the number of registered server instances)
IWMSRSRS	New parameter SPECIFIC	(1)	Returns capacity for the registered server relative to all other registered servers of the same type in the sysplex. Considers: <ul style="list-style-type: none"> <li>▪ Goal Achievement (PI)</li> <li>▪ Queue Time for Enclaves</li> <li>▪ Health Indicator</li> </ul>
IWMSRSRG	New Parameter HEALTH		Allows the server which registers to provide a health indicator from ok=100 to not ok=0. The factor is considered as part of the weight. IWMSRSRG HEALTH can be updated by the server at any time
IWM4HLTH	New Service		For address spaces which are not registered and which want to set a health status. This status is factored into IWM4SRSC return data
IWM4SRSC	New service  New parameter ABNORM_COUNT	(1)	Returns capacity for another address space to which the request is provided by the registered server. Considers <ul style="list-style-type: none"> <li>▪ Goal Achievement (PI)</li> <li>▪ abnormal termination rate expressed as the number of abnormal terminations (as passed to WLM by the IWMRPT interface) per 1000 total terminations.</li> </ul>

(1) Goal achievement is derived from the service class the working is running in. Can be an enclave service class, a transaction manager service class or the service class to which the address space is classified too

## Sysplex Routing: Enhancements for IWMSRSRS Service

- **New Function code „SPECIFIC“**
- **Weight calculation: Product of four factors**
  - System Utilization Factor:
    - Same as the resulting system weight for old „SELECT“ function (as described in the earlier slides)
  - PI Factor:
    - This gives an indication of how good this server, respective the work that is related to this server, is achieving its goals as defined in the active WLM policy.
  - Queue Time Ratio:
    - If the server owns independent enclaves, the ratio of queue time to elapsed time of those enclaves
  - Health Indicator:
    - A registered server can use the IWMSRSRG HEALTH= parameter to inform WLM about additional states which are unknown to WLM but should be factored into the returned weight
  - Remark
    - If multiple servers are registered on the same system, the weight is divided by the number of the servers.

## Sysplex Routing: New IWM4SRSC Service

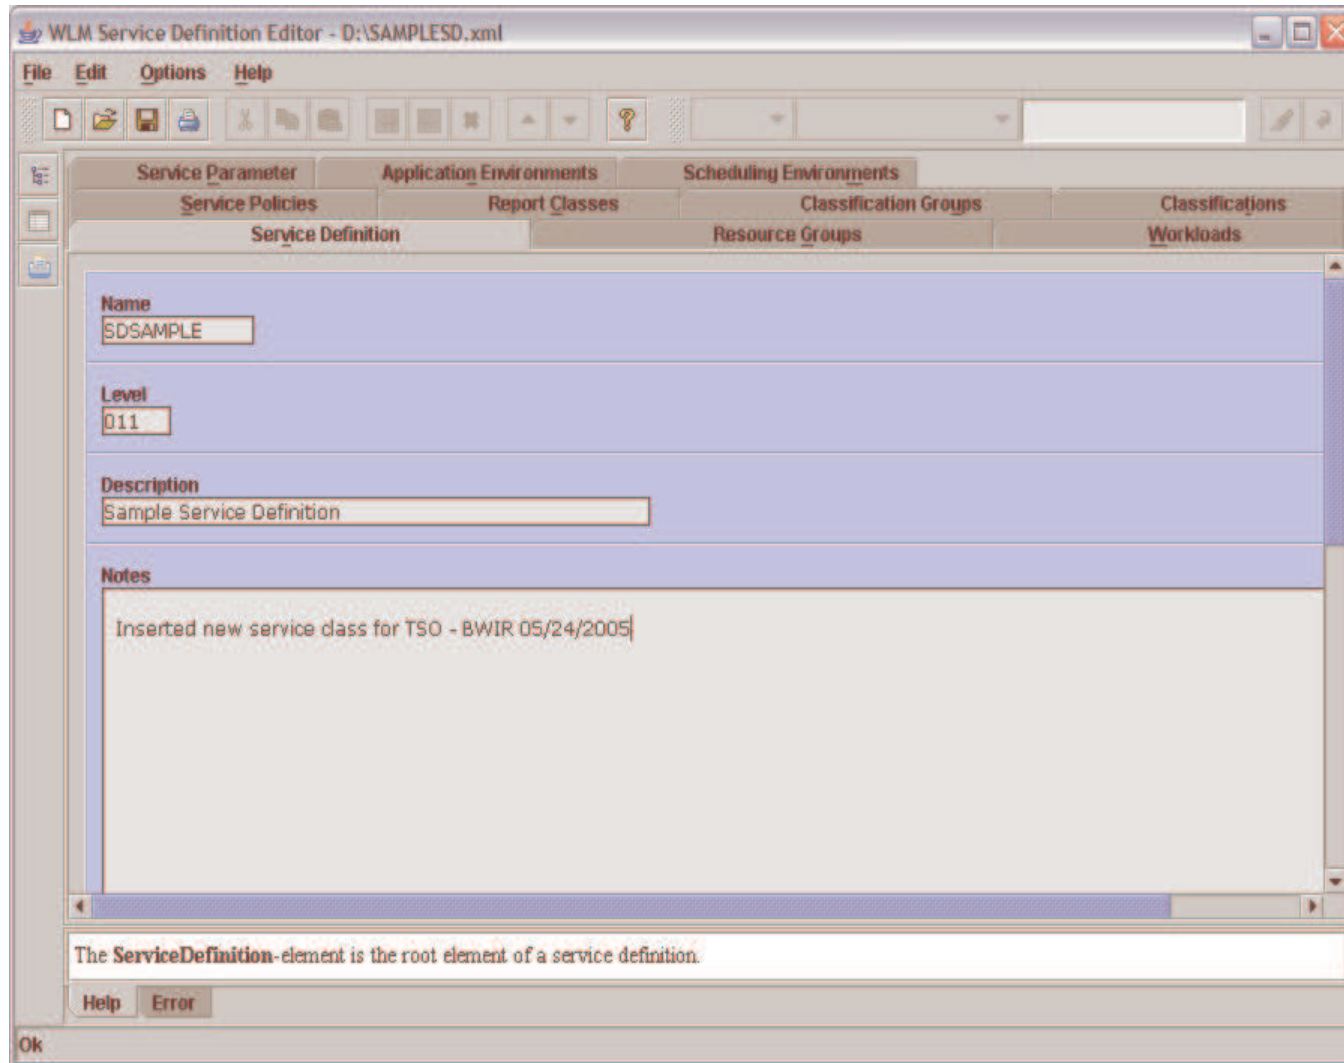
### ■ Characteristics

- Provides routing recommendation for a given server
- System scope
- No registration of servers necessary
- The intention of this service is to obtain routing recommendations for work which is passed through the registered server (e.g. TCP/IP) to a server which consumes/processes the work
- Input: STOKEN that identifies the address space of a server
- Output: Recommendation of how good this server is suitable to receive work → the weight
  - A value between 1 and 64

### ■ Weight calculation: Product of two (four) factors, scaled by 64

- The PI Factor (Performance Indicator Factor)
  - This gives an indication of how good this server, respective the work that is related to this server, is achieving its goals as defined in the active WLM policy.
- The Importance factor
  - This is a measurement of how much CPU Capacity is displaceable by work of the server's importance, respective the work that is related to this server
- The health factor
  - This is an indicator an address space can set via the new IWM4HLTH service
- Abnormal termination rate

## User Friendly Interface for WLM Administrative Application



- **Java application**
- **Interface to WLM ISPF datasets**
- **Different views to show relations between policy elements**
- **Real-time error checking**
- **Real-time help**
- **No popup menus**
- **Tabular input processing**

# User Friendly Interface for WLM Administrative Application

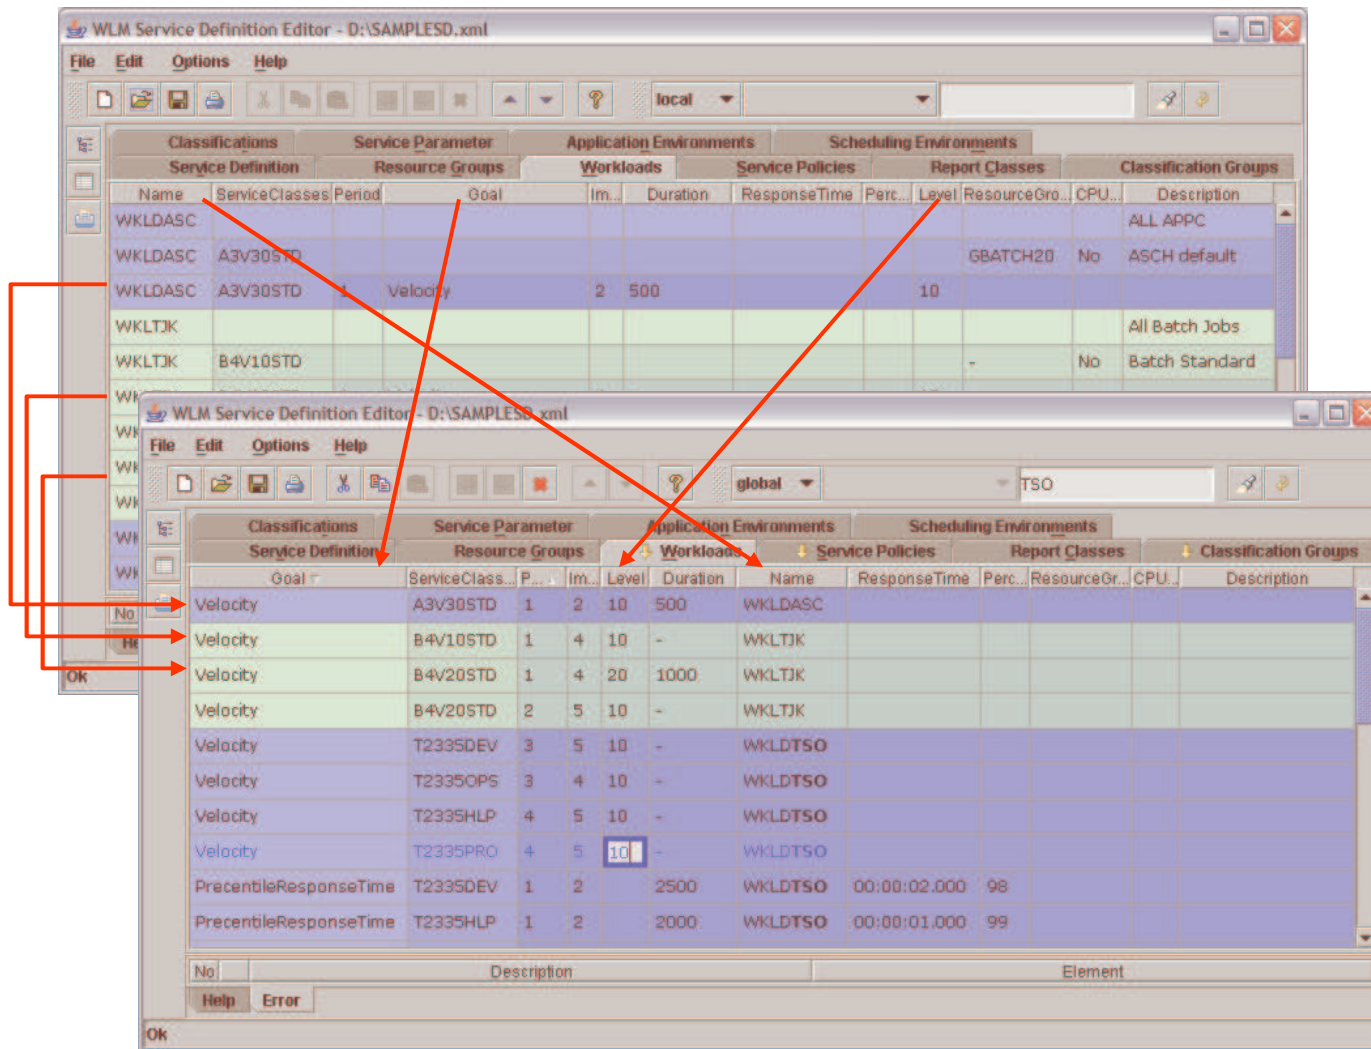
## Example: Workload and Service Class Panel/View

Service Definition		Resource Groups		Workloads		Service Policies		Report Classes			
Name	ServiceClasses	Period	Goal	Im...	Duration	ResponseTime	Perce...	Level	ResourceGr...	CPU...	Description
WKLDASC											ALL APPC Transaction
WKLDASC	A3V30STD								GBATCH20	No	ASCH default Service Class
WKLDASC	A3V30STD	1	Velocity		2 500			10			
WKLDASC	A3V30STD	2	Velocity		2 -			8			
WKLTJK											All Batch Jobs
WKLTJK	B4V10STD								-	No	Batch Standard VEL 10 IMP 4
WKLTJK	B4V10STD		Velocity		4 -			10			
WKLTJK	B4V20STD								-	No	Batch Standard VEL 20 IMP 4
WKLTJK	B4V20STD	1	Velocity		4 1000			20			
WKLTJK	B4V20STD	2	Velocity		-			10			
WKLDTSO											ALL TSO USERIDS
WKLDTSO	T2335DEV								-	No	Developer (Standard) TSO
WKLDTSO	T23		PercentileResponseTime		2 2500	00:00:02.000		98			
WKLDTSO	T23		AverageResponseTime		3 300000	00:00:20.000		95			
WKLDTSO	T23		PercentileResponseTime		5 -			10			
WKLDTSO	T23								-	No	Production TSO Helpers
WKLDTSO	T23		PercentileResponseTime		2 2000	00:00:01.000		99			
WKLDTSO	T23		PercentileResponseTime		3 10000	00:00:01.000		99			

- Information is represented as tables
- Direct manipulation of tables
- Display of relations between policy elements
- Real-time error checking
- Context-sensitive help

# User Friendly Interface for WLM Administrative Application

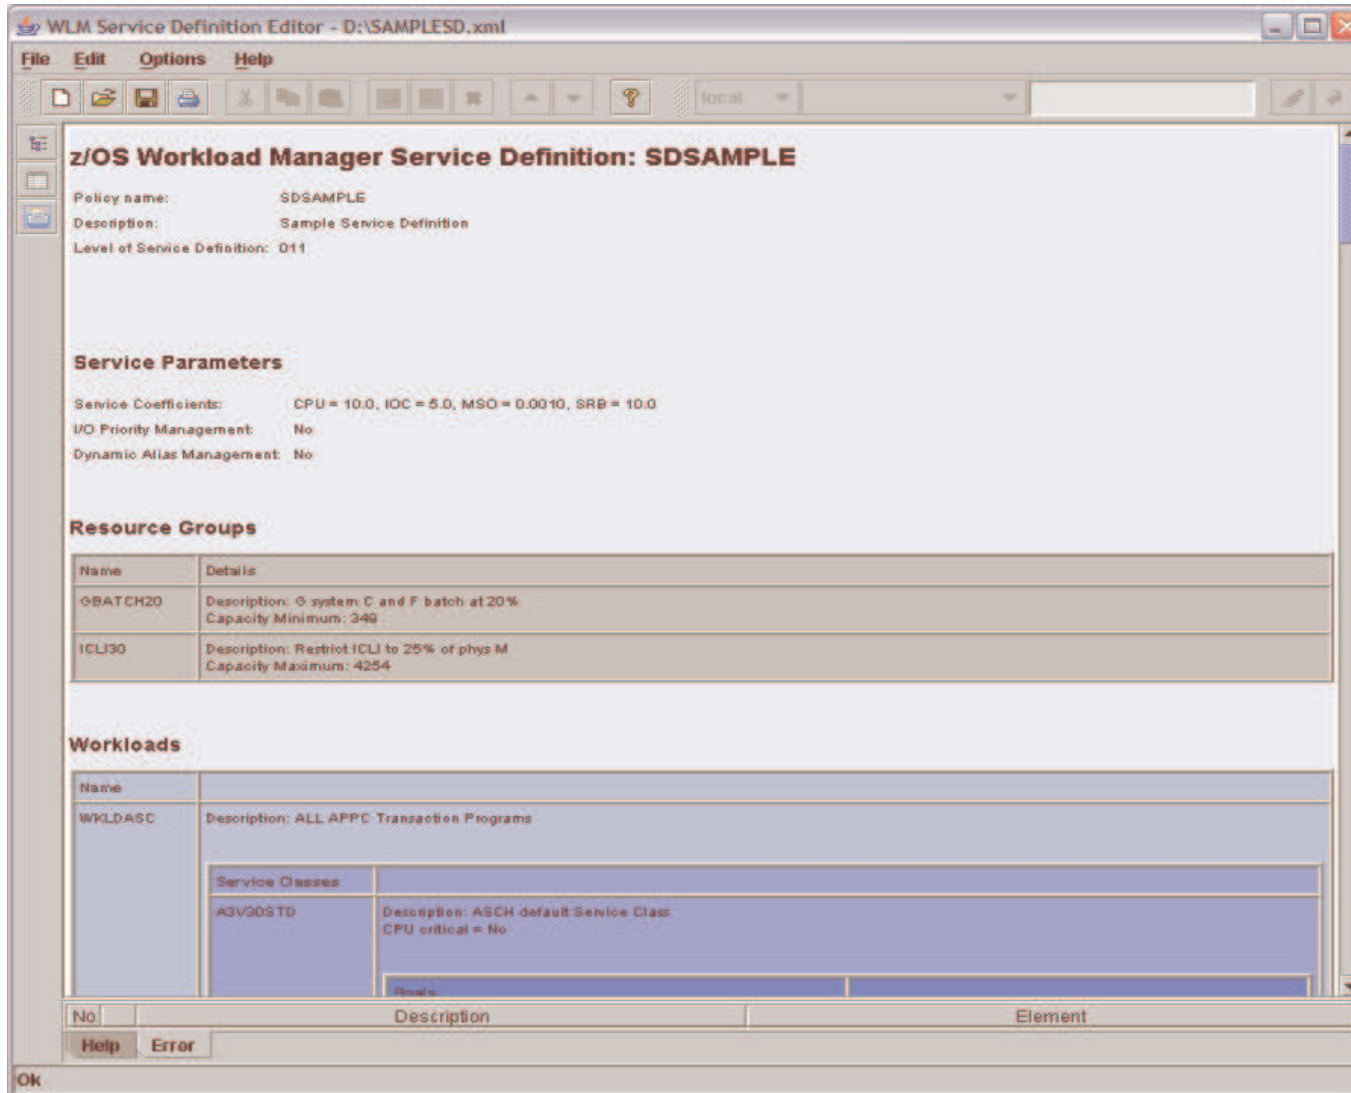
*Example: Easy generation of new views*



- Table columns can be moved as needed
- Table rows can be sorted as needed
- Rearranged tables can be edited
- Search function for arbitrary text within tables

# WLM Administrative Application

## Example: Print View



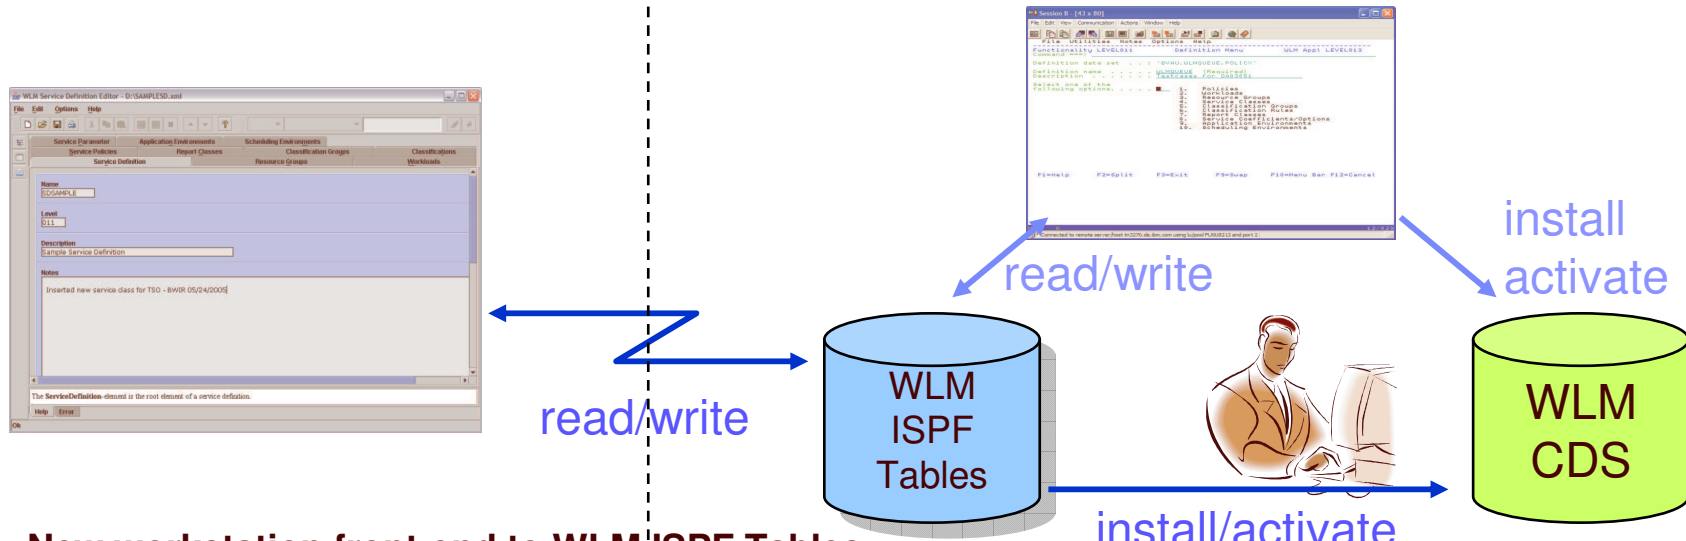
- Service Definition is converted to HTML-Document
- In order to print HTML-Document the configured WWW-Browser of the Operating System is launched



# WLM Administrative Application

*Link to existing WLM Administrative Application*

Workstation z/OS



- **New workstation front-end to WLM ISPF Tables**
  - FTP connection with automatic upload and download
  - FTP connection profiles
- **Policy install and activate via operator command**

**Delivery:**

- As a tool that can be downloaded via the WLM homepage <http://www.ibm.com/servers/eserver/zseries/zos/wlm/tools/>
- Expected availability: 8/30/2006

## Fairly new ITSO Redbook

*System Programmer's Guide to: Workload Manager (SG24-6472)*



— Content:

1. Introduction
2. How WLM works
3. WLM functions
4. Implementation and workload classifications
5. Batch considerations
6. TSO, STC, and APPC workloads
7. DB2 workload considerations
8. WebSphere Application Server workload considerations
9. UNIX System Services considerations
10. Transactional workload considerations

— Download via:

<http://www.redbooks.ibm.com/abstracts/sg246472.html?Open>

## WLM Documentation

### Documentation



z/OS MVS Planning: Workload Management (SA22-7602)

<http://publibz.boulder.ibm.com/epubs/pdf/iea2w151.pdf>

z/OS MVS Programming: Workload Manager Services (SA22-7619)

<http://publibz.boulder.ibm.com/epubs/pdf/iea2w252.pdf>

Redbook – System Programmer's Guide to: Workload Manager (SG24-6472)

<http://www.redbooks.ibm.com/abstracts/sg246472.html?Open>

Redbook – z/OS Intelligent Resource Director (SG24-5952)

Redbook – Effective zSeries Performance Monitoring Using RMF (SG24-6645)

### Internet Links



WLM

<http://www.ibm.com/s390/wlm>

SRM

<http://www.ibm.com/s390/srm>

RMF

<http://www.ibm.com/s390/rmf>

WSC

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

Software Pricing

<http://www-1.ibm.com/servers/eserver/zseries/swprice/>

IRD

<http://www-1.ibm.com/servers/eserver/zseries/ird/>

Redbooks

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

## Additional Optimization and Management bits and bites

### z/OS 1.8

- **WLM Tape I/O priority**
  - SRM calculates a static I/O priority for all ASs and enclaves for tape devices, to be used when no dynamic I/O priority has been assigned
- **WLM delays mapped to ARM services**
- **GRS CNS processing**
  - Now you can specify the contention notifying system (CNS) for GRS Star. Available also for z/OS V1.7 with APAR OA11382.
- **Improved OAM BLOB support**
  - OAM DB2 Binary Large Object Support enabling objects larger than 32 KB using the binary large object (BLOB) data type. Coexistence support for earlier release levels to coexist in an OAMplex.
- **SCLM improvements**
  - Improved performance of the SCLM Library utility (option 3.1). Reduced number of ISPF service calls and caching of certain data,
  - Member level locking avoids loss of updates that can occur when 2 or more users update the same member at the same time
    - Reduced amount of DASD space required to store listings from SCLM since they can now be compressed

# z/OS 1.8

## Helping meet challenges of on demand business

### Improving Usability and Skills

More Health Checks and framework improvements, HCM, ISPF, InfoPrint Central, & RRS improvements, Extensions to the IBM Configuration Assistant for z/OS Comm Server, EE usability improvements, ...

### Integrating new Applications and Supporting Industry and Open Standards

Unicode improvements; inittab support; LE improvements, new LDAP server, z/OS XML System Services, BPXBATCH, CIM, XPLINK, Networking APIs, XL C/C++ enhancements, ...

### Extending the Network

Sysplex partitioning support, Dynamic DNS registration, JES3 NJE via TCP/IP, TN3270 and telnet improvements, IP filtering, IKE, IPsec, WTS, Samba, improved sysplex failure recovery, REXX FTP interface, ...

### Scalability & Performance


Support for up to 4 TB memory, New zIIP processor, CFRM performance improvements, GRS ENQ limits, Parallel VARY ONLINE, z/OS UNIX file descriptors, Device Groups, DADSM/CVAF Rapid Index Rebuild, z/OS USS asynch socket fast path, 64-bit PDSE support ...

### Improving Availability

Master console elimination, Logstream rename, test and prod logstream separation, Recoverable BRML, GDPS enhanced recovery, Fast replication, HyperSwap trigger, z/OS UNIX latch contention detection, ...

### Optimization and Management

WLM, zAAP Stage 2, Sysplex Distributor, Tape I/O priority, JES2 balancing of WLM-managed initiators, GRS CNS processing, Improved OAM BLOB support, SCLM improvements, Group Capacity Limit ...



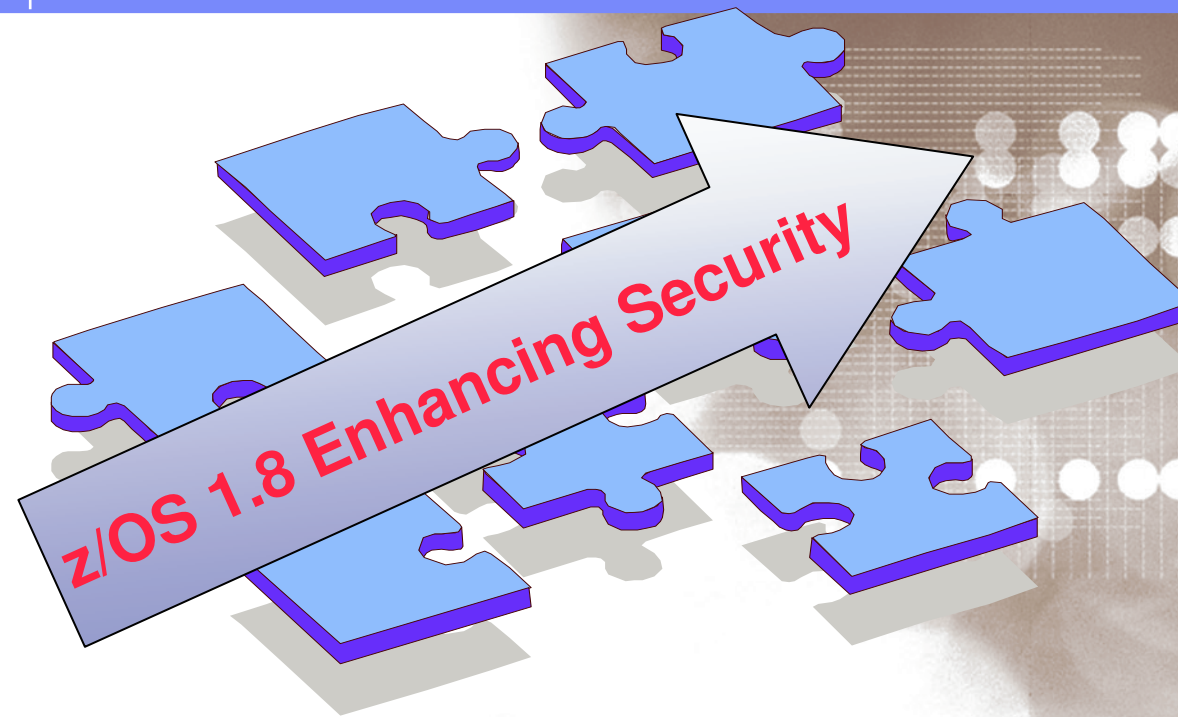
### Enterprise-Wide Roles

for eWLM, DFSMSrmm CIM agent, RMFeServer OS Monitoring Stage II, WLM delay services mapped to ARM ...

**Enhancing Security**  
**PKI extensions, Distributed Identity Support, RACF password phrase support, Tape data set protection, 128-bit AES for IPsec, SAF identity tokens, RACF virtual key rings, IDS configuration, ...**



LSU Danmark



Danmark LSU October 2006

*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation

## z/OS 1.8 Security Enhancements

- **Enhancements to RACF profiles:**
  - RACF database enhancements to allow a template block to be continued into another block
  - New class attribute for disallowing generic profiles in a class
  - IRRDPI00 LIST command granularity
- **IRRUT200 and IRRUT400 utility updates:**
  - Improvements will help prevent database corruption
- **Distributed identity support:**
  - SAF Identity Token support
  - EIM/TAM schema convergence
- **Password phrase support**
  - 14-100 character password phrases
  - Vastly larger namespace
- **Virtual key-ring support**
  - All certificates owned by the same user ID can be in a virtual key-ring
  - No need to manually create the virtual key-ring
  - Can help simplify administration for SSL applications such as FTP



## z/OS 1.8 Security Enhancements...

- **PKI Extensions:**
  - SCEP support (programmatic certificate request support for network devices)
  - Multiple CA support (in a single image)
- **Support for defining IDS policy in a file**
  - In addition to via LDAP
- **Improved tape data set security administration**
  - Can use DATASET class without activating TAPEVOL or TAPEDSN
  - Can specify that all data sets on the same tape should have common authorization
- **IPSEC support for 128-bit AES**
- **EAL4+ certification for z/OS 1.7 with the optional RACF feature:**
  - For Controlled Access Protection Profile (CAPP) and Labeled Security Protection Profile (LSPP)
- **PKDS Key Management SPE**



## z/OS 1.8 Security Enhancements


### *RACF DB Enhancement: Template Expansion*



- **There is no way to disallow generic profiles permanently in z/OS 1.8**
- **IRRDPI00 LIST command has voluminous output**
- **The RACF USER profile template is almost full (approaching 4K Limit).**
  - Template for each profile (group, user, dataset, general resource) fits in one 4K block
  - In order to add new RACF database fields to the user profile the template must expand into another block
- **With z/OS 1.8 RACF Initialization and utility processing is updated to allow for any template to expand more than one block (4K)**
  - This means that new fields can be added to RACF DB in future
- **Any application which reads and processes the RACF database templates directly may be affected**
  - **Support is transparent to applications which use intended interfaces to process RACF database fields**
    - RACROUTE, ICHEINTY / ICHEACTN / ICHETEST, RACF commands (ADDUSER, RDEFINE, etc), RACF callable services (R\_admin, ck\_access, etc)
- ▶ Roll back support to z/OS **V1R4, V1R5, V1R6**, and **V1R7** with APAR **OA12443**

# Security RACF DB Enhancement


## *Generic Profiles*



- No way to disallow Generic profile processing for install defined class
  - With z/OS 1.8 a RACF class can be defined with an attribute that prevents generic profiles from being created in that class
- New ICHERCDE keyword: **GENERIC=ALLOWED | DISALLOWED**
  - Following IBM defined classes are updated with **GENERIC=DISALLOWED**
    - CDT, KERBLINK, REALM, SECLABEL, and SECLMBR
- When sharing DB with a lower level system make sure you:
  - Always administer a dynamic class that disallows generics from a system running z/OS 1.8 or higher.
  - Always administer profiles in dynamic classes where generics are disallowed from systems running z/OS 1.8 and higher.

