


ibm.com



Parallel Sysplex 2001 Frank Kyne - ITSO Poughkeepsie



Redbooks

International Technical Support Organization

© Copyright IBM Corp. 2000

General

Agenda:

- Start 09:00, Lunch about 12:00, Finish about 17:00
- Breaks every 90 minutes or so, for 15 minutes

Questions?? Please ask as I go along. Also, if you can't understand my strange accent, please let me know!

Handouts will be available in PDF format - please send me a note at kyne@us.ibm.com if you want a set. Also, your handouts vary slightly from the presentation as I am constantly changing the presentation as new information becomes available - sorry!

PLEASE fill out the evaluation forms. Especially, please fill in how I can make the workshop more valuable to you.



ibm.com/redbooks

© 2001 IBM Corporation

Agenda

Quick review of z/OS 1.1, 1.2. and 1.3 announcements

Review of sysplex enhancements over last 3-4 years

IRD Enhancements in z/OS 1.2 plus IRD implementation overview

Workload charging relationship to IRD and configuration suggestions

Lunch

Improvements in Sync/Async processing

System-Managed CF Structure Duplexing

msys for Operations

GDPS Overview and latest enhancements

Goodie bag



ibm.com/redbooks

© 2001 IBM Corporation



z/OS Release Summaries



© Copyright IBM Corp. 2000

z/OS V1R1

SNA MCS Support

WLM LPAR CPU Management

Dynamic Channel-path Management

Channel Subsystem I/O Priority Queueing

msys for Setup



ibm.com/redbooks

© 2001 IBM Corporation

z/OS V1R2

Enhancements to Intelligent Resource Director for non-z/OS LPARs

System-Managed CF Structure duplexing

msys for Operations

Redesign of sysplex tape sharing support

Extensions to Sysplex Distributor function

Extensions to msys for Setup



ibm.com/redbooks

© 2001 IBM Corporation

z/OS V1R3

Removal of WLM Compat mode

DFSMSHsm Common Recall Queue

DFSMS Record Level Sharing enhancements

- Support for System-Managed Structure Duplexing for Lock structure
- Enhanced use of CF for caching RLS data sets

DFSMS Transactional VSAM Services

Non-disruptive changes to Logger logstream attributes

msys for Setup enhancements



ibm.com/redbooks

© 2001 IBM Corporation

CF Level 10

Available on 9672 G5/G6 and z900 GA2

Provides support for System-Managed CF Structure Duplexing

Increases buffer space for List and Cache structures (even if not using Duplex support) by 768K so those structures may need to be re-sized

- New allocations may fail if structure size is too small
- New allocations may work, but have poor performance
- Rebuilds from CFLEVEL 9 should work, but make sure APAR OW43778 is applied

Recommend to recheck all structure sizes before migrating to CFLEVEL 10 - use CFSizer, available at:

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



ibm.com/redbooks

© 2001 IBM Corporation

Sysplex enhancements and features review



© Copyright IBM Corp. 2000

Sysplex review

Every year we speak about what was announced over the last year, and then we get questions about what was announced 3 years ago!

This year I decided to go back and summarize the new capabilities delivered over the last 3-4 years.

You are paying for all these whether you use them or not - it is your decision whether to maximize value for money or not...

This covers new features and also touches on availability recommendations and is based on a case study done with a large IBM customer

OS/390-supplied exploiters

| Feature | Exploiting? |
|--|-------------|
| OPERLOG - used for sysplex-wide problem determination - can be used together with syslog | |
| LOGREC - sysplex-wide problem recording medium that never fills | |
| Sysplex tape sharing - free cross-system tape unit sharing, if you can live with current lack of cross-sysplex support | |
| GRS Star - higher performance, superior availability and recoverability compared to GRS Ring. Considered a pre-req for large sysplexes | |
| RACF sysplex data sharing - improved performance, easier multi-system administration | |
| JES2 Checkpoint in CF - equivalent-to-better performance, especially with mixed CPC sizes, and elimination of Reserves | |
| Enhanced Catalog Sharing - performance equivalent to non-shared DASD environment | |
| Sysplex HFS sharing - improved flexibility and availability for end-users and system programmers | |
| Symbolic Alias Facility - swap between subsystem software releases with no JCL changes, renaming, or library copying | |
| Use of system symbols in Parmlib, Proclib, VTAM, TCP, NetView, OPC, AOC, Clists to minimize multi-system maintenance overhead | |



ibm.com/redbooks

© 2001 IBM Corporation

DFSMS-supplied exploiters

| Feature | Exploiting? |
|--|-------------|
| DFSMSshm Secondary Host Promotion | |
| DFSMSshm Single GRSplex - removes false contention when there are >1 HSMplexes in the same sysplex | |
| DFSMSshm RLS for CDSs - provides improved performance for CDS-intensive activities | |
| DFSMSshm Multiple Address Space HSM - lets you break HSM tasks across multiple address spaces, running higher priority tasks (like recall) in one address space, and lower priority tasks in another | |
| OAM sysplex support - ability to access OAM data from any system in the sysplex, and remove OAM affinity to just one system | |
| | |
| | |
| | |



ibm.com/redbooks

© 2001 IBM Corporation

XCF considerations

| Feature | Exploiting? |
|---|-------------|
| XCF structures in Coupling Facilities (2) every structure should be defined as both PATHIN and PATHOUT This is <u>SO</u> much easier than using CTCs | |
| CDSs on dedicated volumes with no RESERVES | |
| CDSs on high performance volumes | |
| Don't back up CDS volumes - could cause RESERVE problems | |
| CDSs should not be larger than necessary - impacts IPL and recovery | |
| Primary, alternate, spare CDSs on separate physical controllers | |
| Automation in place to add a spare CDS if primary or alternate is lost | |
| Primary sysplex, CFRM, and LOGR CDSs on different volumes | |
| Separate RACF profile to protect CDSs | |
| Use XISOLATE (WSC FLASH 10080) program to monitor for single points of failure available from ftp://ftp.software.ibm.com/s390/mvs/tools | |
| Allocate spare CDS at the same time as primary and alternate, to ensure they are all at the same functionality level | |



ibm.com/redbooks

© 2001 IBM Corporation

XCF considerations (cont)

| Feature | Exploiting? |
|--|-------------|
| Always use the POPULATECF flavor of the SETXCF command to repopulate a CF | |
| Implement automation to act on potential hang situations (messages IXL040E and IXL041E). | |
| Update automation/Operator's procedures based on APAR OW44231 - see WSC FLASH 20006 (changes processing following loss of Sysplex Timer signal). | |
| Monitor for, and apply HIPERs for XCF (5752SCXCF) and XES (5752SCIXL) | |
| Check for NEW FUNCTION APARs for XCF and XES | |
| If XCF message rates above 1000/sec, see Setting Up a Sysplex for tuning recommendations | |
| Use enhanced D XCF,PI,STRNM=nnnn command to get response time information for XCF paths (reqs APAR OW37261) | |
| Consider automation to check that each structure is in the first CF in its preference list - first check that all CFs are available! | |
| Apply OW49867 if using FICON Native or CSS I/O Priority Queueing | |



