



LSU Sweden

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

Nordic

Large Systems Update Seminar

October 2006

Uno Bengtsson

IBM Server Technology Group (STG)

e-Mail: uno.bengtsson@se.ibm.com

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

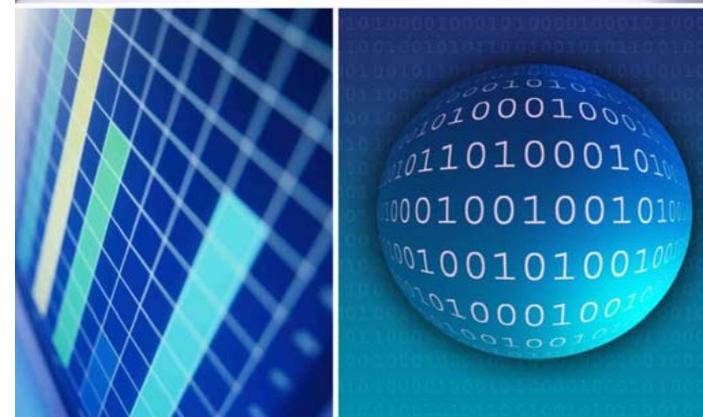
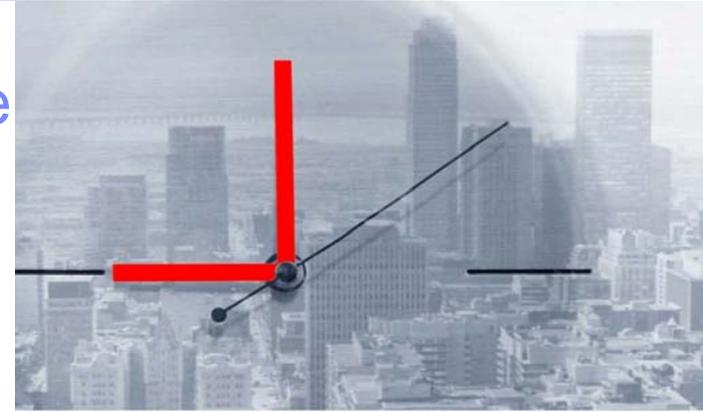


LSU Sweden



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



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."

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**
 - Joint IBM and SHARE event



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

Strategic Investment

System z9

–\$1.2 billion

–5,000 tech
professionals

\$40M



\$100M simplification





LSU Sweden



z/OS 1.7 (3Q 2005) news

Console Restructure Stage 1BOld 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** Addressspace
 - 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.

Old news

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/3rd 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)

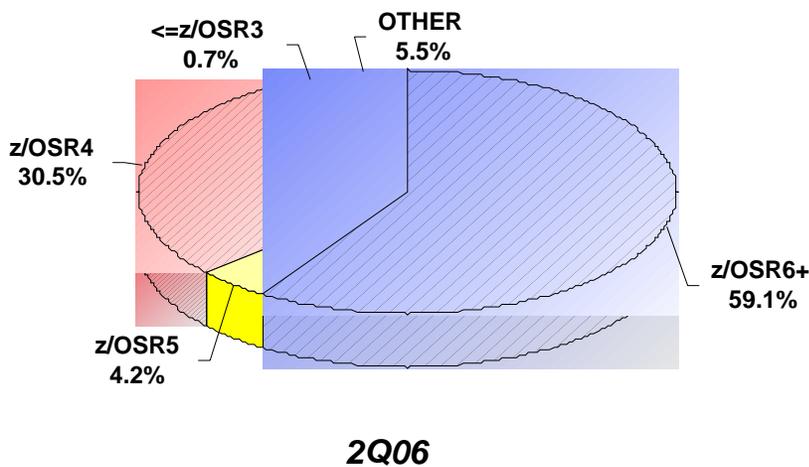
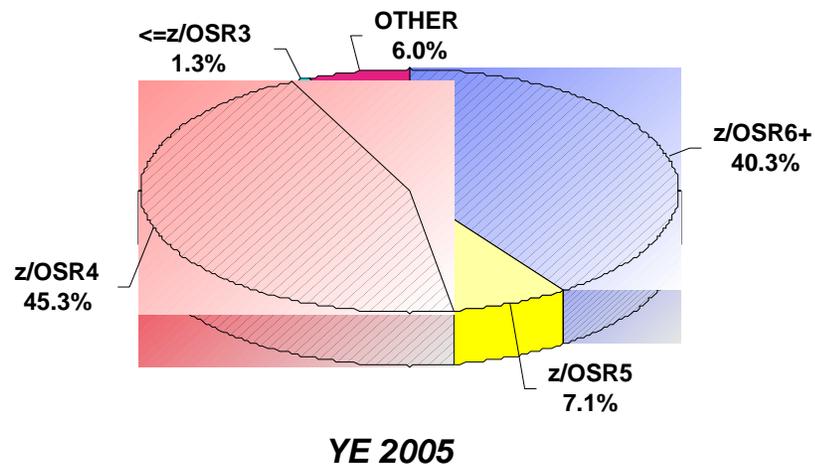
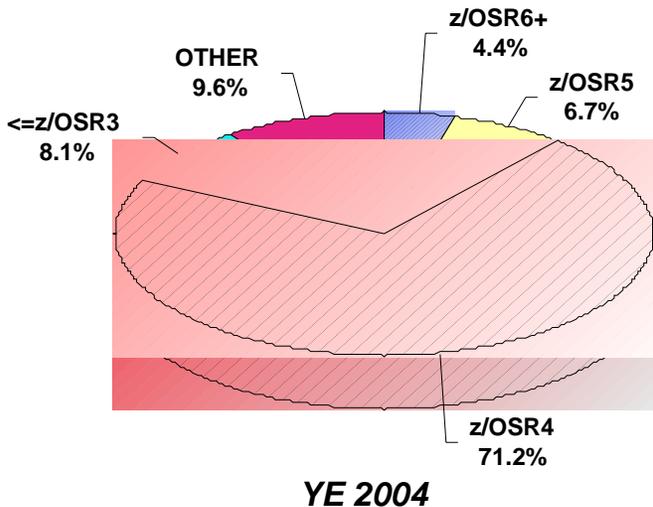


LSU Sweden



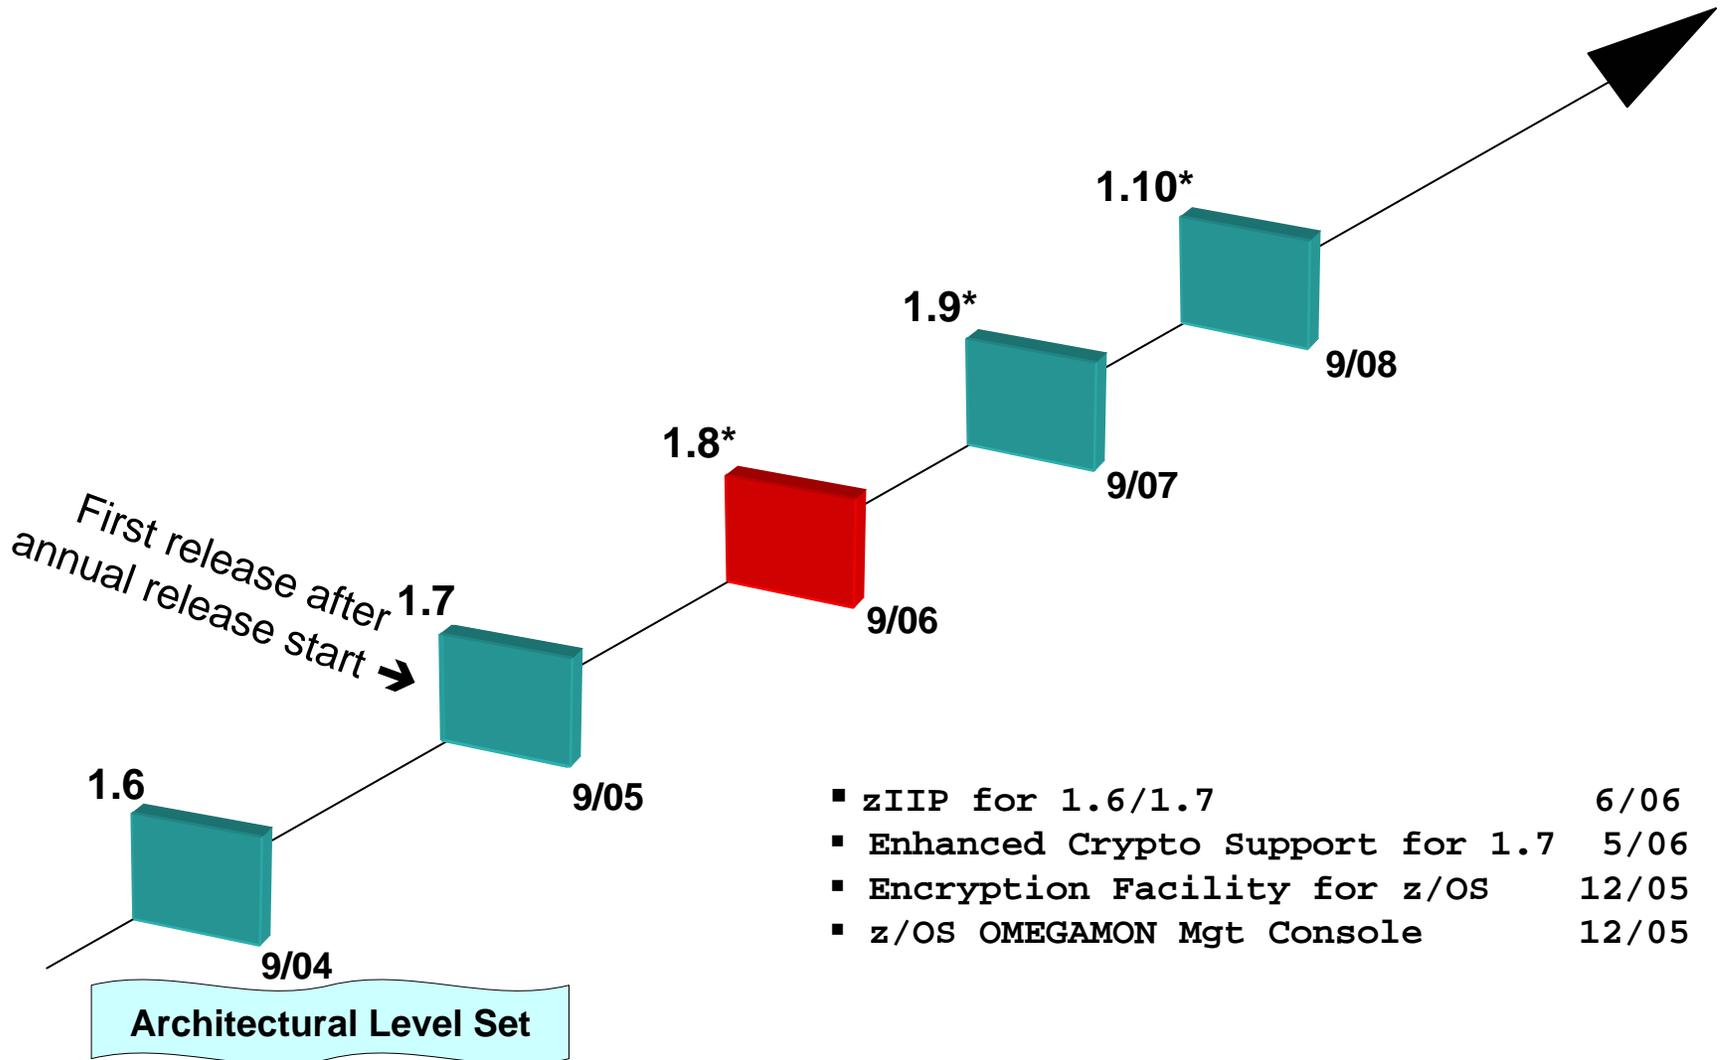
z/OS Release Distribution 2Q 2006

Over 93% on supported releases



**% 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



		G5/G6 Multiprise® 3000	z800	z890	z900	z990	z9-109	DS8000 DS6000	End of Service	Coexists with z/OS...	Planned Ship Date
z/OS	1.1	X	X		X				3/04	1.4	
	1.2	X	X	X ^C	X	X ^C			10/04	1.5	
	1.3	X	X	X ^C	X	X ^C			3/05	1.6	
	1.4	X	X	X ¹	X	X ¹	X ¹	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 coexist with z/OS 1.9

x^C - **Compatibility support only**

x¹ - **z990 compatibility or exploitation feature required**

←→
3 year support,
3 consecutive
Release coexistence
↑↓

* 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 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
- 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 for 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, ...

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 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 ...

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, ...