## Security RACF DB Enhancement

### *IRRDPI00 enhancements*



- IRRDPI00 lists every field definition in the RACF DB.
  - This causes thousands of lines of output
- With this enhancement you will have a possibility to specify a more granular listing
- IRRDPI00 LIST command has new keywords to specify profile type, segment name, and field name
  - IRRDPI00 LIST [ (*profile-type* [ *segment-name* [ *field-name* ] ])]
    - Example: To list all fields in the OMVS segment of the USER profile, issue:  
**IRRDPI00 LIST(USER OMVS)**
    - Example: To list the HOME keyword in the OMVS segment of the USER profile, issue:  
**IRRDPI00 LIST(USER OMVS HOME)**

## Security RACF DB Overwrite Issue – The problem


### *Utility Enhancement*



- **Both IRRUT200 and IRRUT400 has enhancements to prevent potential Database corruptions.**
- **Database corruption resulting from running IRRUT200**
  - IRRUT200 copies/verifies datasets on a 1-1 basis. Used to backup datasets.
    - Corruption may occur when DD SYSRACF (source) equals DD SYSUT1 (target).
    - Corruption may also occur if both DDs are specified and SYSUT1 is an in-use active RACF dataset.
- **Database corruption resulting from running IRRUT400**
  - IRRUT400 copies datasets on an X-to-Y basis. Used to shrink a database from X to X-n datasets. Used to expand a database from X to X+n datasets. Also used to copy X to X datasets across DASD types. All rebuild the index.
    - Corruption may occur when DD INDDx (source) equals DD OUTDDx (target).
    - Corruption may also occur if OUTDDx is an in-use active RACF dataset
  - **Copying active primary datasets to active backups may not result in mirror images.**
    - The process is multi-step and **unserialized**.
      - Today one must use IRRUT200 to copy the active primary to a dataset with the same name as the active backup, but on a different volume. Then **RVARY INACT** the active backup and **uncatalog** it. Then **catalog** the new copy. Then **RVARY ACTIVE** the backup, picking up the new copy.
      - **If updates to the primary occur after the IRRUT200 copy function, then by the time the backup is activated, it is no longer an exact copy of the primary**

# Security RACF DB Overwrite Issue – And a Solution


## *Utility Enhancement*



- **The new safety features for both IRRUT200 and IRRUT400 are similar.**
- **Both will take volume labels into account.**
- **IRRUT200**
  - when DD SYSRACF (source) equals DD SYSUT1 (target) issue message IRR62073I, terminate utility with RC=12
  - when both DDs are specified and SYSUT1 (target) is an active RACF dataset, issue message IRR62072I, terminate utility with RC=12
- **IRRUT400**
  - when INDDx (source) = OUTDDx (target) issue message IRR65041I, terminate utility with RC=16
  - when OUTDDx (target) is an active RACF dataset, issue message IRR65040I, terminate utility with RC=16

## Security RACF DB Overwrite Issue – And a Solution...

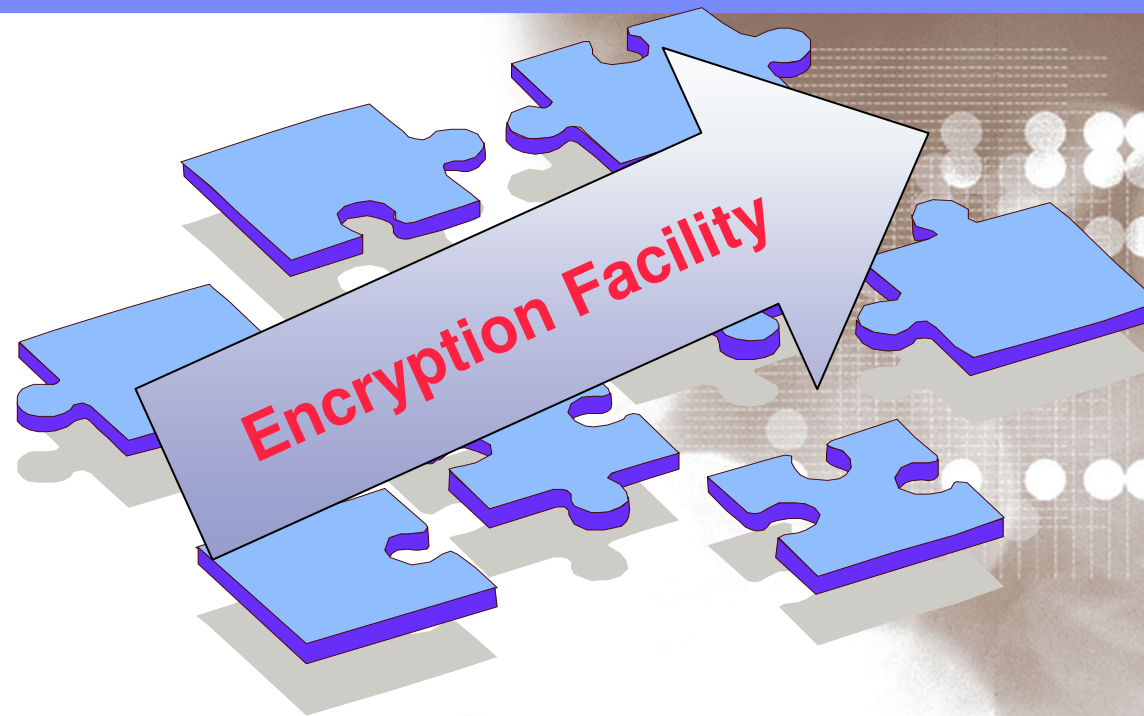
### *Utility Enhancement*



- **The new IRRUT200 synchronized copy:**
  - First run IRRUT200 in the usual **verify mode** to ensure that the active primary dataset is **not damaged**.
  - Run IRRUT200 as if you were doing a copy (DDs for both SYSRACF and SYSUT1), but add new keyword **PARM=ACTIVATE**
    - no verification is performed for PARM=ACTIVATE, SYSIN and SYSPRINT are ignored
  - **SYSUT1 (target)** must be the **in-use inactive backup** associated with the SYSRACF (source) specified in-use active primary
    - the existing IRRUT200 rule of copying only between similar device types still applies
  - **IRRUT200** will get **exclusive serialization**. After the copy completes, an internal RVARY ACTIVATE against the SYSUT1 (target) specified backup will be done. Serialization will be released.
    - In addition to the IRRUT200 messages, RVARY messages will now be found in IRRUT200's SYSUT2 DD.
    - There will be no RVARY password prompt.



## LSU Danmark



Danmark LSU October 2006

*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation

## Encryption of Data

### *A Business Imperative*

- **Businesses are proactively focusing on securing customer and business data**
  - Increasing regulatory requirements driving need for security of data for audit and compliance
  - Recent events highlight impacts caused by loss/theft of removable data
  - Requirements for tighter security driving need for encryption of data





## Security Headlines Daily

*Is Anything More Important to the Success and Survival of Your Business?*

More Than 90% Of Companies Regularly Expose Employee And Customer Data<sup>1</sup>

FBI – Businesses Reluctant To Report Cyber Attacks<sup>2</sup>

One In Four Identity-Theft Victims Never Fully Recover<sup>3</sup>

PCI: Card Associations Unite to Fight Fraud With Collaborative Standard<sup>4</sup>

<sup>1</sup> Reconnex Insider Threat Index August 2005

<sup>2</sup> 2005 CSI/FBI Computer Crime and Security Survey

<sup>3</sup> Nationwide Mutual Insurance Co. Survey July 2005

<sup>4</sup> Green Sheet Inc. August 2005 Issue 2



## Regulatory and Compliance Considerations


- Gramm-Leach-Bliley Financial Services Modernization Act (GLBA)
- Sarbanes Oxley (SOX)
- European Union Data Protection Directive (EUPA)
- International IT Security Standard (ISO 17799)



## Potential Costs of A Security Breach

- \$ Cost of research and recovery
- \$ Cost to notify customers
- \$ Lost customers/business
- \$ Problem solution or remediation
- \$ Claims from trusted vendors and business partners

**\$\$ Damage to brand image**



# The Power of Mainframe Encryption

*Helping to reduce risk across your value-net*




**Helping to protect data over the Internet**



**Helping to protect data leaving your enterprise\***

- Customer objectives:**
- Only intended party is allowed to decrypt
  - Availability of the keys and decryption services when you need them



**Centralized Key Management**



**Helping to protect archived data\***

- IBM Encryption Facility for z/OS planned GA dates:
  - Encryption Services – 28 Oct, 2005
  - DFSMSdss Encryption - 2 Dec, 2005

# IBM Encryption Facility for z/OS, 1.1

Licensed Program Product  
MSU-based pricing\*

Runs on the following servers: System z9 109 (z9-109), or equivalent  
zSeries z900 or z990, or equivalent  
zSeries z800 or z890, or equivalent

Requires: z/OS V1.4 or higher      z/OS.e V1.4 or higher

## Feature: *Encryption Services*

Optional Priced Feature

## Feature: *DFSMSdss Encryption*

Optional Priced Feature

## Encryption Facility Client


Web download

- Supports encrypting and decrypting of data at rest (tapes, disk)
- Supports either Public Key/Private keys or passwords to create highly-secure exchange between partners
- Java technology-based code that allows client systems to decrypt and encrypts data for exchange with z/OS systems
- Allows encryption and compression of DUMP data sets created by DFSMSdss
- Supports decryption and decompression during RESTORE

## The Mainframe

### *A History of Enterprise Security*

- **Hardware Cryptography: 1970**
- **RACF: controls access to resources and applications – 1976**
- **Key management built into operating system (ICSF) – 1991**
- **Security Applications: Tivoli & Leading Software Vendors**
- **Intrusion Detection Services (IDS): 2001**
- **PKI: create digital certificates & act as Certificate Authority (CA) – 2002**
- **Multilevel Security (MLS): 2004**
- **Encryption Facility for z/OS: 2005**
- **Encrypting Tape Drive TS1120: 2006**




## Future Directions –

### *Extending Encryption to IBM TotalStorage*

▪ **Statement of Direction:**

- To address customers' growing concern with data security, IBM is announcing a statement of direction for the development, enhancement and support of encryption capabilities within storage environments such that the capability does not require the use of host server resources.
- This includes the intent to offer, among other things, capabilities for products within the IBM TotalStorage portfolio to support outboard encryption and to leverage the centralized key management functions planned for z/OS ICSF.



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

# Tape encryption with System z *in the enterprise*




Help secure data from theft or compromise


## Why z/OS centralized key management?

- Can help to protect and manage keys
  - Highly secure and available key data store
  - Long term key management
  - Disaster recovery capabilities
- Single point of control
- Over a decade of production use


## Encryption Facility for z/OS, V1.1



- Flexible options for business partner exchange
- Partners can encrypt and decrypt using no-charge JAVA client
- Supports public key or password based exchange



## Plans for encryption in IBM System Storage 2H2006\*



- Highly secure tape library
- High performance archive encryption
- Transparent to existing processes and applications
- Can help provide audit compliance

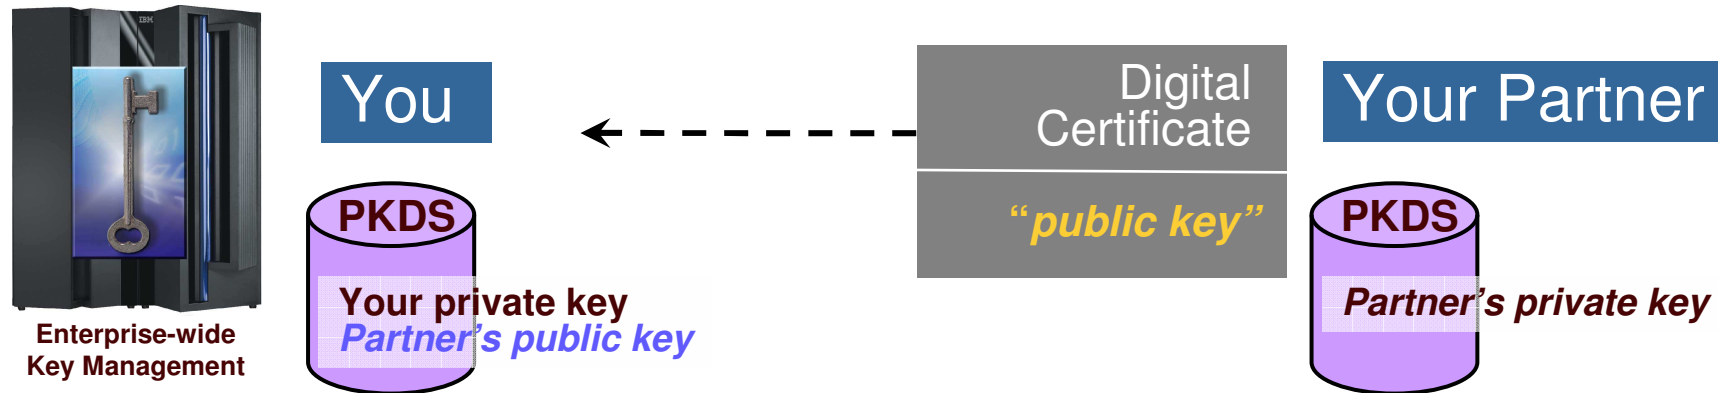
Enterprise scope



## Establishing a Trusted Exchange with Your Partners

### Key Exchange –

- Digital Certificates or passwords can be used to identify and authenticate



*“I know you by your Public Key, so I can create a file that only you can read.”*

OR

*“I know you by your password.”*

### Options for Partners to acquire a Digital Certificate:

1. z/OS customer can generate a certificate for the partner
  - z/OS can be a Digital Certificate Authority using z/OS PKI Services
2. Partner may already have a Digital Certificate
3. Partner may use third party Digital Certificate Authority

# Encrypting with Key Management



Enterprise-wide Key Management

You

Encryption Facility feature:  
*Encryption Services*

Invoked via JCL

- 1) Generate data key for either AES128 or TDES encryption
  - random for RSA key management, or
  - derived from password via PKCS#12
- 2) If RSA key mgmt:
  - encrypt data key with RSA public key, and
  - store encrypted data key in the file header
- 3) Compress the data
- 4) Encrypt the data using the data key
- 5) Send file to partner

Your Partner

*Encryption Facility Client*

*Partner decrypts file*

- 1 Decrypt data key with RSA private key, or derive data key from the password
- 2 Decrypt the data

If z/OS site: can use *Encryption Facility for z/OS* or *Encryption Client (Java code)*

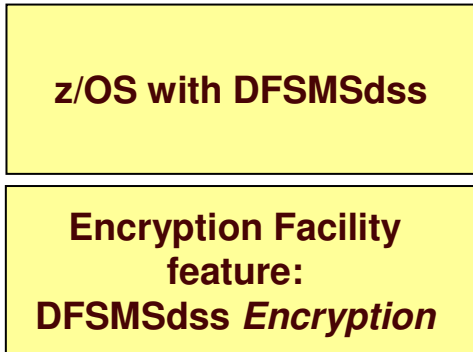
If non-z/OS: uses *Encryption Client (Java code)*

\* Optionally leverages piping functions in z/OS Unix Systems Services to help reduce elapsed time for large datasets

# Archival Encryption with DFSMSdss and Key Management

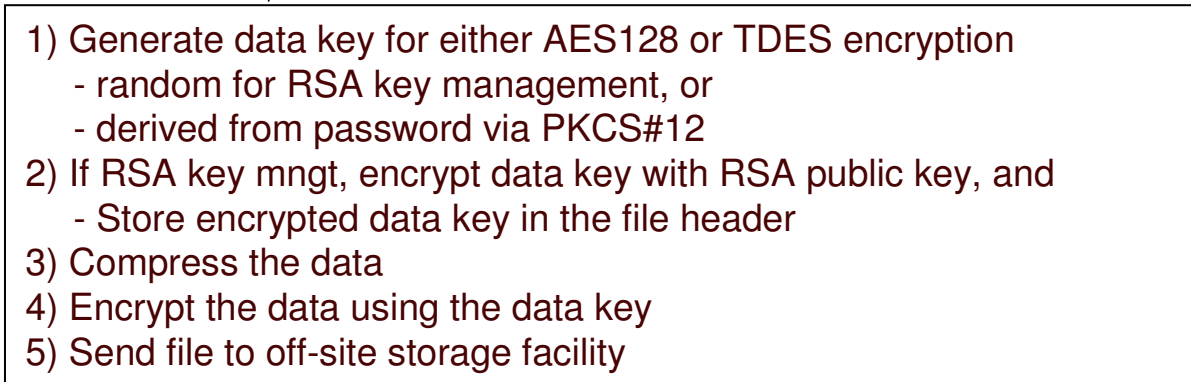


Enterprise-wide Key Management



Three options for invoking encryption in DFSMSdss:

1. Additional Keyword in DUMP and RESTORE commands
2. DFSMShsm
3. ISMF panels



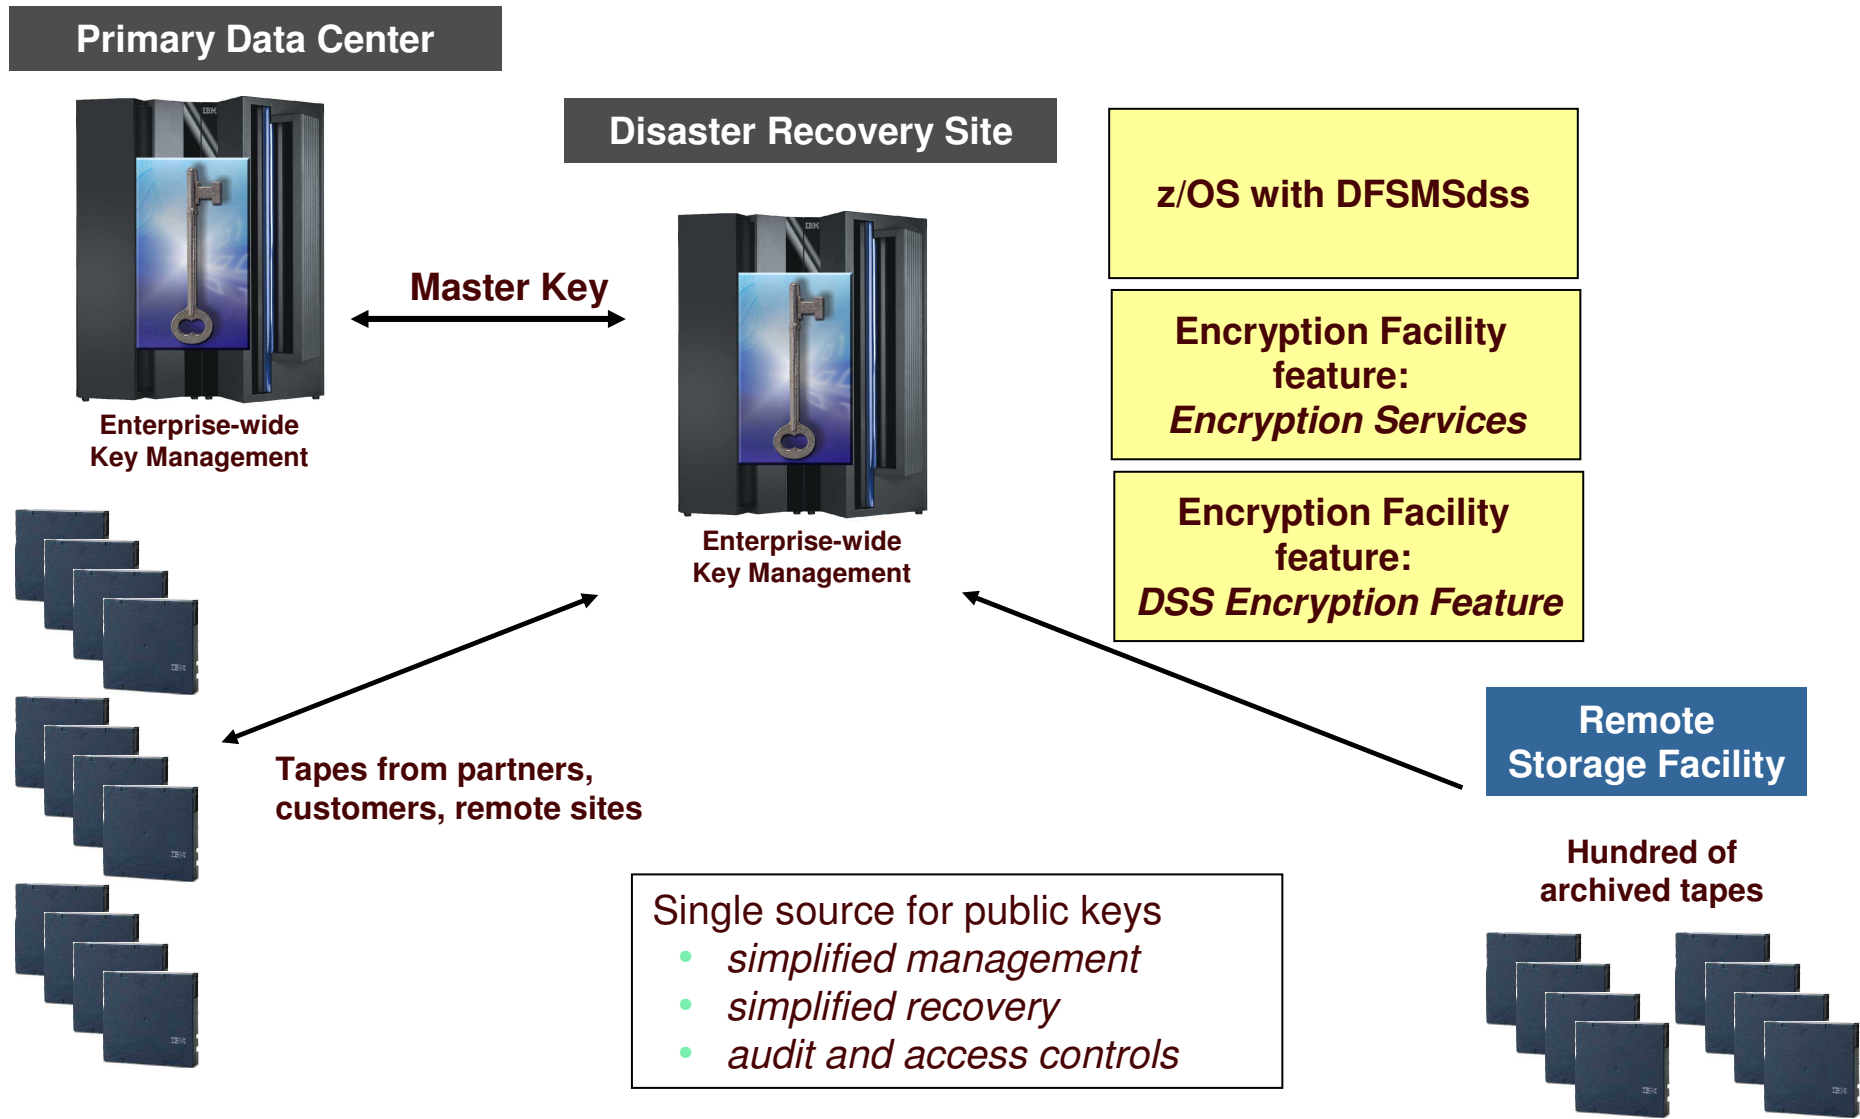
Archived tapes

Estimated\* resulting tape volumes:

Tapes without Encryption	With Encryption No Compression	With Encryption and Hardware Compression
X	2-3X	1-1.5X


\* These measurements are examples of the maximum compression achieved in a lab environment with no other processing occurring and do not represent actual filed measurements. Further, the type of data will affect the hardware and software compression capabilities. Details available upon request.

# Recovery with Centralized Key Management



## Security

### PKDS Key Management SPE



- **z/OS Encryption Facility (EF) became available 4Q2005**
  - Requires exchanging public keys with business partners to enable data encryption
  - No native PKDS key management functions in ICSF
    - End users need to generate key pairs and extract public keys
    - EF requires certificate/key management support in the Security Manager (RACF, ACF/2, etc) to do the above
  - Procedure to exchange keys between mixed shops difficult
  - ICSF for z/OS 1.8 updated with new panels to:
    - Generate/delete public and private key pairs in the PKDS
    - Export the public half as a self-signed x.509 certificate
    - Import a public key from a partner's certificate into the PKDS
- **Function contained in the base of HCR7731 (V1R8)**
- **Also rolled back to earlier releases as an SPE**
  - APAR Number - **OA15156**
  - Releases
    - HCR770A, HCR770B, HCR7720, and HCR7730

# z/OS 1.8

## Helping meet challenges of on demand business

### Improving Usability and Skills

More Health Checks and framework improvements, HCM, ISPF, InfoPrint Central, & RRS improvements, Extensions to the IBM Configuration Assistant for z/OS Comm Server, EE usability improvements, ...

### Integrating new Applications and Supporting Industry and Open Standards

Unicode improvements; inittab support; LE improvements, new LDAP server, z/OS XML System Services, BPXBATCH, CIM, XPLINK, Networking APIs, XL C/C++ enhancements, ...

### Extending the Network

Sysplex partitioning support, Dynamic DNS registration, JES3 NJE via TCP/IP, TN3270 and telnet improvements, IP filtering, IKE, IPSec, WTS, Samba, improved sysplex failure recovery, REXX FTP interface, ...

### Scalability & Performance

Support for up to 4 TB memory, New zIIP processor, CFRM performance improvements, GRS ENQ limits, Parallel VARY ONLINE, z/OS UNIX file descriptors, Device Groups, DADSM/CVAF Rapid Index Rebuild, z/OS USS asynch socket fast path, 64-bit PDSE support ...

### Improving Availability

Master console elimination, Logstream rename, test and prod logstream separation, Recoverable BRLM, GDPS enhanced recovery, Fast replication, HyperSwap trigger, z/OS UNIX latch contention detection, ...

### Optimization and Management

WLM, zAAP Stage 2, Sysplex Distributor, Tape I/O priority, JES2 balancing of WLM-managed initiators, GRS CNS processing, Improved OAM BLOB support, SCLM improvements, Group Capacity Limit ...

### Enhancing Security

PKI extensions, Distributed Identity Support, RACF password phrase support, Tape data set protection, 128-bit AES for IPSec, SAF identity tokens, RACF virtual key rings, IDS configuration, ...



### Enterprise-Wide Roles

WLM/EWLM Service Class Correlation, zAAP Reporting for eWLM, DFSMSrmm CIM agent, RMFeServer OS Monitoring Stage II, WLM delay services mapped to ARM ...

## Enterprise-Wide Roles

- EWLM/WLM service class correlation
  - WLM accepts classification when same-named SC passed from EWLM
  
- New DFSMSrmm CIM agent support
  - Support creation, change, and deletion of volumes and data sets
  - Complements the query and display functions in z/OS 1.7
  
- zAAP reporting for EWLM
  
- Infoprint Central GUI improvements
  - Real-time status information display
  - Change online/offline status, reset a printer
  - Stop a print job without canceling it
  - Auditability enhancements
  - TRACEROUTE support



# z/OS 1.8

## Helping meet challenges of on demand business

### Improving Usability and Skills

More Health Checks and framework improvements, HCM, ISPF, InfoPrint Central, RRS improvements, Extensions to the IBM Configuration Assistant for z/OS Comm Server, EE usability improvements, ...

### Integrating new Applications and Supporting Industry and Open Standards

Unicode improvements; inittab support; LE improvements, new LDAP server, z/OS XML System Services, BPXBATCH, CIM, XPLINK, Networking APIs, XL C/C++ enhancements, ...

**Extending the Network**  
**Sysplex partitioning support, Dynamic DNS registration, JES3 NJE via TCP/IP, TN3270 and telnet improvements, IP filtering, IKE, IPSec, WTS, Samba, improved sysplex failure recovery, REXX FTP interface, ...**

### Scalability & Performance

Support for up to 4 TB memory, New zIIP processor, CFRM performance improvements, GRS ENQ limits, Parallel VARY ONLINE, z/OS UNIX file descriptors, Device Groups, DADSM/CVAF Rapid Index Rebuild, z/OS USS asynch socket fast path, 64-bit PDSE support ...

### Improving Availability

Master console elimination, Logstream rename, test and prod logstream separation, Recoverable BRLM, GDPS enhanced recovery, Fast replication, HyperSwap trigger, z/OS UNIX latch contention detection, ...

### Optimization and Management

WLM, zAAP Stage 2, Sysplex Distributor, Tape I/O priority, JES2 balancing of WLM-managed initiators, GRS CNS processing, Improved OAM BLOB support, SCLM improvements, Group Capacity Limit ...

### Enhancing Security

PKI extensions, Distributed Identity Support, RACF password phrase support, Tape data set protection, 128-bit AES for IPSec, SAF identity tokens, RACF virtual key rings, IDS configuration, ...