ibm.com/redbooks

© 2001 IBM Corporation

CFRM considerations

| Feature | Exploiting? |
|---|-------------|
| System-Managed Rebuild enabled (reqs OS/390 2.8 and CFLEVEL 8) | |
| Structure Full Monitoring enabled and automated (alerts) | |
| Auto Alter enabled for supporting structures, plus monitor for messages so that CFRM policy can be adjusted accordingly | |
| All structures have at least 2 CFs in preference list | |
| Naming convention used so that policy changes can be easily backed out | |
| Use appropriate structure sizes - validate with CFSizer | |
| Suggest SIZE = 2 * INITSIZE | |
| See APAR OW46406 for discussion about MINSIZE values | |
| Never use REBUILDPERCENT > 1 | |
| Apply OW43778 - for persistent structure sizes across rebuilds | |
| Always update the CFRM policy any time a structure size is changed using SETXCF ALTER or Auto Alter | |
| | |



ibm.com/redbooks

© 2001 IBM Corporation

CF considerations

| Feature | Exploiting? |
|---|-------------|
| MUST have at least 2 CFs, even for Resource Sharing | |
| If doing data sharing (prior to SM Duplexing in z/OS 1.2), at least one CF must be failure-isolated and 'category 3' structures should be in that CF See WSC FLASH 98029 | |
| Each CF should have enough storage to hold all structures | |
| Each CF with 1 CP should not run at >25-30% CPU utilization. If there is more than 1 CP, higher utilization is possible without impacting performance | |
| If possible, production CFs should have 2 (or more) CPs | |
| Should be two links from every production CF to every connected image | |
| Use battery backup and UPS when doing data sharing | |
| Never enable Dynamic CF Dispatching for a production CF | |
| Ensure APAR OW43778 is applied - even if you are already on OS/390 2.8 and CFLEVEL 8 | |
| Use highest speed links possible, especially for structures with large data transfers | |



ibm.com/redbooks

© 2001 IBM Corporation

Logger considerations

| Feature | Exploiting? |
|---|-------------|
| Try to ensure staging data sets are not used - this infers that the CFs are failure isolated and LOGR policy is set up correctly | |
| Ensure LOGR CDS is at HBB6603 level (D XCF,C,TYPE=LOGR) | |
| Use data class with 24K CI Size for all offload data sets | |
| Have <i>at least</i> 2 logstreams (connected to different systems) in each LOGR structure | |
| For CICS DFHLOG, monitor offload data sets - if sequence number is large, adjust HIGHOFFLOAD and LOWOFFLOAD or structure size | |
| Also for DFHLOG, monitor frequency of LOG TAIL DELETION (shows up in CICS syslog) - should not be more frequent than 1/minute | |
| Put logstreams with similar average buffer size and request rates in the same structure. Do <i>not</i> mix idle and busy logstreams in same structure | |
| Get your kids to analyze the SMF 88 reports using the IXGRPT1 program - to see how tough life can be! | |
| | |



ibm.com/redbooks

© 2001 IBM Corporation

SFM considerations

| Feature | Exploiting? |
|--|-------------|
| USE SFM. Single change that would have eliminated the largest number of multi-system outages | |
| If the policy is defined and started, SFM can: | |
| <ul style="list-style-type: none"> ✍ Partition a system out of the plex if status update is missing ✍ Partition a system out of the plex if XCF signalling to the system is lost ✍ Assist with V XCF,xxx,OFFLINE processing ✍ Control structure rebuild processing based on REBUILDPERCENT | |
| Don't set INTERVAL in COUPLExx - let system determine it dynamically | |
| You must have a CF to be able to use SFM's system isolation function | |
| Set CLEANUP interval to 30 seconds - should be sufficient for planned shutdowns | |
| Do <i>not</i> specify PROMPT in the SFM - use ISOLATETIME(0) instead to get sick system immediately removed from sysplex | |
| Don't specify OPNOTIFY- let it default to the calculated INTERVAL plus 3 seconds | |
| | |



ibm.com/redbooks

© 2001 IBM Corporation

Automatic Restart Manager considerations

| Feature | Exploiting? |
|--|-------------|
| Use ARM to quickly restart DB2 (using RESTART LIGHT option if DB2 V7) to release retained locks | |
| Use ARM to restart CICS and CPSM regions on the same image as the associated DB2 following a system failure | |
| Use ARM to restart a CICS Server region (Temp Storage, CF Data Table, or Named Counter Server) after a failure (reqs CICS 2.1) | |
| Use ARM to automatically restart NetView, especially if NetView is used to drive automation | |
| Use ARM to restart VTAM and TCP/IP following a failure | |
| Use ARM to restart IMS regions (Control region, Common Queue Server, FDBR) following a failure | |
| Use ARM to restart IRLM following a failure | |
| Use ARM to restart MQ Series following a failure | |
| Use ARM to restart RRS | |
| Use ARM to restart SDSF Server address space | |
| Use ARMWRAP program to add ARM support to other jobs/tasks | |



ibm.com/redbooks

© 2001 IBM Corporation

Network/sysplex considerations

| Feature | Exploiting? |
|---|-------------|
| VTAM over XCF | |
| VTAM Generic Resources for: CICS DB2 IMS TSO NetView Access Services | |
| VTAM Multi Node Persistent Sessions for: CICS IMS Rapid Network Recovery | |
| TCP/IP over XCF | |
| TCP/IP dynamic XCF definitions | |
| TCP/IP Virtual IP Addressing | |
| TCP/IP VIPA Dynamic Takeover | |
| TCP/IP VIPA Dynamic Takeback | |
| TCP/IP Sysplex Distributor (See SG24-5235-02) | |
| TCP/IP WLM/DNS support | |



ibm.com/redbooks

© 2001 IBM Corporation

CICS considerations

| Feature | Exploiting? |
|--|-------------|
| CICS Shared Temporary Storage in CF | |
| CICS Named Counter Server in CF | |
| CICS Global ENQ/DEQ | |
| CICS Data Tables in CF | |
| CICS VSAM Record Level Sharing | |
| CICSplex System Manager - for CICSplex management and dynamic transaction routing | |
| Use CICS Transaction Affinities Utility or CICS Interdependency Analyzer to identify, remove, or manage transaction affinities | |
| CICS TOR defined as VTAM Generic Resource | |
| | |
| | |



ibm.com/redbooks

© 2001 IBM Corporation

Other MVS-related exploiters

| Feature | Exploiting? |
|--|-------------|
| WLM-Managed Batch Initiators | |
| WLM Resource Affinity Scheduling | |
| OPC Hot Standby Controller | |
| OPC Interface to WLM to alter WLM Service Class for selected jobs/applications | |
| Sysplex enhancements in System Automation for OS/390 (See APAR OW39485) | |
| GRS Enhanced Contention Analysis (enhanced D GRS command) | |
| SNA Console support - full function MVS console (except NIP support) from any SNA or IP screen | |



ibm.com/redbooks

© 2001 IBM Corporation

Fix time



ibm.com/redbooks

© 2001 IBM Corporation



IRD Enhancements and Implementation Considerations



© Copyright IBM Corp. 2000

Intelligent Resource Director

Brief review of IRD functions

Implementation considerations for each function

z/OS 1.2 enhancements



ibm.com/redbooks

© 2001 IBM Corporation

IRD - CSS IO Priority Queueing

Already have priority-based queueing on UCB and in ESS

CSS IOPQ adds priority-based queueing in the channel subsystem

Works with DCM to ensure that the 'right' user gets the additional capacity when more paths are added to an LCU