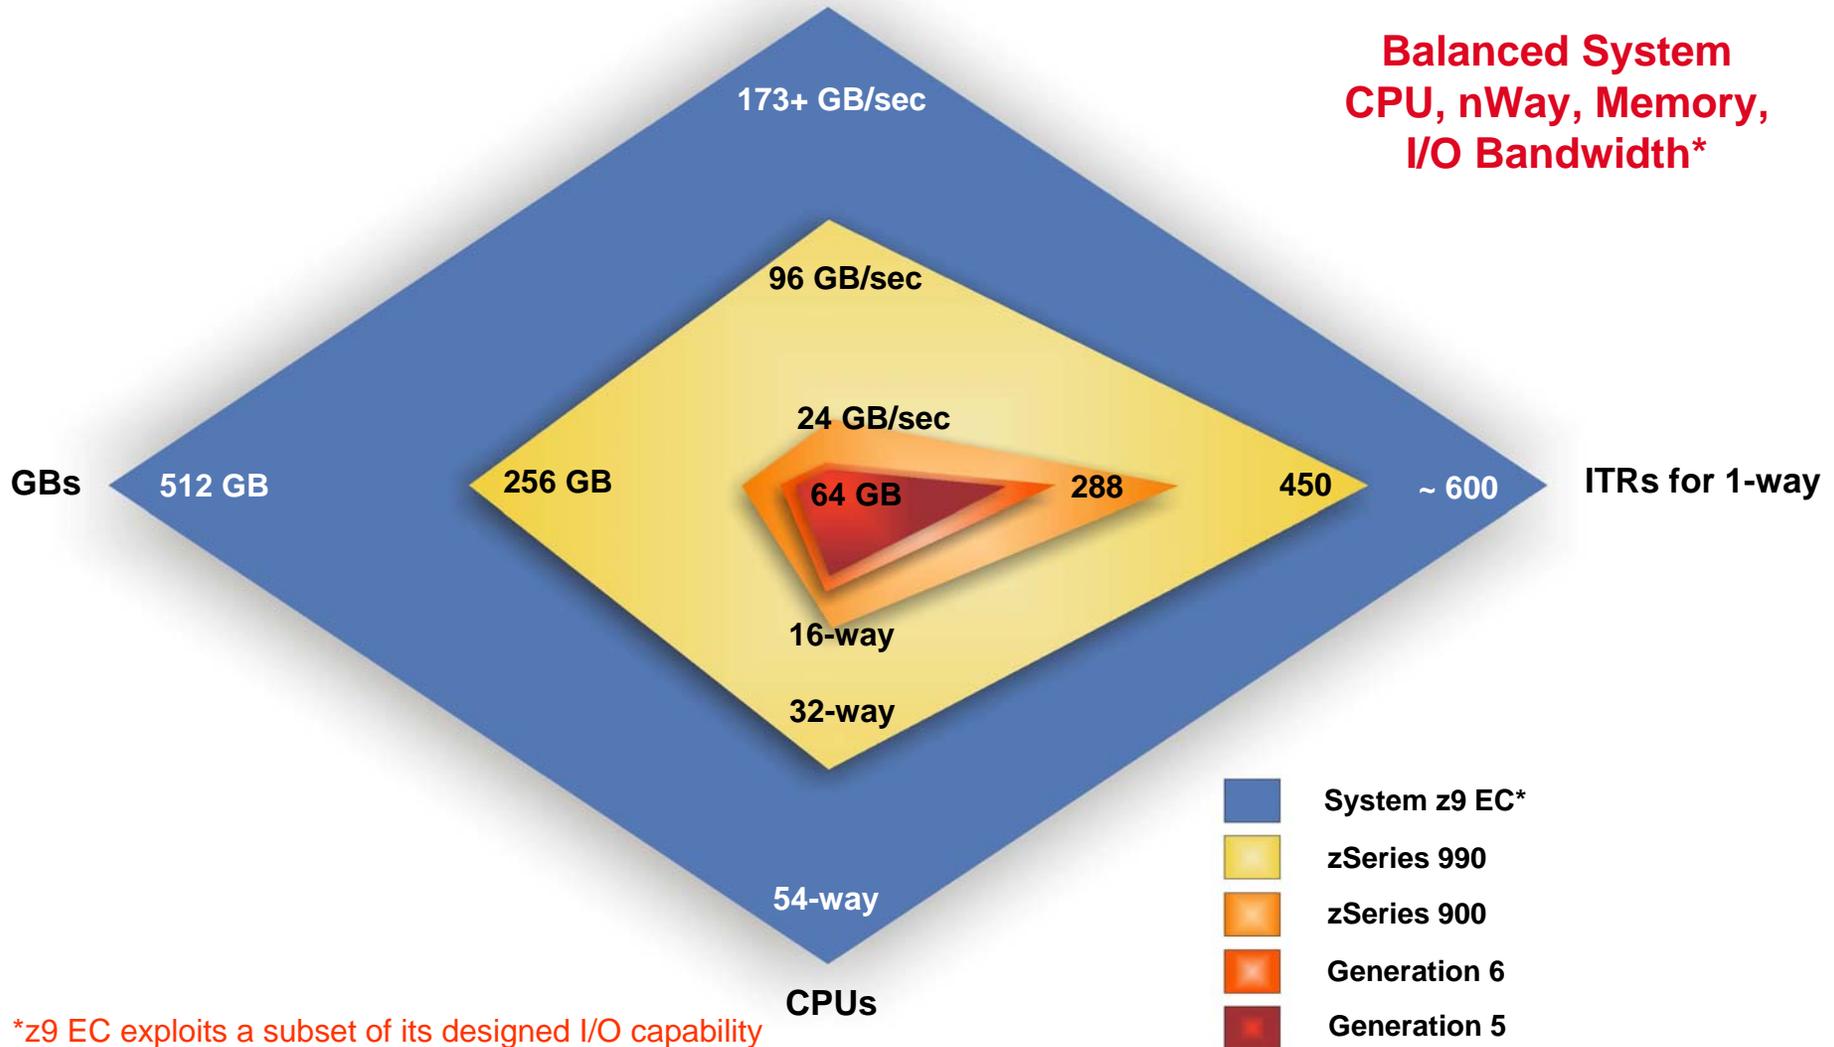
...

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 250k; can specify higher max than defaults
 - **New maximum is 2G**
 - New ISGADMIN authorized interface, T GRS command, GRSCNFxx parmlib support for setting higher maximums
- GRS 31-bit constraint relief
 - **Move star-mode global QCBs and QCBS control blocks above 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: 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 FILEPROCMAX 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**

z/OS Scalability

more performance: LE, VARY processing, PDSE, CFRM

- 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
- LE support for sequential data sets larger than 64K tracks and for VSAM extended addressability data sets
 - QSAM support in LE for C/C++ programs (using noseek)
 - Support for ESDS, KSDS, RRS extended format data sets with extended addressability
- PDSE exploitation of 64-bit addressing
 - New SMS init 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
- CFRM performance enhancements Stage 2
 - Designed to allow more systems, structures, and connectors to be added **without availability impacts**



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

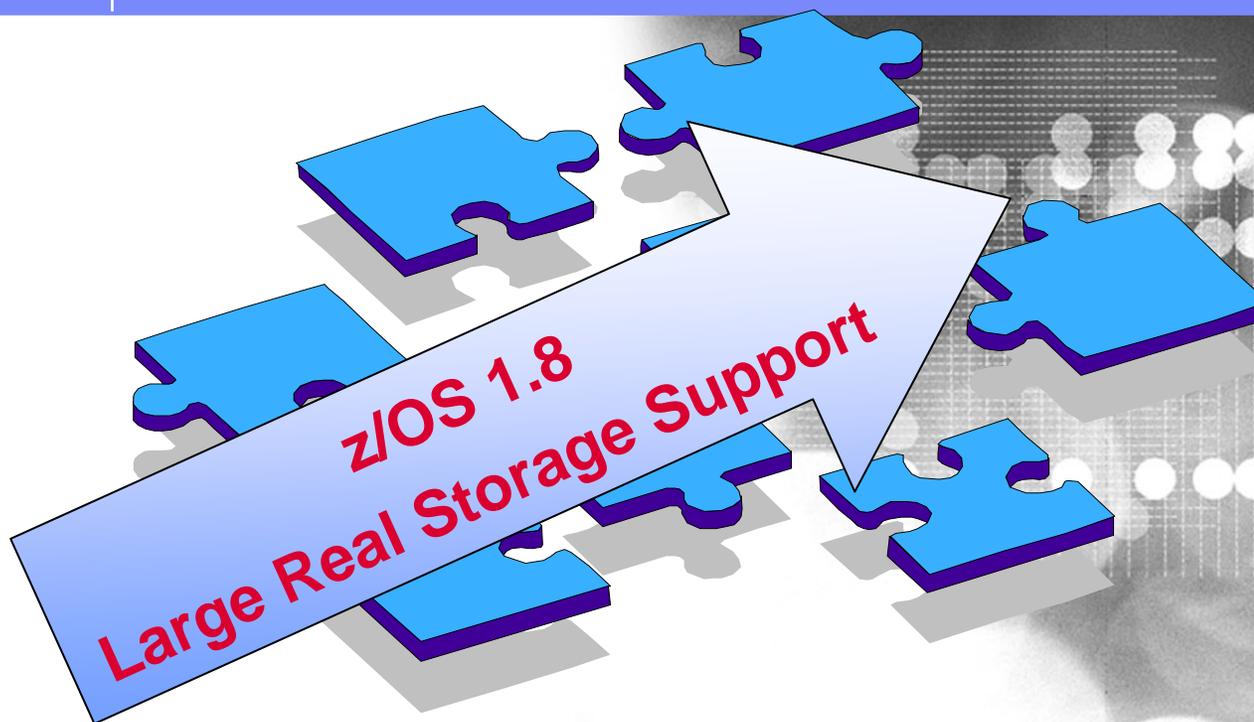
Stage 2 – z/OS 1.8

- **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.



LSU Sweden

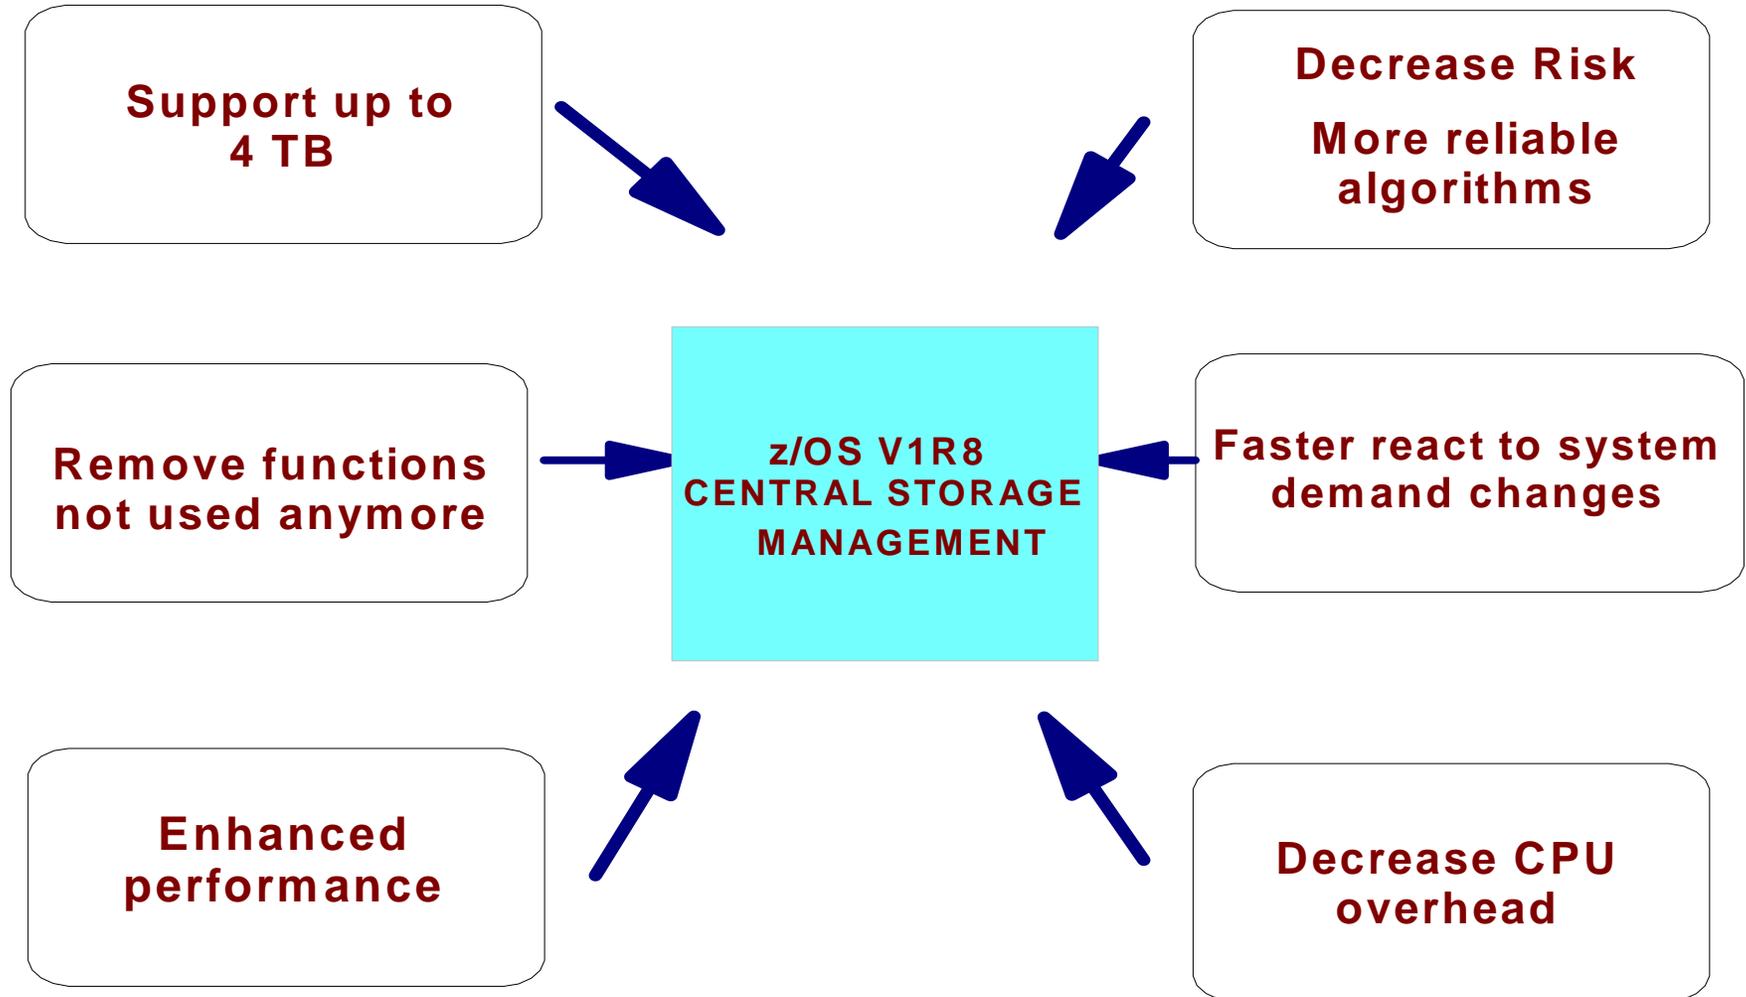


z/OS 1.8 RSM/SRM Support for Large Real storage *reasons why?*

- Support up to 4TB 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**.
- Pageable Storage Shortages
 - Used to be resolved **by physical swaps**
 - Often RSM **spin loops during swap-in**
 - Are now resolved by **frame exchanges**
- Physical Swap Processing
 - Physical **Swaps** to Auxiliary storage **eliminated**