### Enterprise-Wide Roles

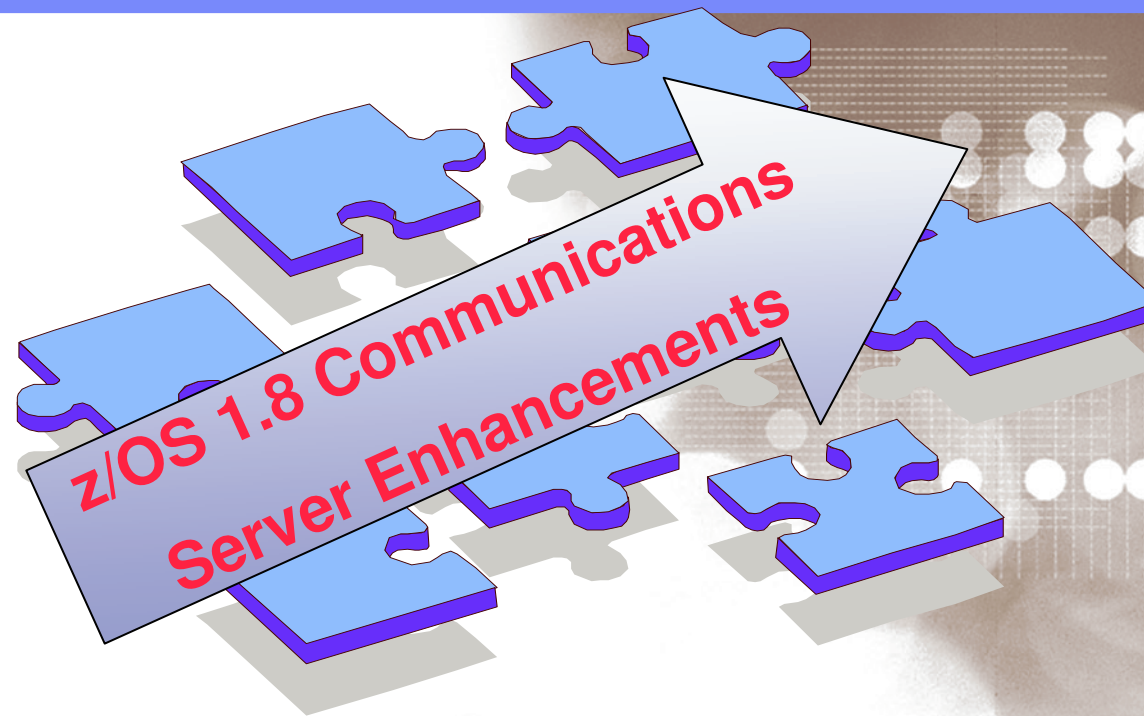
WLM/EWLM Service Class Correlation, zAAP Reporting for eWLM, DFSMSrmm CIM agent, RMFeServer OS Monitoring Stage II, WLM delay services mapped to ARM ...







LSU Danmark



Danmark LSU October 2006

*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation

## Communications Server Support

### *Recent history: z/OS 1.4-1.5*

- z/OS 1.4
  - Sysplex wide security associations
  - Sysplex wide source VIPA
- z/OS 1.5
  - Increased VIPADISTRIBUTE ports (4 -> 64)
  - Increased maximum DVIPAS per stack (256 -> 1024)
  - Sysplex Distributor enhancements
    - timer based "stickiness" for IP address, DVIPA, and ports
    - round robin distribution
  - Enhanced DVIPA activation in a sysplex

## Communications Server Support

### *Recent history: z/OS 1.6-1.7*

#### ■ z/OS 1.6

Extended configuration control for outbound connections

- outbound connections use same IP address as inbound
- job specific source IP addressing

Self monitoring for improved availability of TCP/IP stacks

Can use sysplex distributor without z/OS stack being a general intermediate routing node

All sysplex functions support IPv6

z/OS Load Balancing Advisor (LBA) - 4Q04

- for outboard load balancers
  - supports all IP based application workloads
  - sysplex scope
  - z/OS 1.4 and above

#### ■ z/OS 1.7

Improved automomics to reduce/eliminate operator actions

Improved usability for DVIPAs

Improved intra-sysplex flexibility for sysplex distributor and DVIPA

Improved sysplex distributor load balancing decisions and end user response time

## Communication Server Sysplex Directions

- Use of “local” vs “remote” connectors in a sysplex results in performance vs. availability tradeoffs
  - Improved Comm Server controls allowing configuring the level of “favor” for local connectors
- Additional (beyond that provided in z/OS 1.6) capability for choosing outbound IP connections
  - Destination based source IP address selection
  - Useful if JOBNAME is unpredictable or if a JOB connects to multiple business partners
- z/OS 1.6/1.7 Comm Server recovers from range of error conditions within Sysplex
  - Add automated recovery from network outage conditions
  - Network interface (active/inactive), 1<sup>st</sup> hop router issues, etc.

## Communication Server Sysplex Directions...

- Connectivity to multiple network security areas (zones)
  - Enhanced structure and function to enable networking Sysplex functions connected to multiple security zones
- Support for Dynamic DNS (DDNS) for dynamic name registration (previously supported by DNS/WLM function)

## z/OS 1.8 Communications Server news

- **VTAM and TCP/IP network sysplex partitioning support**
  - Can define multiple logical networks within a sysplex
- **Dynamic DNS registration/deregistration**
- **IBM Configuration Assistant for z/OS Communications Server GUI**
  - Introduced in z/OS 1.7 as the Network Security Configuration Assistant
  - Extended in z/OS 1.8 to help with IDS and QoS configuration
- **IPv6 support improved**
- **WTS (Windows Terminal Server) support**
- **SMB support for Linux Samba clients**
- **REXX interface designed to allow calling FTP programmatically**
  - Sample program included
- **JES3 support for NJE via TCP/IP (delayed availability)\***
  - New **NJERMT** parameters
  - New **NETSERV**, and **SOCKET** statements.
  - Planned availability 1H2007

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

## z/OS 1.8 Communications Server news...

- **Communications Server callable APIs:**
  - Support for dropping multiple TCP connections or UDP endpoints
  - Enhancing the Enterprise Extender network management interface to allow CP name wildcarding for HPR connection requests
- **TN3270 connection management:**
  - New API designed to allow for the retrieval of performance data for TN3270 server sessions
  - Add TCP connection IDs to TN3270 SMF records
- **Telnet enhancements:**
  - Ability to specify TN3270 servers should automatically clean up hung SNA Telnet sessions when a new TN3270 connection is initiated
  - Support for the resolution of system symbols for Telnet unformatted system services message (USSMSG) processing.
  - USS table assignment from the LU exit designed to provide more flexibility in assigning a USS table based on client criteria.



- LSU Specific Network presentation will cover much more details



# z/OS 1.8

## Helping meet challenges of on demand business

### Improving Usability and Skills

More Health Checks and framework improvements, HCM, ISPF, InfoPrint Central, & RRS improvements, Extensions to the IBM Configuration Assistant for z/OS Comm Server, EE usability improvements, ...

### Scalability & Performance


Support for up to 4 TB memory, New zIIP processor, CFRM performance improvements, GRS ENQ limits, Parallel VARY ONLINE, z/OS UNIX file descriptors, Device Groups, DADSM/CVAF Rapid Index Rebuild, z/OS USS asynch socket fast path, 64-bit PDSE support ...

### Improving Availability

Master console elimination, Logstream rename, test and prod logstream separation, Recoverable BRLM, GDPS enhanced recovery, Fast replication, HyperSwap trigger, z/OS UNIX latch contention detection, ...

### Integrating new Applications and Supporting Industry and Open Standards

Unicode improvements; inittab support; LE improvements, new LDAP server, z/OS XML System Services, BPXBATCH, CIM, XPLINK, Networking APIs, XL C/C++ enhancements, ...



### Optimization and Management

WLM, zAAP Stage 2, Sysplex Distributor, Tape I/O priority, JES2 balancing of WLM-managed initiators, GRS CNS processing, Improved OAM BLOB support, SCLM improvements, Group Capacity Limit ...

### Enhancing Security

PKI extensions, Distributed Identity Support, RACF password phrase support, Tape data set protection, 128-bit AES for IPSec, SAF identity tokens, RACF virtual key rings, IDS configuration, ...

### Extending the Network

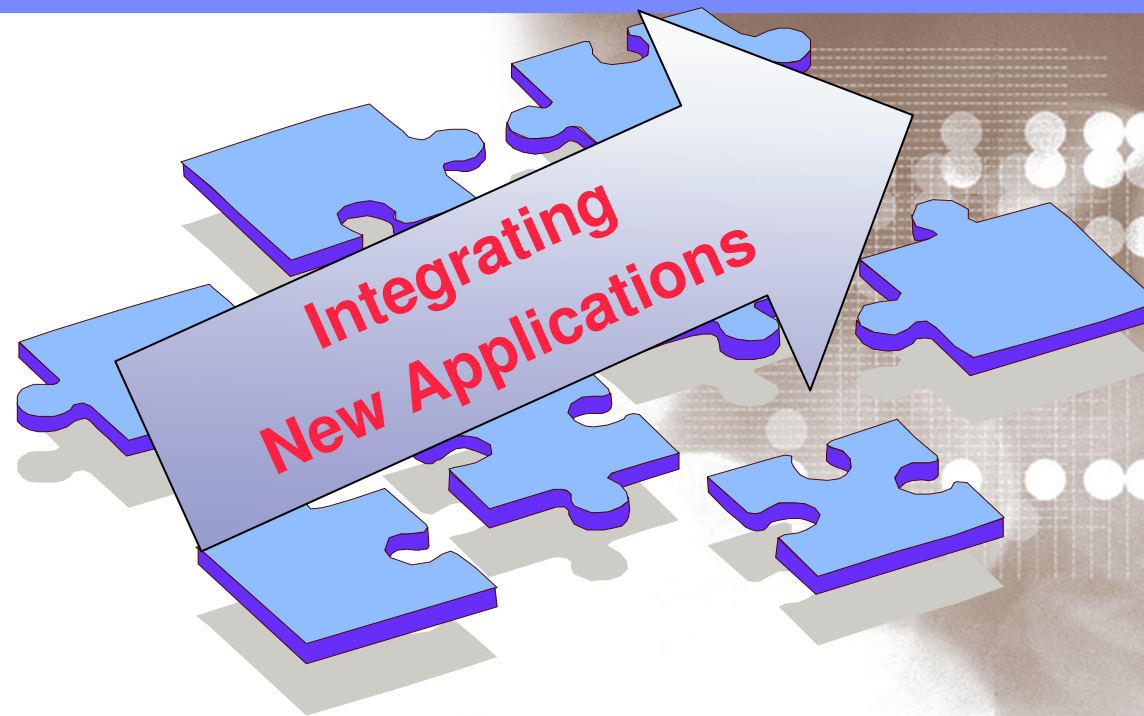
Sysplex partitioning support, Dynamic DNS registration, JES3 NJE via TCP/IP, TN3270 and telnet improvements, IP filtering, IKE, IPSec, WTS, Samba, improved sysplex failure recovery, REXX FTP interface, ...

### Enterprise-Wide Roles

WLM/EWLM Service Class Correlation, zAAP Reporting for eWLM, DFSMSrmm CIM agent, RMFeServer OS Monitoring Stage II, WLM delay services mapped to ARM ...



## LSU Danmark



Danmark LSU October 2006

*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation


## z/OS 1.8 Enhancements at a Glance

### *Integrating New Applications*

- NEW LDAP server
- BPXBATCH enhancements
- z/OS XML System Services
- Unicode improvements
- USS `/etc/inittab` support
- XPLINK enhancements



LSU Danmark



Danmark LSU October 2006

*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation

## z/OS 1.8 LDAP Enhancements at a Glance

*Significantly improved performance, availability, scalability and auditing*

- Better performance
  - The z/OS LDAP V1.8 server is rewritten:
    - from C++ to C, which includes a restructure of both front and back end.
    - The SDBM backend is modified to use the enhanced SAF R\_Admin interface.
    - A file based back end, LDBM, provides better performance as **all entries are cached in storage for quick retrieval.**
- Better availability
  - The z/OS LDAP server will now provide **ARM, TCP/IP restart, Dynamic network interface management, and sysplex support for LDBM**
- **Scalability and constraint relief**
  - The z/OS LDAP server is **64 bit capable** for all backends **except for TDBM** because DB2 does not support 64 bit in the 1.8 time frame.
- **Auditing support**
  - z/OS LDAP 1.8 provides new **SMF 83 audit records** for the LDAP server.

## z/OS 1.8 LDAP Enhancements at a Glance...

*Significantly improved RAS, ease of use, cross-platform consistency and security*


- Additional RAS support
  - LDAP 1.8 provides 2 Health Checks. Ability to trace and invoke dumps in LDAP
- Ease of use support
  - LDAP server provides a file-based backend, LDBM, eliminating need for DB2. Also the change log (GDBM) no longer uses DB2. Instead it is file based.
  - Both of these supports make it easier for you to setup and deploy z/OS LDAP.
- Cross-platform consistency
  - z/OS 1.8 LDAP Client and Server provide additional functions that the current IBM Tivoli Directory Server already has.
- Security enhancements
  - z/OS 1.8 LDAP eliminates requirement for OCSF for cryptographic support. The functions provided by OCSF are now contained in z/OS LDAP.

**New ITDS LDAP server **NOT** available at z/OS 1.8 GA**  
Planned availability on z/OS 1.8 in **1H 2007**

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



LSU Danmark



Danmark LSU October 2006

*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation

## z/OS 1.8 BPXBATCH enhancements

### *longer input parameters and MVS DS output support: description*

- Current limitations in BPXBATCH input and output
  - **make it more difficult to create and maintain BPXBATCH jobs** compared to other batch utilities.
- BPXBATCH Enhancement enables you to:
  - Specify much longer *input* parameters
    - **Up to 64,536 chars** as opposed to old **100 chars limit**
    - **Up to 32.754 chars for TSO Commands**
  - Be able to have data from the standard *output* and standard error streams **to MVS Data sets** (**instead of just z/OS Unix Files**)





## z/OS 1.8 BPXBATCH Enhancements

### *longer input parameters and MVS DS output support: implementation*

- Larger Input Parameter Strings:
  - New **STDPARM DD** to allow specification of UNIX file or MVS DS to contain input parameters up to **65,536** in length
  - Traditional MVS input data sets can be used for **STDPARM**
    - **SYSIN**, Sequential, PDS/PDSE member
  - Specification of **STDPARM overrides** usage of **PARM=** on BPXBATCH invocation
  - **BPXBATCH TSO** Command enhanced to allow up to **32,754** length input string on command invocation
- **STDOUT/STDERR** to MVS Data Sets
  - DDs will now support specification of traditional MVS data sets:
    - Allows output from BPXBATCH jobs to be directed to **SYSOUT**, **Sequential** or **PDS/PDSE** Members
    - Data sets should have **LRECL** that will allow all lines to fit or **truncation** will occur



## LSU Danmark



Danmark LSU October 2006

*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation

## z/OS 1.8 XML Services

### Overview – why is this new component needed?


- z/OS XML System Services is a **NEW** component of z/OS 1.8
  - XML is important to z/OS
  - New workloads are often XML based
  - Databases (i.e. DB2) provides integrated XML Support
  - Platforms, middleware, applications must handle XML efficiently
  - XML parsing overhead threatens to become serious cost problem
  - Imbedding, or pre-req. of a parsing package not feasible
  - Need a good place to implement platform-unique optimization
  - Open software APIs are often not stable
  - Need to support z/OS operating environments



## z/OS 1.8 XML Services

### Overview – and what are the benefits?

- Using z/OS XML, you can expect these **benefits**:
  - Parse XML documents into an easily digested binary form
  - Verify “well-formedness” of XML documents
  - Have large documents be parsed in segments.
  - Use application-specific exits for storage management and string identifiers
  - Query the encoding of XML documents
  - Obtain high performance and high availability in z/OS environments designed for enterprise class z/OS applications with a need to handle XML





LSU Danmark



Danmark LSU October 2006  
*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation


## z/OS 1.8 Unicode Improvements

- "Unicode on Demand" introduced in z/OS 1.7
  - Support dynamic loading of tables when running in **TCB mode only**.
  - SRB mode environments **cannot dynamically add tables** to storage.
    - All tables that are loaded in storage are always **page fixed**.
    - Tables loaded in Common Area Data Space - CADS). **You do not have a choice**.
- Changes are made to z/OS 1.8 Unicode support:
  - Allow adding tables when running both **TCB mode and SRB mode**.
  - Provide an **option to page fix** the tables when loaded into storage.
- The previous support based on the Unicode **3.0.1 Standard**
- z/OS 1.8 Provide support for later Unicode Standards in Normalization service
  - **Unicode 3.2.0**
  - **Unicode 4.0.1**
  - **Unicode 4.1.0**





## LSU Danmark




## z/OS 1.8 USS etc/inittab support

- In z/OS 1.8 the `/etc/inittab` file is the same as used on other UNIX platforms.
- `/etc/inittab` allows you to:
  - Identify system processes that can be started at system initialization;
  - To identify processes that can be restarted automatically when they unexpectedly end.
- If `/etc/inittab` exists in your system, it is used instead of the `/etc/rc` files.





## LSU Danmark



Danmark LSU October 2006

*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation

## XPLINK – *What is it? And what is it good for?*

*... implemented at the System z Hardware level*

- **Extra Performance Linkage (XPLINK) is a call linkage between programs that has the potential for significant performance increase in case of frequent calls between small functions or subprograms.**
- **The C/C++ subroutine linkage on z/OS cannot be considered state-of-the art with respect to performance.**
  - It represents a disproportionate percentage of total execution time, higher yet for C++ than for C due to the many, typically small, functions.
  - Depending on the style of programming, the total prolog/epilog cost may reach a double digit percentage even for C, and thus represents a significant potential for further program optimization.
- **XPLINK significantly speeds up the linkage for C and C++ routines by using a downward-growing stack and by passing parameters in registers. It includes support for reentrant and non-reentrant code, for calls to functions in DLLs, and compatibility with old code.**
- **With XPLINK, the linkage and parameter passing mechanisms for C and C++ are identical. If you link to a C function from a C++ program, you should still specify extern C to avoid name mangling.**
- **The primary objective of XPLINK is to make subroutine calls as fast and efficient as possible by removing all nonessential instructions from the main path.**

## XPLINK Enhancements Overview

### ■ XPLINK Improvements

- Improve execution environment between XPLINK and non-XPLINK programs,
  - LE now allows a **function pointer to be passed from a non-XPLINK compiled program to an XPLINK compiled program without restrictions**
- Improved XPLINK transition tracing
  - **XPLINK greatly increases the performance** of C/C++ applications
  - **Transitions** between **XPLINK** and **non-XPLINK** code **can reduce** its effectiveness.
  - Often **difficult/impossible to determine where transitions occur** in the application
  - With this enhancement a new **trace levels** added to the existing LE TRACE run-time option to trace these transitions

## Application Integration

### New z/OS XML System Services component of z/OS

- Expected to help satisfy requirements for high performance or unique environmental XML non-validating parsing
- Support for cross-memory and SRB modes
- Initial support for HLASM; IBM plans to add C/C++ high-level language in a future release
- Also planned to be made available on z/OS 1.7\*

### New CIM version, designed to include

- An upgrade of the CIM Server Runtime Environment to V2.5.1 of OpenPegasus from the Open Group
- CIM Schema upgrade to 2.9
- Additional resource instrumentation
- Support for Embedded Objects, Events (CIM Indications), HTTP Chunking, and the capability to run CIM providers in a separate address space
- Command line interface for running CIM Client requests against the CIM Server
- Designed to provide security, reliability, and scalability improvements for CIM Server



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

## Application Integration: New LDAP server (1H2007)\*

- Designed to improve performance, scalability, auditability, availability, and ease of use, and planned to include:
  - z/OS UNIX file-based backend for small- and medium-sized directories with memory caching for all directory entries, to help simplify setup and operation of small- and medium-sized directories
    - In addition to existing support for DB2-based backstore
  - ARM and TCP/IP restart support
  - Parallel Sysplex support designed to enhance synchronization of LDAP servers within a sysplex and allow a sysplex group to replicate with other LDAP servers
  - SMF records designed to improve LDAP auditability.



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

## Application Integration

- Unicode collation tables:
  - New needed by middleware applications added
  - Based on the Unicode Collation Algorithm (UCA), which conforms to the Unicode 4.0 standard
- Improvements for Language Environment's Preinitialized Environments for Authorized Programs:
  - Support for changing the number of system-managed environments dynamically
  - Improved failure data collection capabilities and more recovery options
- inittab support in z/OS UNIX for Java recovery
  - Provide de facto standard recovery methods
- XPLINK transition tracing
  - Can help optimize application performance
- Additional locale Euro support for EU countries
  - Czech Republic, Estonia, Hungary, Latvia, Lithuania, Poland, Slovakia, Slovenia



## Application Integration

- RMF eServer OS monitoring Stage 2 for z/OS
  - A WBEM profile based on the IETF SLP protocol, intended to make self-discovery of the CIM-based monitoring service easier
  - Support for CIM indications, designed to enable exploiters using standard CIM client applications to subscribe to asynchronous events
  - Expose some additional existing RMF data to CIM
- BPXBATCH improvements
  - STDIN, STDOUT support for sequential & SYSOUT data sets, PDS(E) members
  - Also available now on z/OS R5 and up with APAR OA11699
- CLIST variable pool and REXX OUTTRAP constraint relief for authorized programs
  - Move data area above the 16MB line
  - Larger storage area available for authorized programs

## Application Integration

- XL C/C++ improvements
  - Some C99 features added for C++ to help improve portability
  - New compiler options, including some for storage use and code optimization
  - Improved interoperability with CICS, Inter-Procedure Analysis (IPA) optimization enhancements, ...
- Support for flockfile() family in Language Environment
  - Implemented in a way intended to be compliant with SUSV3 standard for these functions
- ISPF EDIT informational messages made available to EDIT macros



## Perl - Practical Extraction and Report Language

### *Perl for z/OS*

- Perl V5.8.7
  - All core functions and libraries supported
  - Pre-configured/compiled
  - SMP/E installed
  - All known defects addressed
  - Runs on z/OS 1.4 and higher
- Feature of IBM Ported Tools for z/OS V1R1M1
  - **Program number: 5655-M23**
  - **COMPID: 5655M2303**
  - **FMID: HPRL110**
- Facilitate migration of web-based applications to z/OS
- Provide a popular tool for new applications

# z/OS 1.8

## Helping meet challenges of on demand business

### Improving Usability and Skills

**More Health Checks and framework improvements, HCM, ISPF, InfoPrint Central, & RRS improvements, Extensions to the IBM Configuration Assistant for z/OS Comm Server, EE usability improvements, ...**

*Integrating new Applications and Supporting Industry and Open Standards*

*Unicode improvements; inittab support; LE improvements, new LDAP server, z/OS XML System Services, BPXBATCH, CIM, XPLINK, Networking APIs, XL C/C++ enhancements, ...*

*Extending the Network*

*Sysplex partitioning support, Dynamic DNS registration, JES3 NJE via TCP/IP, TN3270 and telnet improvements, IP filtering, IKE, IPSec, WTS, Samba, improved sysplex failure recovery, REXX FTP interface, ...*

*Scalability & Performance*

*Support for up to 4 TB memory, New zIIP processor, CFRM performance improvements, GRS ENQ limits, Parallel VARY ONLINE, z/OS UNIX file descriptors, Device Groups, DADSM/CVAF Rapid Index Rebuild, z/OS USS asynch socket fast path, 64-bit PDSE support ...*

*Improving Availability*

*Master console elimination, Logstream rename, test and prod logstream separation, Recoverable BRLM, GDPS enhanced recovery, Fast replication, HyperSwap trigger, z/OS UNIX latch contention detection, ...*

*Optimization and Management*

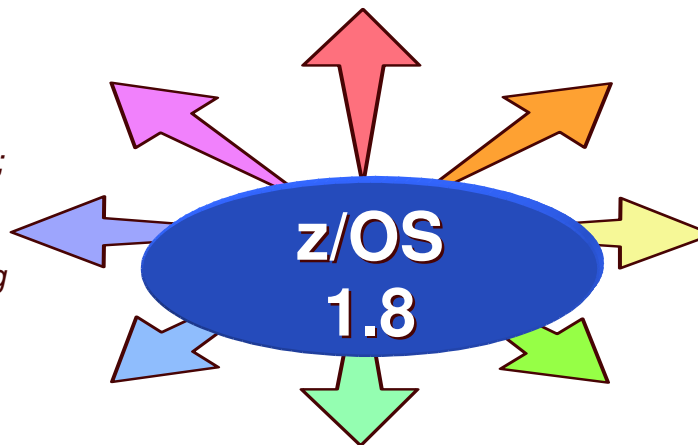
*WLM, zAAP Stage 2, Sysplex Distributor, Tape I/O priority, JES2 balancing of WLM-managed initiators, GRS CNS processing, Improved OAM BLOB support, SCLM improvements, Group Capacity Limit ...*

*Enhancing Security*

*PKI extensions, Distributed Identity Support, RACF password phrase support, Tape data set protection, 128-bit AES for IPSec, SAF identity tokens, RACF virtual key rings, IDS configuration, ...*

*Enterprise-Wide Roles*

*WLM/EWLM Service Class Correlation, zAAP Reporting for eWLM, DFSMSrmm CIM agent, RMFeServer OS Monitoring Stage II, WLM delay services mapped to ARM ...*



## z/OS 1.8 Simplification *IBM Health Checker for z/OS*

### Value

- Configuring for best practices
  - Helping to avoid outages
- Checks against active settings
- Notifies when exceptions found
- Runs on all supported releases of z/OS



### Latest Enhancements

- In z/OS 1.8, the Health Checker Framework is enhanced and will provide following:
  - Enabled for NLS
  - PARMLIB Check Definition
    - Define/add a check via a PARMLIB definition instead of needing a program that does the definition
  - Enhanced Policy Support
    - Support multiple policies
  - Verbose command output
    - Check output be as verbose as possible
  - Parameter Parser
    - Provide a service usable by check routines to parse parameters

<sup>1</sup> Majority of these checks are available on prior releases

## z/OS 1.8 Simplification

### **IBM Health Checker for z/OS – New Checks**

- **Communications Server checks to alert you to non-recommended TCP/IP configuration values**
- **TCPIP Checks**
  - `CSTCP_TCPMAXRCVBUFRSIZE_tcpipstackname`
  - `CSTCP_SYSTCPIP_tcpipstackname`
- **VTAM Check**
  - `CSVAM_CSM_STG_LIMIT`
- **RACF Checks**
- **RRS check to highlight when a configuration varies from the IBM recommendation**
- **ASM**
  - Check for unused PARTE entries
  - Check for COMMON page data set
  - Check the current TOTAL AUX SLOT usage
- **LOGGER (OA15593) – z/OS 1.7 and z/OS 1.8**
  - `CHECK (IBMIXGLOGR, IXGLOGR_STRUCTUREFULL)`
  - `CHECK (IBMIXGLOGR, IXGLOGR_STAGINGDSFULL)`
  - `CHECK (IBMIXGLOGR, IXGLOGR_ENTRYTHRESHOLD)`

## z/OS 1.8 Simplification

### *IBM Health Checker for z/OS – GRS Check*

- GRS will provide GRS\_GRSQ\_SETTING, which will examine the current GRSQ setting.
  - IBM recommends having a GRSQ setting of CONTENTION
    - Which may significantly reduce dump time
- IBM Development has provided two sample checks:
  - **Download website**
    - [ibm.com/server/eserver/zseries/zos/integst/](http://ibm.com/server/eserver/zseries/zos/integst/)
  - Select IBM health Checker for z/OS: Sample Checks

## z/OS 1.8 Simplification

### *Resource Recovery Services Orderly Shutdown*


- RRS currently must be cancelled to come down. Cancel processing results in one or more RRS abends during the shutdown process
  - Allow “clean” shutdown of RRS
  - Allows an installation to shutdown RRS without experiencing abends
  - New shutdown command - **SETRRS SHUTDOWN**
  - **ATR104I SHUTDOWN REQUEST WAS RECEIVED FOR RRS**
  - **ATR105I RRS SHUTDOWN REJECTED, RRS SHUTDOWN IS ALREADY IN PROGRESS**
  - **ATR106I AN UNEXPECTED ERROR OCCURRED DURING RRS SHUTDOWN PROCESSING. RRS CANCEL COMMAND IS ISSUED**
- Install toleration APAR OA15144 on all 1.4 to 1.7 systems

## z/OS 1.8 Ease of Use

- IBM Configuration Assistant for z/OS Communications Server GUI improvements:
  - Support for QoS and IDS policy configuration
  - Designed to allow you to configure IPsec, Application Transparent TLS, QoS, and IDS policy using a consistent user interface
  
- Support for changing file format (extattrs) from the z/OS UNIX shell

## z/OS 1.8 Ease of Use

- ISPF improvements
  - Display and process z/OS UNIX files and directories
  - Data Set List (OPT3.4) support for SuperC and Search-for
  - Better member list searching
  - Autodiscovery of IP address for ISPF workstation GUI
  - Improved alias processing for copy, move, delete
  - Tape data set support in DSLIST
- SDSF display of zAAP utilization
- RMF postprocessor & Monitor III improvements
- HCM enhancements:
  - Literally, too many to list!
  - Some are: IODF import/export; saving layouts for modeling; Copy, Add Like, and Repeat support; and IODF compare





## New z/OS Management Console

### IBM OMEGAMON z/OS Management Console V1.1

#### Capabilities

- Graphical, web-based console for z/OS
- Status of z/OS sysplexes and systems displayed using Tivoli® Enterprise Portal, plus:
  - Event notification
  - Drill-down to problem details and expert advice
- Integrated z/OS Health Checker reports
- Easy upgrade to comprehensive Tivoli Monitoring Services products
- Foundation for the future




- Available at no charge to z/OS customers
- Supported on z/OS 1.4 and above

[www.ibm.com/zseries/zos/zmc](http://www.ibm.com/zseries/zos/zmc)

**Hundreds of orders  
and downloads!**

## Continued investment in skills for IBM Mainframe community



### ■ Enhancing mainframe skills today

- Training for customer "in-house" skills
- Global road shows and workshops assisting in customer system designs
- Significant investment in Technical Skills Vitality initiative and certifications

### ■ Advancing toward goal of 20,000 skilled mainframe professionals in marketplace by 2010

- Explosive growth in IBM Academic Initiative for System z; over 200 participating institutions
- Educating university and college students worldwide


### ■ 2005 IBM Student Mainframe Contest

- Over 700 students from 77 schools!
- <http://www.developer.ibm.com/university/students/contests/mainframecontest.html>
- Student comments:
  - I have submitted Parts 1 and 2 and am eager to continue. Lets do this :D -Chris Peterman
  - I don't know if I'll make it any further- schoolwork is eating up the time I should be investing in mainframe...but I'm enjoying it more than I can admit in public. -Beth Boose
  - I'm one of the "no experience required" guys the contest was advertised for and the whole thing is really interesting to me. I really had a blast with parts 1 and 2 and am looking forward to logging in to part 3. - Brian Capps
  - I enjoyed this contest very much, makes me appreciate my elders...kidding. Thanks for the great contest, - Walter Weiss

### ■ Draft of z/OS Basics available:

<http://publibz.boulder.ibm.com/zoslib/pdf/zosbasic.pdf>

<http://publibz.boulder.ibm.com/zoslib/pdf/zosbasic.pdf>



# Gateway to the mainframe community ecosystem

**High Availability Center**

[http://www.ibm.com/servers/eserver/design\\_center/highavailabilitycenter/](http://www.ibm.com/servers/eserver/design_center/highavailabilitycenter/)

**Design Centers**

[http://www.ibm.com/servers/eserver/design\\_center](http://www.ibm.com/servers/eserver/design_center)

**Executive Briefing Centers**

<http://www.ibm.com/servers/eserver/briefingcenter>

**Mainframe Magazine**

[http://www.ibm.com/servers/eserver/zseries/eserver\\_magazine.html](http://www.ibm.com/servers/eserver/zseries/eserver_magazine.html)

**Benchmark Centers**

[http://www.ibm.com/servers/eserver/customer\\_benchmarks/](http://www.ibm.com/servers/eserver/customer_benchmarks/)

**iSource – customized weekly e-news**

<https://isource.ibm.com/world/>

**The Mainstream**

**System z9 and zSeries software newsletter**

<http://www.ibm.com/software/os/zseries/mainstream>

**Mainframe Favorites**

**Utilize the Mainframe Community Web Portal**

[http://www.ibm.com/servers/eserver/zseries/about/charter\\_community.html](http://www.ibm.com/servers/eserver/zseries/about/charter_community.html)

**Mainframe News**

**Skills Portal (Enhance customer skills)**

<http://www.ibm.com/servers/eserver/zseries/about/charter/skills.html>

**System z Technical Support**

<http://www.ibm.com/servers/eserver/support/zseries>

**Redbook™ Customer Residency**

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

**Consultants/System Integrators**

<http://www.ibm.com/easyaccess/integration>

**System z Innovation Center**

<http://www.developer.ibm.com/vic/hardware/>

**IBM Academic Initiative for Mainframe**

<http://www.ibm.com/university/zseries>

**Mainframe Co-op Program**

[email:inivprg@us.ibm.com](mailto:inivprg@us.ibm.com)

## “New Face of z/OS” Messages

- IBM is investing in mainframe simplification and modernization:
  - System Management: We’re making z/OS systems easy to set up, operate, and administer
  - Application Development: We’re delivering a modern, easy to use development environment.
- IBM is partnering with clients and schools to build the pool of mainframe-savvy IT professionals:
  - IBM Academic Initiative for System z
  - z/OS Basic Skills Information Center



## z/OS System Management Strategy

*Goal: Make z/OS servers easy to deploy, administer, and service.*

<b><i>Today</i></b>	<b><i>Tomorrow</i></b>
Manual, time-consuming tasks	Highly automated, efficient tasks
Multiple UIs and manuals often required to perform a single task	A modern web UI for performing tasks end-to-end, with integrated task guidance;
Years of z/OS experience required	Designed with novices in mind; but experts will benefit

### Over the next five years, IBM plans to significantly simplify:

- **System health monitoring with event analysis and problem management**
- **System installation and configuration (for instance, network and hardware configuration)**
- **Workload management**
- **Storage and data management**

## z/OS 1.8 Simplification Highlights

- IBM Health Checker for z/OS enhancements – more checks, easier way to write health checks
- IBM OMEGAMON z/OS Management Console 4.1.0 (available 12/06) – more monitoring workspaces
- Major Hardware Management Console enhancements
- IBM Configuration Assistant for Communications Server – support for QoS and IDS configuration has been added
- Infoprint Central enhancements
- ShopzSeries Version 10.2 – an easy and fast way to order System z software over the Internet
- Application development simplification through products such as WebSphere Developer for zSeries and z/OS support for open standards such as Common Information Model (CIM) and XML
- Plus enhancements to ISPF, SDSF, RMF, and RRS



See the z/OS 1.8 [announcement letter](#) for details.