No visible effect on overall response times - however, high importance I/Os should be faster, low importance will be slower (assuming there is queueing for a channel) - average should be unchanged

Works for ALL device types (even 3420!) and all channel types except FICON Bridge (even parallel) - not just DASD or ESS

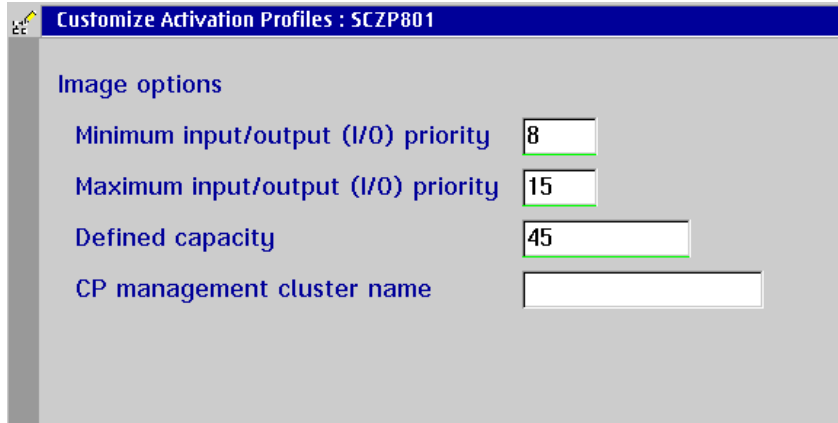


ibm.com/redbooks

© 2001 IBM Corporation

IRD - CSS IOPQ Implementation

- ✓ Specify range of priorities in each image profile



Customize Activation Profiles : SCZP801

Image options

| | |
|-------------------------------------|---------------------------------|
| Minimum input/output (I/O) priority | <input type="text" value="8"/> |
| Maximum input/output (I/O) priority | <input type="text" value="15"/> |
| Defined capacity | <input type="text" value="45"/> |
| CP management cluster name | <input type="text"/> |

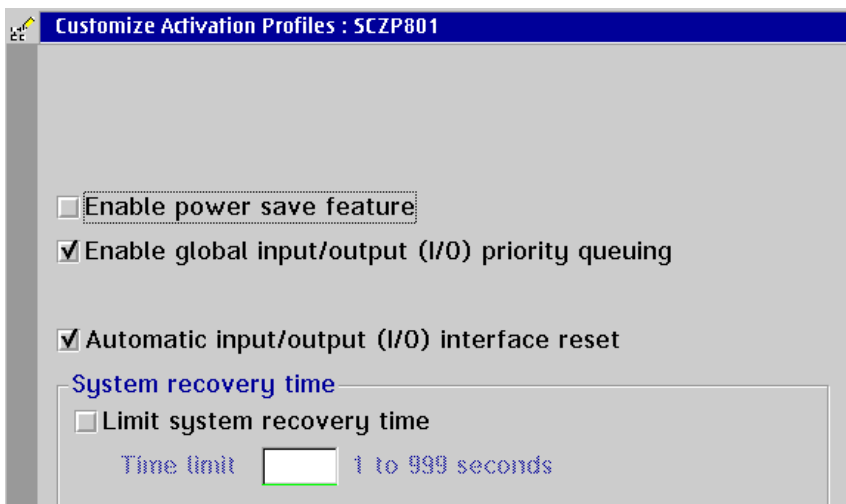


ibm.com/redbooks

© 2001 IBM Corporation

IRD - CSS IOPQ Implementation

- ✓ Enable I/O Prty Queueing at CPC level in Reset profile



Customize Activation Profiles : SCZP801

☐ Enable power save feature

☒ Enable global input/output (I/O) priority queueing

☒ Automatic input/output (I/O) interface reset

System recovery time

☐ Limit system recovery time

Time limit 1 to 999 seconds



ibm.com/redbooks

© 2001 IBM Corporation

IRD - CSS IOPQ Implementation

Remember that once enabled, priorities apply to ALL LPARs - z/OS and others

Recommend:

- Range of 8-15 for Production z/OS LPARs
- Range of 10-10 to 13-13 for production non-z/OS LPARs depending on importance of work in LPAR
- Range of 3-10 for Development z/OS LPARs
- Range of 0-4 for test and sysprog LPARs
- **MUST** be in WLM Goal Mode to exploit range. Non-Goal mode z/OS and all other LPARs will use lowest value in the range

Display status with D WLM,IRD command

No externals in RMF



ibm.com/redbooks

© 2001 IBM Corporation

IRD - CSS IOPQ Implementation

Implementation steps:

- Set priority ranges in Image profile in HMC
- Enable in the CPC Reset profile in HMC (remember, it is enabled or disabled at the CPC, not the LPAR, level)
- Enable dynamically using the "Enable I/O Priority Queueing" icon on the HMC
- Monitor effect in RMF Workload Reports




Enable I/O
Priority Queueing



ibm.com/redbooks

© 2001 IBM Corporation

IRD - CSS IOPQ Implementation



Environment:

TSO and CICS/DB2 DASD behind same CU
TSO and CICS/DB2 DASD on same volume

EMIFed channels to CU CSS contention
Large IOP queues

Test: Run with/without CSS Priority queuing

Expected: Decreased resp. time for High Imp work
Increased resp. time for Medium Imp work




ibm.com/redbooks

© 2001 IBM Corporation

IRD - CSS IOPQ Implementation

For samples with similar IOP queue lengths



- Response time for more important work improves (at the expense of less important work)



ibm.com/redbooks

© 2001 IBM Corporation

IRD - WLM LPAR CPU Mgmt Implementation

To use WLM LPAR CPU Management, you must be running z/OS 1.1 or later, in Goal mode, in an LPAR Cluster, on z900

Gives WLM the ability to move weight between the LPARs in an LPAR Cluster, based on Importance and PI of workloads in each LP

Also provides ability for WLM to vary logical CPs on and offline based on capacity requirements