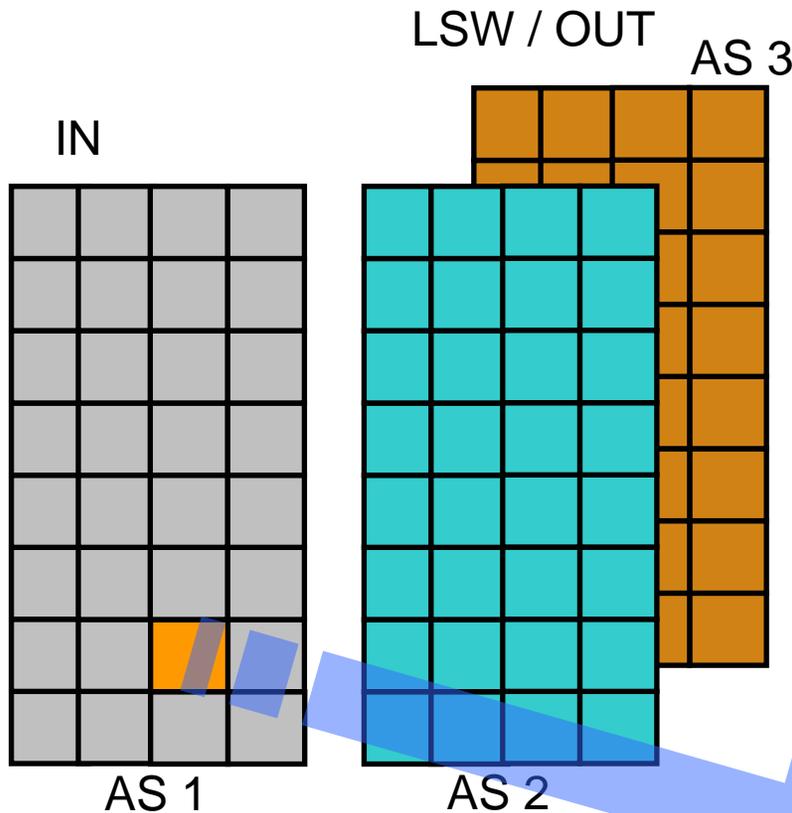
z/OS 1.8 RSM/SRM Support for Large Real storage

...and the expected benefits

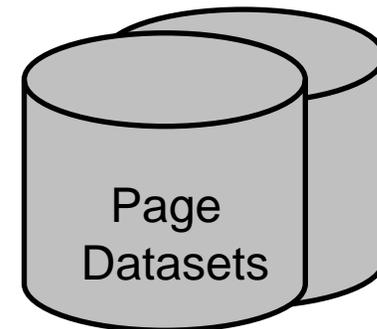


Page Replacement

prior to z/OS 1.8



- Today's LRU algorithm
 - Implemented by keeping an Unreferenced Interval Count (UIC)
 - Available Frame Queue (AFQ) Low/OK thresholds
- UIC update process runs **periodically** updating the UIC **of each in-use frame, costs CPU and holds locks**
- Stealing** starts when below **AFQ Low** and continues until **AFQ OK** reached
- Stealing happens on an **AS basis**
 - Oldest frames are stolen first



UIC Calculation

a long and tedious process performed irrespective of storage demand

Runs every 10 Seconds!

Look at each frame and calculate a UIC

Pageable Frame Queue of AS-1



Pageable Frame Queue of AS-2



Look at each frame and calculate a UIC

Pageable Frame Queue of AS-3



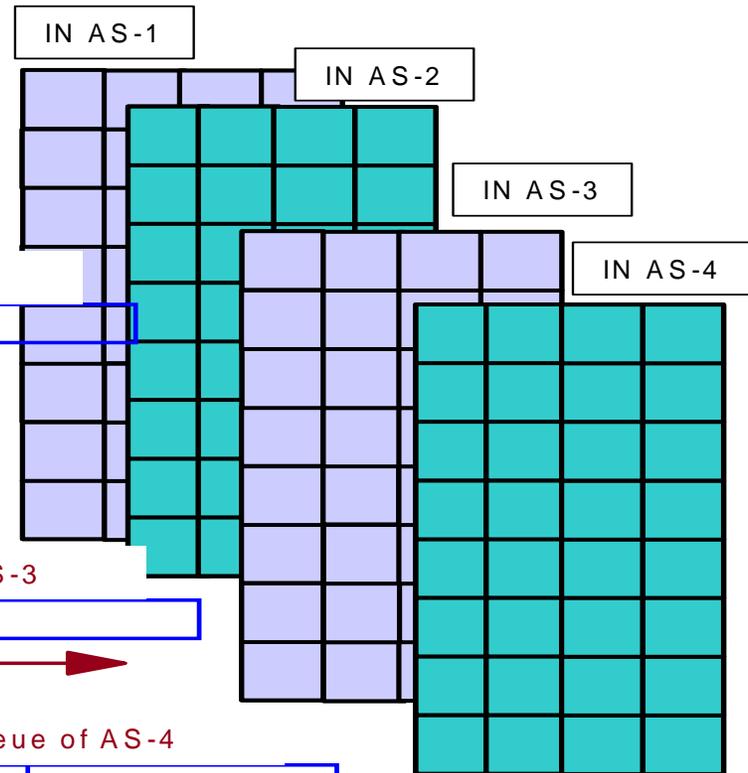
Look at each frame and calculate a UIC

Pageable Frame Queue of AS-4



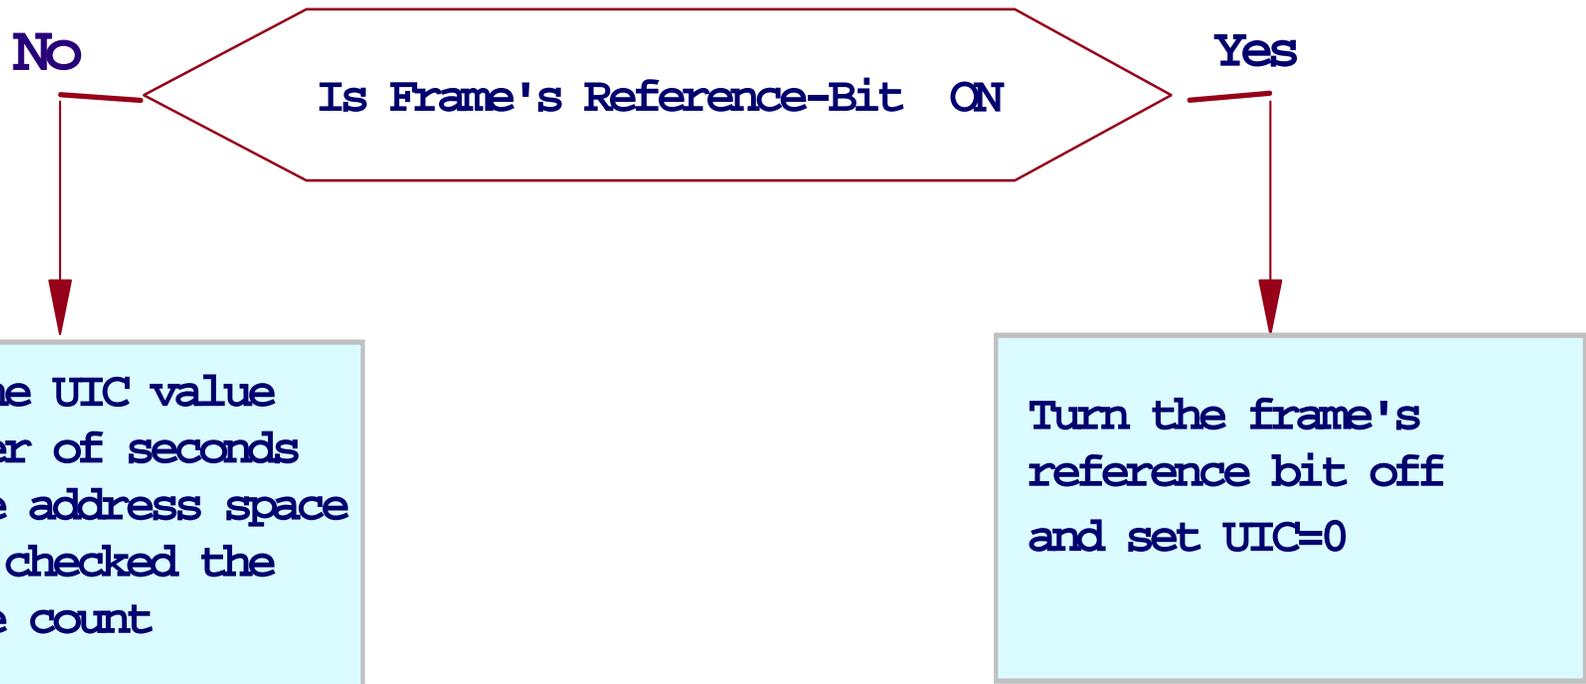
Look at each frame and calculate a UIC

LOOK in EVERY IN-AS !



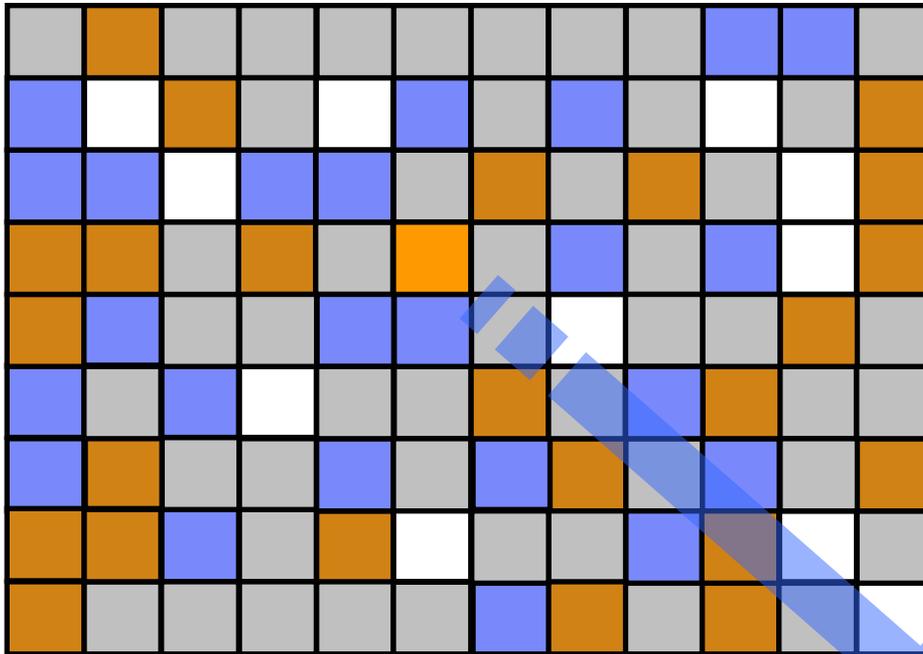
...

UIC calculation details the “old” way



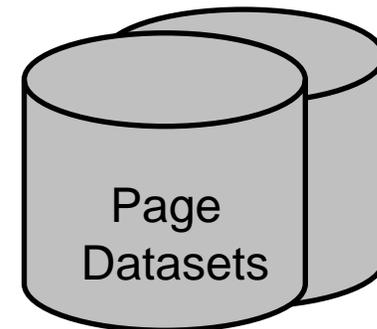
Page Replacement

z/OS 1.8 and beyond



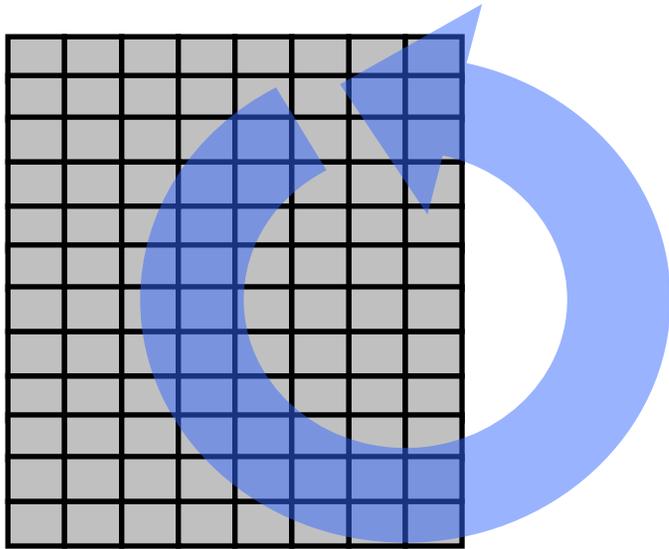
Why Change?