## IBM OMEGAMON z/OS Management Console

- Available at no charge for z/OS 1.4 and above; easy upgrade to comprehensive Tivoli Monitoring products
- For monitoring the health of z/OS sysplexes and systems; replaces complex DISPLAY commands; integrated Health Checker reports
- Powerful features:
  - Event notification
  - Drill-down to problem details and expert advice
- Version 4.1 adds new workspaces for monitoring sysplexes, LPAR clusters, and UNIX processes



Version 4.1  
available  
12/06

Learn more at: <http://www.ibm.com/servers/eserver/zseries/zos/zmc/>

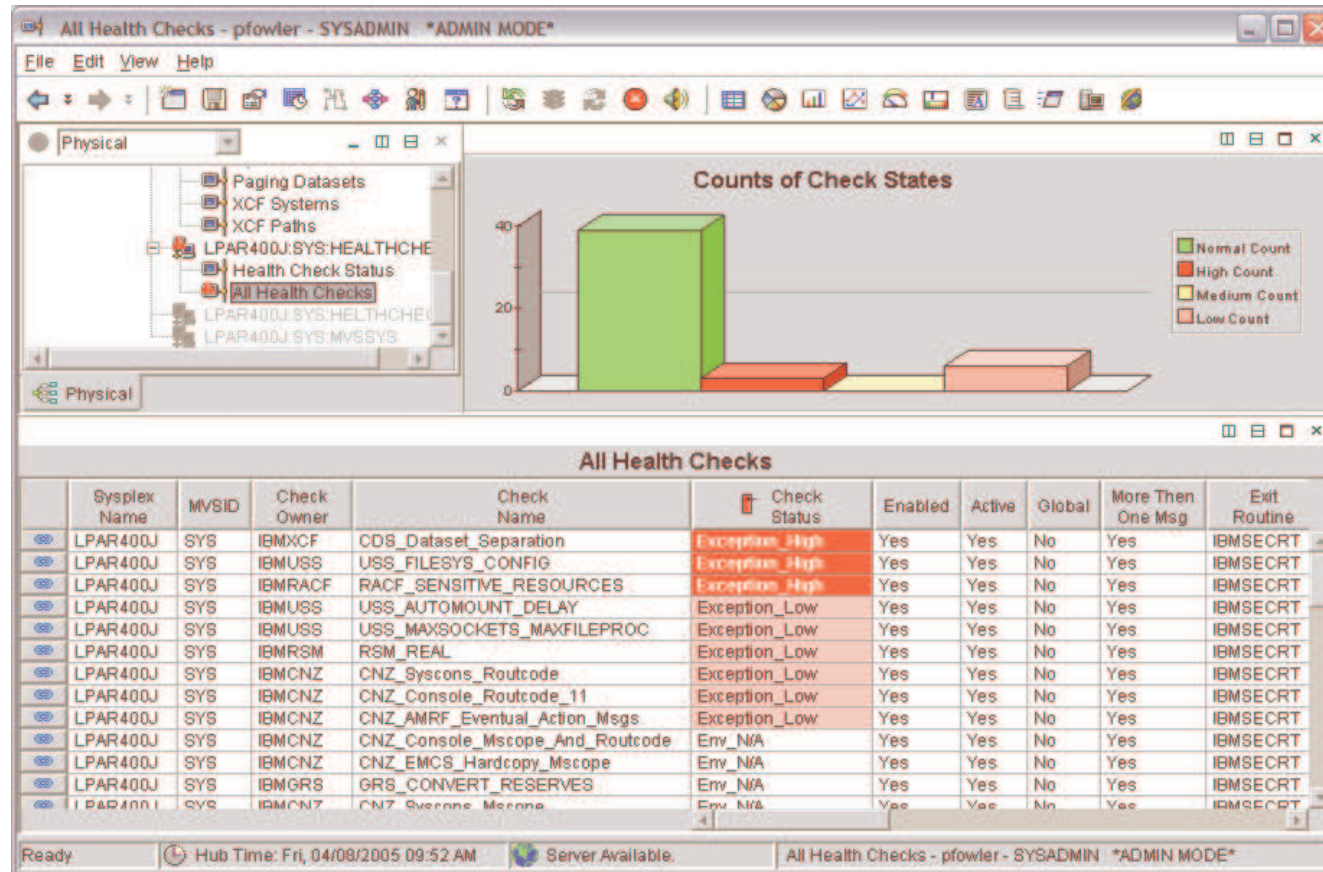


## z/OS Health Checks in zMC 1.1.1/4.1.0

The Health Checker examines configuration values for deviations from best practices.

Health check exceptions are highlighted in the zMC.

IT staff can open a report for details and recommended corrective actions.



**Health Checker integration is a unique feature of the IBM OMEGAMON z/OS Management Console (zMC)**

LPAR Clusters Data - RGATSKI - SYSADMIN

File Edit View Help

View: Physical

Logical Processor Counts

Physical & Special CPs

LPAR Clusters Information

LPAR Clusters Information

Cluster Name	LPAR Name	LCPs Online	LCPs Offline	LCPs Reserved	LCPs Standby	CP Status	LPAR Status	CPC Model#	SA	Timestamp
LPAR400J,0960.2064	CANSP23	2	0	1	0	SHR	Active	2064-109		08/29/06 11:2
LPAR400J,0960.2064	CANSYSG	3	0	1	0	SHR	Active	2064-109		08/29/06 11:2
N/A	RALNSCF9	1	0	0	0	DED	Active	2064-109		08/29/06 11:2
N/A	RALNS0	2	0	0	0	SHR	Active	2064-109		08/29/06 11:2
LPAR400J,0960.2064	CANSP11	2	0	1	0	SHR	Active	2064-109		08/29/06 11:2
N/A	TIVVM2	5	0	2	0	SHR	Active	2064-109		08/29/06 11:2
LPAR400J,0960.2064	CANSP13	1	0	1	0	SHR	Active	2064-109		08/29/06 11:2
LPAR400J,0960.2064	CANSP22	3	0	1	0	SHR	Active	2064-109	050960	08/29/06 11:2
N/A	RALHCD	1	0	0	0	SHR	Active	2064-109	050960	08/29/06 11:2
N/A	PHYSICAL	9	0	0	0	SHR	NotActive	2064-109	050960	08/29/06 11:2

Hub Time: mar., 08/29/2006 11:31 AM Server Available LPAR Clusters Data - RGATSKI - SYSADMIN

LPAR Clusters View in zMC 4.1


## IBM Configuration Assistant for z/OS Communication Server

- Accelerates the implementation new network features
- Version 1.8 adds support for setting up Quality of Service (QoS) and Intrusion Detection Services (IDS) policies
- Helps customers make the right configuration decisions for their environment
- Ready to take it for a spin?  
[Click here](#) to download it!



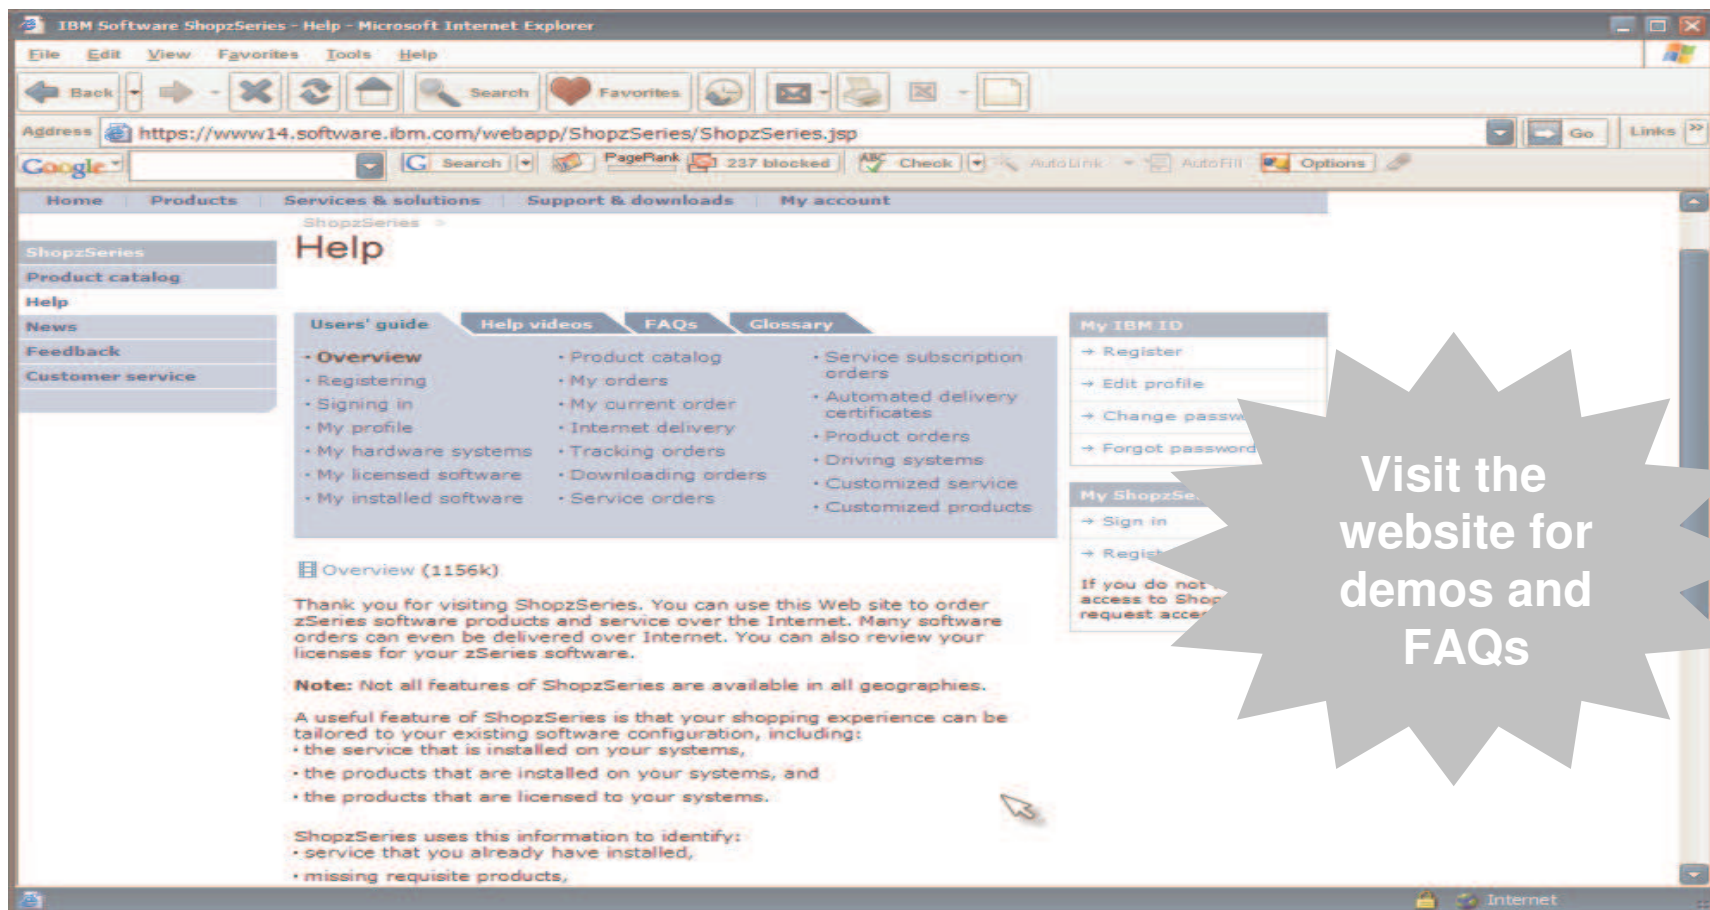
## Hardware Configuration Manager

- Lets you easily navigate through configuration diagrams and make changes
- 1.8 adds many new ease of use enhancements including:
  - **Wizards that simplify complex configuration tasks**
  - **Improved configuration reports**
  - **Access to RMF Monitor III reports which speeds up the detection and resolution of performance bottlenecks**
  - **Ability to import and export I/O Definition Files (IODFs) – another productivity saving**



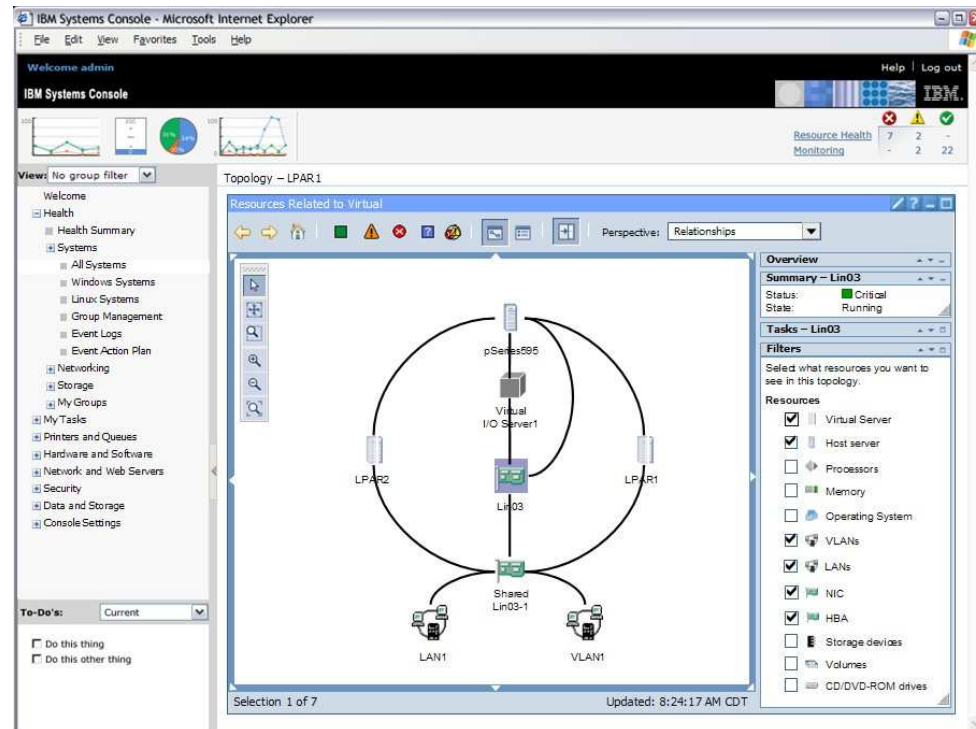
ShopzSeries - <http://www.ibm.com/software/shopzseries>

- Easy and fast way to plan and order upgrades and service over the Web.
- Supports all products offered in CBPDO and ServerPac



## z/OS Management Console Future Directions\*

- IBM is moving to a common web console solution for administering servers – including z/OS, z/VM, and Linux on System z.
- This console will be an integration point for all management tasks.
- Tasks such as health monitoring and software updates will have a common look and feel across server platforms.
- The underlying system management infrastructure will be open so that vendors can also exploit it. (This is the role of the z/OS Common Information Model (CIM) element in z/OS 1.8).

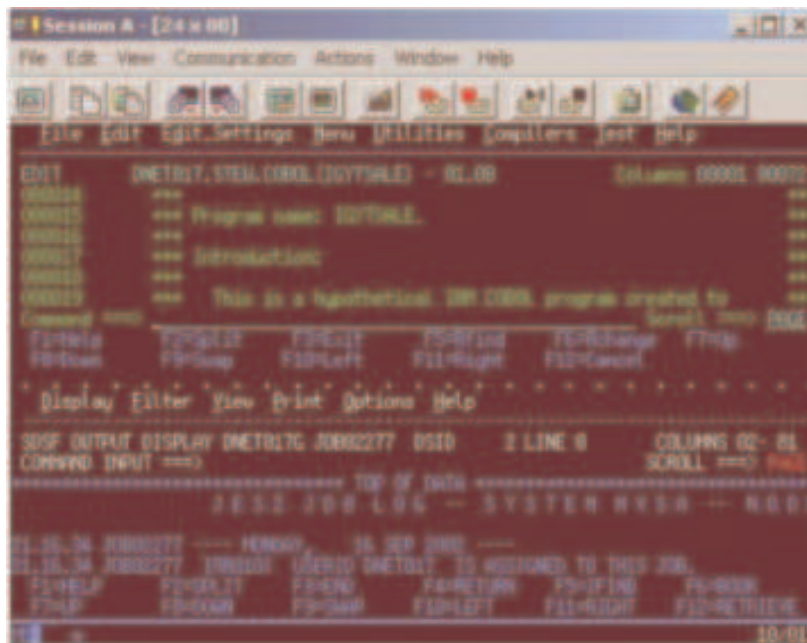


Future web console

\* IBM development plans are subject to change or withdrawal without further notice.

## z/OS Application Programmer's Experience Today

Scenario: Correcting an error in a COBOL source file



The screenshot shows the ISPF editor window titled 'Session A - [24 x 80]'. The menu bar includes 'File Edit View Communication Actions Window Help'. The main display area shows a COBOL source file with the following content:

```
EDIT DNETB17C.STEEL (DBL (DVTYSALE) - 01.06) Column 0001 0072
000100 ***
000101 *** Program name: DVTYSALE.
000102 ***
000103 *** Introduction:
000104 ***
000105 *** This is a hypothetical DBL DBR program created to
000106 ***
Command ***
F1=Help F2=List F3=Exit F4=Print F5=Refresh F6=Top
F7=Down F8=Swap F9=Left F10=Right F11=Cancel
Display Filter View Print Options Help
SDSF OUTPUT DISPLAY DNETB17C JOB02277 0510 2 LINE 8 COLUMN 02- 84
COMMAND INPUT ***
***** TOP OF DATA *****
JOB02277 LOG - SYSTEM MSGS - 0001
01.05.04 JOB02277 --- HENNY 16 SEP 040 ---
01.05.04 JOB02277 1000000 USERID DNETB17C ASSIGNED TO THIS JOB.
F1=HELP F2=LIST F3=END F4=RETURN F5=PAGE F6=JOB
F7=UP F8=DOWN F9=SWP F10=LEFT F11=RIGHT F12=RETRIEVE
10-015
```

1. Start 3270 emulator
2. Logon to the z/OS system
3. Navigate to the dataset and member using ISPF
4. Select the member for editing
5. Locate the line in the source code
6. Change the source code and save the member
7. Switch to the JCL that is used to submit and schedule the compile job
8. Submit the JCL job
9. Switch to SDSF to monitor the job and review the output
10. Repeat 5-9 until program runs correctly

## Strategic Environment for Application Developers

- **WebSphere Developer for zSeries (WDz) provides a modern environment for developing:**
  - Dynamic Web applications including Java™ and Java 2 Enterprise Edition
  - Traditional COBOL and PL/I applications
  - High level Enterprise Generation Language (EGL) applications
  - Web services to integrate these applications together
  
- **WDz offers advantages over traditional environment:**
  - Single environment for programming tasks: design, code, debug from the same UI
  - Productivity features like syntax-checking, JCL generation, visual debugging
  - Integrated user assistance



[Click for Details](#)



The screenshot shows the IBM Rational Software Development Platform interface. The main window displays a COBOL program being edited. The program text is as follows:

```

Line 93      Column 1      Insert
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
000093      A000-MAIN-DRIVER.
000094      *
000095      DISPLAY "***** START OF CALLIVP1 *****".
000096
000097      PERFORM B100-CALL.
000098
000099      IF TEST-STATUS = "Y"
000100      DISPLAY "***** CALLIVP1 SUCCESSFUL *****"
000101      ELSE
000102      DISPLAY "+++++ CALLIVP1 FAILED +++++".
000103      STOP RUN.
000104
000105      *
000106      B100-CALL.
  
```

The interface includes a 'Remote Systems' tree on the left showing the connection to 'DEMOMVS' and various file types (JES, USS Files, MVS Files). A 'Properties' window at the bottom left shows 'Remote Scratchpad' with fields for 'Line count', 'Title length', and 'Title start'. The status bar at the bottom indicates 'z/OS LPEX Editor is active'.

### WDz's z/OS Perspective

- System name is DEMOMVS
- You can work with z/OS datasets, UNIX files, and JES jobs from the UI
- This example shows a COBOL program being edited

## z/OS Basic Skills Information Center

- Information for “on-the-job” training
- Educational topics plus interactive exercises

The screenshot shows the IBM z/OS Basic Skills Information Center website. At the top, there is a navigation bar with the IBM logo, a search bar, and links for 'United States [change]' and 'Terms of use'. Below this is a secondary navigation bar with links for 'Home', 'Products', 'Services & solutions', 'Support & downloads', and 'My account'. The main content area is titled 'z/OS basic skills information center' and includes a list of topics with brief descriptions and links to HTML or PDF versions. A right-hand sidebar contains sections for 'Related links', 'IBM Academic Initiative', 'Podcast', and 'z/OS poster'.

# Sample Topic in z/OS Information Center

The screenshot shows the IBM z/OS Information Center interface. The top navigation bar includes links for Home, Products, Services & solutions, Support & downloads, and My account. A search bar is present with a 'GO' button and a 'Search scope' dropdown set to 'All topics'. The left-hand 'Contents' pane shows a tree view with 'z/OS concepts' expanded to 'Basic shared DASD'. The main content area is titled 'z/OS concepts' and contains the following text:

A basic shared DASD environment is illustrated in [Basic shared DASD](#). The figure shows z/OS images, but these could be any earlier version of the operating system. This could be two LPARs in the same system or two separate systems; there is absolutely no difference in the concept or operation.

The diagram below illustrates this concept. It shows two 'zSeries (or LPAR)' boxes, each containing a 'z/OS' image with 'channels'. These channels are connected to two 'control unit' boxes. Each control unit is further connected to a set of DASD storage devices (represented by blue cylinders). A note on the right states: 'Real system would have many more control units and devices'.

## z/OS Basic Skills Info Center: Sample Course

**IBM. z/OS Basic Skills - The mainframe's processors**

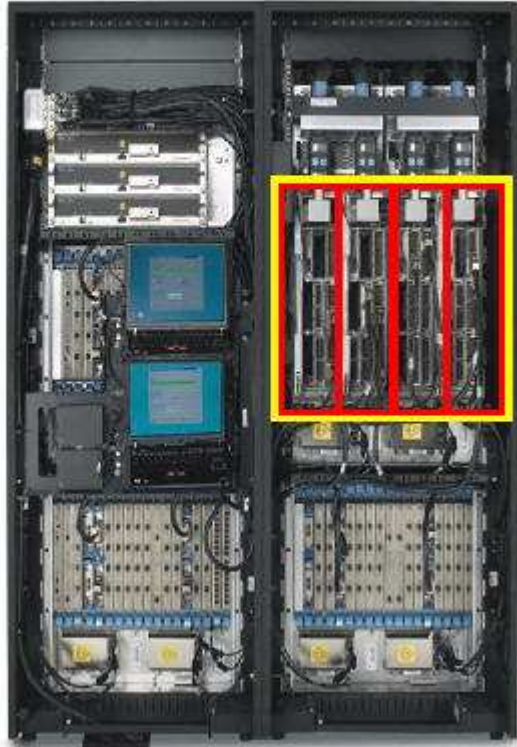
*Main menu > CPC cage*

**The central processor complex**

Mainframes have one or two metal frames that contain specialized cages, as well as other physical elements. This diagram shows the interior front view of an IBM System z9 Enterprise Class (EC) model that has two frames. The z9 is slightly larger than a household refrigerator.

The central processor complex, or CPC, resides in its own cage inside the mainframe, and consists of one to four book packages. Just like its personal-computer counterpart, the motherboard or system board, each book consists of processors, memory, timers, and I/O connections.

These collections of hardware parts are called "book packages" because you can slide them in or out of the CPC cage almost as easily as you can slide a book on or off a bookshelf.



## IBM Academic Initiative for System z


Develop and manage a worldwide program to build skills for the mainframe community.

Assist customers in locating students educated on the mainframe and with establishing relationships at schools.

Market and assist faculty worldwide to include Large Systems Thinking into curricula.

Create innovative campaigns and initiatives to generate interest and excitement around the mainframe.

Encourage and assist university ambassadors and anyone in the mainframe community who wants to help.