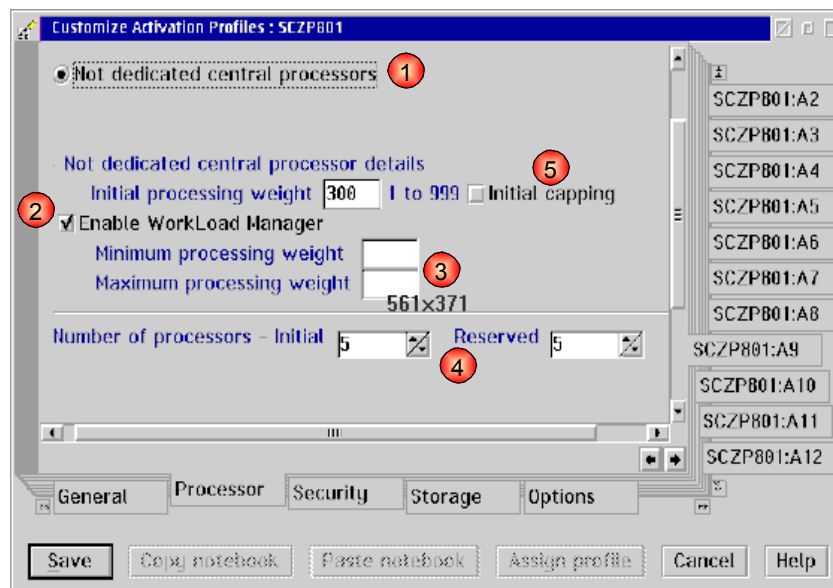


ibm.com/redbooks

© 2001 IBM Corporation

IRD - WLM Weight Mgmt Implementation

- ✓ Shared CPs ①, WLM Enabled ②, min&max weight ③, Initial and Reserved CPs ④, no initial capping ⑤



ibm.com/redbooks

© 2001 IBM Corporation

IRD - WLM Weight Mgmt Implementation

- ✓ Can display and modify Weight Mgmt info at CPC level...

| Logical Partition | Active | Defined Capacity | Current Weight | WLM Managed | Initial Processing Weight | Minimum Processing Weight | Maximum Processing Weight | Initial Capping | Current Capping |
|-------------------|--------|------------------|----------------|-------------------------------------|---------------------------|---------------------------|---------------------------|--------------------------|-----------------|
| A1 | Yes | 0 | 300 | <input type="checkbox"/> | 300 | 10 | 0 | <input type="checkbox"/> | No |
| A2 | Yes | 0 | 300 | <input type="checkbox"/> | 300 | 10 | 0 | <input type="checkbox"/> | No |
| A3 | Yes | 0 | 300 | <input type="checkbox"/> | 300 | 10 | 0 | <input type="checkbox"/> | No |
| A4 | Yes | 0 | 300 | <input type="checkbox"/> | 300 | 10 | 0 | <input type="checkbox"/> | No |
| A5 | Yes | 0 | 300 | <input type="checkbox"/> | 300 | 10 | 0 | <input type="checkbox"/> | No |
| A6 | Yes | 0 | 300 | <input type="checkbox"/> | 300 | 10 | 0 | <input type="checkbox"/> | No |
| A7 | Yes | 0 | 300 | <input type="checkbox"/> | 300 | 10 | 0 | <input type="checkbox"/> | No |
| A8 | Yes | 0 | 300 | <input checked="" type="checkbox"/> | 300 | | | <input type="checkbox"/> | No |
| A9 | Yes | 0 | 300 | <input checked="" type="checkbox"/> | 300 | | | <input type="checkbox"/> | No |
| A10 | Yes | 0 | 300 | <input checked="" type="checkbox"/> | 300 | | | <input type="checkbox"/> | No |
| A11 | Yes | 0 | 300 | <input type="checkbox"/> | 300 | 10 | 0 | <input type="checkbox"/> | No |
| A12 | Yes | 0 | 300 | <input type="checkbox"/> | 300 | 10 | 0 | <input type="checkbox"/> | No |

Buttons: Save to profiles, Change running system, Save and change, Reset, Cancel, Help



ibm.com/redbooks

© 2001 IBM Corporation

IRD - WLM Weight Mgmt Implementation

- ✓ Can display Vary CPU ① and Weight Management ② status in each LPAR...

```

D WLM,IRD
IWM059I 18.12.19 WLM DISPLAY 663
OPTIONS
  VARYCPU ENABLED: YES ①
  CPU MANAGEMENT ENABLED: YES ②
  CHANNEL SUBSYSTEM PRIORITY ENABLED: YES
WLM CPU MANAGEMENT STATUS
  CPU MANAGEMENT ACTIVE ②
DCM STATUS
  DCM ACTIVE
WLM LPAR CLUSTER DATA
  SYSPLEX NAME: SANDBOX
  WLM LPAR CLUSTER STRUCTURE: SYSZWLM_0ECB2064
  SYSTEM PARTITION MVS CAPABILITY CONNECT
  NAME IDENTIFIER LEVEL LEVEL STATUS
  SC64 009 013 011 CONNECTED
  SC63 008 012 012 CONNECTED
  SC65 010 012 012 CONNECTED
  
```

Annotations:

- ① In IEAOPTxx
- ② On HMC
- ② Active in z/OS



ibm.com/redbooks

© 2001 IBM Corporation

IRD - WLM Weight Mgmt Implementation

VARY CPU Management:

- Can be switched on/off in IEAOPT member
- Is designed to balance efficiency and flexibility
- Will *not* take logical CPs offline until CPU consumption increases from low levels (immediately after IPL, for example)
- Will always attempt to give each LPAR spare capacity at least equivalent to 1 logical CP
- Will start with Initial number of logical CPs as defined in Image Profile and never go higher than that (unless you vary on Reserved CPs manually)
- Recommend to set Initial number=physically installed shared CPs for production LPARs. Initial + Reserved should = maximum number of CPs that can be installed on the CPC - allows you to manually bring more CPs online without a POR or even an IPL



ibm.com/redbooks

© 2001 IBM Corporation

IRD - WLM Weight Mgmt Implementation

WLM LPAR Weight Management:

- Moves weight between LPARs in an LPAR Cluster - total Cluster weight never changes
- Uses similar methodology to Dispatching Priority management:
 - Only 1 donor and 1 receiver per adjustment
 - Uses 'net value' calculation described in 2.2.1.8 in SG24-5326 (WLM Exploitation)
- Generally only 1 adjustment per minute (unless higher importance SCP needs help)
- Adjustment is a percentage of total LPAR Cluster weight - amount can vary depending on how much help the receiver needs
- Adjustments are prompted by an important SCP having a PI (sysplex *or* local) greater than 1



ibm.com/redbooks

© 2001 IBM Corporation

IRD - WLM Weight Mgmt Implementation

WLM LPAR Weight Management recommendations:

- Start off with similar minimum and maximum weights, increasing gap with experience
- Target should be no min or max weight specification, BUT beware LPARs where utilization increases dramatically very quickly - WLM may not be able to react quickly enough, so set weights accordingly.
- Consider the importances of the workloads in all the LPARs in the LPAR Cluster - would they work if you merged all the work into a single system? Could be a concern, especially if LPAR Cluster contains production and non-production systems
- Make sure all the service in the IRD bucket (2064DEVICE Subset IRD) is applied to pick up latest algorithm fine tuning




ibm.com/redbooks

© 2001 IBM Corporation

IRD - WLM Weight Mgmt Implementation

Steps:

- Define WLM LPAR Cluster structure (SYSZWLM_2064xxxx)
- Define initial, min, and max weights on HMC (start with +/- 10% of Initial Weight)
- Ensure LPARs to be managed have shared CPs
- Define more logical CPs than you need, and let WLM manage
- Enable Workload Management in Image profile on HMC for each LPAR to be managed
- Dynamically enable Workload Management for each LPAR as you are ready for it to be IRD-managed  Change LPAR Controls
- Monitor with RMF - use LPAR Cluster reports and Workload Reports



ibm.com/redbooks

© 2001 IBM Corporation

IRD - WLM Weight Mgmt Implementation

| L P A R C L U S T E R R E P O R T | | | | | | | | | | | | | |
|-----------------------------------|-----------|---------|----------------------|-----|-----|-----------------|-------|-------|----------------------------------|--------|--------------|-------|--|
| z/OS V1R2 | | | SYSTEM ID SC64 | | | DATE 09/10/2001 | | | INTERVAL | | | | |
| | | | RPT VERSION V1R2 RMF | | | TIME 15.40.00 | | | CYCLE 1 | | | | |
| ----- WEIGHTING STATISTICS ----- | | | | | | | | | ----- PROCESSOR STATISTICS ----- | | | | |
| --- DEFINED --- | | | | | | | | | --- NUMBER --- | | -- TOTAL% -- | | |
| CLUSTER | PARTITION | SYSTEM | INIT | MIN | MAX | AVG | MIN % | MAX % | DEFINED | ACTUAL | LBUSY | PBUSY | |
| SANDBOX | A10 | SC65 | 300 | 0 | 0 | 222 | - | - | 12 | 2.0 | 46.14 | 13.18 | |
| | A12 | WTSCVMT | 100 | 0 | 0 | 100 | - | - | 2 | 2.0 | 79.04 | 22.58 | |
| | A8 | SC63 | 300 | 0 | 0 | 157 | - | - | 12 | 2.0 | 45.07 | 12.88 | |
| | A9 | SC64 | 300 | 0 | 0 | 520 | - | - | 10 | 4.0 | 27.84 | 15.91 | |
| TOTAL | | | 1000 | | | | | | 36 | | 198.1 | 64.55 | |
| | | | | | | | | | | 4 | 2.98 | 1.70 | |
| WTSCPLX1 | A1 | SC52 | 300 | | | 300 | | | 4 | 2 | 5.47 | 1.56 | |
| | A2 | SC53 | 300 | | | 300 | | | 2 | 2 | 5.43 | 1.55 | |
| | A5 | SC42 | 300 | | | 300 | | | 2 | 2 | 5.85 | 1.67 | |
| | A7 | SC66 | 300 | | | 300 | | | 2 | | | | |
| TOTAL | | | 1200 | | | | | | 10 | | 19.73 | 6.49 | |

* Defined Processors is Initial + Reserved

Actual processors is average online over interval



ibm.com/redbooks

© 2001 IBM Corporation

IRD - WLM Weight Mgmt Implementation

To monitor effect in RMF:

- Use OVW options to define variables, one line/interval
 - Keywords WACTL, NLACTL
 - Labels - up to 8 characters
 - For example:
 - ▶ OVW(AWGTSYSA(WACTL('A01')))) /* actual weight for SYSA */
 - ▶ OVW(LPSYSA(NLACTL('A01')))) /* actual LPs for SYSA */
- Can produce report (OVERVIEW(REPORT)), record format for use by a spreadsheet (OVERVIEW(RECORD)) or both. If you want record format, the output data set must be DCB=(RECFM=VB,LRECL=32756,BLKSIZE=32760).



ibm.com/redbooks

© 2001 IBM Corporation

IRD - WLM Weight Mgmt Implementation

RMF PP JCL

```
//SYSIN DD *
DATE(09182001,09182001)
ETOD(1200,1400)
PTOD(1200,1400)
RTOD(1200,1400)
STOD(1200,1400)
SUMMARY(INT,TOT)
SYSOUT(X)
OVERVIEW(REPORT)
OVW(AWGTSC64(WACTL('A9')))
OVW(AWGTSC65(WACTL('A10')))
OVW(ALPSSC64(NLACTL('A9')))
OVW(ALPSSC65(NLACTL('A10')))
/*
```

LPAR name

RMF OVERVIEW REPORT

```
z/OS V1R1          SYSTEM ID SC63          START 09/18/2001
CONVERTED TO z/OS V1R2 RMF          END 09/18/2001

NUMBER OF INTERVALS 12          TOTAL LENGTH OF INTERVALS 02.00.00
DATE    TIME    INT    AWGTSC64  AWGTSC65  ALPSSC64  ALPSSC65
MM/DD HH.MM.SS MM.SS
09/18 12.00.00 09.59          297          297          5.0          2.0
09/18 12.10.00 09.59          284          297          5.0          2.0
09/18 12.20.00 09.59          270          297          5.0          2.0
09/18 12.29.59 10.00          270          282          5.0          2.0
09/18 12.40.00 10.00          270          269          5.0          2.0
09/18 12.50.00 10.00          270          269          5.0          2.0
```

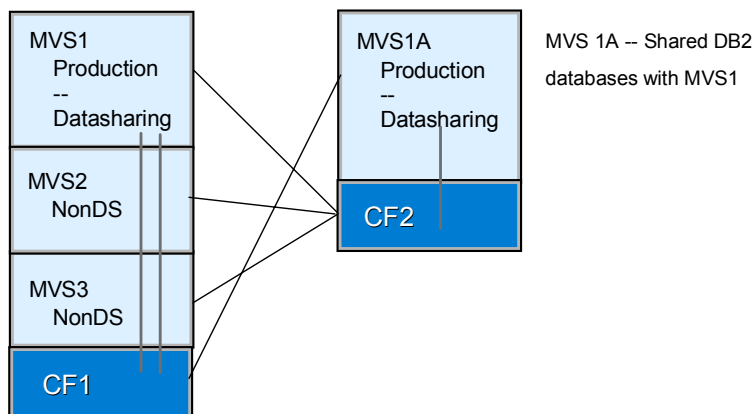


ibm.com/redbooks

© 2001 IBM Corporation

IRD - WLM Weight Mgmt Implementation

Add another CPC (z900 with 4 CPs) and CF with 1 ICF




Test: Increase MVS2 non-DS workload (TSO)
Result: MVS1 workload moves to MVS1A, thrupt constant
MVS2 gets more capacity, thrupt increases



ibm.com/redbooks

© 2001 IBM Corporation

IRD - WLM Weight Mgmt Implementation



ibm.com/redbooks

© 2001 IBM Corporation

IRD - DCM Implementation

Dynamic Channel-path Management:

- Does **not** move channels between LPARs
- Does add or remove channels from the configuration for a DCM-capable control unit - similar to a systems programmer modifying HCD and doing dynamic activates - it is not just varying paths on and offline
- Must be z/OS on a z900. Goal mode is optional, and CPC can be in basic or LPAR mode. Control Unit *must* be attached via ESCON Director
- Works with ESS and RVA, and some OEM DASD - must be ESCON or FICON-Bridge attached




ibm.com/redbooks

© 2001 IBM Corporation

IRD - DCM Implementation

- Configuration today, without DCM





ibm.com/redbooks

© 2001 IBM Corporation

IRD - DCM Implementation

- Configuration with DCM



ibm.com/redbooks

© 2001 IBM Corporation

IRD - DCM Implementation

DCM recommendations:

- Use DCM to cater for the difference between average and peak workloads
- Configuration *must* be symmetric - all LPARs in the LPAR Cluster *must* have access to the same devices through the same number of shared channels
- Must have more than one managed LCU behind an ESCON director for DCM to be effective
- DCM specifically addresses the components of PEND time, so select LCUs with high PEND time (remember, management is at the LCU, not physical CU, level)
- Spread managed channels across processor cages and SAPs, and across multiple ESCON directors so DCM can select path that will provide the best availability



ibm.com/redbooks

© 2001 IBM Corporation

IRD - DCM Implementation

DCM restrictions:

- CU must be ESS or RVA or supported OEM
- ESCON or FICON Bridge (FCV) only
- LCU must be attached to managed channels via a switch
- Managed channels cannot be shared outside the LPAR Cluster
- LCU load should not increase dramatically in very short time - makes a maximum of one adjustment every 10 seconds
- All 'affected' LCUs must also support DCM
- Cannot mix FC and FCV/ESCON on a managed LCU
- 'Affected' control units *also* cannot have mix of FC and FCV/ESCON

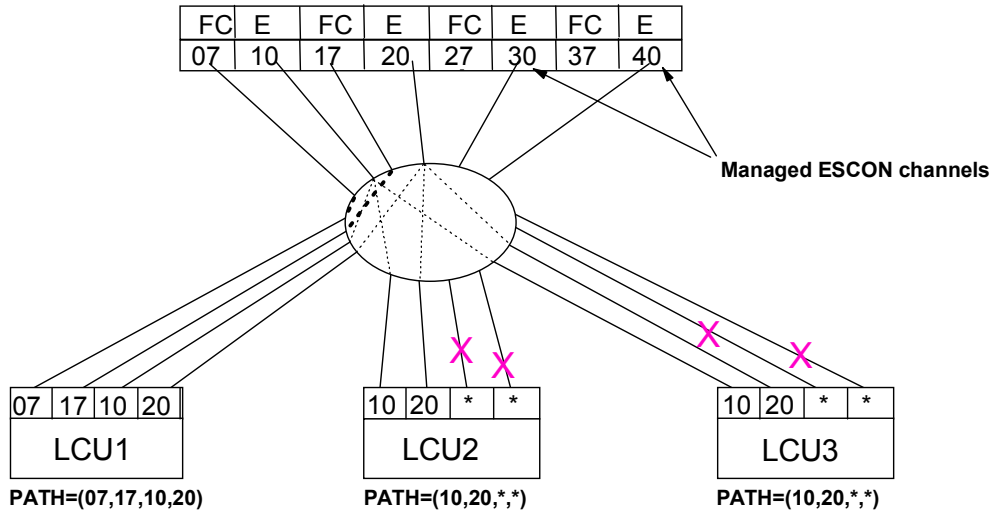


ibm.com/redbooks

© 2001 IBM Corporation

IRD - DCM Implementation

- What is an "Affected" control unit?



Because CHPIDs 10 and 20 share a control unit with FICON channels, no changes can be made that will affect those channels, so LCUs 2 and 3 will not be managed, even though they are only connected to ESCON channels and are defined in HCD with managed paths.



ibm.com/redbooks

© 2001 IBM Corporation

IRD - DCM implementation

```

Goto  Filter  Backup  Query  Help
----- Add Channel Path -----

Specify or revise the following values.

Processor ID . . . . . FK1          Dynamic CHPID Management Example
Configuration mode : LPAR

Channel path ID . . . . . 82  +
Number of CHPIDs . . . . . 1
/ Channel path type . . . . . CNC  +
Operation mode . . . . . SHR  +
Managed . . . . . Yes  + I/O Cluster: PRDPLEX
Description . . . . . Managed channel

/ Specify the following values only if connected to a switch:

Dynamic switch ID . . . . . 01  (00 - FF)
Entry switch ID . . . . . 01  +
Entry port . . . . . 82  +

F1=Help    F2=Split    F3=Exit    F4=Prompt    F5=Reset    F9=Swap
F12=Cancel

96  CNC  SHR  04    04 82    Yes
97  CNC  SHR  04    04 83    Yes
F1=Help    F2=Split    F3=Exit    F4=Prompt    F5=Reset    F7=Backward
F8=Forward    F9=Swap    F10=Actions    F11=Add    F12=Cancel    F13=Instruct
F20=Right    F22=Command
    
```

Mode can be Basic or LPAR

MUST be defined as shared if in LPAR mode

Defines a managed CHPID

For managed CHPIDs, you MUST specify the sysplex that can share this CHPID.

This MUST be specified

These are optional but highly recommended



ibm.com/redbooks

© 2001 IBM Corporation

IRD - DCM implementation

Steps:


- Identify candidate LCUs
- Install all service identified in IRD subset of 2064DEVICE bucket
- If member of multisystem sysplex and in LPAR mode, define WLM LPAR Cluster structure
- Control unit ports must be defined in HCD for all ESCON directors
- Define set of channels as managed channels
 - If using existing channels, must remove from ALL LCUs in HCD first
- Define managed ports on selected control units
- Educate ops and sysprogs in new and changed commands
- If managed channels and managed control units are available at IPL, system will automatically come up with DCM active if this is the first system in the LPAR Cluster. If not the first, it comes up in the same mode as existing systems.



ibm.com/redbooks

© 2001 IBM Corporation

IRD - DCM implementation



All Response times improved when DCM activated



ibm.com/redbooks

© 2001 IBM Corporation

IRD - Documentation

z/OS MVS Planning: Workload Management, SA22-7602

Redbook z/OS Intelligent Resource Director, SG24-5952

WLM Web site:

- <http://www.ibm.com/servers/eserver/zseries/zos/wlm/documents/ird/ird.html>

IRD Bucket - 2064DEVICE subset IRD



ibm.com/redbooks

© 2001 IBM Corporation

IRD - Enhancements

z/OS 1.2 - WLM LPAR CPU Management for non-z/OS LPARs



ibm.com/redbooks

© 2001 IBM Corporation

IRD - non-Z/OS support



Provides ability to include non-z/OS images in LPAR Cluster and have WLM manage the weights of those LPARs

Requires at least one system in the LPAR Cluster to be at z/OS 1.2 or later

Must be on GA2-level z900 (Driver 3C or later)

New support in WLM - new subsystem type (SYSH) - single period with velocity goal

Must be in LPAR mode using shared CPs, not capped, not IFLs, and with WLM CF Structure

Need enabling PTF for APAR OW50221 (planned 10/01)

Only weight management - no logical CP management



ibm.com/redbooks

© 2001 IBM Corporation

IRD - non-z/OS support



Supported operating systems:

- Linux (must be Kernel 2.4 or above)
- z/VM (3.10 and above)
- VSE
- OS/390 V2R6 and later that are in a different sysplex
 - Specifically checks and excludes OS/390 systems in the same sysplex
- Systems must pass their sysname to the HMC



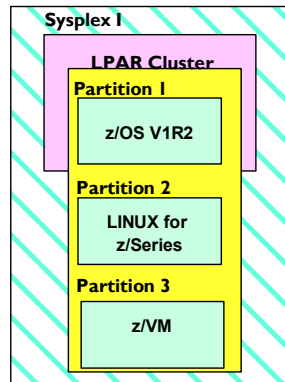
ibm.com/redbooks

© 2001 IBM Corporation


IRD - non-z/OS support



z/900 CPC



z/900 CPC



ibm.com/redbooks

© 2001 IBM Corporation

IRD - non-z/OS support



Assigning a system name:

- In OS/390, use normal methods (IEASYMxx, IEASYSxx)
- In VM, system name is defined in SYSTEM CONFIG file
- To give Linux system a sysname, issue:
 - `insmod hwc_cpi.o system_name=xxxxxxx`
 - `hwc_cpi.o` command is in directory
`/lib/modules/2.4.5-0tape-dasd/kernel/drivers/s390/char/`



ibm.com/redbooks

© 2001 IBM Corporation

IRD - non-z/OS LPAR Management



- Assign cluster name in HMC image profile

Customize Image Profiles: SCZP801:A12

Image options

Minimum input/output (I/O) priority

Maximum input/output (I/O) priority

Defined capacity

CP management cluster name

General Processor Security Storage Options Load Crypto PCI Crypto

Save Copy notebook Paste notebook Assign profile Cancel Help

ibm.com/redbooks

© 2001 IBM Corporation

IRD - non-z/OS support



- System and cluster name show in HMC

SCZP801:A12 Details

Instance information

Status: **Operating**

Group: CPC Images

SysPlex Name:

Lockout disruptive tasks: ☐ Yes ☒ No

Activation profile: A12

Last used profile: not set via Ac

Operating System: **WTSCVMT**

CPU LPAR Cluster Name: **SANDBOX**

Last task information

Task name: Customize/Delete Activation Profiles

Task status:

Acceptable status

☒ Operating -

☐ Not Activated -

☐ Not Operating -

☐ Power save -

☐ Exceptions -

☐ Status check -

ibm.com/redbooks

© 2001 IBM Corporation

IRD - non-z/OS support



- Define workload in WLM - can be any name you want

Workload Notes Options Help

Create a Workload

Command ==> _____

Enter or change the following information:

Workload Name IRDSYSH_ (Required)

Description IRD-managed non-z/OS LPARS _____



ibm.com/redbooks

© 2001 IBM Corporation

IRD - non-z/OS support



- Define Service Class(es)

Service-Class Notes Options Help

Create a Service Class

Row 1 to 1 of 1

Command ==> _____

Service Class Name IRDSYSHI (Required)

Description High priority non-z/OS LPARs_

Workload Name IRDSYSH_ (name or ?)

Base Resource Group _____ (name or ?)

Cpu Critical NO (YES or NO)

Specify BASE GOAL information. Action Codes: I=Insert new period,
E=Edit period, D=Delete period.

---Period--- -----Goal-----

Action # Duration Imp. Description

i_

***** Bottom of data *****



ibm.com/redbooks

© 2001 IBM Corporation

IRD - non-z/OS support



- Select Goal Type - must be Velocity

```
Service-Class  Notes  Options  Help
-----
Choose a goal type for period 1
-----
C 3_ 1. Average response time
D   2. Response time with percentile
W   3. Execution velocity
B   4. Discretionary
C
F1=Help      F2=Split      F5=KeysHelp
F9=Swap      F12=Cancel
S
E

---Period---  -----Goal-----
Action #  Duration  Imp.  Description
i
***** Bottom of data *****
```



ibm.com/redbooks

© 2001 IBM Corporation

IRD - non-z/OS support



- Specify Execution Velocity and importance - must be j period

```
Execution velocity goal

Enter an execution velocity for period 1

Velocity . . . 40 (1-99)

Importance . . 2 (1=highest, 5=lowest)
Duration . . . (1-999,999,999, or
               none for last period)

F1=Help      F2=Split      F5=KeysHelp  F9=Swap
F12=Cancel
```



ibm.com/redbooks

© 2001 IBM Corporation

IRD - non-z/OS LPAR Management



- Assign Service Class to LPARs using new qualifier of SY or SYG (based on SYSNAME) or sysplex name (PX)

```

Subsystem-Type  Xref  Notes  Options  Help
-----
Create Rules for the Subsystem Type          Row 1 to 1 of 1
Command ===> _____ SCROLL ===> PAGE

Subsystem Type SYSH (Required)  Fold qualifier names?  Y (Y or N)
Description . . . IRD-managed non-z/OS LPARs _____

Action codes:  A=After      C=Copy      M=Move      I=Insert rule
               B=Before     D=Delete row  R=Repeat    IS=Insert Sub-rule
                                           More ===>

-----Qualifier-----
Action  Type      Name      Start      Service      Report
-----
_____ 1  px_____ WTSCPLX1 _____ IRDSYSHI  IRDSYSH_
_____ 2  sy_____ WTSCVMT_ _____ IRDSYSHI  IRDSC50_
*****
***** BOTTOM OF DATA *****

```



ibm.com/redbooks

© 2001 IBM Corporation

IRD - non-z/OS support



Implementation steps:

- Define WLM LPAR Cluster structure (if not already defined)
- Add definitions to WLM - workload, service class, report class, classification rules. After this, only z/OS 1.2 systems can change or activate WLM policy. Install and activate policy from z/OS 1.2 system
- Change image profile to add cluster name
- Monitor velocity in RMF - should have separate report class (and probably separate service class) for each non-z/OS LPAR to be managed
- Fine tune WLM definitions
- Enable Workload Management for the selected LPAR dynamically and in image profiles



ibm.com/redbooks

© 2001 IBM Corporation

IRD - non-z/OS support



How it works:

- z/OS asks LPAR for information about all the partitions in the same LPAR Cluster
 - This information includes the system name, CPU utilization and waiting-for-CPU numbers
- WLM calculates velocity, and thereby Performance Index (PI) for each partition
- Based on the importance of the work in each partition, and the PI of that work, WLM may adjust weight of a given partition (just one partition per adjustment), using a z/OS or non-z/OS LPAR as the donor
- Adjustment may be made by any z/OS 1.2 system in the LPAR Cluster - there is not a single 'managing' image



ibm.com/redbooks

© 2001 IBM Corporation

IRD - non-z/OS support

- New reporting in RMF Workload Activity report

| W O R K L O A D A C T I V I T Y | | | | | | | | | |
|-----------------------------------|--------------|---------------------------------------|-------------------------|---------------|------------------------------|------------|------|--|--|
| z/OS V1R2 | | SYSPLEX SANDBOX | | | DATE 09/11/2001 | | | | |
| | | RPT VERSION V1R2 RMF | | | TIME 12.30.00 | | | | |
| REPORT BY: POLICY=DAYTIME | | WORKLOAD=IRDSYSH | | | SERVICE CLASS=IRDSYSH | | RESO | | |
| number of managed non-z/OS LPARs | | Time since z/OS started managing LPAR | | | CRITICAL =NONE | | | | |
| | | | | | DESCRIPTION =High priority S | | | | |
| TRANSACTIONS | TRANS.-TIME | HHH.MM.SS.TTT | --DASD I/O-- | ---SERVICE--- | | | | | |
| AVG 143.50 | ACTUAL | 0 | SSCHRT 0.0 | IOC 0 | A | | | | |
| MPL 143.50 | EXECUTION | 0 | RESP 0.0 | CPU 0 | T | | | | |
| ENDED 0 | QUEUED | 0 | CONN 0.0 | MSO 0 | T | | | | |
| END/S 0.00 | R/S AFFINITY | 0 | DISC 0.0 | SRB 0 | S | | | | |
| #SWAPS 0 | INELIGIBLE | 0 | Q+PEND 0.0 | TOT 0 | R | | | | |
| EXCTD 0 | CONVERSION | 0 | IOSQ 0.0 | /SEC 0 | I | | | | |
| AVG ENC 0.00 | STD DEV | 0 | | | | | | | |
| REM ENC 0.00 | | | | | | | | | |
| MS ENC 0.00 | | | | | | | | | |
| PER IMPORTANCE | PERF | --TRANSACTIONS-- | -----RESPONSE TIME----- | | | | | | |
| | INDX | -NUMBER- | ---- | GOAL----- | ---- | ACTUAL---- | | | |
| 1 1 | 411 | 0 0 | | | | | | | |