- Disruption of the UIC update process **intolerable as amount of RS increases**
- Enhanced paging algorithm:
 - Runs only when **AFQ needs** to be **replenished**
 - Stealing** happens on a **global basis**
 - By exchanging frames



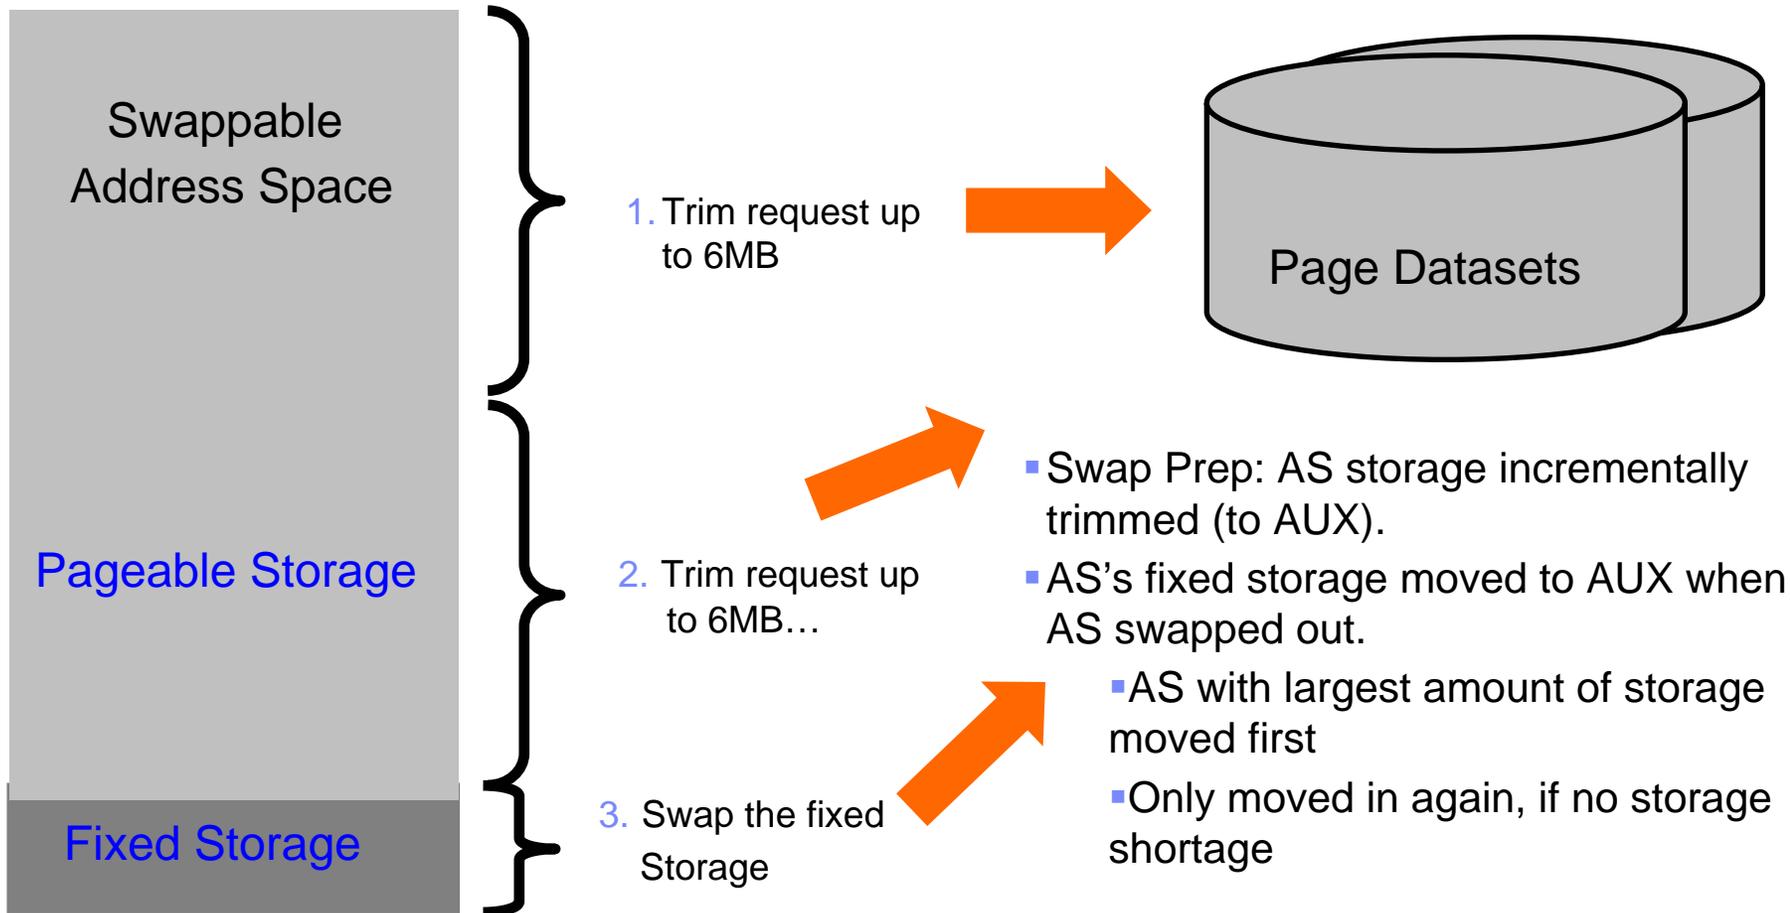
UIC Calculation

z/OS 1.8 and beyond



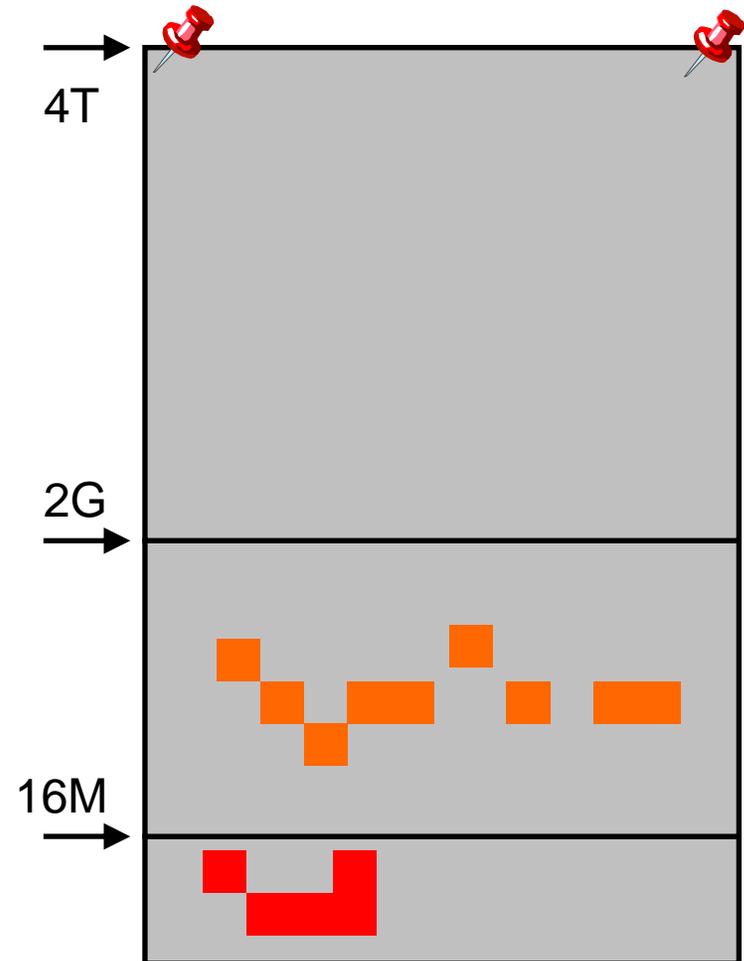
- Page replacement algorithm enhanced to efficiently process large amounts of RS
 - z/OS 1.8 defines UIC as a single walk through the whole storage in secs.
 - As a result the UIC values in RMF reports vary from **0–65535 (18 hrs)**
 - in the past the UIC value range was **0-2540.**
 - **High UIC value => less contention for storage**
 - **Very low UIC => storage constrained**
-
- 3 UICs displayed by performance monitors:
 - **Current UIC** (calculated every sec)
 - **Minimum UIC** (last walk through)
 - **Maximum UIC** (last walk through)

Physical Swap - pre z/OS 1.8



Pageable Storage Shortages – z/OS 1.8 ...

- Pre-1.8 pageable storage shortages were **resolved via physical swaps**
- z/OS 1.8 **“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 exchanged with frames above 16M or 2G preferably
 - Message IRA40xx lists the five largest users of fixed frames in the shortage area and criticality of shortage



Migration & Coexistence Considerations RSM/WLM

- Analyze your pre-1.8 system and pay attention to **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 ASes on OUT Ready or OUT Wait queue, you should consider increasing the amount of real storage




```

      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

-----
      IN      IN      OUT      OUT      LOGICAL  LOGICAL  BATCH  ADDRESS SPACE TYPES
      READY   IN      READY   WAIT   OUT RDY  OUT WAIT  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
-----
  
```

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

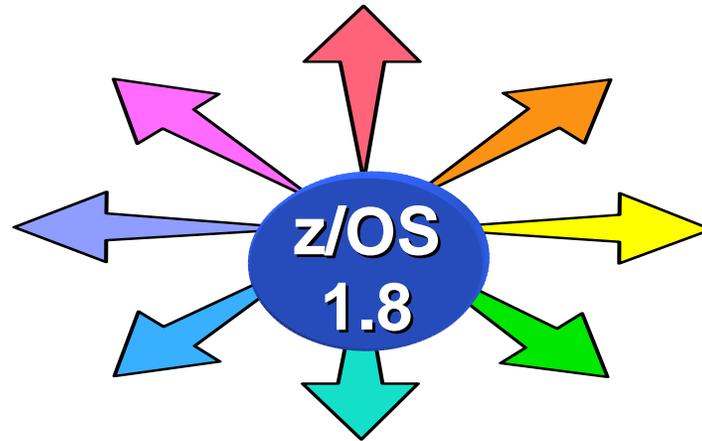
Sysplex partitioning support, Dynamic DNS registration, JES3 NJE via TCP/IP, TN3270 and telnet improvements, IP filtering, IKE, IPsec, WTS, Samba, improved sysplex failure recovery, REXX FTP interface, ...

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 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 ...

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, ...

Some z/OS 1.8 Availability Enhancements at a Glance

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



Console Restructure - Summary



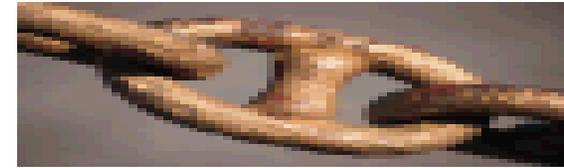
Old news

- Stage 1a (z/OS 1.5) solved msg. delivery problems
 - **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):
 - Internal RAS / Problem determination enhancements
 - **Ability to delete EMCS consoles**
 - Operational changes in preparation for stage 2 (z/OS 1.7 and **1.8**)
- Stage 2 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
 - Future z/OS release; no “rollback” planned

z/OS 1.8 Console Enhancements

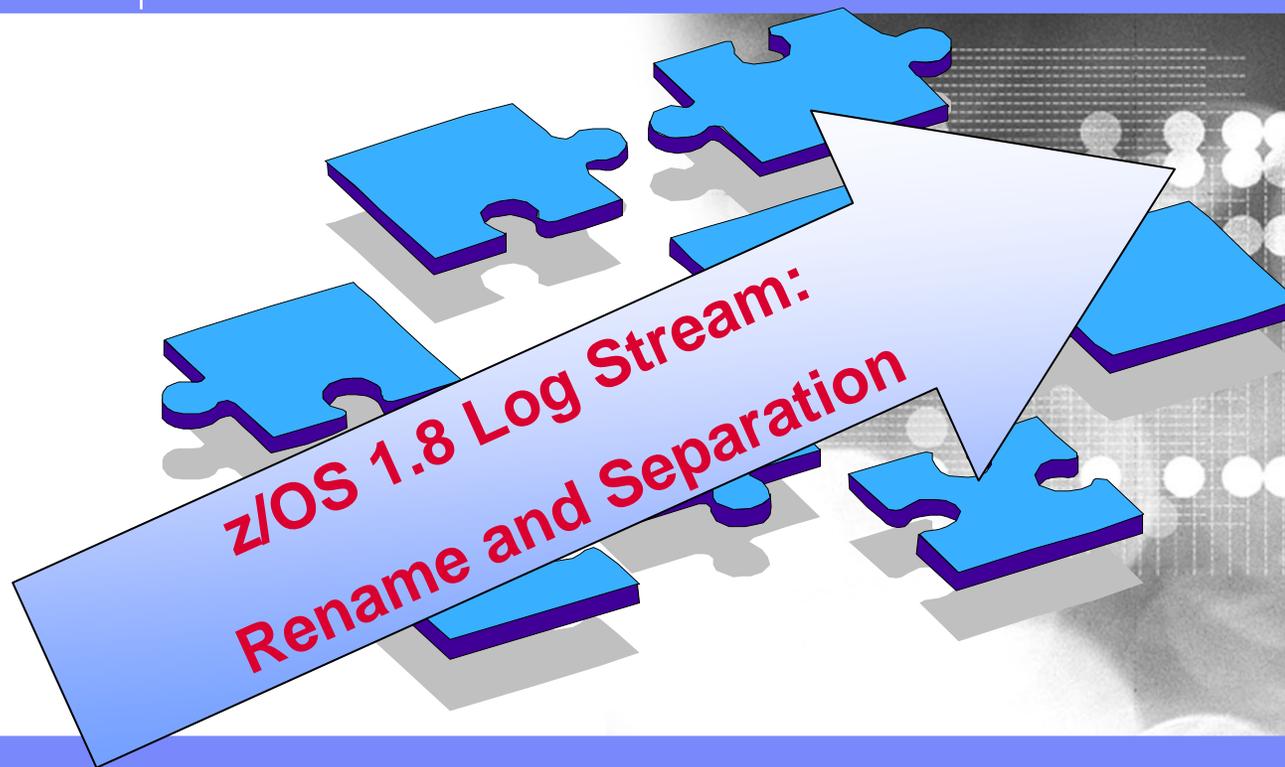
Master Console Elimination, 1-byte consoles now history

- Master Console elimination
 - **Removes a single point of failure for improved availability**
 - 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





LSU Sweden



z/OS 1.8 System 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 System Logger 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 Sweden

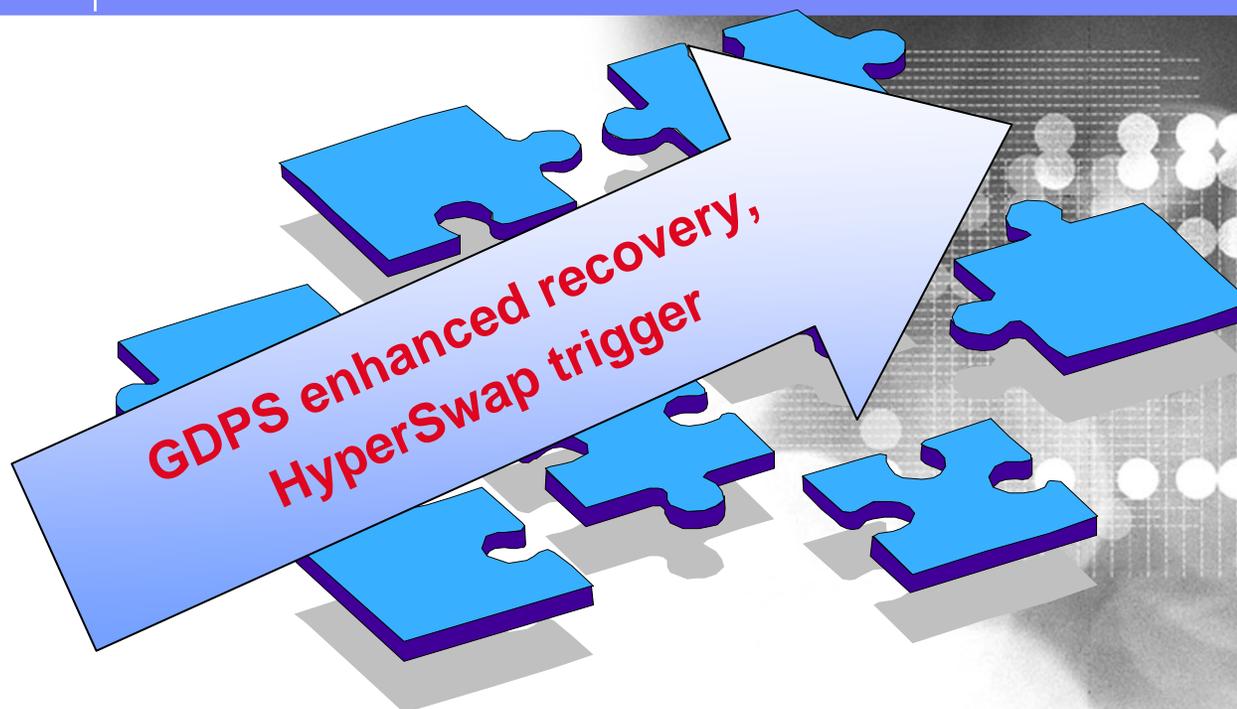


z/OS 1.8 Recoverable Byte Range Lock manager (BRLM)