**For more information, see:**

[ibm.com/ict09002c/university/scholars/products/zseries/](http://ibm.com/ict09002c/university/scholars/products/zseries/)

## Major Initiative: System z Mastery Exam in Fall 2006

- **Measure and validate mainframe skills and knowledge**

- *Introduction to the Mainframe: z/OS Basics*

- **Panel of subject matter experts**

- IBM (authors and instructors)
  - College and university faculty (mainframe educators)
  - Mainframe customers

- **Worldwide proctored exam**  
(Thomson Prometric, Pearson VUE, Prime)



THOMSON  
PROMETRIC

PEARSON  
VUE

- **Results recorded in IBM certification database**

- **Goal: Increase the chances for success within an organization's mainframe community**


- **Academic Initiative**

- Qualifies students to submit their resumes to the Student Opportunity System
  - Database (Accessible by customers)
  - Leverage the value and importance of System z courses in academia

## Student mainframe contest



- Completed first contest in North America, Fall 2005
  - Over 700 students enrolled from 85 schools
  - First ever “remote hands-on” contest by IBM
    - Each student got an ID and access to the mainframe
  - Levels of challenges with prizes awarded at every level
    1. (Hours) Breaking the Ice
    2. (Days) Practical Experience
    3. (Months) Real World Challenges
  - Top 5 students received awards and a trip to the Poughkeepsie Lab in March
- 2006 Fall Semester US/Canada, Europe, 2007 – Brazil and Europe



Anna Kuriger  
St. Ambrose University

“z/OS has blown me away in terms of polish and usability compared to MVS.”  
– student, Michigan State

“I’m enjoying it more than I can admit in public.”  
– student, Rutgers University

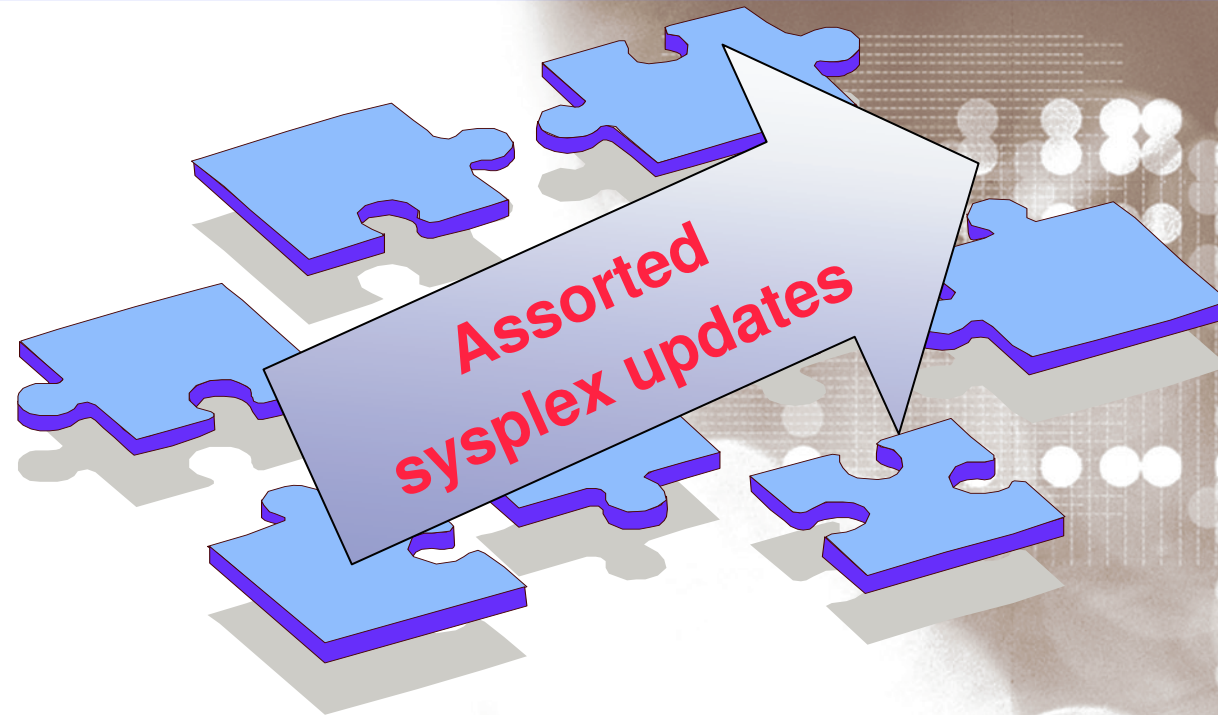
## Academic Initiative Status

Skills Metrics	YE 2003	YE 2004	YE 2005	2006 (4Q) Actual to date
Students educated or being educated by the AI	--	n/a	n/a	15,750
Schools registered in AI -900% growth in 2 years - Over half of schools are outside the US	24	70	213	263
AI Initiated Customer/School partnerships	1	n/a	17	43
Complete AI Courses	3	6	12	13 (3 more planned in 2006)
AI University Hubs	1	4	5	6
AI Contests / Schools / Students	0	0	1/85/ 725	2 contests in progress
AI IBM Ambassadors	0	0	195	234





## LSU Danmark



Danmark LSU October 2006

*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation

## Recent Sysplex updates at a glance

### *z/OS 1.7 time frame (last year)*

- Faster and more efficient HW
  - Faster z9 senders, faster z9 CFs, faster Links (Internal)
  - Statement of intent on Sysplex Timer Links using existing CF links (STP)
- Scalability, performance, tips and techniques
- CF Levels 12, 13 and 14 revisited
- IFLs, zAAPs, ICF new processor pooling options
- Bit bucket
  - Parallel Sysplex aggregation update
  - XRC+

## Coupling Facility Control Code – Level 15?

- Yes, we are working on new CFCC function but too early to discuss specifics

## CF Level summary

*CF Level 14*■ **CF Level 14**

- ✓ z990/z890/z9 - not available on z800/z900
- ✓ Implements a re-write of the CFCC Dispatcher
  - The enhanced dispatcher deliver the most benefits for environments that are XCF-signaling intensive and have a large number of CPs in the CF
  - CFCC did not process requests in FIFO order. Gave preference to starting new work rather than finishing existing suspended work. Especially in a System Managed Duplexed environment, starting out of FIFO order can impact performance. Now it is treated as true ready queue
  - **CFCC Level 13 and lower reports time when it WAS NOT REALLY BUSY as busy time - known as “false utilization”**
- ✓ Requires z/OS APAR OA08742 - Our testing showed an insignificant increase in structure sizes with a CF LEVEL 14

## CF Level Summary

### *CF Level 12 and CF Level 13*



#### ■ **CF Level 12**

- ✓ Number of concurrent tasks increased from 16 to 48
- ✓ Support for Message Time Ordering
- ✓ 64 bit support
- ✓ Support for System Managed CF Duplexing
- ✓ CF Level 12 provides batch CF GBP request enhancements for DB2 V8


#### ■ **CF Level 13**

- ✓ Enhancements to Read Castout Class (RCC) command
- ✓ Ability to install patches and new CF Levels without requiring a POR (z890/z990/z9) Requires z/OS APAR OA01517
- ✓ Minimal increase in size of some cache structures, no impact on list or lock structure

# ETR Network Limitations – Today

## *Distance, time gap and timer limitations*

- Distance between 9037 Sysplex Timers **cannot exceed 40 km**
  - Requires intermediate site for second timer if data centers more than 40 km apart
- “**Best case**” messaging times over ICB links (8 us) approaching “**Worst case**” TOD synchron. between Processors stepping to 9037s 40 km apart (4 us)
- SOD - IBM intends to **withdraw 9037 Model 2** during 2006



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

—————> Active ETR link  
 - - - - -> Alternate ETR link

## Server Time Protocol (STP) overview

### *preview*

- Capability for multiple System z servers to maintain time synchronisation with each other:
  - **Does not require the 9037** Sysplex Timer if all servers STP capable
  - **Timing information transmitted over ISC-3 links (Peer mode), ICB-3 and ICB-4 links.** Supports a multi-site timing network of **up to 100 km**
  - **May reduce the cross-site connectivity** required for a multi-site Parallel Sysplex
  - **Can coexist with an External Time Reference** (ETR) network (9037 based)
  - Designed to allow use of dial-out time services to set the time to international time standard (UTC) as well as adjust to UTC
- Planned to be available as a feature on z9-EC, and an RPQ on z990 and z890
- Prerequisites:
  - z9-EC HMC and SE Code load
  - z/OS V1.7 and above

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

## Server Time Protocol (STP) overview

### *description*

- Designed to provide capability for multiple servers to maintain time synchronization with each other and form a Timing Network
- Server-wide facility planned to be implemented in System z9, z990, z890 LIC
  - Single view of “time” to LPAR
  - LPAR can virtualize this view of time to individual partitions (for example z/OS)
  - STP not available on z900, z800 or 9672 Gx servers
- Message-based time synchronization protocol
  - Based on Network Time Protocol (NTP) – an industry standard
  - Time message information transmitted over all defined Coupling Facility Links
    - ISC-3 links (Peer mode), ICB-3 and ICB-4 links
  - NOT standard NTP
- Supports a multi-site Sysplex of at least 100 km
  - Does not preclude going to longer distances in the future
- Allows use of dial-out time services to set time to international time standard (UTC) to within +/- 100 ms as well as adjust to UTC on a periodic basis

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



## Server Time Protocol (STP) overview

### *configurations: mixed and STP-only CTN*

- Mixed Coordinated Timing Network (CTN)
  - Allows co-existence of ETR and STP networks
  - 9037 provides the primary time reference for the CTN
  - At least one server in Mixed CTN capable of both:
    - Time coordination using ETR (attach to 9037)
    - Time coordination using STP
  - 9037 console continues to be used for all timing related functions
  - HMC must be used for CTN ID initialization and modification
- STP-only CTN
  - All servers configured to be part of STP network
  - 9037s not used as a timing source for any server in timing network
  - HMC must be used for
    - 9037 console functions related to timing
    - New STP related functions (CTN initialization; Stratum 1 configuration definition etc.)
    - Dialing out to Time services providing accurate time (External Time Source – ETS)

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

## IBM STP Implementation Assistance Program (IAP)

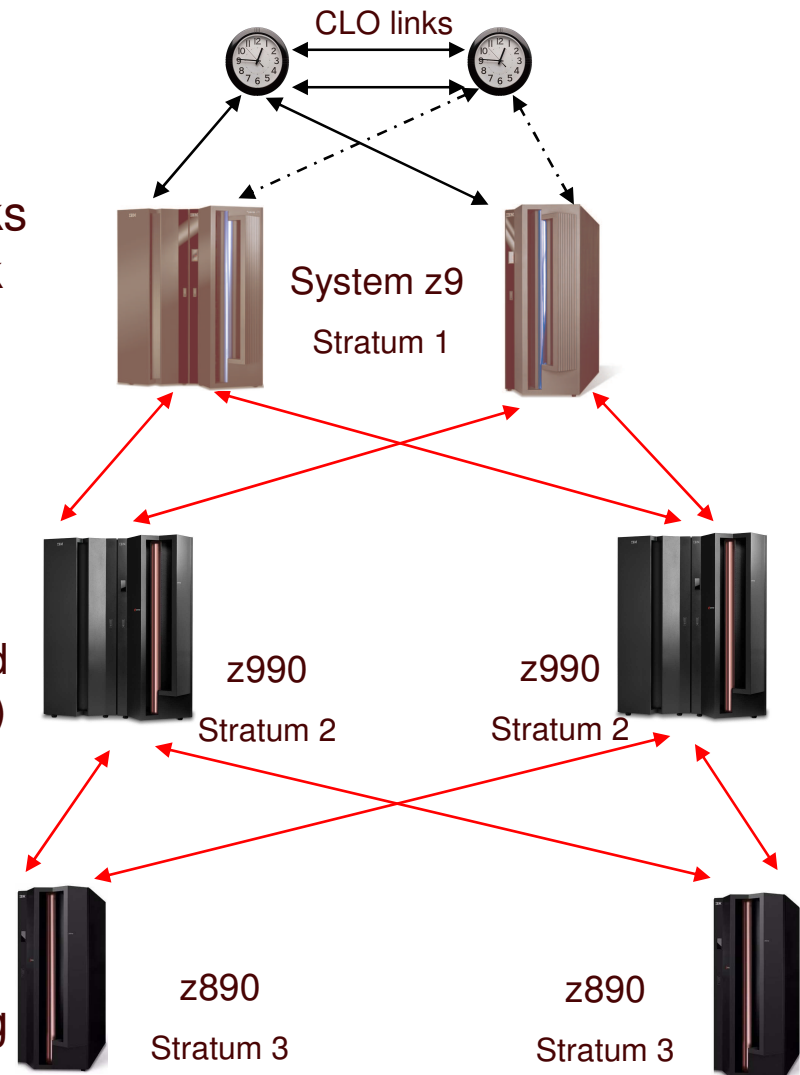
- Objective
  - Accelerate the adoption of STP with IBM assistance
- Assistance planned
  - Consultation
  - Review of migration plans
  - Technical support
- Availability dates
  - October 2005 (z990, z890)
  - 1Q06 (z9-EC)

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

# Server Time Protocol (STP) overview

## Mixed Coordinated Timing Network (CTN)

- Allows co-existence of ETR and STP networks
- 9037s are “root” for both ETR and STP networks
- Need at least 1 server in Mixed Timing Network capable of:
  - Time coordination using ETR (to 9037) and time coordination using STP message-based protocol
- Any STP capable server (z9, z990, z890) attached to 9037 can be a Current Time Server (Stratum 1) for the STP network
  - As long as one server in the CTN is still connected to the ETR, there is no Current Time Server (CTS) configured. Configuring a CTS removes connectivity of all Servers to the ETR
  - Backup time server can be Stratum 2 or 3
    - Stratum 2 uses Stratum 1 as clock source
    - Stratum 3 uses Stratum 2 as clock source
- 9037 console continues to be used for all timing related functions of the Mixed Timing Network
- HMC must be used for CTN configuration



Maximum Stratum level for STP is 3. Stratum 3 not recommended

## Changes to DST in the USA - 2007

*may affect clients with world wide operations*

- **Flash10483**
- **The new USA Energy Policy Act of 2005, signed by the President in August 2005 has changed Daylight Saving Time (DST) in the USA, beginning in 2007**
- **DST Prior to 2007**
  - DST began - 2:00AM first Sunday of April
  - DST ended - 2:00AM last Sunday of October
- **DST in 2007**
  - DST begins - 2:00AM second Sunday of March
  - DST ends - 2:00AM first Sunday of November
- **IBM 9037 Sysplex Timer – users will be responsible to schedule DST offset at the appropriate time at the 9037 console**
- **STP – IBM is planning on making code changes to support new times**  
[ibm.com/support/techdocs](http://ibm.com/support/techdocs)



## zSeries z9 Coupling Technology

### **CBU for ICF engines**

- **All ICFs, zAAPs, and IFLs managed in separate PR/SM “pools”**
- **CFCC Level 14**
  - Dispatcher Restructure
  - DB2 castout performance enhancement
  - Etc.
- **z9....**
- **64 Bit Architecture**
- **Dedicated or Shared CPs**
- **Up to 60 LPs**
- **Up to 512 GB of storage**
- **zSeries Peer Channels (only)**
  - InterSystem Channels-3 (ISC3)
  - Integrated Cluster Bus-3 (ICB3)
  - Integrated Cluster Bus-4 (ICB4)



## z9-109 CF Link Connectivity (Reference)

Connectivity Options	z9-109 ISC-3	z9-109 ICB-3	z9-109 ICB-4
<b>z900/z800 ISC-3</b>	2 GB/s Peer Mode	N/A	N/A
<b>z9-109/z990/z890 ISC-3</b>	2 GB/s Peer Mode	N/A	N/A
<b>z900/z800 ICB-3</b>	N/A	1 GB/s Peer Mode	N/A
<b>z9-109/z990/z890 ICB-3</b>	N/A	1 GB/s Peer Mode Recommendation: Use ICB-4	N/A
<b>z9-109/z990/z890 ICB-4</b>	N/A	N/A	2 GB/s Peer Mode

Compatibility mode (sender/receiver) ISC-3 not supported

## Configuring CF Links (Reference)

Server	IC	ICB-4	ICB-3	ICB	ISC-3	Max # Links
z800	32	-	5 6 (0CF)	-	24	26 + 32
z900-100 CF	32	-	16	16	32 42 w/ RPQ	64
z900	32	-	16	8 16 w/ RPQ	32	64
z890	32	8	16	-	48	64
z990	32	16	16	8	48	64
z9-109	32	16	16	-	48 Peer Mode Only	64

## System z9 and zSeries – Statements of Direction\*

- **Fulfilled SODs (from May, 2003 and April, 2004); support withdrawn**
  - ICB-2
  - ISC-3 G5/G6 ISC-2 compatibility mode
- **z9-109 will be the last system to support ICB-3s Coupling Links**
  - ICB-3 links provide connectivity between z9-109, z900 and z800 servers.
- **Withdrawal from Marketing – 9037 Model 002**
  - IBM intends to withdraw the Sysplex Timer Model 2 (9037-002) from marketing in 2006.
  - A Sysplex Timer or equivalent (Server Time Protocol) is a mandatory HW requirement for a sysplex consisting of more than one server.
  - All strategic planning should include a plan to migrate from a Sysplex Timer **Model 2** network to an STP network, when it becomes available
- **Withdrawal from Marketing – 9036 Model 003 (RPQ 8K1919)**
  - IBM withdrew from marketing the 9036 M3 effective 2Q, 2006. The 9036-003 extends the distance between 2 Sysplex Timers (9037-002s) or extends the distance between a Sysplex Timer and a server by providing a multimode fiber to single mode fiber conversion.

\*All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represents goals and objectives only.



# DB2 Data Sharing Directions

- V7

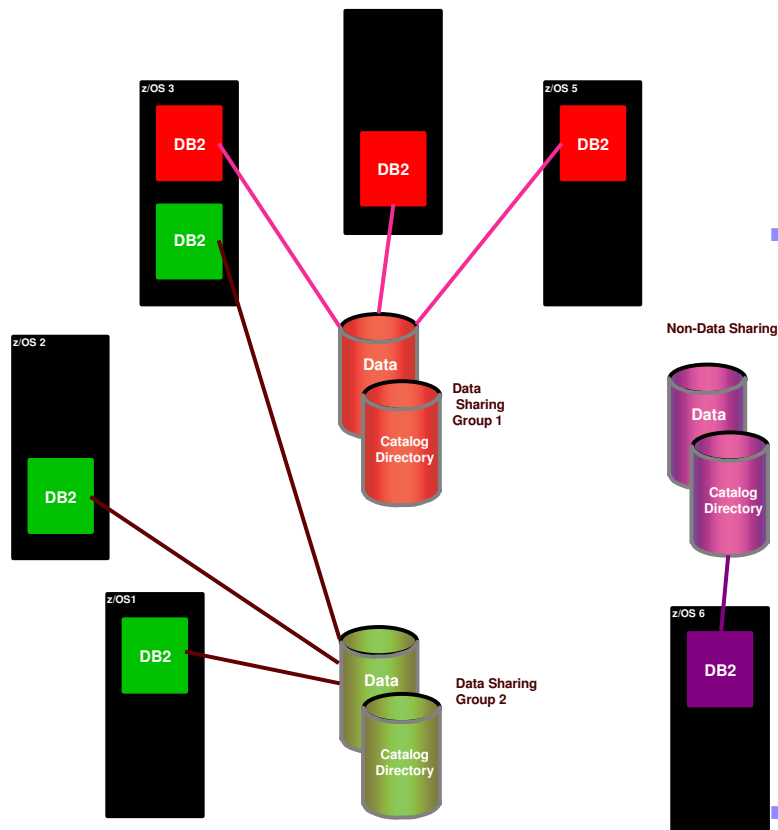
- Restart "Light"
- N, N-2 Migration/Fallback Support
- Structure Rebuild Enhancements
- Group attach enhancements
- Persistent CF structure sizes
- System Managed Duplexing of Lock and SCA Structures

- V8

- Batched GBP writes and cast-outs
  - improved performance especially for batch updates
- Reduced global contention for tablespace L-locks
  - improved performance especially for OLTP
- More efficient index split processing
- Less Disruptive LPL Recovery
- Restart "Light" Recovery Enhancements
- Immediate Write PH1 becomes system default
- Enhanced DISPLAY GBPOOL output
- DPSIs for more efficient member affinity routing
- DDF sub-groups

- **V9 data sharing enhancements**

- Restart Enhancements**
- Performance – indexing & batch enhancements**
- Logging enhancements**



## IMS Sysplex Manager – New!

### Real-time management of the IMS Sysplex Environment

- Single system image for management of IMS in a Sysplex
- Provides local and aggregate view of data
- Simplified User Interface (TSO/ISPF) to reduce operational complexity
- Structured displays of IMS related address spaces, CF structures and resources
- Recorded system exception and console alerts
- IMS type 1 (“/”) local and global command support
- IMS type 2 integrated SPOC support
- IMS Data Sharing support
- IMS Shared Message Queues and CQS support
- IMSplex Operations Manager, Resource Manager and Structure Call Interface support
- Support IMS DB/TM, DBCTL, and DCCTL for IMS v8 and later
- See <http://www-306.ibm.com/software/data/db2imstools/imstools-library.html> for more information

## ICSF Parallel Sysplex Support for Shared CKDS

### ■ Problem

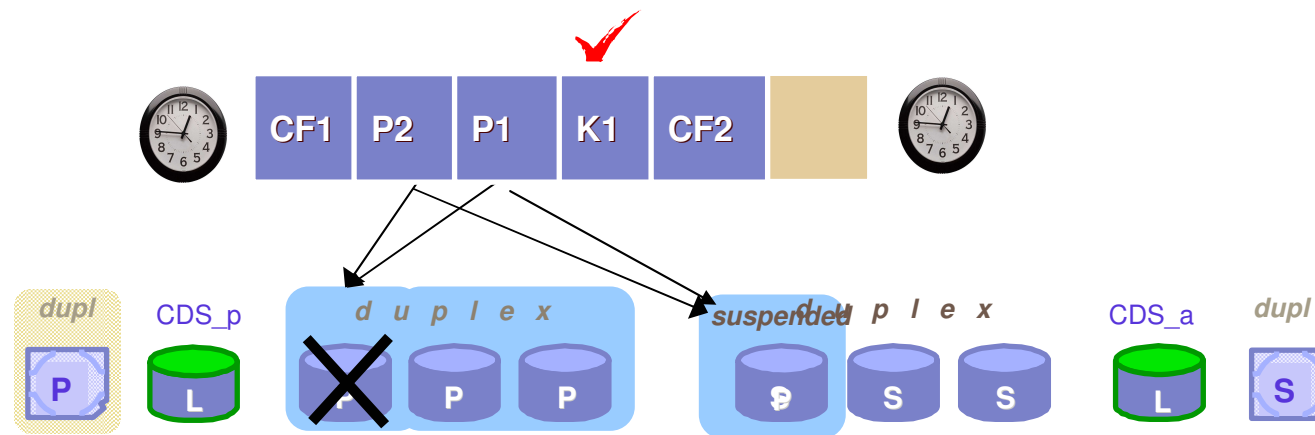
- The CKT (in storage copy of Cryptographic Key Dataset (CKDS)) is only updated on the system that initiated the change.
- When CKDS is shared across a Parallel Sysplex, the change is not propagated to the CKTs on the other systems.
- Change propagated manually or by running a batch job
- However, if not propagated, **units of work routed to the systems that did not see the change may fail**

### ■ Solution

- **Sysplex wide consistent updates to the CKDS (and CKT) using XCF signaling**
- **Sysplex wide updates triggered by new SYSPLEXCKDS option in ICSF installation options dataset**
- Was 05 Web deliverable for installation on z/OS 1.6 and 1.7; rolled into future z/OS release.

## GDPS/PPRC and GDPS/PPRC HM HyperSwap Extensions

- IOS timing trigger
  - Autonomic detection of “Soft Failures” to trigger HyperSwap
  - Based upon customer defined I/O timing thresholds
- Secondary PPRC device failures
- Unplanned HS IEA491E PRIMARY-DEVICE-FAILURE trigger
- Keep current master and secondary PPRC volumes in same site
- Reduce alias bind CPU consumption stage 1 / no GDPS change required
- Detect CC and XRC sessions

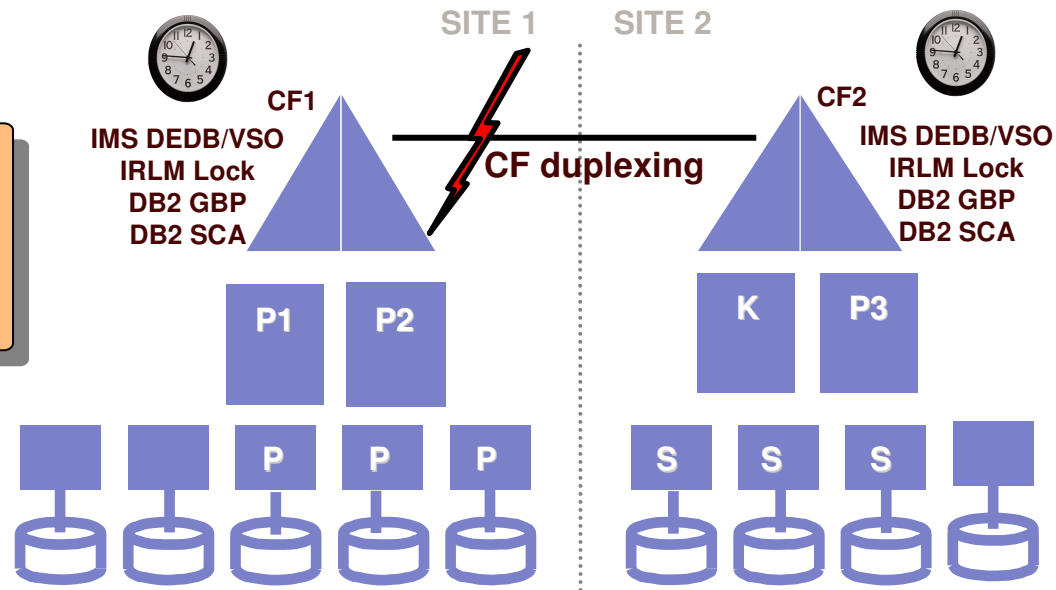


\* All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represents goals and objectives only.

# GDPS/PPRC enhanced recovery support

for GDPS V3.3

Should we keep Site1 or Site2 instance?



- Avoids potentially long running recovery procedures to restore data
- Provides time-consistency between CF structures and disk
- Requires Freeze=STOP policy

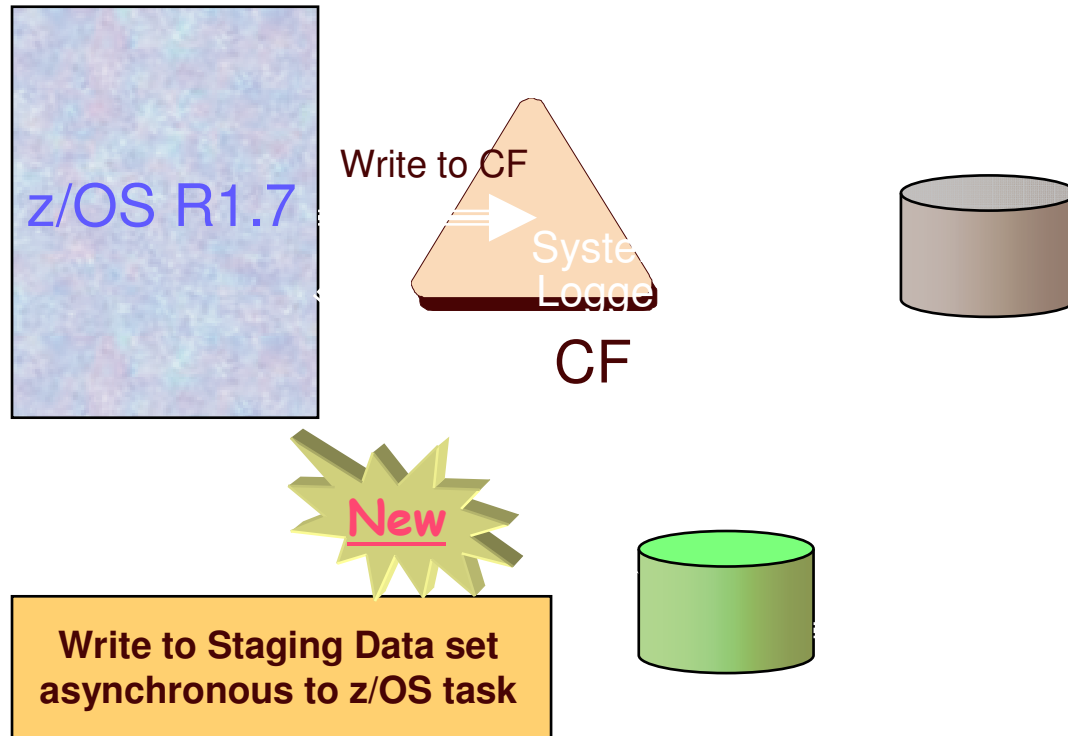
•No special recovery actions (e.g., GRECP) required  
 •Facilitates faster application restart (may improve RTO)  
 •Can provide consistent recovery time

## GDPS/XRC V3.3 Enhancements

- Write Pacing
- Scalability Enhancements
  - Parallel Execution
  - Mapping macro for Flash Copy
  - Improved panel responsiveness
- XRC+ Volume Management
- Standard Action Panel Usability Enhancements

# GDPS Support for “XRC+”

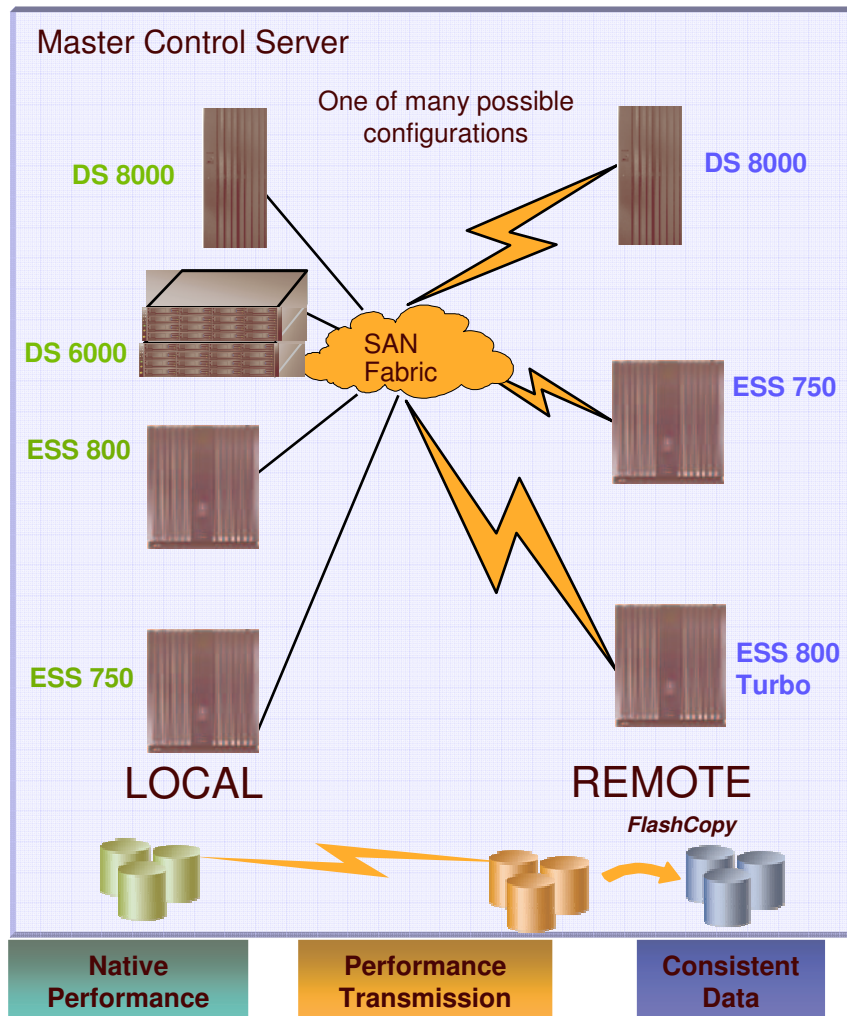
*Performance Enhancement for z/OS R1.7 (3Q05)*



- **Designed to provide better support for remote mirroring of high volume logging applications**
- **Only with XRC**
- **GDPS support will include the options in the setup for volume management.**

**Improved throughput for high volume logging applications**

# IBM TotalStorage Global Mirror



■ **Designed to Provide:**

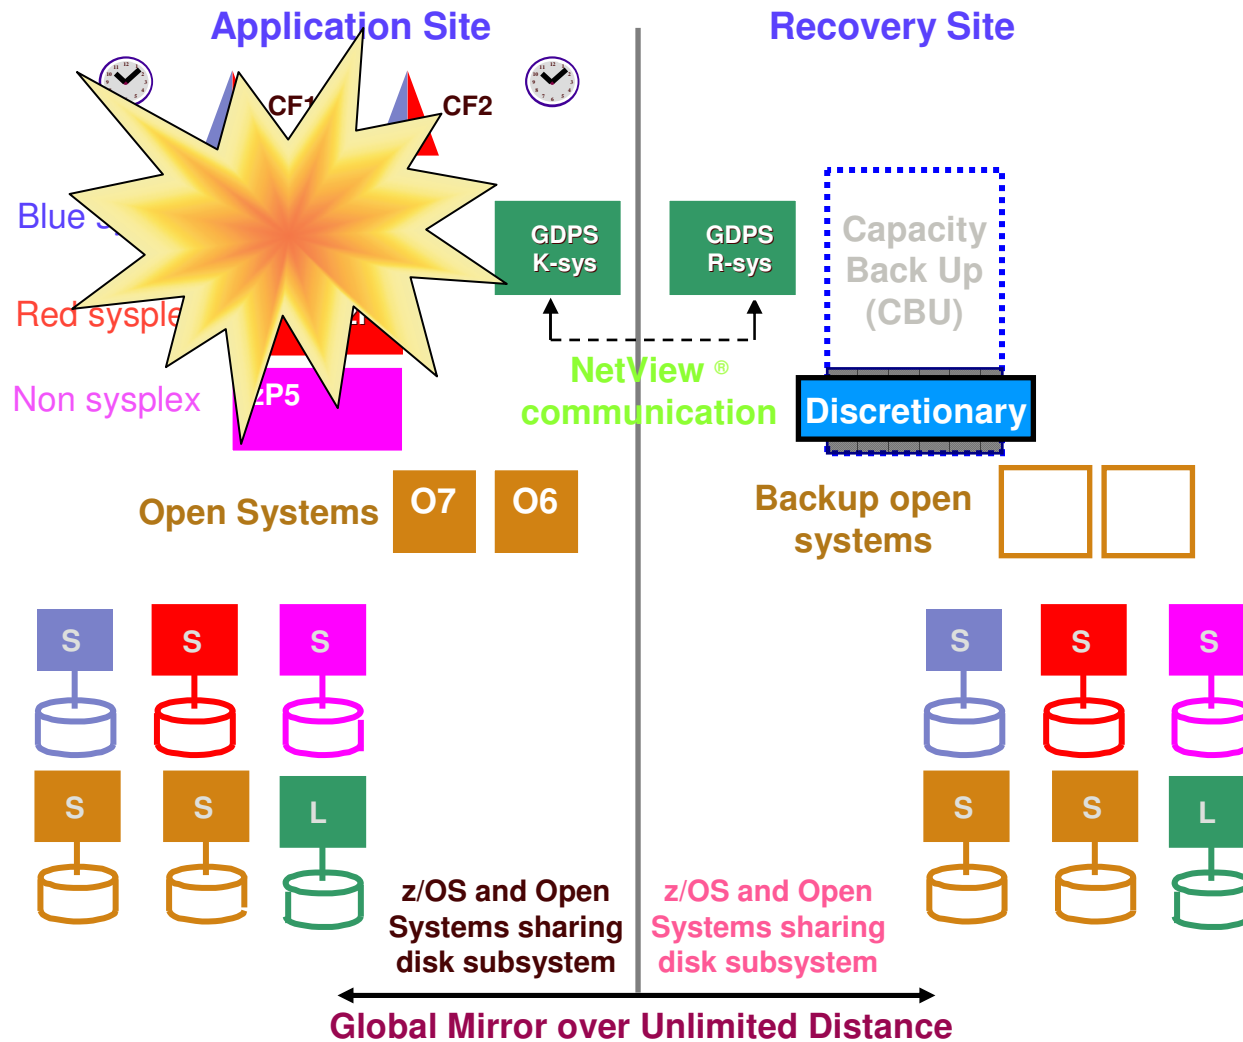
- **Global Distance:** Two-site, unlimited distance, data consistent asynchronous disk mirroring
- **Heterogeneous:** Data can span System z9, zSeries and open systems data, and can contain a mix of System z9, zSeries and open systems data
- **Scalability:** Consistency Group supported across up to 17 total IBM TotalStorage Enterprise Storage Server® (ESSs) in Global Mirror session (with RPQ)
- **Flexibility:** Many possible configurations
- **Application Performance:** Native
- **Mirroring Performance:** Two ESS Fibre Channel disk mirroring links per ESS sufficient for almost all workloads

■ **Intended Benefits**

- **Autonomic:** No active external controlling software required to form consistency groups
- **Saves cost:** No server cycles required to manage consistency groups
- **Lowers TCO:** designed to provide improved performance, global distances, and lower costs

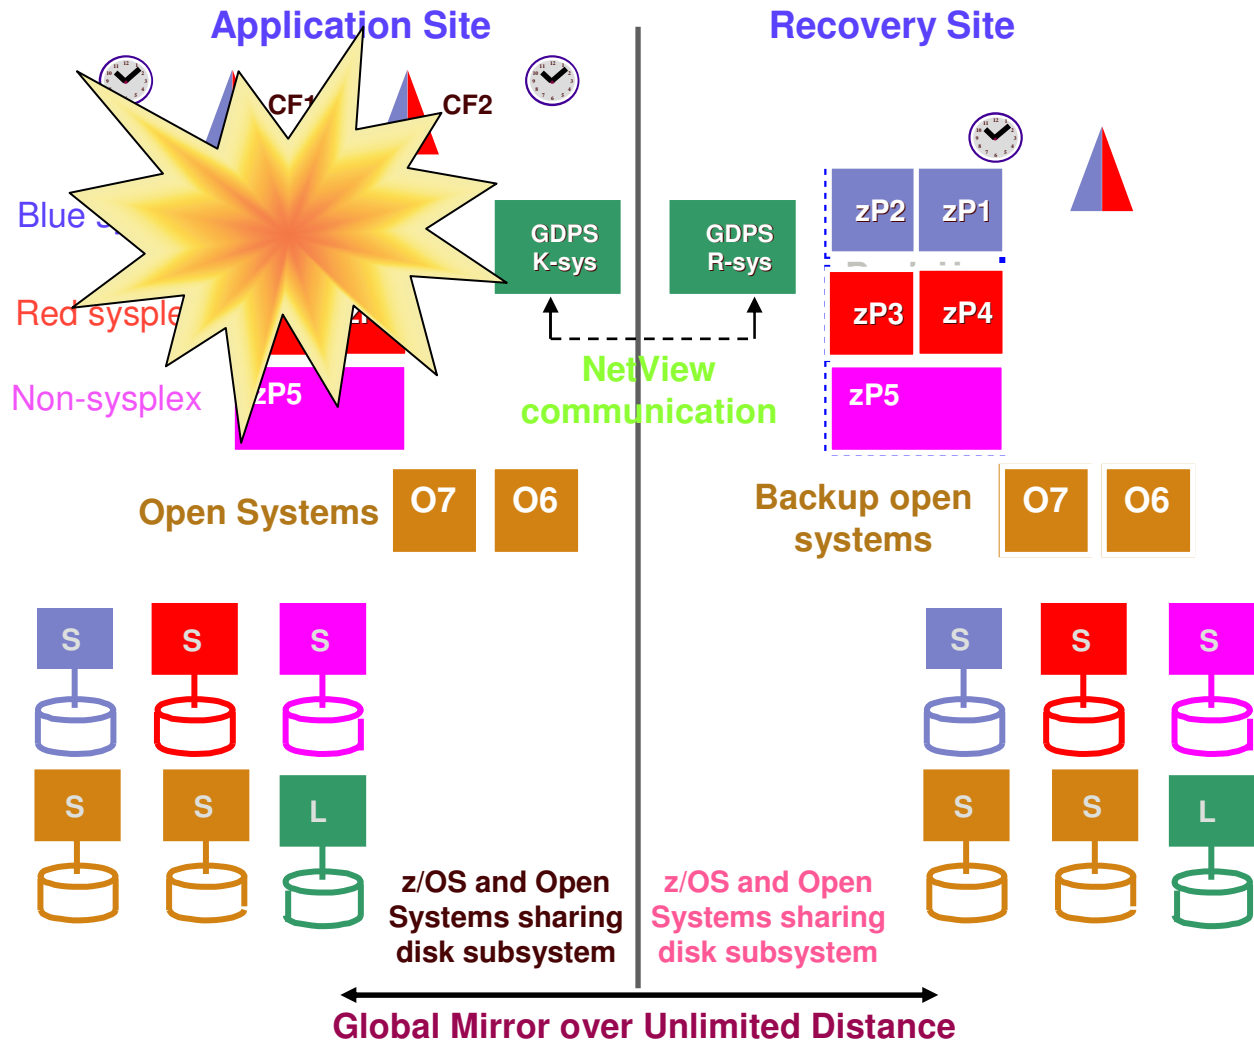


# GDPS/Global Mirror



- Application site can have single z/OS systems, Open systems, systems in a Sysplex
- All data (z/OS and Open Systems) is planned to be mirrored using Global Mirror
- K-sys activities**
  - Manages multiple Global Mirror sessions
  - Sends device info, scripts, alerts to R-sys
- R-sys activities:**
  - Secondary disk recovery, CBU activation, activate backup LPARs, IPLs systems.


# GDPS/Global Mirror



- **RTO < 2 hours**
- **RPO < 5 seconds**
  - (depends on bandwidth)



LSU Danmark




Danmark LSU October 2006

*Henrik Thorsen*



11/5/2007

© 2006 IBM Corporation

# RMF Architecture Overview

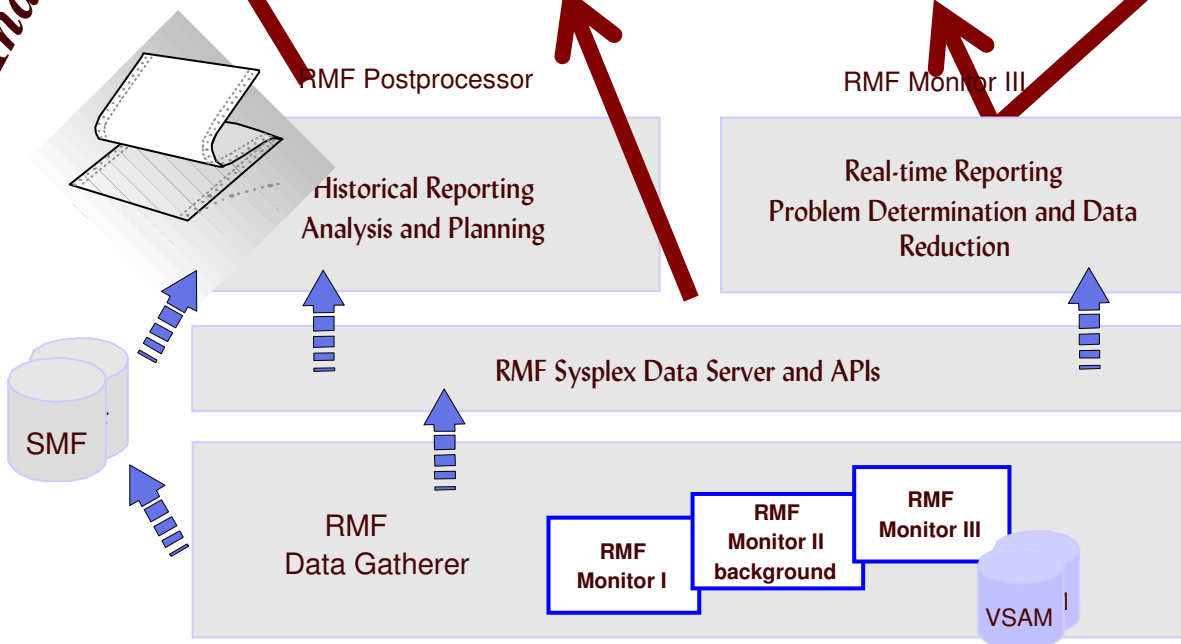


## Snapshot Reporting



*Long-term Analysis*

*Online Monitoring Asynchronous Events (CIM)*



## z/OS 1.7 Timeframe News: *Recent RMF enhancements*



- zFS support
  - Monitor III zFS reports
- Disk Space Monitoring
  - Monitor III disk space reports
- CPU activity enhancements
  - z9 separate processor pool
  - RMF Address Space Queue depth enhancement
- SPEs and RMF Distributed Data Server (DDS) extensions
  - Monitor III now collects MASTER AS User/LSQA/SWA/229/230 page information (no Report displays it)
  - Monitor III Common Storage report enhanced to show unallocated common storage
  - RMF Monitor III Data Portal preview

## z/OS 1.8 Timeframe News: *Recent HW support*



- System z HW support
  - System z Specialty Engines (zIIPs, zAAPs, ICFs, IFLs)
  - z9 separate processor pool reporting
  - Support for >128 GB
    - New memory management (UIC calculation)
    - More precise and less CPU consuming UIC calculation
- Disk Space Monitoring
  - Monitor III disk space reports
- SPEs and RMF Distributed Data Server (DDS) extensions
  - Monitor III now collects MASTER AS User/LSQA/SWA/229/230 page information (no Report displays it)
  - Monitor III Common Storage report enhanced to show unallocated common storage
  - RMF Monitor III Data Portal preview

## IBM System z Specialty Engines



WebSphere software

Information Management software

- **z9 Integrated Information Processor (zIIP) is another speciality engine that allows to execute zIIP eligible workload – that's currently DB2 star schema parallel queries and some DB2 utilities. RMF reports zIIPs as „IIPs“.**  
**For zIIPs, you need a z9 EC or BC server together with z/OS V1.6 or higher.**
- **System zAAPs are attractively priced processors limited to execute z/OS Java workload. In z/OS V1.6/V1.7, RMF reported them under the old name „IFA“ – Integrated Facilities for Applications. For z/OS V1.8, the abbreviation „AAP“ is used (to be more consistent with the naming of the new zIIPs).**
- **In order to use zAAPs, you need z/OS 1.6, IBM SDK for z/OS Java 2 Technology Edition 1.4, and for WAS workload WAS V5.1 or higher on z890, z990 or z9-109 systems.**
- **On z890, please note zAAPs always run at full speed. To make CPU numbers comparable, RMF introduces normalization. Older Speciality Engines – IOPs/SAPs, ICFs and IFLs – still supported by RMF. Similar things can happen for zIIPs on z9 BC, so there is a normalization factor to accommodate that.**

## zIIP Usage Projection

- **zIIP on CP values are filled in RMF and SMF if IEAOPT parameter PROJECTCPU = YES**
- **For more detailed information on zIIPs, please refer to**
  - zIIP Overview <http://www-03.ibm.com/systems/z/ziip/>
  - zIIP FAQ <http://publibz.boulder.ibm.com/zoslib/pdf/zosr6r7ziip.pdf>



# zAAP and zIIP: Monitor I Workload Activity Report

- ▶ **The Resource Consumption section of the WLMGL report is extended by following zAAP and zIIP related fields:**
  - ▶ AAP / IIP                                  zAAP / zIIP service time in seconds
  - ▶ APPL% AAP / IIP                          Percentage of CPU time on zAAP/zIIP processors used by eligible workloads
  - ▶ APPL% AAPCP / IIPCP                      Percentage of CPU time used by zAAP/zIIP eligible workload executed on regular CPs
- ▶ **The meaning of APPL% changes. It is the percentage of CPU time used by transactions running on standard CPs. Thus, the field is renamed from APPL% to APPL% CP**
  - ▶  $APPL\% \text{ CP} = (TCB + SRB + RCT + IIT + HST - AAP - IIP) \times 100 / \text{interval length}$
- ▶ **TCB is the task control block time spent on regular CPs as well as IFA processors**
- ▶ **If no zAAPs / zIIPs are configured, N/A is shown for the new values.**

```

----- SERVICE CLASS PERIODS
-
-REPORT BY: POLICY=DRDAIC1    WORKLOAD=DB2    SERVICE CLASS=DB2ADRS    RESOURCE GROUP=*NONE    PERIOD=1    IMPORTANCE=1
-
-TRANSACTIONS    TRANS-TIME HHH.MM.SS.TTT    --DASD I/O--    ---SERVICE---    SERVICE TIMES    ---APPL %---    PAGE-IN RATES
-AVG    3.00    ACTUAL    0    SSCHRT 585.3    IOC    300191    TCB    18.6    CP    10.68    SINGLE    0.0
-MPL    3.00    EXECUTION    0    RESP    1.0    CPU    547934    SRB    11.1    AAPCP 0.00    BLOCK    0.0
-ENDED    0    QUEUED    0    CONN    0.7    MSO    0    RCT    0.0    IIPCP 0.01    SHARED    0.0
-END/S    0.00    R/S AFFIN    0    DISC    0.1    SRB    329049    IIT    1.5    HSP    0.0
-#SWAPS    0    INELIGIBLE    0    Q+PEND    0.2    TOT    1177K    HST    0.0    AAP    0.00    HSP MISS    0.0
-EXCTD    0    CONVERSION    0    IOSQ    0.0    /SEC    4024    AAP    0.0    IIP    2.81    EXP SNGL    0.0
-AVG ENC    0.00    STD DEV    0    IIP    8.2    EXP BLK    0.0
-REM ENC    0.00    ABRPTN 1341    EXP SHR    0.0
-MS ENC    0.00    TRX SERV 1341
-
-GOAL: EXECUTION VELOCITY 80.0%    VELOCITY MIGRATION:    I/O MGMT 85.3%    INIT MGMT 85.7%
-
RESPONSE TIME EX    PERF    AVG    ----- USING% -----    EXECUTION DELAYS %    ----- DLY%--
-SYSTEM    VEL%    INDX    ADRSP    CPU    AAP    IIP    I/O    TOT    I/O    AAP    IIP    CPU    UNKN    IDLE
-HZ    --N/A--    85.7    0.9    3.0    3.3    0.0    10.9    16.9    5.4    3.0    0.1    1.4    1.0    0.0    63.5
    
```

# zAAP and zIIP: CPU Activity Report

```

-
                                C P U   A C T I V I T Y

-CPU  2094   MODEL  716   H/W MODEL  S18

-----CPU----- ONLINE TIME  LPAR BUSY      MVS BUSY      CPU SERIAL  I/O TOTAL      % I/O INT
- NUM  TYPE  PERCENTAGE  TIME PERC     TIME PERC     NUMBER        INTERRUPT RATE  HANDLED
-  0   CP   100.00      5.46         5.41         396AAD        32.99           0.27
-  1   CP   100.00      5.58         5.54         396AAD        35.66           0.27
-  2   CP   100.00      5.34         5.29         396AAD        44.29           0.22
-CP   TOTAL/AVERAGE      5.46         5.41         113.0         0.25

-  3   AAP  100.00      0.34         0.22         396AAD
-AAP  AVERAGE      0.34         0.22

-  4   IIP  100.00      20.34        20.22        396AAD
-  5   IIP  90.00      23.88        23.00        396AAD
-IIP  AVERAGE      22.29        21.99
    
```

- ➔ CPU section is grouped per processor type – CP, AAP, IIP, IFL, ICF
- ➔ New type column indicates to which pool a processor belongs
- ➔ The I/O related columns are only available for CPs, not for System z speciality engines
- ➔ A TOTAL/AVERAGE line is printed per processor pool

# zAAP and zIIP: Partition Data Report

```

PARTITION DATA REPORT
MVS PARTITION NAME          TRX2          NUMBER OF PHYSICAL PROCESSORS          19
IMAGE CAPACITY              77              CP              15
NUMBER OF CONFIGURED PARTITIONS 60              AAP              1
WAIT COMPLETION             NO              IFL              0
DISPATCH INTERVAL          DYNAMIC          ICF              1
                               IIP              2
----- PARTITION DATA ----- -- LOGICAL PARTITION PROCESSOR DATA -- -- AVERAGE PROCESSOR UTILIZATIO
          ----MSU----  -CAPPING-- PROCESSOR-  ----DISPATCH TIME DATA---- LOGICAL PROCESSORS  --- PHYSICA
NAME      S  WGT  DEF  ACT  DEF  WLM%  NUM  TYPE  EFFECTIVE  TOTAL  EFFECTIVE  TOTAL  LPAR  MGMT
TRX2      A   10   77   10  NO   0.0   3.0  CP   00.03.10.650 00.03.16.598    5.30   5.46   0.03
...
*PHYSICAL*
TOTAL                                           00.07.21.739    2.30
-----
TRX2      A   10           1  AAP  00.00.01.776 00.00.04.059    0.15   0.34   0.19
TRX1      A   50           1  AAP  00.00.00.667 00.00.01.568    0.06   0.13   0.08
*PHYSICAL*
TOTAL                                           00.00.48.459    4.04
-----
TRX2      A  100           2  IIP  00.18.02.906 00.18.05.033   88.24  88.42   0.18
*PHYSICAL*
TOTAL                                           00.00.54.798    4.57
-----
TOTAL                                           00.18.02.906 00.18.59.832    4.74
    
```

# zAAP and zIIP: Overview Control Statements

- IIPSRV                      zIIP service units per second
- IIPCPSRV                 zIIP service units on standard CPs per second
- IIPSEC                     zIIP service time in seconds
- IIPNSEC                  zIIP service time in seconds (normalized)
- IIPCPCSEC                zIIP service time in sec. spent on standard CPs
- APPLIIP                  zIIP application execution time %
- APPLIPCP                 zIIP on CP application execution time %
- IIPDLYP                  zIIP delay %
- IIPUSGP                  zIIP using %
- IPCUSGP                 zIIP on CP using %

Similar overview control statements are also defined for zAAPs, just with AAP instead of IIP

OVW(CPUIBUSY(CPUIBSY))  
 OVW(ISECCBS(IFASEC(S.CBSRV.1)))  
 OVW(ICPSCBS(IFACPSEC(S.CBSRV.1)))  
 OVW(APPLCBS(APPLPER(S.CBSRV.1)))  
 OVW(IAPPCBS(APPLIFA(S.CBSRV.1)))  
 OVW(ICPACBS(APPLIFCP(S.CBSRV.1)))  
 OVW(IUSGCBS(IFAUSGP(S.CBSRV.1)))  
 OVW(ICPUCBS(IFCUSGP(S.CBSRV.1)))  
 OVW(IDLYCBS(IFADLYP(S.CBSRV.1)))

```

RMF OVERVIEW REPORT
z/OS V1R6                SYSTEM ID WEBD          START 07/28/2004-17.10.00  INTERVAL 00.07.05
                           RPT VERSION V1R5 RMF              END   07/28/2004-17.40.00  CYCLE 1.000 SECONDS

NUMBER OF INTERVALS 4          TOTAL LENGTH OF INTERVALS 00.40.00

DATE    TIME    INT    CPUIBUSY  ISECCBS  ICPSCBS  APPLCBS  IAPPCBS  ICPACBS  IUSGCBS  ICPUCBS  IDLYCBS
MM/DD  HH.MM.SS MM.SS
07/28  17.10.00 10.00    53.5    1050.8   18.6     280.8    175.1     3.1     2.5     0.0     0.8
07/28  17.20.00 10.00    48.1     925.7   18.9     268.6    154.3     3.1     2.0     0.0     0.8
07/28  17.30.00 09.59     8.8     162.4    2.0      70.8     27.1     0.3     0.1     0.0     0.0
07/28  17.40.00 10.00     5.0      86.8    1.2      40.6     14.5     0.2     0.0     0.0     0.0
    
```

## zAAP and zIIP: Monitor III Adaptations

- ➔ Existing CPU utilization values in SYSINFO report (*CPU Util%*, *MVS Util%*, *Appl%*, *EAppl%*) do not consider processor time spent on zAAPs or zIIPs
- ➔ Processor Using and Delay samples consider regular CPs and zAAP/zIIP processors
- ➔ AAP/IIP using and delay included in Using%, Delay% and Workflow%
- ➔ ExecVelocity% in SYSSUM report contains AAP/IIP using and delay samples
- ➔ Minor changes only in SYSINFO and ENCLAVE report

# zAAP and zIIP: Proc report

```

- RMF V1R8 Processor Delays          Line 1 of 138
- Command ==>                        Scroll ==> HALF

-Samples: 60      System: MVSI  Date: 10/31/06  Time: 09.10.00  Range: 60 Sec

-
-
- Jobname      Service  CPU  DLY  USG  EAppl  -----  Holding Job(s) -----
-              CX Class  Type  %    %    %      % Name    % Name    % Name
-
- WWS7         O  OMVS   CP   11  46  59.4   9 *ENCLAVE  7 DBS3DIST  7 WSP1S2F
- WSP1S2FS     SO  WASCR  CP    4   4  42.5   2 DBS3DIST  2 WWS7     2 VTAM44
-              AAP      6   0  98.4   6 *ENCLAVE
- WSP1S6FS     SO  WASCR  CP    0   0   5.3
-              AAP      6   0   7.7   6 *ENCLAVE
- DBS3DBM1    S  DB2HIGH CP    2   6   0.8   2 XCFAS     2 DBS3DIST  2 WSP1S2F
- WSP1S6F     SO  WASCR  CP    0   2   1.9
-              AAP      2   2   0.7   2 *ENCLAVE
- U078069     O  OMVS   CP    2   4   1.2   2 WWS7     2 DBS3DIST  2 U078069
- WSP1S4F     SO  WASCR  CP    0   0   0.1
-              AAP      2   0   0.4   2 WSP1S6F
- U078068     O  OMVS   CP    2   0   0.2   2 XCFAS     2 WWS7     2 *ENCLAVE
- DBS3DIST    SO  DB2HIGH CP    0  78 111.0
-              IIP      0   2  21.3
- XCFAS       S  SYSTEM  CP    0  28  24.1
- TCP/IP      SO  SYSSTC  CP    0  22  16.1
- VTAM44      S  SYSSTC  CP    0  19  14.5
- WSP1S2F     SO  WASCR  CP    0  15  14.0
    
```

- ➔ Separation of CP / zAAP / zIIP pools for usings, delays, holding jobs
- ➔ Appl% removed

# zAAP and zIIP: Monitor III SYSINFO Report

```

RMF VIR8  System Information

Samples: 120      System: TRX2  Date: 03/14/06  Time: 13.46.00  Range: 120  Sec

Partition:  TRX2  2094 Model 716      Appl%:      65  Policy: SVPOL1
CPs Online:  3.0  Avg CPU Util%:  4    EAppl%:     66  Date:   02/22/06
AAPs Online: 1.0  Avg MVS Util%:  4    Appl% AAP:  8   Time:   13.30.55
IIPs Online: 1.0  Appl% IIP:  29

Group      T WFL --Users--  RESP TRANS -AVG USG-   -Average Number Delayed For -
          %  TOT  ACT   Time  /SEC  PROC  DEV   PROC  DEV  STOR  SUBS  OPER  ENQ
-----
*SYSTEM    83  103   1      0.13  0.1  0.5    0.0  0.1  0.0  0.0  0.0  0.0
*TSO              5   0      0.13  0.0  0.0    0.0  0.0  0.0  0.0  0.0  0.0
*BATCH              1   0      0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  0.0
*STC         83   95   1      0.00  0.1  0.5    0.0  0.1  0.0  0.0  0.0  0.0
*ASCH              0   0      0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  0.0
*OMVS              2   0      0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  0.0
*ENCLAVE       0  N/A      N/A  0.0  N/A    0.0  N/A  0.0  N/A  N/A  N/A
IFASOAK  W         2   0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  0.0
IFASOAK  S         2   0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  0.0
          1         2   0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  0.0
SYSTEM   W      83  95   1  .000  0.00  0.1  0.5    0.0  0.1  0.0  0.0  0.0  0.0
SYSSTC   S      79  73   0  .000  0.00  0.0  0.1    0.0  0.0  0.0  0.0  0.0  0.0
          1      79  73   0  .000  0.00  0.0  0.1    0.0  0.0  0.0  0.0  0.0  0.0
    
```

➔ CPU%, MVS%, Appl%, EAppl% are exclusively related to CPs as before

# zAAP and zIIP: Monitor III SYSINFO Report

```

RMF Enclave Classification Data

Details for enclave ENC00003 with token 000000B0 0000008A.
Press Enter to return to the Report panel.

- CPU Time -          -zAAP Time -          -zIIP Time -
Total 26.78           Total 0.00              Total 6.33
Delta 22.50           Delta 0.00              Delta 1.01

State  --- Using% ----  ---- Execution Delays% ----  IDL  UNK
Samples CPU AAP IIP I/O CPU AAP IIP I/O STO CAP QUE
592    11 0.0 1.0 0.0  88 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.3


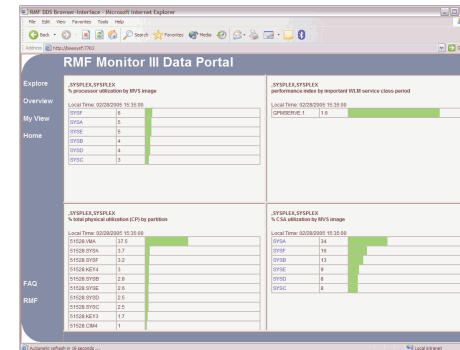
Classification Attributes:
More: +
Subsystem Type: DDF      Owner:  ENCLAVE1      System: RMF3
Accounting Information  :
Q123ERF7
    
```

- Enclave Details pop-up panel extended with information about zIIP usage
- CPU time includes time on the zAAP as well as the zIIP
- zIIP time consumed by enclave (Total)
- zIIP time consumed in reported Mon III range (Delta)
- Percentage of zIIP using and delay samples



# zAAP and zIIP: New DDS Metrics

- **New metrics for resource SYSPLEX**
  - % total utilization (zIIP/zAAP) by partition
  - % total physical utilization (zIIP/zAAP) by CPC
  - % total LPAR management time (zIIP/zAAP) for PHYSICAL by CPC
  - % zIIP/zAAP by job
  
- **New metrics for resource CPC**
  - % effective physical utilization (zIIP/zAAP)
  - % total physical utilization (zIIP/zAAP)
  - % total LPAR management time (zIIP/zAAP) for PHYSICAL
  - % effective physical utilization (zIIP/zAAP) by partition
  - % total physical utilization (zIIP/zAAP) by partition
  
- **New metrics for resource PROCESSOR**
  - % zIIP / %zAAP
  - % zIIP/zAAP on CP
  - # zIIPs/zAAP online
  - % zIIP/zAAP by enclave
  - % zIIP/zAAP delay by enclave
  - total / delta zIIP/zAAP seconds by enclave
  - % zIIP/zAAP by job / service or report class
  - % zIIP/zAAP by service or report class period



## Current System UICs (new with z/OS 1.8)

- With z/OS SRM V1.8, the algorithm to calculate the UIC has been changed. The new algorithm is much faster and more precise. With the new algorithm, there is also a set of new UIC based metrics.
- The old metrics are still computed based on the new UIC metrics (essentially, capped at 2540 maximum), so existing applications won't break. The "old" UICs as externalized in z/OS 1.8 SMF records are not wrong either, however the new metrics are more precise.