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**

LSU Sweden



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, time-consuming 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**

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

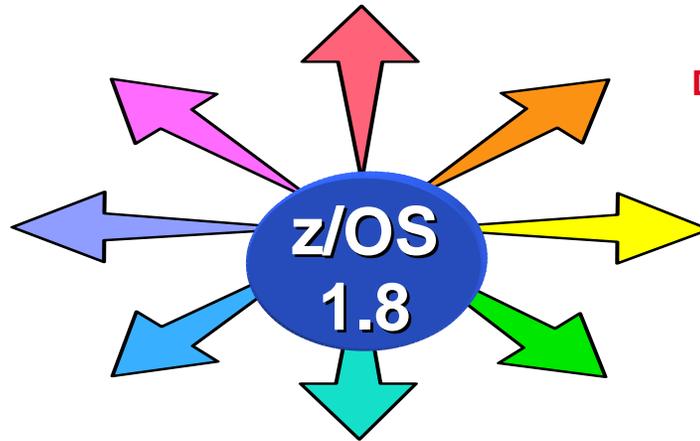
WLM, zAAP Stage 2, Sysplex Distributor, Tape I/O priority, JES2 balancing of WLM-managed initiators, GRS CNS processing, Improved OAM BLOB support, SCLM improvements, Group Capacity Limit ...

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, ...

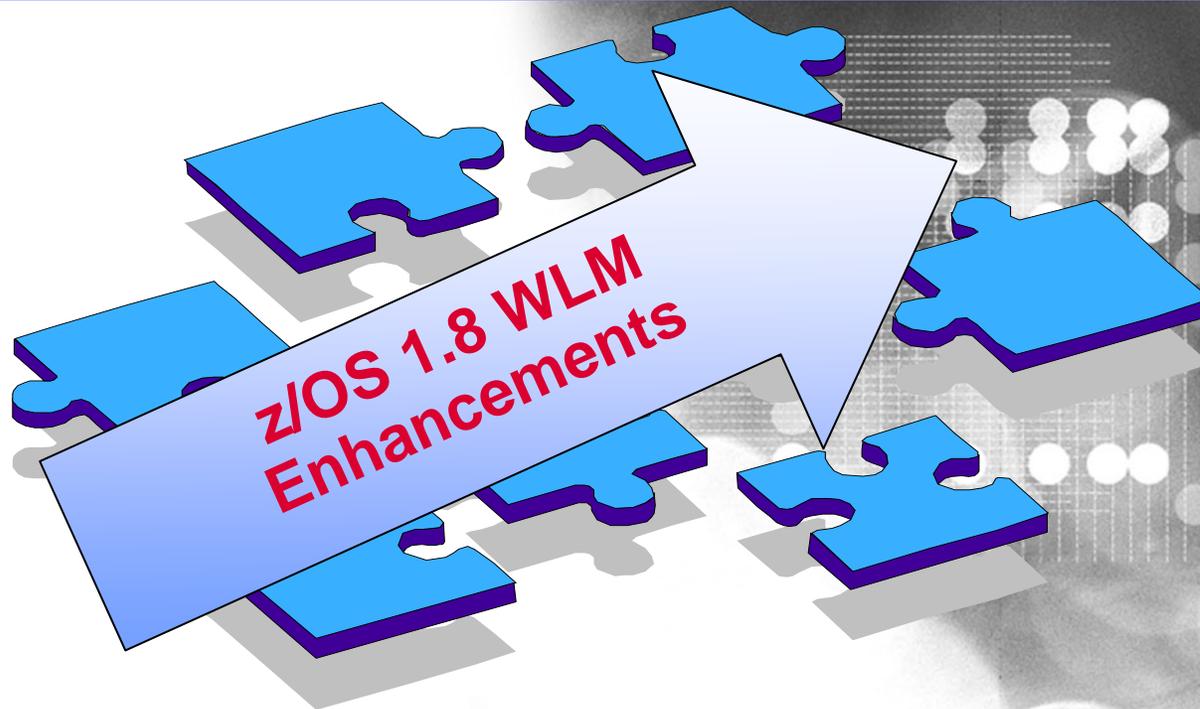
Enterprise-Wide Roles

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





LSU Sweden



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 again)
- New Resource Group types
- Routing enhancements
- User Friendly Interface for WLM Admin Application
- WLM support for DB2 Buffer Pool automatic adjustment *
 - To be picked up by future release of DB2
- WLM/IOS/DASD (DS8000) HyperPAV co-operation preview *

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

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 if the zAAPs become unable to process all queued work Notice: regular CPs are asked for help therefore not all CPs may process zAAP work at the same time
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”**

IFAHONORPRIORITY Processing the “old” way

- IFACROSSOVER=YES **is needed** to specify IFAHONORPRIORITY=YES
 - When IFAHONORPRIORITY=YES was specified:
 - At each dispatch event, general processors evaluated the zAAP eligible work and the priority of jobs is considered
 - When the highest priority unit of work is zAAP eligible work, the general purpose CPs dispatched it - **This results in more zAAP work run in general purpose CPs although zAAPs did not need any help**
- When there were no zAAPs online **you may need to set both parameters to YES**

IFAHONORPRIORITY Processing the “new” way

- IFACROSSOVER=YES is no longer needed to specify IFAHONORPRIORITY=YES
 - They are not related to each other any longer
- If you specify IFAHONORPRIORITY=YES (the default) in the IEAOPTxx parmlib member
 - Indicates that standard CPs may execute both Java and non-Java work in priority order
 - If zAAP processors are unable to execute all zAAP-eligible work
- If you specify IFAHONORPRIORITY=NO:
 - Then zAAP-eligible work can execute on standard CPs, but at a lower priority than non-Java work

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

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.**

Preview: WLM Group Capacity Limit on IBM System z9

- IBM plans to enable an **LPAR group capacity** limit on System z9 and later allowing you to manage a group such that capacity used by the LPARs will not exceed a capacity limit.

- You define: `group_name` & `group_limit`:
 - On HMC/SE
 - **Group Limit** independent of **Defined Capacity**. Both limits can be defined and both work together
 - 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; 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 LPARs and LPARs with **wait completion YES** ignored

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 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.

JES2 Improved Batch Initiator Balancing

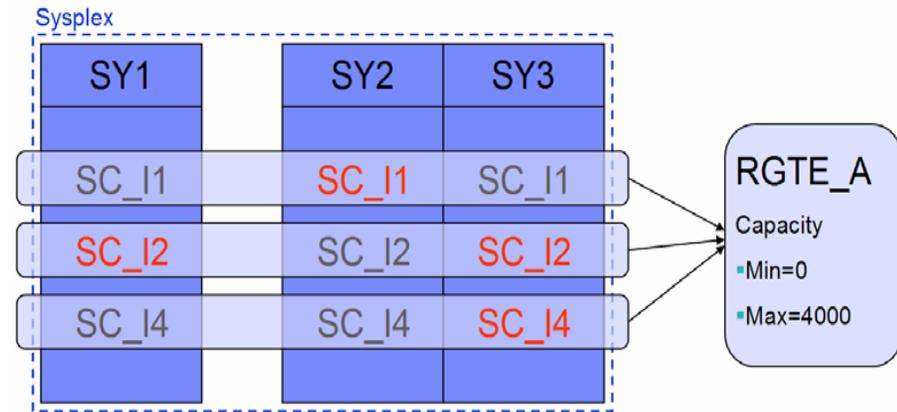
Problem prior to z/OS 1.8

- 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

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



Existing WLM Resource Group Concept...

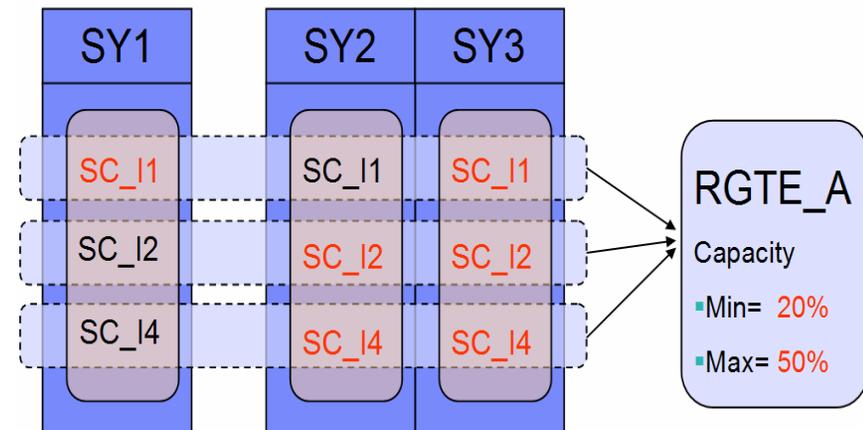
potential skewed CPU usage across systems

- WLM exchanges information in the sysplex about the CPU consumption of all Service Classes in the Resource Group
- WLM attempts to distribute the available capacity to the service classes based on their goal achievement
- This behavior can cause a Service Class to get a **very high part of CPU usage on some systems, and less or no capacity on others in the same sysplex**
- The existing Resource Group (now type 1) concept is:
 - Sysplex-wide defined in service units
 - Sysplex-wide managed

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
 - # Logical CP 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



Sysplex Routing: Enhancements

z/OS 1.4 and z/OS 1.7 enhancements



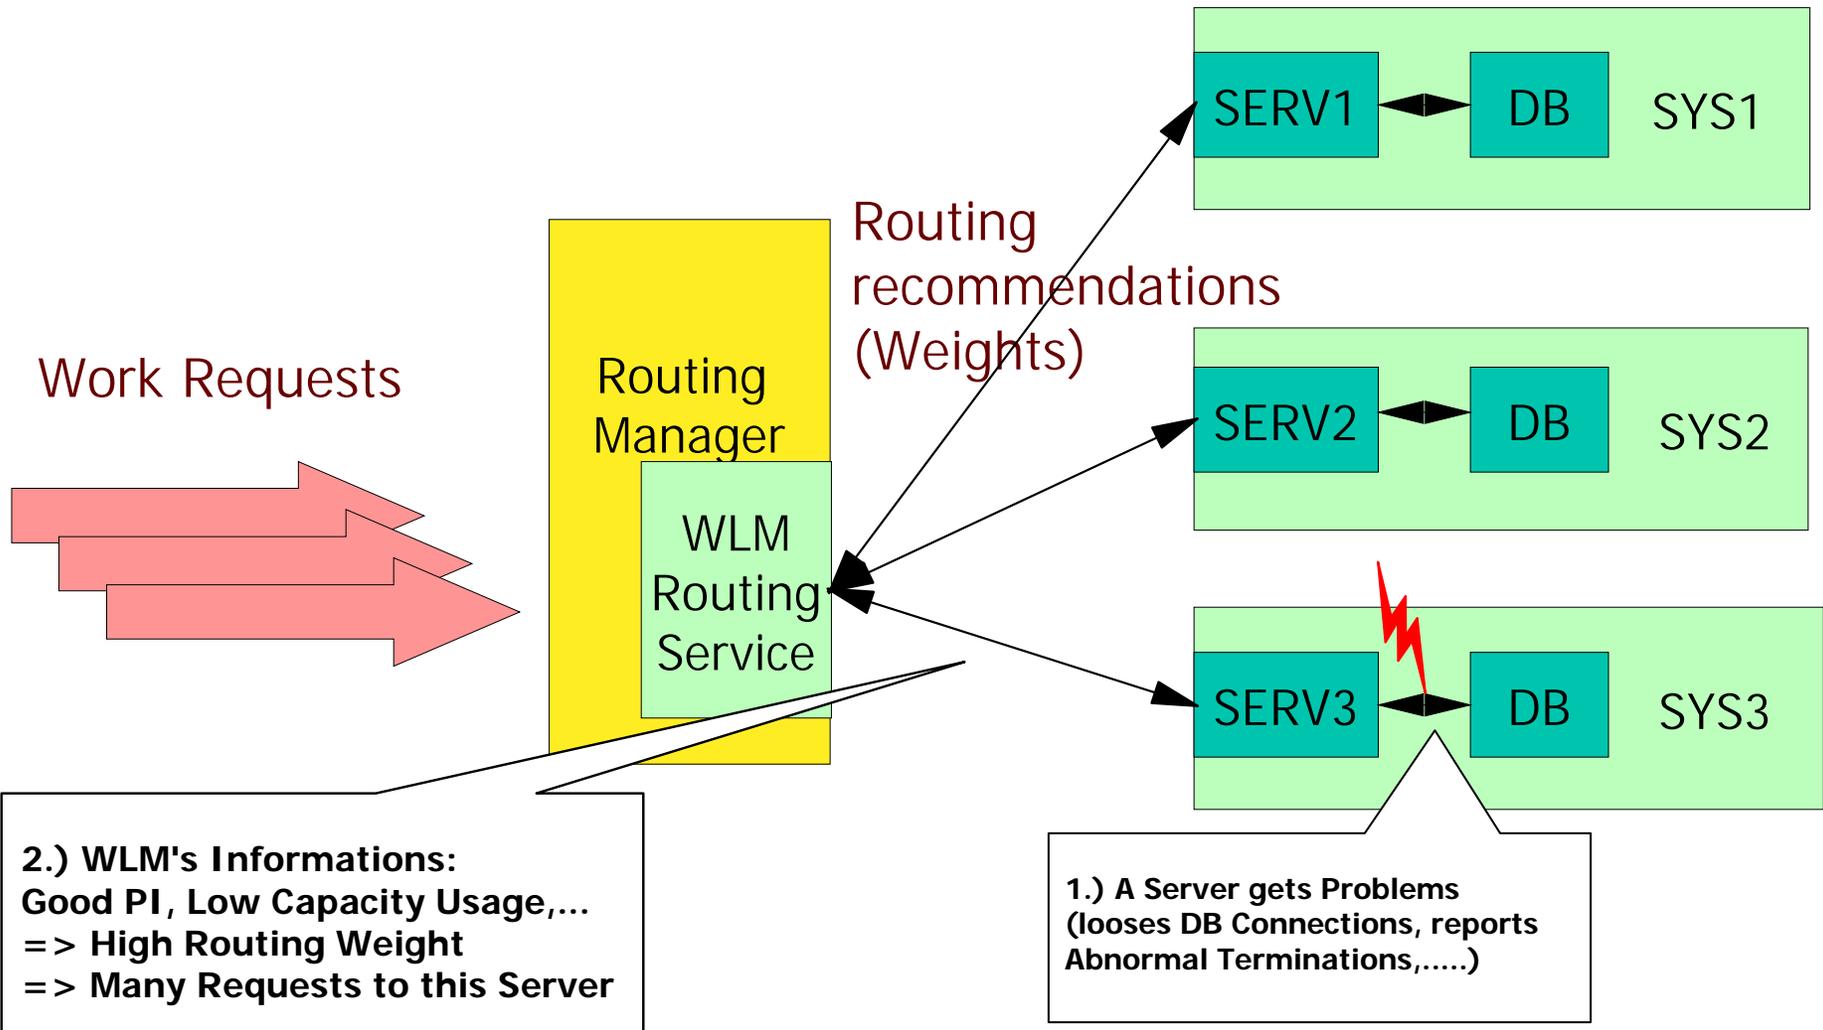
Old news

- **Current recommendation was 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** May be exploited by DB2 and others in the future
- **TCP/IP only registers the stack and does not know how 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**
- **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**
- 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 **OA1006 on z/OS 1.4 and above**

Storm drain issue

- In some situations, a server with WLM routing recommendations starts processing the incoming requests **with error codes**
 - Each request **finishes very fast**
 - This leads to a situation in which this server has a **good PI and its LPAR usage is low**
 - When this occurs, WLM could **calculate a high weight** as a routing recommendation for that server
 - This would cause the routing service **to send many requests to that server**

Sysplex Routing: z/OS 1.8 Enhancements



Sysplex Routing: z/OS 1.8 Enhancements

WLM is better informed

- As WLM is not aware of abnormal conditions of a server, **it can make bad recommendations**
- New decision made by servers:
 - Inform WLM about abnormal conditions and WLM takes abnormal conditions into account as follows:
 - **Health value of a server**
 - **Abnormal termination rate of a server**
- By taking the abnormal termination rate and the health value of a server into account in routing recommendations, **the routing recommendations are more adequate in abnormal situations**

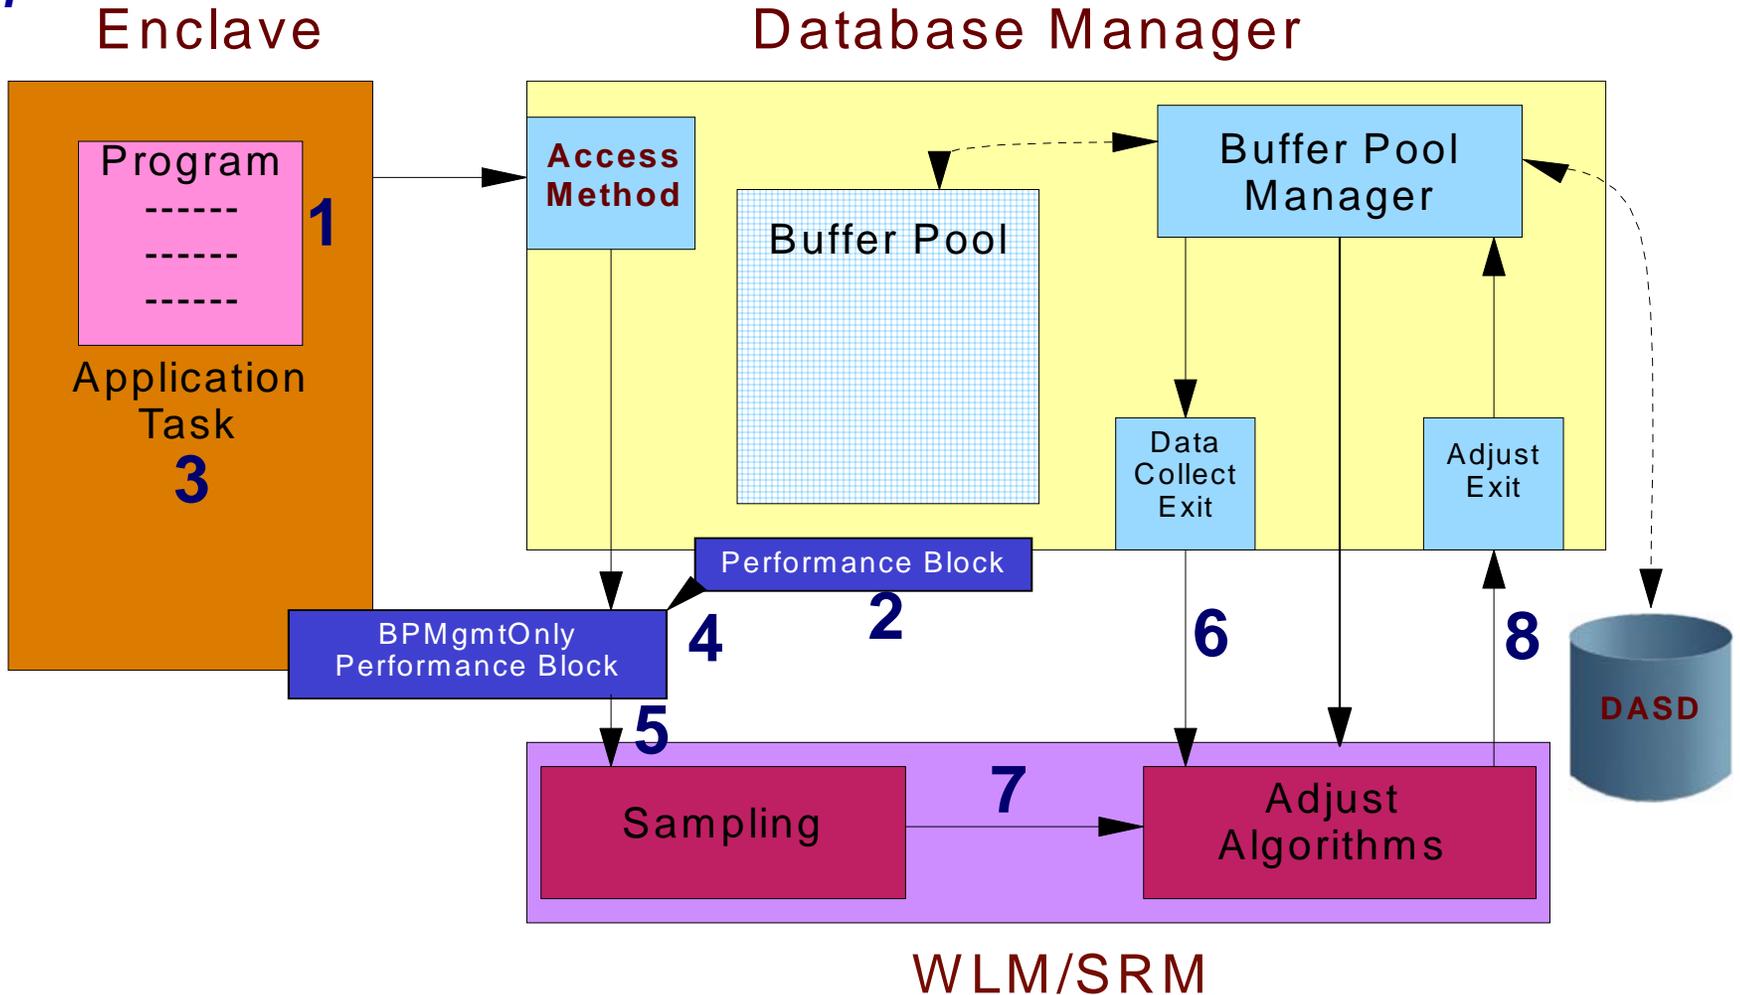
Planned exploitation

- The planned exploiters of these enhancements are:
 - DB2 V8 and up *
 - IMS *
- **Current exploiters**
 - CICS
 - Communication Server for z/OS 1.8
- For more detail information about services and enhancements see
 - z/OS: Planning Workload Manager, SA22-7602
 - z/OS: Programming: Workload Management Services, SA22-7619.

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

WLM support for DB2 Buffer Pool automatic adjustment

*preview **



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

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
 - Reduced DASD space required for SCLM listings due to compression

WLM z/OS 1.8 I/O Priority for Tape

- **SRM provides two I/O priorities for use by the (IOS)**
 1. Channel Path I/O Priority
 2. DASD I/O Priority
- **Both of them are dynamically calculated**
- **Cases when SRM does not provide any dynamic I/O priority for Tape I/O**
- **New I/O priority scheme for tape based on a static mapping of the importance**
- **The default for I/O priority management is no, which sets I/O priorities equal to dispatching priorities**
- **If you specify yes, workload management sets I/O priorities in the sysplex based on goals**

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

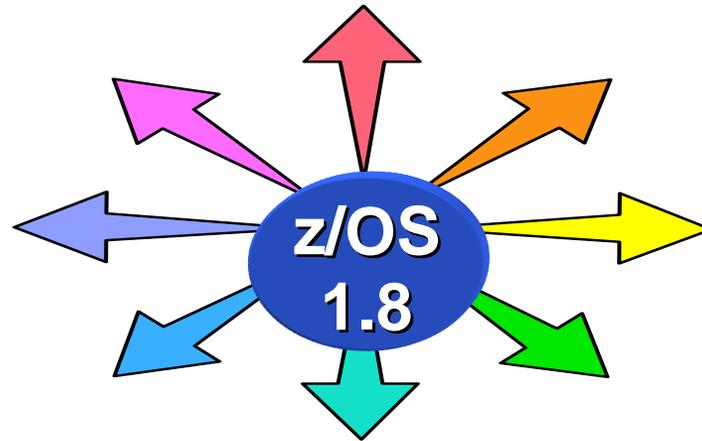
WLM, zAAP Stage 2, Sysplex Distributor, Tape I/O priority, JES2 balancing of WLM-managed initiators, GRS CNS processing, Improved OAM BLOB support, SCLM improvements, Group Capacity Limit ...

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

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

- zAAP reporting for 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

- Infoprint Server 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

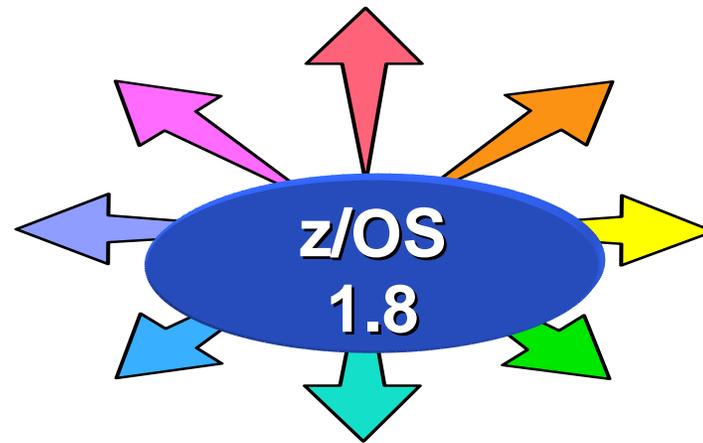
Support for up to 4 TB memory, New zIIP processor, CFRM performance improvements, GRS ENQ limits, Parallel VARY ONLINE, z/OS UNIX file descriptors, Device Groups, DADSM/CVAF Rapid Index Rebuild, z/OS USS asynch socket fast path, 64-bit PDSE support ...

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 ...

**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 ...

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

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

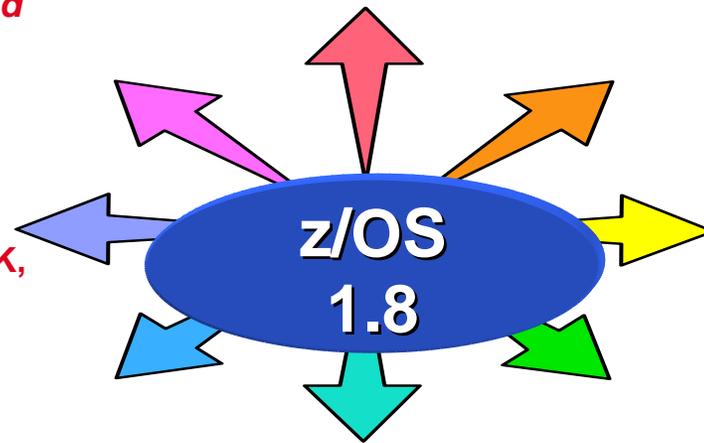
WLM, zAAP Stage 2, Sysplex Distributor, Tape I/O priority, JES2 balancing of WLM-managed initiators, GRS CNS processing, Improved OAM BLOB support, SCLM improvements, Group Capacity Limit ...

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 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

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.