## What RMF gets from SRM: Highest System UIC (MVCSTCRI) and Current System UIC (MCTCurSystemUIC)

- Highest System UIC: that's the "old" UIC, maximum of 2540
  - For compatibility, *SMF71LIC*, *SMF71HIC*, *SMF71ACA* etc. are still filled by RMF, just capped at 2540.
  - Current System UIC: "new" UIC, maximum of 65535 (that's 18 hours)

## RMF Report Changes

- RMF replaces the Highest System UIC by the Current System UIC in following reports:
  - Header area of any Monitor II report
  - Mon II SRCS – Central Storage/Processor/SRM report
  - Mon II SPAG – Paging Activity report
  - Mon III STORS – Storage Delay Summary report
  - Mon III STORR – Storage Resource Delays report
  - Postprocessor Paging report – Central Storage section
- RMF Overview Conditions AVGHUIC and MAXHUIC are based on the new Current System UIC, too

## Monitor II (Snapshot) SRM report

RMF - SRCS Central Storage / Processor / SRM

CPU= 2/ 2 **UIC= 10K** PR= 0 System= SYSF Total

TIME	AFC	HI UIC	SQA F	LPA F	LPA FF	CSA F	L+C FF	PRI FF	LSQA CSF	LSQA ESF	CPU UTL	IN Q	OUT LOG	OUT RQ	OUT WQ
13:12:53	294K	9922	9.5K	4.6K	80	5.1K	365	7434	26K		2	51	47	0	47
13:12:55	294K	9923	9.5K	4.6K	80	5.1K	365	7434	26K		2	52	46	0	46
13:16:55	294K	9996	9.5K	4.6K	80	5.1K	365	7428	26K		2	52	46	0	46
13:16:56	294K	9997	9.5K	4.6K	80	5.1K	365	7428	26K		2	52	46	0	46
13:16:57	294K	9998	9.5K	4.6K	80	5.1K	365	7428	26K		2	51	47	0	47
13:16:59	294K	9999	9.5K	4.6K	80	5.1K	365	7428	26K		2	51	47	0	47
13:17:00	294K	10K	9.5K	4.6K	80	5.1K	365	7431	26K		3	51	47	0	47

- Width of **UIC field in Status Area** and **HI UIC column** is limited to four digits
- Current System UIC values >9999 are displayed like nnK

# Postprocessor Paging Activity Report

```

                P A G I N G   A C T I V I T Y

z/OS V1R8          SYSTEM ID SYSF          DATE 02/28/2006          INTERVAL 14.59.999
                   RPT VERSION V1R8 RMF    TIME 15.30.00          CYCLE 1.000 SECONDS

OPT = IEAOPT00    MODE = ESAME           CENTRAL STORAGE MOVEMENT RATES - IN PAGES PER SECOND
-----
HIGH UIC (AVG) = 60333    (MAX) = 65535    (MIN) = 50888

                WRITTEN TO          READ FROM          *--- CENTRAL STORAGE FRAME COUNTS ---*
                CENTRAL STOR        CENTRAL STOR        MIN           MAX           AVG

HIPERSPACE  RT           0.00           0.00           2             2             2
PAGES

VIO         RT           0.00           0.00           0             0             0
PAGES
    
```



- The HIGH UIC fields MIN, MAX and AVG do no longer format the old UIC fields SMF71LIC, SMF71HIC and SMF71ACA
- The new Current System UIC fields SMF71UAC, SMF71UHC and SMF71ULC
- ➔ Highest values possible change from 2540 to 65535

## RMF z/OS 1.8 Part 2

### *Software*

- Group Capacity Limits
- zFS support
- Disk Space Monitoring
- CPU activity enhancements
- WLM report class periods

# Group Capacity Limits: PP CPU Report

```

GROUP CAPACITY REPORT


GROUP-CAPACITY  PARTITION  SYSTEM      -- MSU --   WGT   -CAPPING--   - ENTITLEMENT -
NAME            LIMIT                                     DEF  ACT      DEF   WLM%   MINIMUM MAXIMUM

GROUP1          200    A           DOMSYS01    0   110    70   NO     0.0     93    200
                B           DOMSYS02    80   70    50   YES    0.0     67    80
                C           DOMSYS02    30   10    30   YES    0.0     30    30
-----
                                TOTAL          190   150
    
```

- This new section in the Postprocessor CPU report lists each group with its defined limit and the LPARs of the group
- MINIMUM ENTITLEMENT reports the guaranteed MSU share the LPAR gets
- MAXIMUM ENTITLEMENT reports the maximum MSU share the LPAR can get



## Group Capacity Limits: PP Partition Data Report



```

PARTITION DATA REPORT

z/OS V1R8          SYSTEM ID SYS1          DATE 30/04/2006          INTERVAL 15.00.999
                    RPT VERSION V1R8 RMF          TIME 13.30.00          CYCLE 1.000 SECONDS

MVS PARTITION NAME      A      NUMBER OF PHYSICAL PROCESSORS 16  GROUP NAME      GROUP1
IMAGE CAPACITY          120          CP      8          LIMIT          200 *
NUMBER OF CONFIGURED PARTITIONS      6          IFA      2
WAIT COMPLETION          NO          IFL      5
DISPATCH INTERVAL      DYNAMIC          ICF      1

```

- The header section of the Partition Data report is extended
  - Group name and limit printed by CPU report generator ERBMFRRCR
- If the Partition belongs to a capacity group, the group name and group capacity limit is printed
- An asterisk following the group limit indicates that the partition is less than four hours a member of the capacity group with regard to the RMF interval start. That is, the complete four hour history of MSU consumption was not available and the group limit could not be guaranteed


## RMF z/OS 1.8 Part 3

### *Infrastructure & Usability Enhancements*

- **DMTF CIM support stage 2:  
Subscription for Events / Indications**
- **RMF Postprocessor Overview  
Duration Periods with up to 100 hour  
intervals**
- **RMF Spreadsheet Reporter support  
for Secure FTP and for Passive FTP  
(useful to co-exist with some  
firewalls)**
- **Linux on System z: RMF-like LPAR /  
Channel / Device metrics**

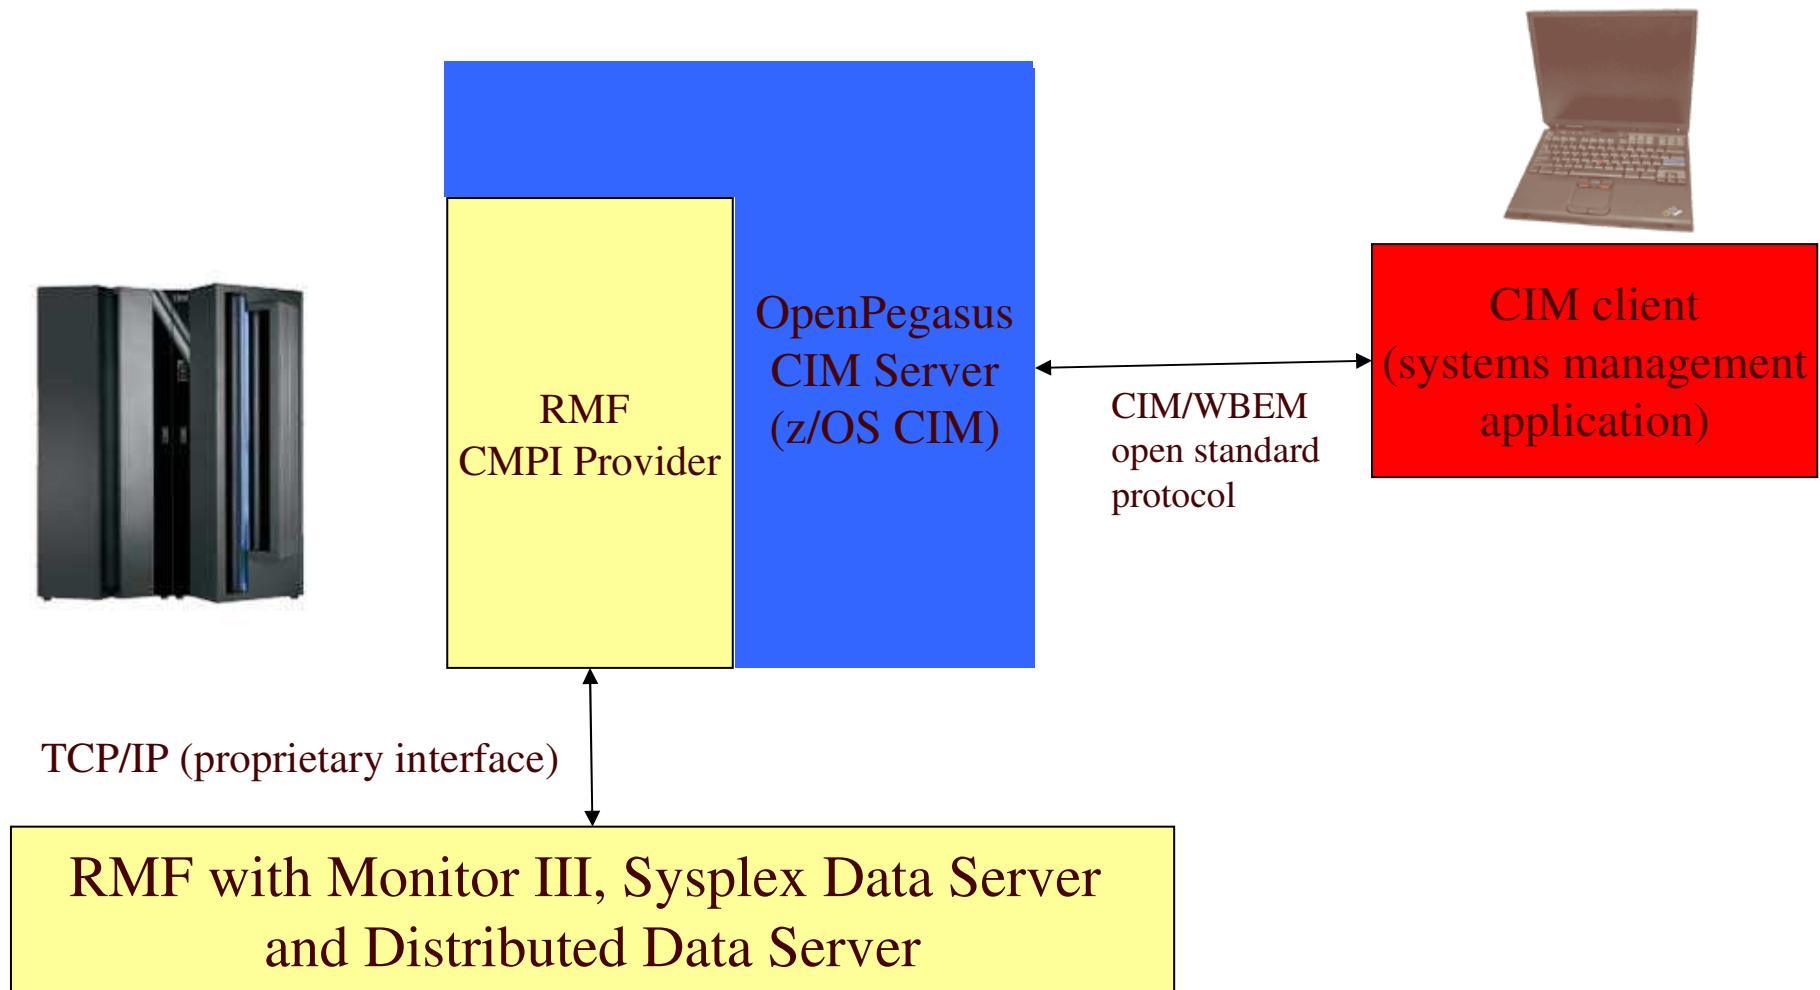
## Tivoli OMEGAMON XE integration

Tivoli software



- Some System z interfaces are expensive to be exploited, so it's much better to gather that data only once per sysplex
- First step since the acquisition of Candle Corp. is that Tivoli OMEGAMON XE will be enabled to use RMF gathered CF data, so we can reduce overall monitoring overhead if using RMF as well as OMEGAMON
- Existing OMEGAMON customers will lose the capability to monitor individual CF paths if using RMF as the data source, on the other hand the gathering interval will be much shorter (100 seconds instead of 5 minutes if running with default settings), so we think the quality of the data will be much better. Nevertheless, it's still optional to have the data gathered by OMEGAMON if you prefer for whatever reason
- RMF DDS Infrastructure has been extended to add HA features if exploited by Tivoli

# CIM System Architecture



## CIM Indications

- Beginning with z/OS RMF 1.8, it is possible to subscribe for asynchronous events, so an indication gets delivered to the destination server whenever a given threshold condition is fulfilled.
- The filter strings use an SQL like query language modified for CIM requirements, that's the new CIM Query Language (CQL) which has been implemented by IBM for the OpenSource project OpenPegasus.

## eServer OS monitoring

- The problem to be solved is to provide a *remote instrumentation interface for performance data* with the following properties:
  - Based on **open standards** – and not only “standards” like *XML* (or *EBCDIC* with *TCP/IP*, ...), but conceptual data model standardized, so interoperable systems management applications can be developed
  - **Common model across eServer** platforms, so you can get the same metric on i5/OS (since V5.3), z/OS, Linux (on zSeries) and potentially other platforms
  - Exploit CIM server infrastructure (new z/OS base component), based on OpenPegasus project
  - **RMF** uses some expensive interface to gather performance data (e.g. for CF data or for volumes). Therefore, it is best practice to **just use this data instead of gathering it again**, with high effort.



## CIM Goals



THE *Open* GROUP

- Seamless integration, componentized software
- Enablement of platform agnostic software / service
- Ease zSeries skill problem by making the platform more common (if possible without loosing zSeries specific value add)
- Embrace open standards
- Support systems management vendors

## Who uses CIM today ?

- **Storage industry (SNIA), SMI-S standard**
- **Microsoft with WMI**
- **Sun with WBEM**
- **Tivoli in ITM**
- **IBM Director**

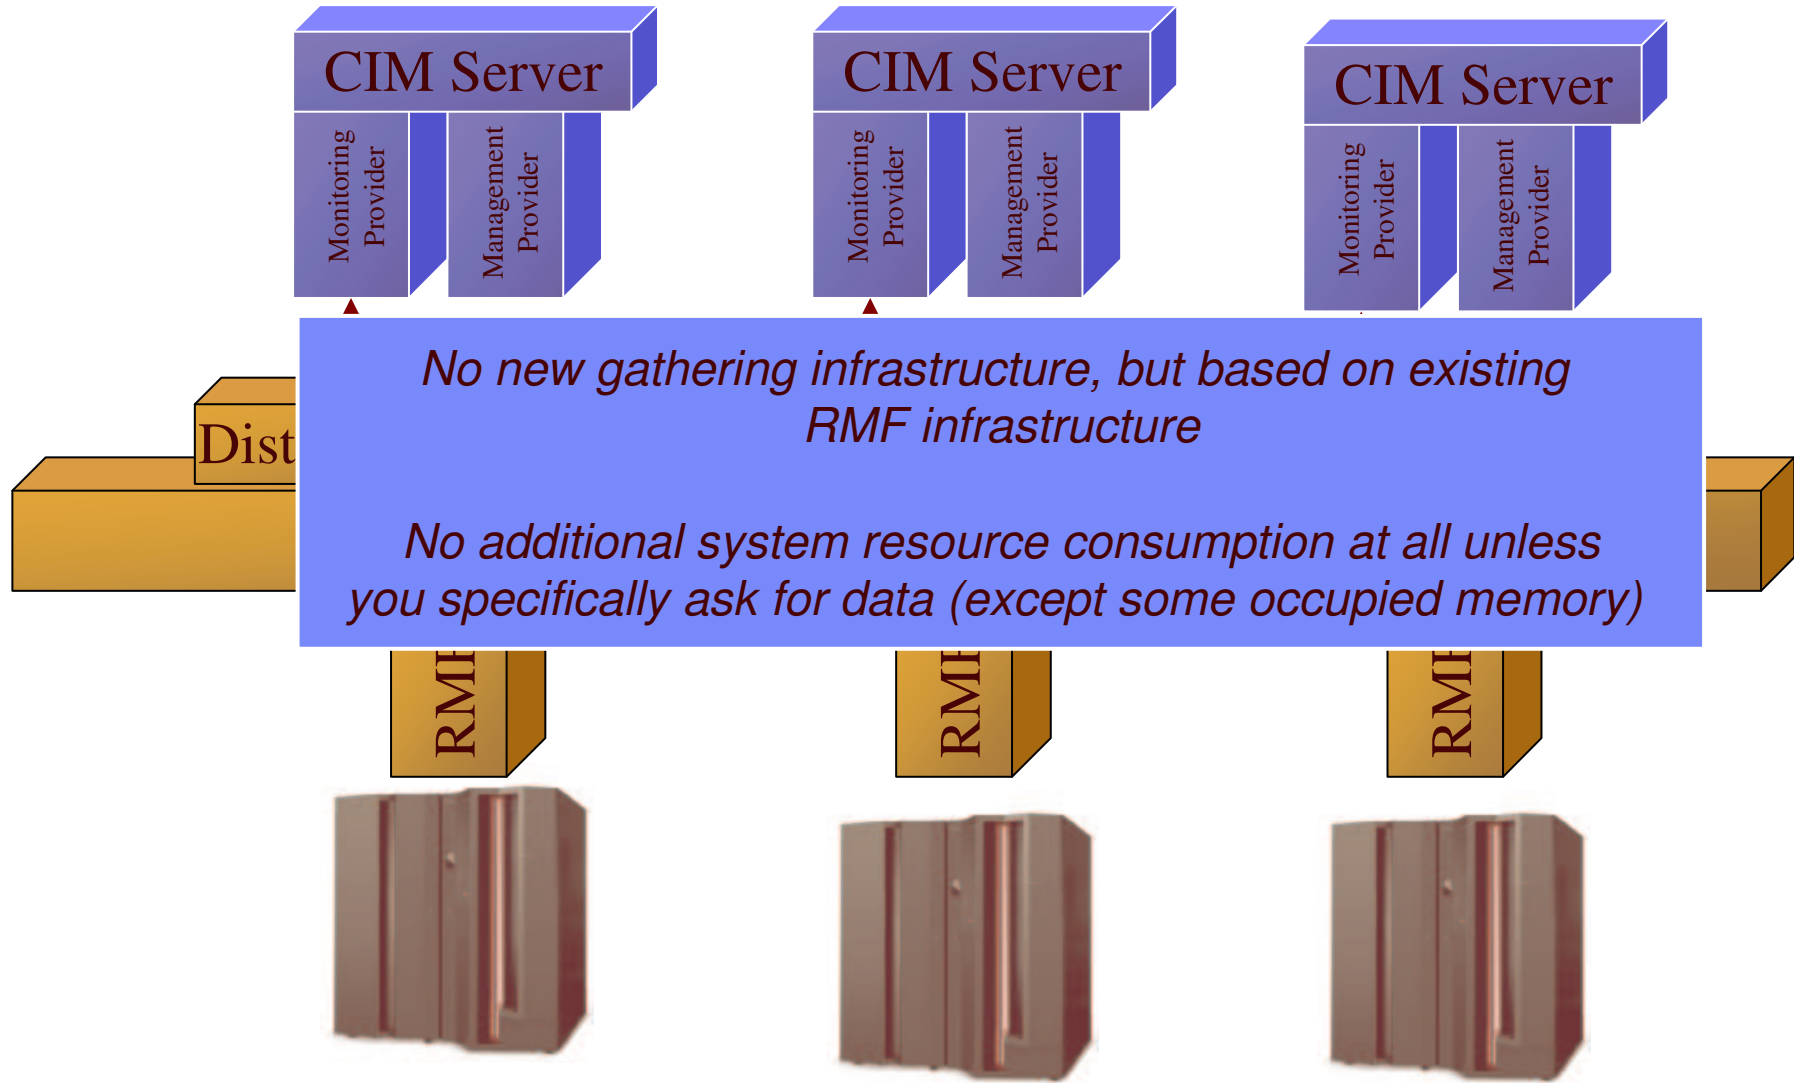
... and many more

**On z/OS with V1R7, it is**

- **The new base component *z/OS CIM*, including support for base systems management - resource classes representing address spaces, z/OS image, etc.**
- *z/OS Communications Server*
- *DFSMS Removable Media Manager* to view, add, change or delete RMM resources
- **and, of course, *z/OS RMF***

**CIM technology is already a significant investment for IBM. Stay tuned!**





## CIM Query Language (CQL)

- CQL (CIM Query Language) is based on ISO SQL-92, W3C XML-Query and WQL draft.
- WQL, WBEM Query Language, was a language draft implemented by Microsoft but never accepted by DMTF; it's deprecated because several limitations and inconsistencies)
- CQL is used for indication filters by z/OS RMF
- Example:

```
• SELECT *  
  FROM CIM_InstModification  
  WHERE SourceInstance ISA IBMzOS_BaseMetricValue  
  AND  
  SourceInstance.IBMzOS_BaseMetricValue::MetricDefinit  
  ionId  
    = '8D2870'
```

## Creation of an Indication in CIM

- One needs to create three object instances in the CIM Server, in PG\_InterOp namespace:
  - One instance of ***CIM\_ListenerDestinationCIMXML*** which describes where the indication should be delivered (TCP/IP address of the CIM client where the CIM listener thread has been started)
  - One instance of ***CIM\_IndicationFilter*** which contains the CIM:CQL query filter string
  - One instance of ***CIM\_IndicationSubscription*** which combines the CIM\_IndicationFilter and the CIM\_ListenerDestinationCIMXML instances

## Recipe to create an GetInstance based metric subscription

- 1. Get some metric you are interested in as usual, like traverse the CIM\_MetricForME association from a resource you are interested in to get all metrics associated, or traverse the CIM\_MetricInstance association from a given metric definition**
- 2. Remember the InstanceId of that instance**
- 3. Create a CIM\_IndicationFilter instance with a CIM:CQL query filter like this:**

```
select * from CIM_InstModification
where SourceInstance isa IBMzOS_BaseMetricValue
and SourceInstance.IBMzOS_BaseMetricValue::InstanceId
    = `#presource=SYSF, *, CENTRAL_STORAGE#aid=8D2EE0`
```
- 4. Create CIM\_ListenerDestination and CIM\_IndicationSubscription instances**
- 5. Start the indication listener at the destination address**

## MetricInstance association traversal

- Alternatively, it's also possible to subscribe for indications whenever any instance of a given metric definition is above a given threshold, like create an indication whenever the ResidentSetSize (WorkingSetSize) of any address space is above 1 GB of memory

- Example CQL filter string:

```
select * from CIM_InstModification
       where SourceInstance isa CIM_BaseMetricValue
       and
SourceInstance.CIM_BaseMetricValue::MetricDefinitionId =
'8D1280'
       and BaseMetricValue > '1073741824'
```

- In order to generate indications based on such a filter, the CIM\_MetricInstance association is traversed. An indication is generated whenever any address space has a ResidentSetSize of more than 1073741824 Bytes (=1 GB)

## Gathering architecture

- CIM metrics data is obtained based on existing RMF Monitor III data.
- This means that you don't add any additional gathering overhead by using the CIM infrastructure. It's not a new monitor, it's just a re-structured data model and standard based protocol for easy exploitation in client applications.
- It may happen that it takes significant time between indication checks (although unlikely). But it may not happen that the indication generation infrastructure consumes significant system resources

## CIM Installation

- The z/OS RMF CIM providers are installed as part of z/OS V1R7 Common Information Model (CIM) base element.
- However, RMF Metrics are only available if RMF is properly started and configured
- z/OS RMF CIM providers are installed as part of the z/OS CIM base element. Depending on your installation, you may need to customize the `cimserver_env` file.

## RMF CIM Environment Variables

- **Those environment variables are set in the file *cimserver.env***
- **RMF\_CIM\_HOST: contains the TCP/IP address of the z/OS MVS image on which the DDS responsible for this system is running**
- **RMF\_CIM\_PORT: TCP/IP port number of the DDS (default: 8803)**
- **RMF\_CIM\_TRACE: trace level of the RMF CIM provider, with 5 giving all information possible**
- **RMF\_CIM\_TRACE\_FILE: file to which the trace data should go**
- **RMF\_CIM\_BENCH: used for performance benchmarks, e.g. can identify the response time of the underlying RMF infrastructure**
- ***RMF\_INDICATION\_POLLING\_INTERVAL***  
**The wallclock time interval between checking the same indication subscription twice should be at least *RMF\_INDICATION\_POLLING\_INTERVAL* seconds**
- ***RMF\_INDICATION\_RESTTIME***  
**Time slept (in second) after checking for an indication; default 1 second**



## CIM SYS1.PARMLIB

- **The systems with RMF CIM providers need to be authorized to use RMF DDS.**
- **Please use *HTTP\_NOAUTH()* in *SYS1.PARMLIB(GPMSRV00)* in order to do so. The following description can be found in the parmlib member:**

### HTTP\_NOAUTH

Host names or TCP/IP addresses that can use the HTTP interface without authentication (user ID/password). Wildcards \* and ? are allowed.

Examples: HTTP\_NOAUTH(sysa.boeblingen.ibm.com)

Default: HTTP\_NOAUTH()

Notes:

1. The host running the RMF CIM provider or RMF LDAP backend must be specified here. More than one HTTP\_NOAUTH statement may be present.
2. If your installation uses a proxy server for http access, you have to specify the hostname of the proxy server to allow access from users who use this proxy server.

## eServer CIM metric

### *PartitionDefinedCapacityUsed*

- **PartitionDefinedCapacityUsed** metric: percentage of CPU capacity actually used by the operating system
- On z/OS, this is calculated by dividing the image capacity (in MSUs) by the actual consumed MSUs
- This metric can give you an impression on what percentage of the resources you gave to an LPAR is actually used
- If running z/OS under z/VM, this metric won't be available

## eServer CIM metric

*OperationalStatus*

- The *OperationalStatus* metric in CIM open standard has several possible states like “*Starting*”, “*Predictive Failure*”, etc. however with z/OS 1.7, only the states “Ok” and “Stressed” can be reported
- On z/OS, the metric is based on “% Workflow” with \*SYSTEM workscope. So the system is reported to be stressed if even workloads with system priority have significant delays. This works even if the customer installation defines WLM goal mode in a very strange way. If system workloads have significant delays, that’s always a problem.

## eServer CIM metric

### *LoadAverage*

- Average in-read queue length, average number of dispatchable units in the in-ready queue of the scheduler, not counting any idle tasks.
- This in relation to the number of processors available can be an indicator on CPU contention. Dispatchable units in the in-ready queue are workloads which are waiting for nothing but CPU resources.

*Note that this metric actually counts all dispatchable units – not only address spaces, but also USS threads, enclaves, etc. So it's a very nice metric to identify CPU contention online.*

## CIM/WBEM References

- General information:
  - <http://www.dmtf.org> (DMTF home page)
  - <http://www.openpegasus.org> (OpenPegasus home page)
  - <http://sblim.sourceforge.net> (related Linux project, sponsored by IBM LTC)
- IBM eServer OS profiles: [http://publib.boulder.ibm.com/eserver/v1r1/en\\_US/index.html](http://publib.boulder.ibm.com/eserver/v1r1/en_US/index.html)
- DMTF:
  - CIM Metrics Model – DMTF White Paper (DSP0141)  
<http://www.dmtf.org/standards/documents/CIM/DSP0141.pdf>
  - CIM Query Language Specification (DSP0202)
  - CIM Indications White Paper (DSP0141)
  - IEEE eTransactions on Network and Service Management Vol. 1 No. 2 (December 2004)  
Alexander Keller et al: “*The CIM Metrics Model: Introducing Flexible Data Collection and Aggregation for Performance Management in CIM*”  
<http://www.comsoc.org/livepubs/etn/public/2004/dec/index.html>
- z/OS RMF Programmer’s Guide (SC33-7994)
- z/OS Common Information Model User’s Guide (SC33-7988)

## RMF Overview Duration Periods

- With the enhanced support in z/OS RMF 1.8 for overview report processing, it is possible to aggregate multiple intervals to so-called *duration periods*, giving you more flexibility than in previous RMF releases. While a single interval is limited to a maximum of 60 minutes, duration periods up 100 hours are now possible.
- Based on this, it's also possible to create aggregated Overview Records for further analysis in RMF Spreadsheet Reporter.
- The RMF Postprocessor now honors duration intervals processing option "DINTV(hhmm)" to specify duration period lengths in hours and minutes. Now, you are able to produce trend reports over long periods of time!

# Overview Duration Periods: Example

- Format of the **LENGTH OF INTERVAL** columns changes from **MM.SS** to **HH.MM.SS** to cover duration intervals up to 100 hours

O V E R V I E W   R E P O R T

ID SCLM                      START 07/01/2006-00.00.00    INTERVAL 24.00.00  
 VERSION V1R8 RMF            END    07/30/2006-23.59.59    CYCLE 1.000 SECONDS

NUMBER OF INTERVALS 30                      TOTAL LENGTH OF INTERVALS 720.00.00

DATE	TIME	INT	CPUBUSY	MVSBUSY	APPLPER	NUMPROC	EXCP	EXCPRT	OCPU1	OCPU2
MM/DD	HH.MM.SS	HH.MM.SS								
07/01	00.00.00	24.00.00	3.5	3.3	4.9	3.0	15880	4.411	0.0	0.0
07/02	00.00.00	24.00.00	3.6	3.4	5.0	3.0	15841	4.400	0.0	0.0
07/03	00.00.00	24.00.00	3.5	3.3	4.9	3.0	15904	4.418	0.0	0.0
07/04	00.00.00	24.00.00	3.5	3.3	4.9	3.0	17945	4.985	0.0	0.0
07/05	00.00.00	24.00.00	3.5	3.3	4.9	3.0	15880	4.411	0.0	0.0
07/06	00.00.00	24.00.00	3.6	3.4	5.0	3.0	15841	4.400	0.0	0.0
07/07	00.00.00	24.00.00	3.5	3.3	4.9	3.0	15904	4.418	0.0	0.0
07/08	00.00.00	24.00.00	3.5	3.3	4.9	3.0	17945	4.985	0.0	0.0

## Spreadsheet Reporter FTP

*Passive FTP, SSL*

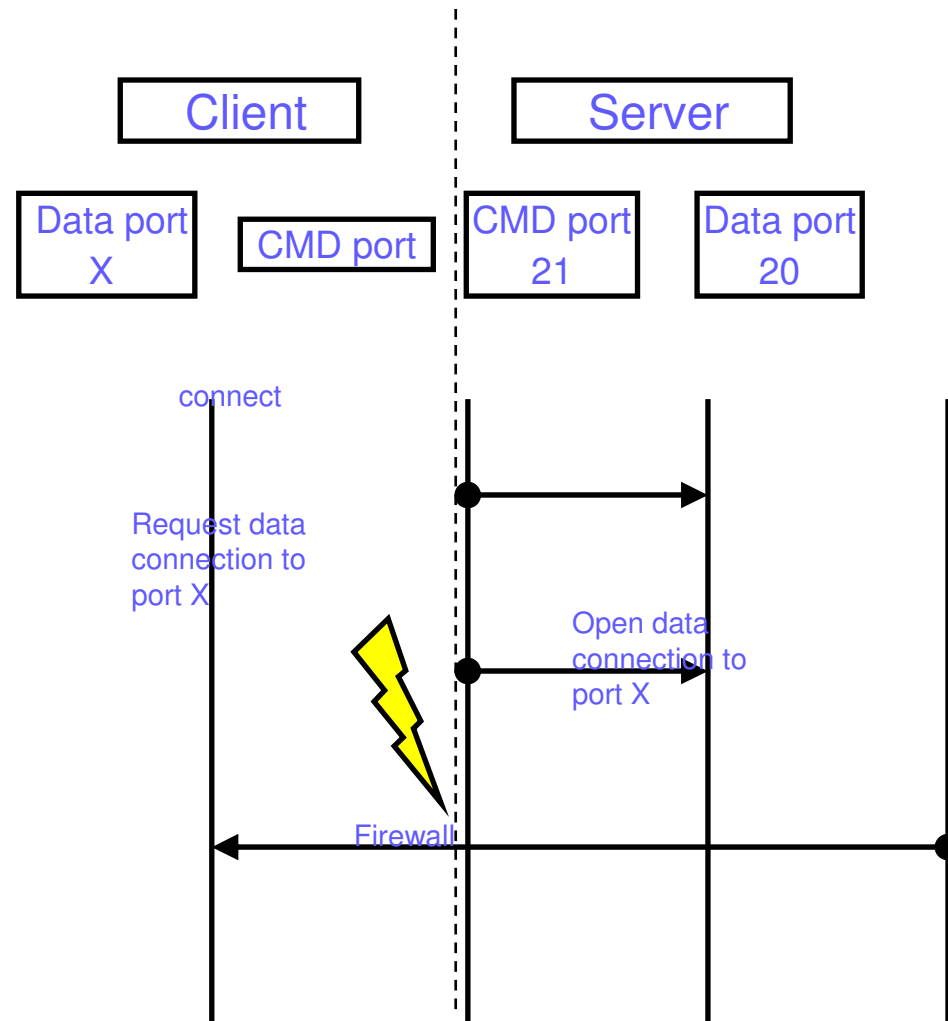
- Spreadsheet Reporter FTP Connection
  - Uses FTP connection to connect to the FTP server on the z/OS system, for example to download Postprocessor Listings
- Today
  - Active FTP connection is supported
- New support for
  - Passive FTP connection
  - Secure Socket Layer (SSL) FTP connection with Server Authentication