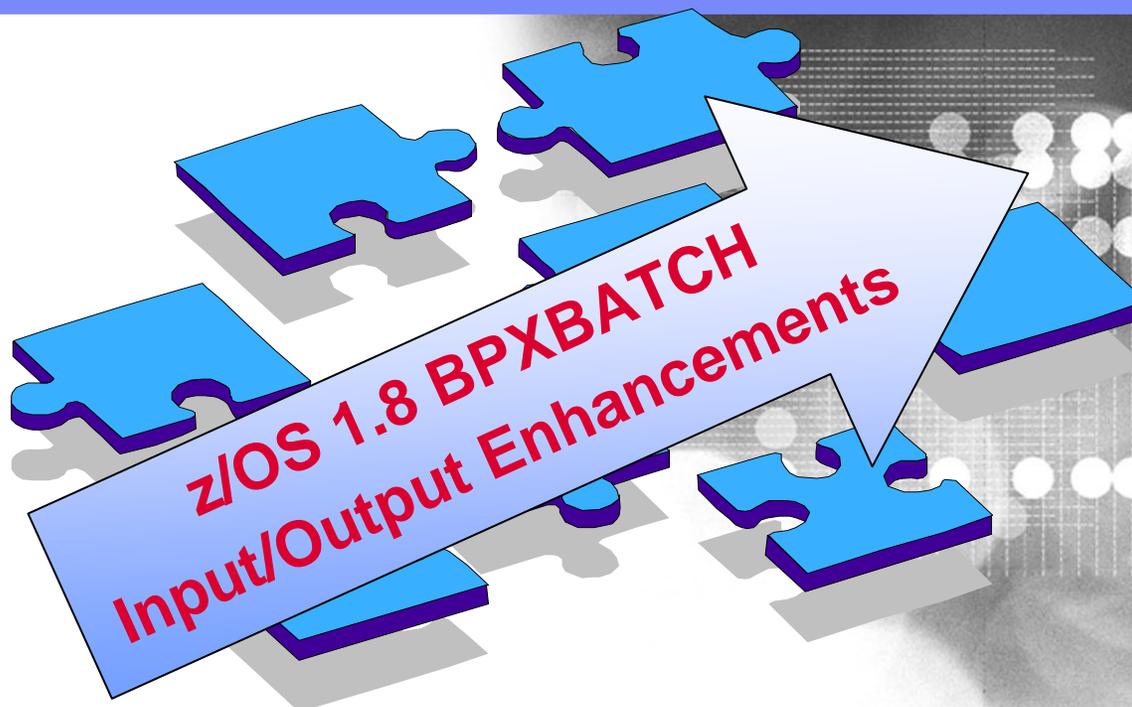
New ITDS LDAP server NOT available at z/OS 1.8 GA

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 Sweden



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 65,536 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)**





LSU Sweden



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





LSU Sweden



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 Sweden



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.

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 lang 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.

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

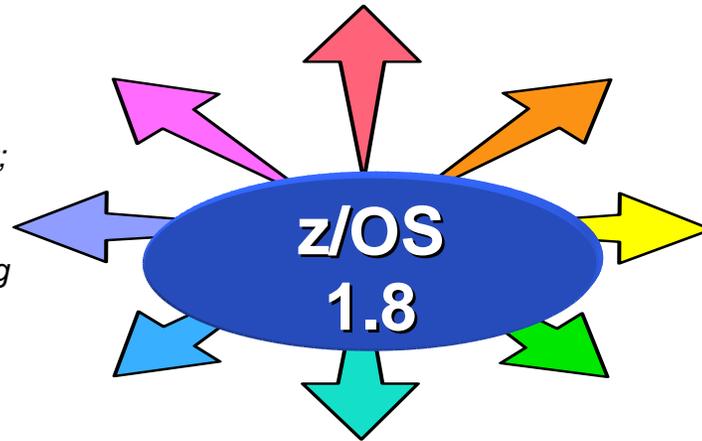
WLM, zAAP Stage 2, Sysplex Distributor, Tape I/O priority, JES2 balancing of WLM-managed initiators, GRS CNS processing, Improved OAM BLOB support, SCLM improvements, Group Capacity Limit ...

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



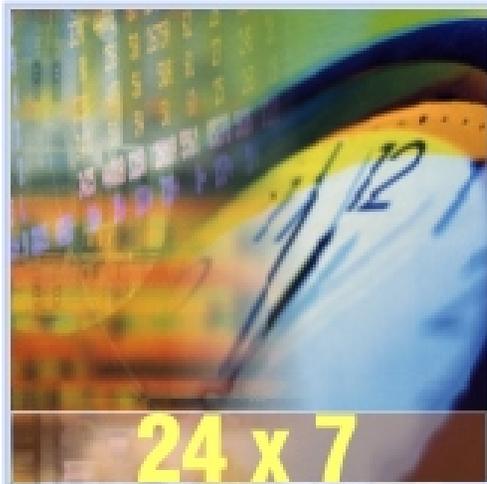
...

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

¹ Majority of these checks are available on prior releases

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

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

**Hundreds of orders
and downloads!**

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

The screenshot shows a Microsoft Internet Explorer browser window displaying a PDF document. The address bar shows the URL: <http://publibz.boulder.ibm.com/zoslib/pdf/zosbasic.pdf>. The browser's menu bar includes File, Edit, View, Favorites, Tools, and Help. The toolbar contains various navigation and utility icons. The left sidebar shows a 'Bookmarks' pane with a tree view of the document's contents, including sections like 'Front cover', 'Contents', 'Preface', and multiple chapters. The main content area displays the title page of the PDF, which features the IBM logo in the top right corner. The title is 'Introduction to the New Mainframe: z/OS Basics'. Below the title, there are three bullet points: 'Basic mainframe concepts, including usage and architecture', 'z/OS fundamentals for students and beginners', and 'Mainframe hardware and peripheral devices'. The background of the title page features a stylized red and white globe graphic. The bottom of the browser window shows the Windows taskbar with the Start button, several open application windows, and the system tray displaying the time as 10:48 PM and the date as 11/17/2006.

“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.

Today	Tomorrow
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; consistent across IBM
Years of z/OS experience required	Designed with novices in mind; but experts will benefit too

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 HW configuration)
- Workload management
- Storage and data management
- Security management

z/OS 1.8 Simplification Highlights

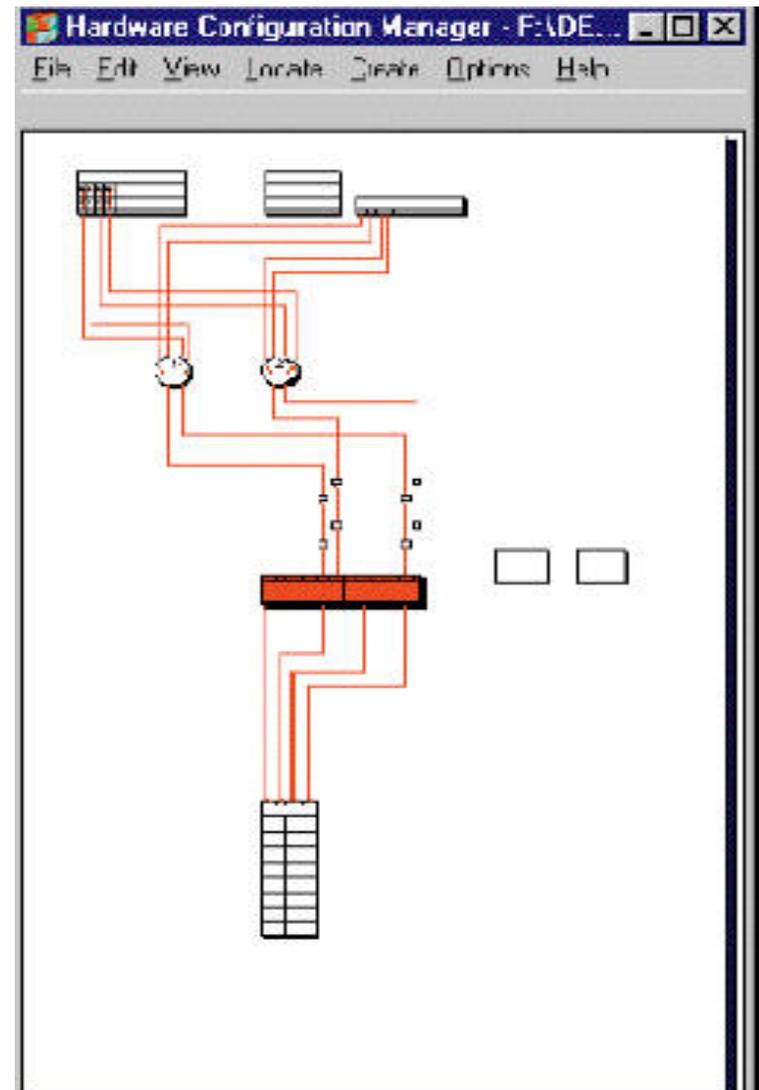
- IBM Health Checker for z/OS enhancements – more checks, easier way to write health checks
- IBM OMEGAMON z/OS Management Console 4.1.0 (available 12/06) – more monitoring workspaces
- Major Hardware Management Console enhancements
- IBM Configuration Assistant for Communications Server – support for 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.

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**

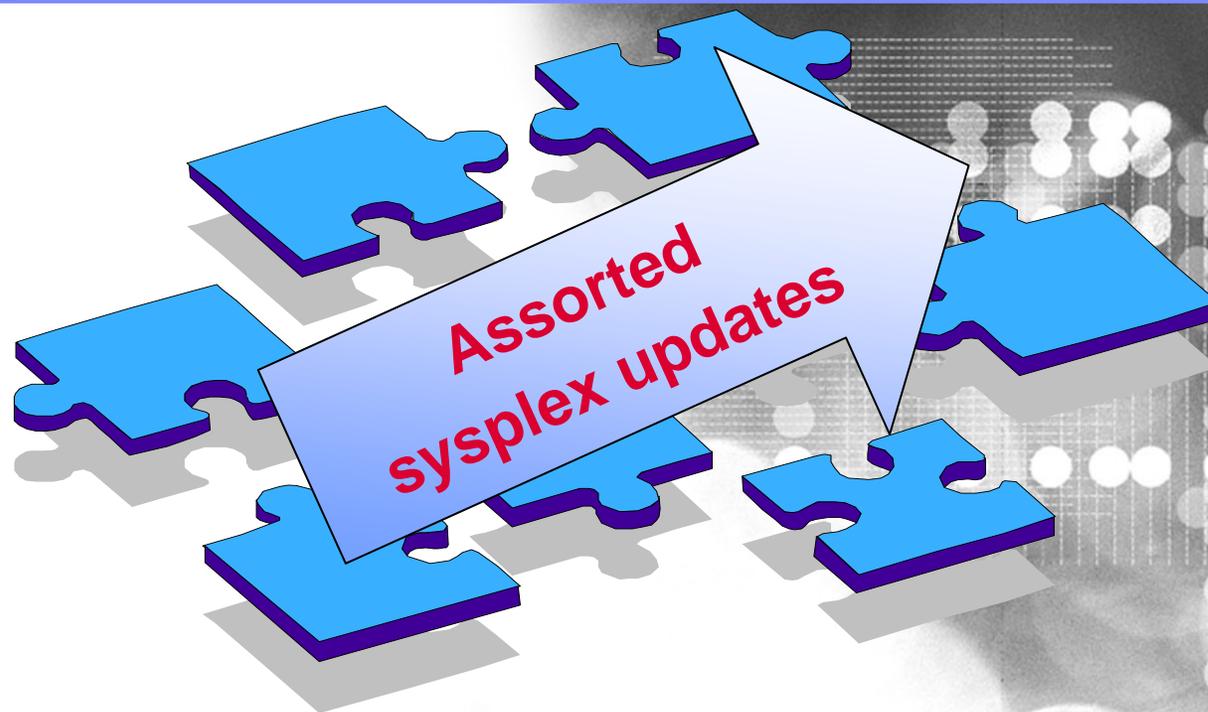


Academic Initiative Status

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 Sweden



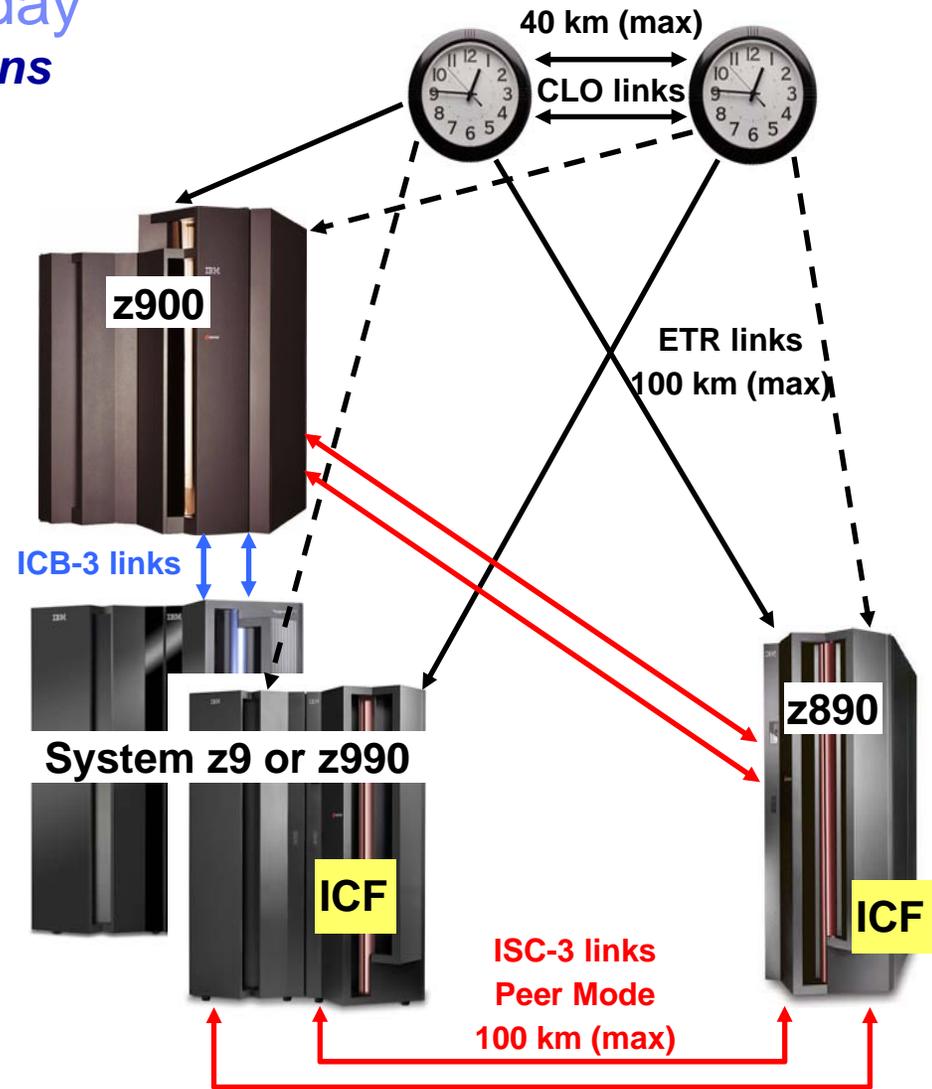
Coupling Facility Control Code – Level 15?

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

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 synch. 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

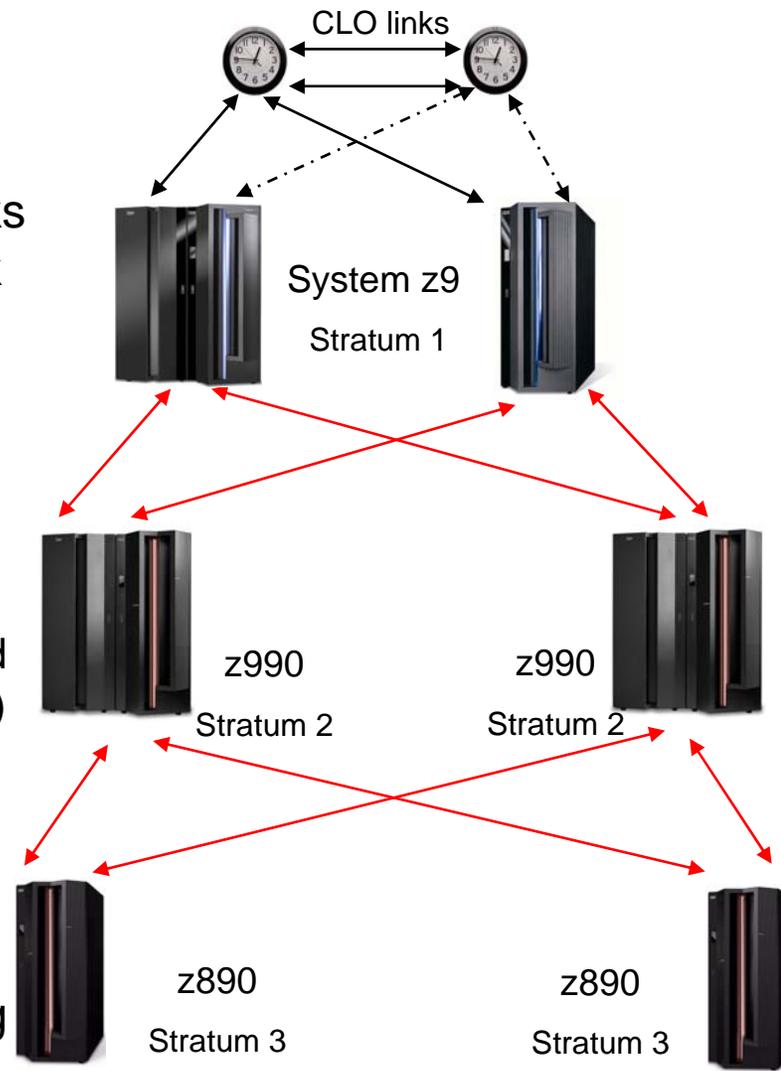
- Capability for multiple System z servers to maintain time synch. with each other:
 - **Does not require the 9037** Sysplex Timer if all servers STP capable
 - **Timing information transmitted over ISC-3 links (Peer mode), ICB-3 and ICB-4 links.** Supports a multi-site timing network of **up to 100 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

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

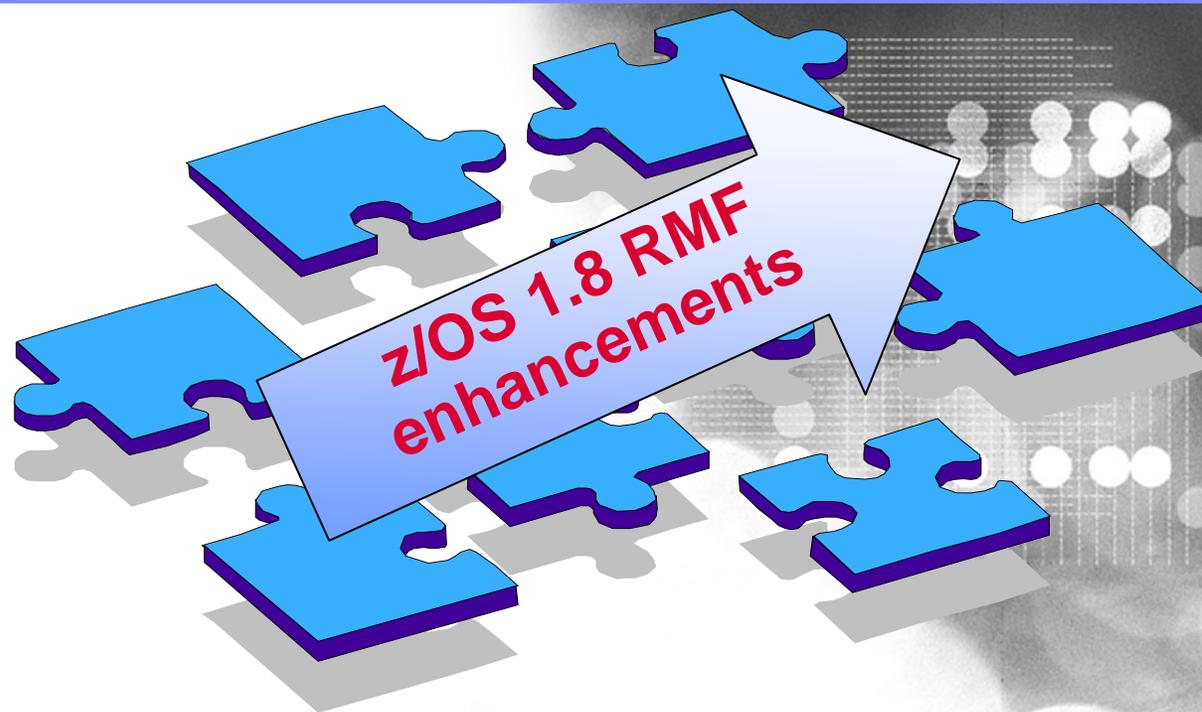
Changes to DST in the USA - 2007

may affect clients with world wide operations

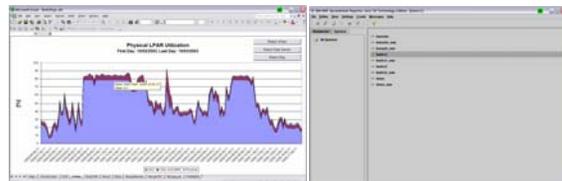
- **Flash10483**
- **The new USA Energy Policy Act of 2005, signed by the President in August 2005 has changed Daylight Saving Time (DST) in the USA, beginning in 2007**
- **DST Prior to 2007**
 - **DST began - 2:00AM first Sunday of April**
 - **DST ended - 2:00AM last Sunday of October**
- **DST in 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



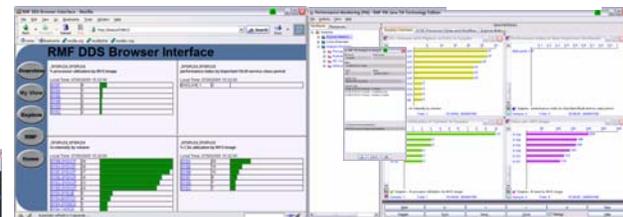
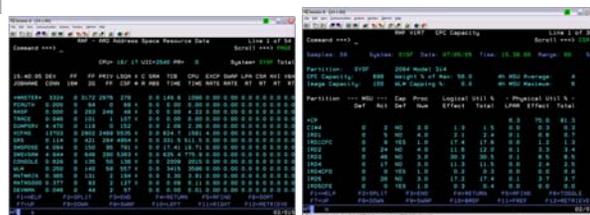
LSU Sweden



RMF Architecture Overview

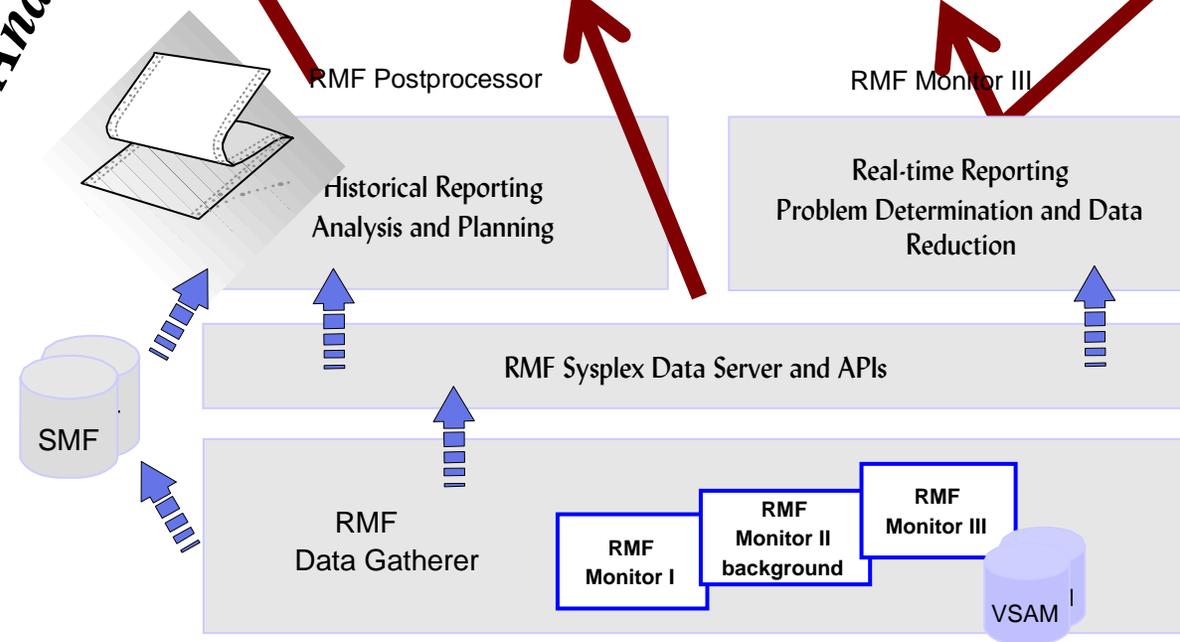


Snapshot Reporting



Long-term Analysis

Online Monitoring Asynchronous Events (CIM)



Monitor II (Snapshot) SRM report




```

RMF - SRCS Central Storage / Processor / SRM

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

    HI  SQA  LPA  LPA  CSA  L+C  PRI  LSQA  LSQA  CPU  IN  OUT  OUT  OUT
    UIC  F    F    FF   F    FF   FF   CSF  ESF  UTL  Q  LOG  RQ  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

Group Capacity Limits: PP CPU Report




```

      G R O U P   C A P A C I T Y   R E P O R T

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
                200    B           DOMSYS02    80   70    50   YES   0.0       67    80
                200    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

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

Meeting the Future challenges of on demand business*

Improving Usability and Skills

Still more health checks, ISPF improvements, DFSMSrmm usability improvements, Improved Configuration Usability for Communications Server

Integrating new Applications and Supporting Industry and Open Standards

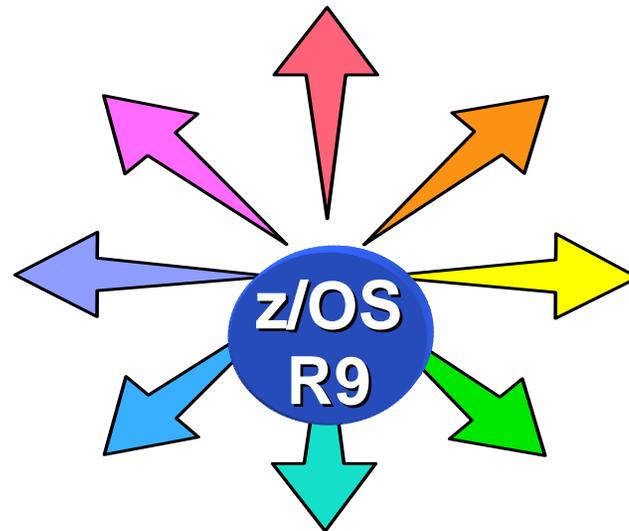
pthread enhancements, crypt and dcrypt support, z/OS XML System Services

Extending the Network

Improved TCP/IP Routing, Centralized Policy Based Networking, Expanded Network Encryption,

Scalability & Performance

64-way support, 64-bit GRS, zFS Sysplex, SMF to Logger, 24-bit VSCR, TSO/E support for large sequential data sets ...



Enterprise-Wide Roles

EWLM end-to-end transaction management, MQ Monitoring, CIM monitoring enhancements

Improving Availability

Improved latch contention detection, JES2 APF Mitigation

Self Managing Capabilities

WLM support for cross-system routing of zAAP workloads, WLM "Trickle" Support, Promotion of canceled jobs

Enhancing Security

Password phrase exploitation, Kerberos AES support, Enhanced CRL support, PKCS#11 support, Java user and group SAF admin classes

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

ITSO Workshops Week 48 - Copenhagen

- <http://www.redbooks.ibm.com/projects.nsf/WorkshopIndex?OpenView&Start=59>
- z/OS Version 1 Release 8, ITSO1DK **Open**
 - Starts 27 Nov 2006 for 1 day in Copenhagen, Denmark
- Parallel Sysplex Update and High Availability Topics, ITSO2DK **Open**
 - Starts 28 Nov 2006 for 1 day in Copenhagen, Denmark
- System z Hardware Update -2006, ITSO3DK **Open**
 - Starts 29 Nov 2006 for 1 day in Copenhagen, Denmark
- System z Security, ITSO4DK **Open**
 - Starts 30 Nov 2006 for 1 day in Copenhagen, Denmark
- System z Networking Technologies Update, ITSO5DK **Open**
 - Starts 01 Dec 2006 for 1 day in Copenhagen, Denmark