## Spreadsheet Reporter FTP

### Active FTP

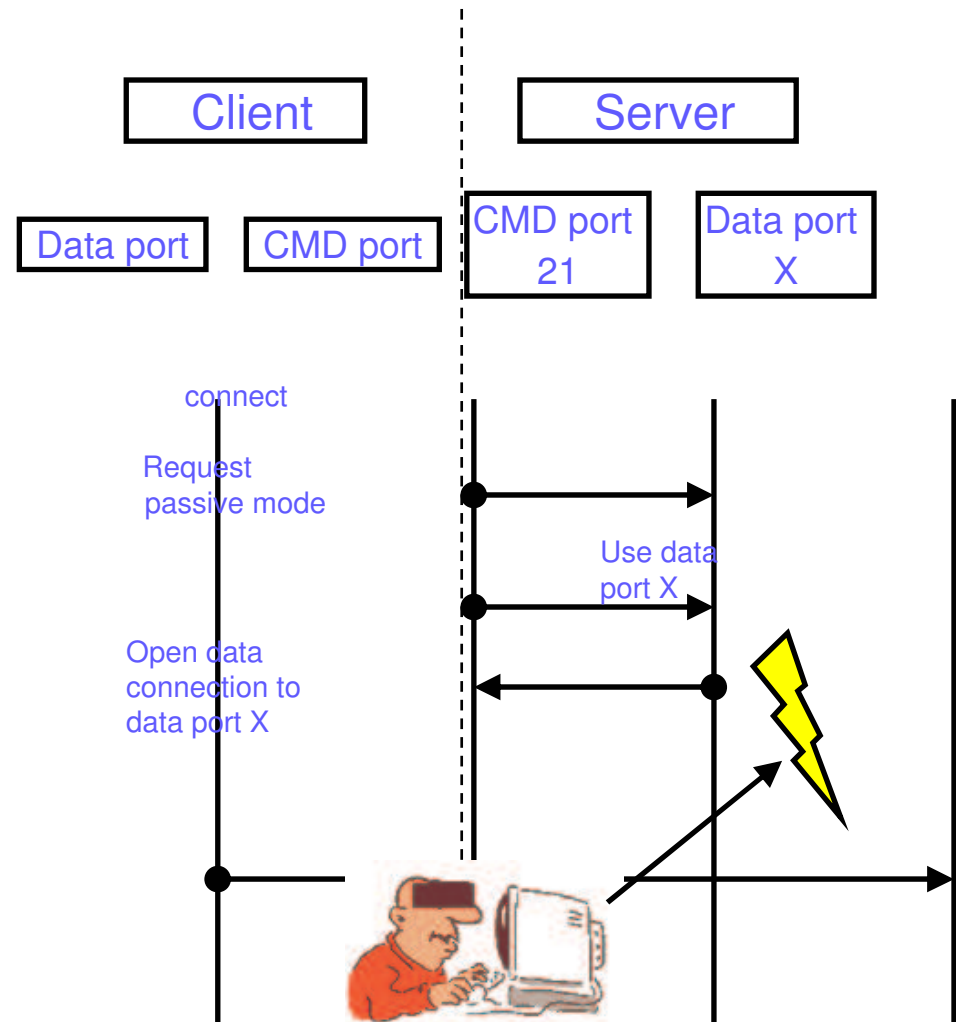
- Flow
  - Client connects to server on the host
  - Client asks the Server to open a data connection to a dataport
  - Client waits and listens on the data port for a connection
  - Server initiates a connection to the client on the dataport
- Firewalls may prevent that if not properly configured



# Spreadsheet Reporter FTP

## Passive FTP Basics

- **Flow**
- **Client connects to server on the host**
  - Client asks the Server to switch to passive mode
  - Server open a dataport, answer the client and wait and listen for connection
  - Client initiates a connection to the server dataport
- **No problem with standard Firewall, but risk that someone hijacks the open connection**

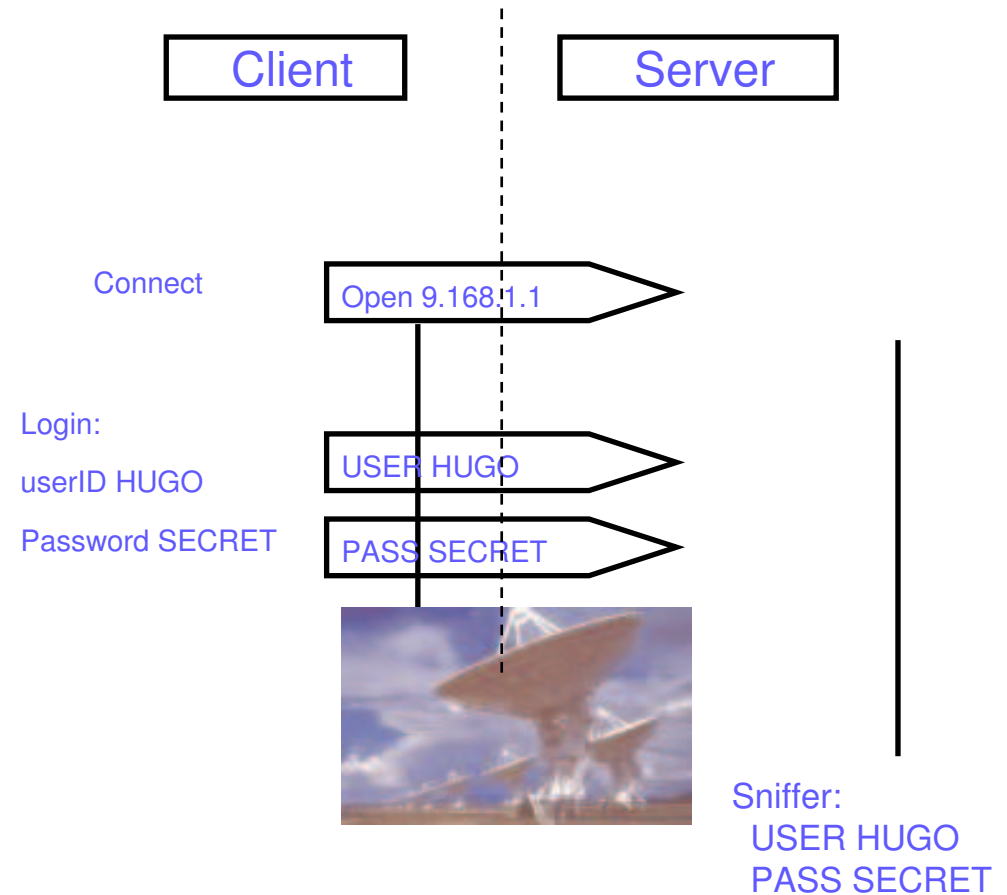


# Spreadsheet Reporter FTP

## Non-secure FTP Basics



- All communication is send as plain text
- UserID and password can be sniffed

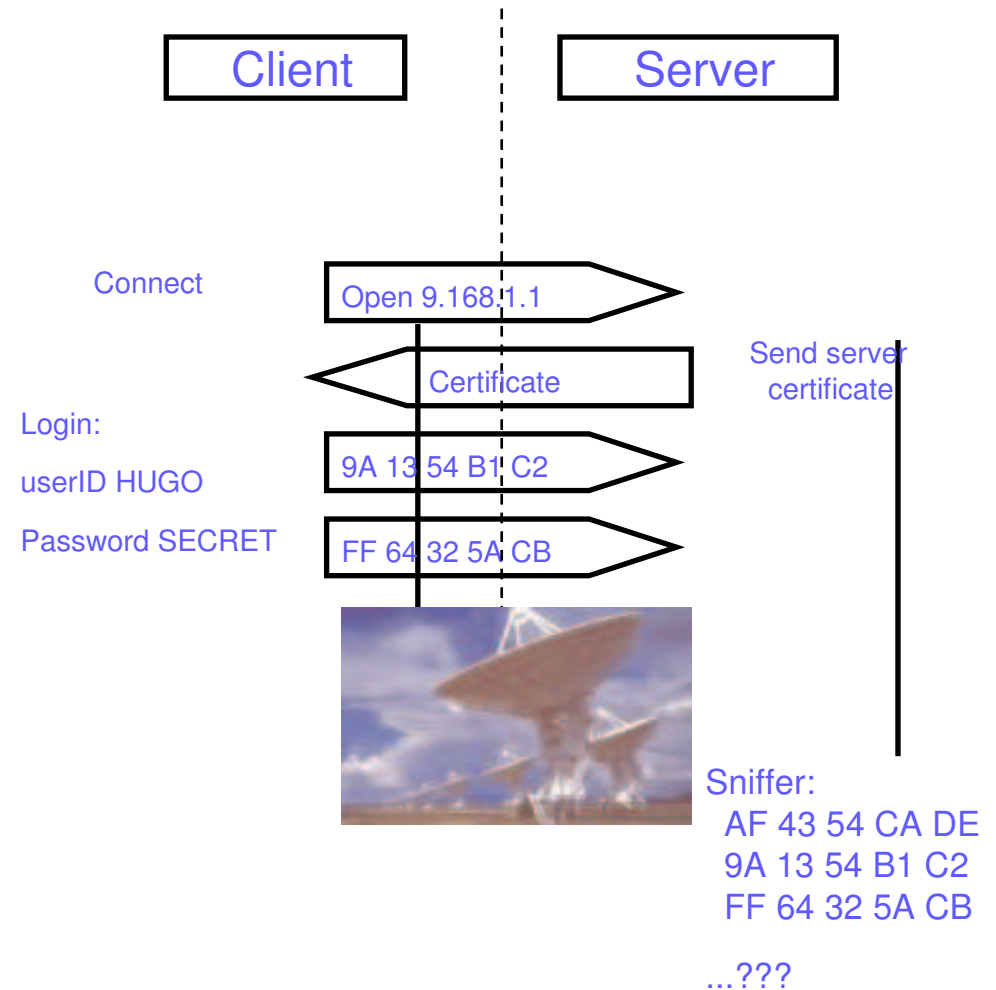


# Spreadsheet Reporter FTP

## Secure FTP Basics



- **Secure FTP flow**
  - After initiating communication, all further communication is crypted
  - Server authenticates itself to the client with an certificate
- **Server Certificate**
  - Certificate is verified:
    - Does the certificate belongs to the host that send it?
    - Is the certificate expired?
  - Client can decide to trust certificate – even if it's invalid (for example it's expired)



## Spreadsheet Reporter FTP

### *Certificate Handling*



- Spreadsheet Reporter certificate errors
  - Unknown certificate
    - The certificate is unknown – the user can decide to trust the certificate.
  - Certificate is expired
    - The certificate valid date is expired – the user can still decide to trust that invalid certificate for the current session.
  - Certificate hostname verification failed
    - A certificate is issued for one host and the hostname is stored in the certificate. If the certificate is now used by another server, the hostname of the server is not the same as stored in the certificate. The user can still decide to trust that invalid certificate for the current session.

# Spreadsheet Reporter FTP

## Activate FTP enhancements



- How to use?
  - On system level in the system profile:
  - FTP Security:
    - None
    - SSL with Server Authentication
  - FTP Mode
    - Active
    - Passive

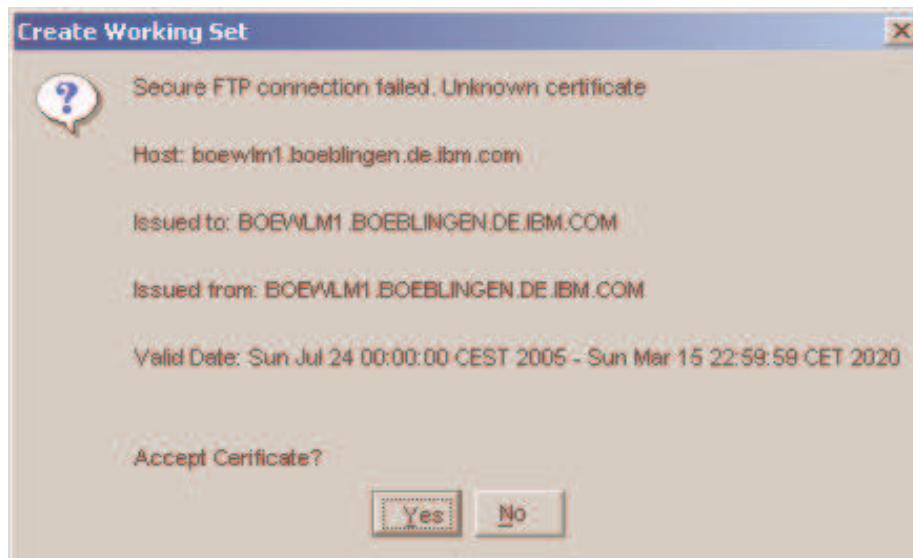
The screenshot shows a 'System Properties' dialog box with the following fields and values:

System ID	SYSA
Hostname	boesysa.boeblingen.de.ibm.com
Dataset HLQ	bhbe
User ID	bhbe
Password	*****
Account	dept007
Jobclass	a
Msgclass	h
FTP Security	SSL Server Authentication
FTP Mode	Passive
OVW	

Buttons: Ok, Cancel

# Spreadsheet Reporter FTP

## Certificate Handling



- **When a certificate error occur, the certificate is displayed, and the user is asked to accept it or not.**
  - No: Connection is aborted
  - Yes: Certificate is trusted and stored
- **Even an accepted Certificate can be invalid**
  - Certificate is expired
  - Hostname does not fit to Certificate
    - ➔ User can still decide to trust the certificate – but only for the current Spreadsheet Reporter session

# RMF DDS Data Portal

The RMF DDS Data Portal is a web browser-based frontend to RMF Monitor III data

- Just use <http://<hostnam>:8803> as URL in Mozilla, Firefox, Internet Explorer or whatever XSL-capable web browser you have
- Same data model which is also used in RMF PM Java client
  - Limitations compared with RMF PM:
  - No Analysis capability
  - Only limited “save” capabilities
  - Advanced features like filtering or work scopes not supported
- Most of the RMF Monitor III data is available, including some hidden fields





## Linux on System z: RMF like metrics

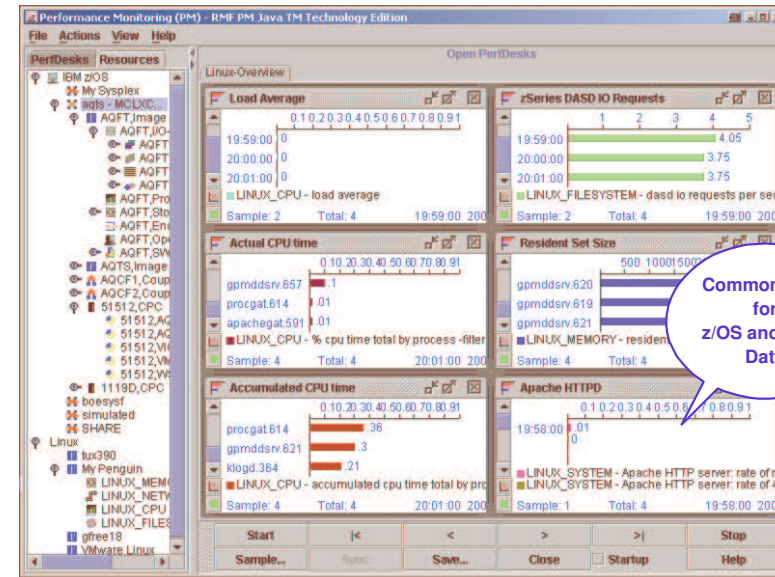
- ... implemented as part of the SBLIM open source project, to be integrated into distros (<http://www.sblim.org>)
- Metrics from the Partition Data Report (same instrumentation source as RMF – *LPAR management time, dispatch time, ...*), Channel metrics (*Partition/Total/Bus utilization, bandwidth, ...*), metrics for FICON attached devices (*Response time, Connect Time, CU Queue Time, Device Utilization, I/O Intensity, ...*)

## Linux on System z: New Metrics

- IBM implemented several System z HW specific metrics for Linux on z.
- The implementation is to be shipped by RedHat and Novell, so you don't have to install anything special on the Linux end points in order to get this data.
- The new metrics are based on System z hardware interfaces. Obviously, they are similar to or identical with related RMF metrics.
- The new metrics allow you to display most of what's displayed in RMF CPC / Partition Data report.
- In addition, a few Channel related metrics and metrics for FICON devices attached to Linux.
- All working if you're running Linux on z in an LPAR, side-by-side with z/OS!

## Linux Performance Measurement

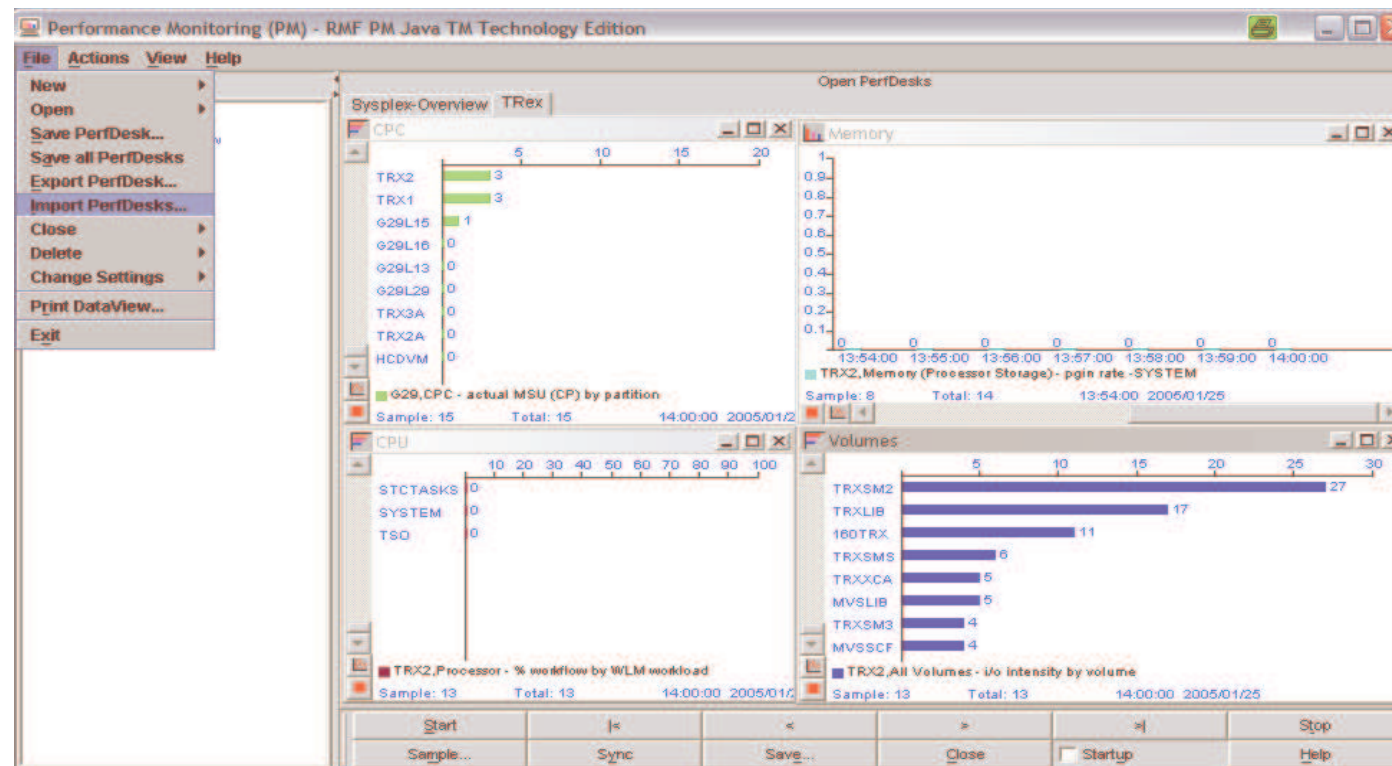
- Historical data collection
- Integration with z/OS RMF and z/VM FCON and z/VM FCON
- Most of the z/OS RMF PM functionality – like customizing performance desktops, saving pre-defined performance analysis scenarios, exporting data to spreadsheet applications, etc – is also available for Linux (for zSeries)



## rmfpms (Linux data gathering) – recent updates

- **New script to automatically start Linux gatherer at Linux guest IPL (boot) time (“`enable_autostart`”); in addition, this script moves rmfpms to `/var/opt/rmfpms` and `/opt/rmfpms` in conjunction with Linux standards and it user user ID nobody for security reasons**
- **New “`delete_old_perfdata`” script to delete old Linux performance data archives**
- **Automatic repository compression now also applied for those customers which did not install a specific `cronjob` as described in the documentation**

## RMF PM: Export/Import PerfDesk functionality




- In order to share PerfDesk definitions with your colleagues, you can now export and import PerfDesks in standard file format
- An updated RMF PM Java client for all z/OS releases and Linux versions in service is available on the RMF home page:  
<http://www-1.ibm.com/servers/eserver/zseries/zos/rmf/>

## New functionality ...

Abstract	APAR	PTF / Release
zIIP support	OA13499	UA26131 (z/OS V1R5) UA26170 (z/OS V1R7)
z9-109 support	OA10346	UA90164 (z/OS V1R2) UA90166 (z/OS V1R5) UA90171 (z/OS V1R7)
zAAP support for SDSF	OA11802	UA90172 (z/OS V1R5) UA90173 (z/OS V1R7)
zAAP support	OA05731 OA07950	UA90081 (z/OS V1R5) UA90142 (z/OS V1R5)
IBM TotalStorage ESS 2105 support	OA04877	UA90077 (z/OS V1R2) UA90078 (z/OS V1R5)
IBM TotalStorage DS8000 and DS6000 series support (D/T2107 and D/T1750)	OA06476 (PE OA10992)	UA18905 (z/OS V1R2) UA18906 (z/OS V1R5)


... and otherwise important APARs ☹️



Abstract	APAR	PTF / Release
<p><b>Incorrect APPL% in WLMGL report</b></p>	<p><b>OA17089 (PE fixing for OA13499)</b></p>	<p><b>Unknown at time of printing</b></p>



## LSU Danmark



Danmark LSU October 2006

*Henrik Thorsen*


11/5/2007

© 2006 IBM Corporation





## LSU Danmark



## z/OS 1.8 Unix System Services

### *File descriptor limit*

- File descriptors are used for all open file system objects, including files, sockets, pipes, terminals, and directories
- Currently a single Unix Process is limited to 128K file descriptors could be a constraint for very large servers
- The limit is raised to **512K** open descriptors
  - Note; the limit was raised from 64K to 128K in 1.6.

## z/OS 1.8 Unix System Services

### *extattr support*

- Currently **no** supported commands to set file format from shell
  - The **extattr** command is enhanced to accept a **-F** option with values consistent with the cp command to indicate the format of the file.
  - **Usage** - `extattr [-F format] <file name>`
    - For ***format*** you can specify:
      - NA - Not specified
      - BIN - Binary data
      - Or the following text data delimiters:**
      - NL – New Line
      - CR - Carriage Return
      - LF - Line Feed
      - CRLF - Carriage Return followed by Line Feed
      - LF CR - Line Feed followed by Carriage Return
      - CRNL - Carriage Return followed by New Line

## z/OS 1.8 Unix System Services

### *BPXBATCH Enhancements*

- Today there is a limitations in BPXBATCH input and output make which makes it more difficult to create and maintain BPXBATCH jobs compared to other batch utilities.
- Using [the BPXBATCH enhancement](#), you can:
  - Specify a much **longer parameter** (64,536 chars vs. the old 100 chars limit)
    - New **STDPARM** DD statement
    - **STDPARM** will override the usage of PARM=
  - Output the data from the standard output and standard error streams to **MVS Datasets** (instead of just z/OS Unix Files)
    - Allows output from BPXBATCH jobs to be directed to **SYSOUT**, **Sequential** or **PDS/PDSE** Members
    - Dataset should have LRECL that will allow all lines of output to fit otherwise truncation will occur.

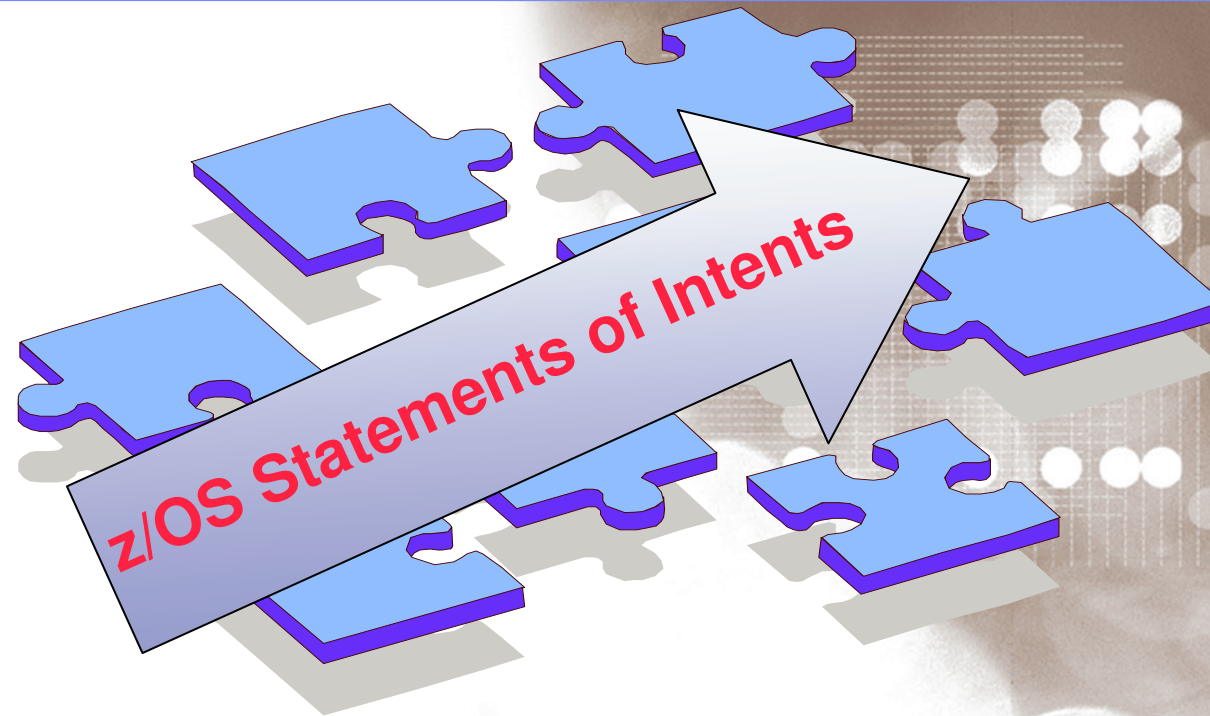
## z/OS 1.8 Unix System Services

*BPXBATCH Enhancements Alternate entrypoints*

BPXBATA2	BPXBATA8
Both are APF Authorized entry points intended for authorized invocations of BPXBATCH	
Both will force a local spawn by setting the inhemustbelocal bit of the inhe to cause it to function in a similar manner to BPXBATSL	
BPXBATA2 runs in Key 2	BPXBATA8 runs in Key 8
PPT Entry is needed to cause BPXBATA2 to startup running	PPT Entry is unnecessary to cause BPXBATA8 to startup running



## LSU Danmark



Danmark LSU October 2006

*Henrik Thorsen*

11/5/2007

© 2006 IBM Corporation

## Statements of Direction

### *Effective with z/OS 1.8 (support withdrawals)*



Old news

- Support for sharing zFS file systems contained in multi-file system aggregates to be shared across a sysplex is removed.
- z/OS 1.8 does not include the Firewall Technologies component of the Integrated Security Services element.
- 1-byte Console IDs are not supported.
- Communications Server no longer supports:
  - Configuration profile block definition statements (ASSORTEDPARMS and KEEPALIVEOPTIONS), The PAGTSNMP subagent, Defining EE TGs with multiple SAP addresses, AnyNet<sup>®</sup>
- Support for the TCP/IP, USS, LE, Parallel Sysplex, ISPF, and RMF plug-ins for msys for Setup is withdrawn. IBM intends to deliver z/OS setup and configuration improvements in the future.
  - DB2 V8 plug-in is not affected
  - TCP/IP plug-in will be made available via the Web and will not require msys for Setup
- The z/OS msys for Operations element is removed.

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

## Statements of Direction ...

*Effective with z/OS 1.8 (prelinker stabilization and Crypto enhancements)*

- IBM intends to stabilize the prelinker
  - The prelinker is designed to be used with the linkage editor
  - Note: The Binder provides additional function, and is *not* stabilized
- IBM plans to enhance the IBM Encryption Facility for z/OS (product 5655-P97) Encryption Services feature to use the OpenPGP standard, RFC 2440.
  - This is expected to allow file exchange with others having OpenPGP-capable software running on their systems.

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



## Statements of Direction ...

### *Beyond z/OS 1.8*

- z/OS 1.8 is planned to be the last release to support APPC communication between HCD and HCM
  - TCP/IP communication remains supported
- z/OS 1.8 is planned to be the last release to include the C/C++ IBM Open Class (IOC) Dynamic Link Libraries (DLLs)
  - You should plan to migrate to the standard C/C++ libraries
  - See the *IBM Open Class Library Transition Guide*
- In a future release of z/OS Comm. Server, support for the in-stack version of the TN3270 Server is planned to be discontinued.
  - The standalone TN3270 server introduced in z/OS 1.6 offers significant advantages, and remains supported.
- In a future release, IBM plans to replace the RMF LDAP backend.

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

## Statements of Direction ...

### *Beyond z/OS 1.8*

- IBM plans to make it possible to define a logical partition (LPAR) group capacity limit on System z9 servers.
- IBM intends to release an X.25 over TCP/IP product that uses the X.25 NPSI enablement interface available on IBM Communication Controller for Linux for System z9 and zSeries V1.2.
- IBM intends to deliver enhancements to the consoles component in the future, not in 2007 as originally planned.
- z/OS 1.8 is planned to be the last release that will include the APPC Application Suite in z/OS Communications Server.
  - Note that APPC itself remains an integral part of z/OS Communications Server's SNA functions. There are no plans to remove APPC from z/OS. For more information, see:  
<http://www.ibm.com/software/network/commserver/zos/>

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

## Statements of Direction .

### *Beyond z/OS 1.8 (network functions)*


- **Layer 3 Virtual MAC for z/OS environments:** When the associated OSA-Express function is available, the z/OS Layer 3 VMAC function will help to simplify the network infrastructure and facilitate load balancing when multiple operating system instances are sharing the same OSA port or MAC address. With Layer 3 VMAC, each operating system instance can now have its own unique "logical" or "virtual" MAC (VMAC) address. Instead of sharing the same physical MAC address, each TCP/IP stack and its associated IP addresses are accessible using their own VMAC addresses.
- **Network Traffic Analyzer:** When the associated OSA function is available, the Network Traffic Analyzer function will help to provide a new diagnostic trace facility, allowing OSA packet trace records to be sent directly to the host operating system.
- **QDIO Diagnostic Synchronization:** When the associated OSA function is available, the QDIO Diagnostic Synchronization function will help provide system programmers and network administrators the ability to coordinate and simultaneously capture both software (z/OS) and hardware (OSA-Express2) traces.

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

## Notable Changes\*

- Effective January 15, 2006, the S/390 Service Update Facility (SUF) will be discontinued.
- Effective March 2006, new ESO and CBPDO physical delivery subscriptions will not be accepted.
- Effective June 2006, CBPDO product orders will include service only for included products.
- Effective June 2006, Service-Only CBPDO (5751-CS3) orders will no longer be accepted.
- Effective September 2006, existing ESO and CBPDO physical delivery subscriptions will be discontinued.

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



Thanks!!  
Hope I see you again next year