ibm.com/redbooks

© 2001 IBM Corporation

IRD - non-z/OS support

- RMF LPAR Cluster report now shows non-z/OS LPARs that are in the LPAR Cluster, with same information as other LPARs

| L P A R C L U S T E R R E P O R T | | | | | | | | | | | | | |
|-----------------------------------|-----------|---------|----------------------|-----|-----|-----|-------|-----------------|---------|--------|----------|-------|--------|
| z/OS V1R2 | | | SYSTEM ID SC65 | | | | | DATE 09/11/2001 | | | INTERVAL | | |
| | | | RPT VERSION V1R2 RMF | | | | | TIME 13.40.00 | | | CYCLE 1. | | |
| ----- WEIGHTING STATISTICS ----- | | | | | | | | | | | | | |
| ----- PROCESSOR STATISTICS ----- | | | | | | | | | | | | | |
| | | | --- DEFINED --- | | | | | --- ACTUAL --- | | | | | --- |
| | | | --- | | | | | --- | | | | | --- |
| CLUSTER | PARTITION | SYSTEM | INIT | MIN | MAX | AVG | MIN % | MAX % | DEFINED | ACTUAL | LBUSY | PBUSY | TOTAL% |
| SANDBOX | A10 | SC65 | 300 | 0 | 0 | 177 | - | - | 12 | 2.0 | 45.23 | 12.92 | |
| | A12 | WTSCVMT | 100 | 0 | 0 | 158 | - | - | 2 | 2.0 | 76.25 | 21.78 | |
| | A8 | SC63 | 300 | 0 | 0 | 507 | - | - | 12 | 2.0 | 89.79 | 25.65 | |
| | A9 | SC64 | 300 | 0 | 0 | 157 | - | - | 10 | 3.0 | 29.82 | 12.78 | |
| ----- | | | | | | | | | | | | | |

157 is approximately 5% of the total weight of the CPC
- IRD will stop using that LPAR as a donor at this point



ibm.com/redbooks

© 2001 IBM Corporation

IRD - non-z/OS support



Example scenario:


- 3 Workloads running on 1 CEC:
 - OLTP: Most important workload. Runs only during day shift
 - Web server running on Linux talking to DB2 on z/OS. Medium importance
 - Batch: Low importance. Always running. Capable of consuming whole CEC if allowed
- CEC divided into 2 partitions
 - Partition 1: z/OS running OLTP, DB2, Batch
 - Partition 2: Web server on Linux
- IRD (Intelligent Resource Director) to manage environment



ibm.com/redbooks

© 2001 IBM Corporation

IRD - non-z/OS support




Initial CPU resource distribution



ibm.com/redbooks

© 2001 IBM Corporation

IRD - non-z/OS support



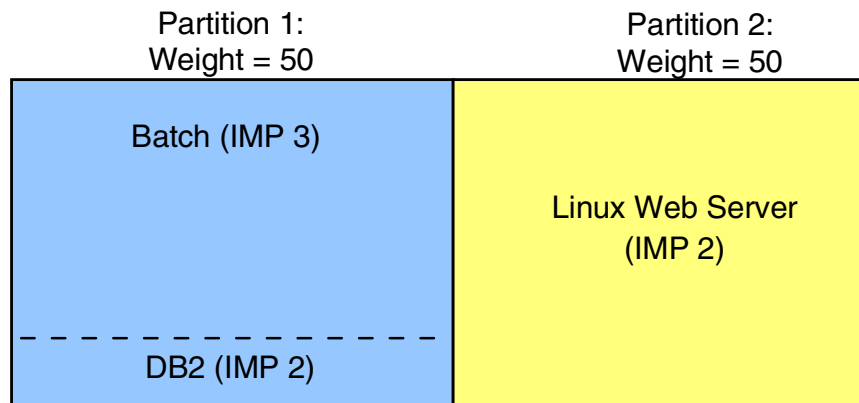
Spike in OLTP work:
Batch squeezed out
Weights adjusted to give more resource to partition 1



ibm.com/redbooks

© 2001 IBM Corporation

IRD - non-z/OS support



OLTP work ends:

Weights adjusted to give more resource to partition 2

Resource given to batch kept constant



ibm.com/redbooks

© 2001 IBM Corporation



IBM®

**Workload Charging - How to give IBM
even less money!**



Redbooks

International Technical Support Organization



© Copyright IBM Corp. 2000

Workload Charging

Planning for WLC

- Get WLC tool to identify 4-hour rolling average for all LPARs
 - http://www.ibm.com/servers/eserver/zseries/wlc_lm/wlctool.html
- Get SCRT tool to identify which VWLC products are running in each LPAR
 - http://www.ibm.com/servers/eserver/zseries/wlc_lm/scrt.html
- Run both tools against (at least) one month's SMF data
- Understand what is running in each LPAR - savings come from not running VWLC products in every LPAR. Note that not every VWLC product creates Type 89 records
- Analyze software inventory - FLWC products provide better value on a small number of large CPCs



ibm.com/redbooks

© 2001 IBM Corporation

Workload Charging

Defined capacity considerations:

- Independent of IRD LPAR CPU Management
 - Can co-exist, but IRD is *not* a pre-req for WLC or Defined Capacity
- Does not require WLM to be in Goal mode
- Defined in Image profile on HMC and dynamically using "Change LPAR Controls" icon
- Reported in RMF Postprocessor LPAR Cluster report (z/OS 1.1), RMF PP Overview report (WCAPPER) (z/OS 1.1), in new 'CPC Capacity' option in RMF Monitor III (1.2+), and in RMF PM (1.2)
- Be careful with CPCs where the total of the defined capacities of the LPARs is greater than the MSU capacity of the CPC
- Be careful when moving from full capacity (defined capacity=0 for all LPARs) to defined capacity (for example, when you install white space)



ibm.com/redbooks

© 2001 IBM Corporation

Workload Charging

Mechanisms for controlling capacity:

- Defined capacity based upon rolling 4-hour average. However until the 4-hour average reaches the defined capacity, this places no limits on CPU consumption
- LPAR Capping. If this option is used, any defined capacity for that LPAR is ignored
- Number of logical CPs. Limiting the number of logical CPs defined for an LPAR will limit how much CPU it can consume
- IRD. Doesn't limit CPU consumption by an LPAR, but attempts to ensure that important workloads are protected. Mutually exclusive with LPAR Capping



ibm.com/redbooks

© 2001 IBM Corporation

Workload Charging

Recommend to use defined capacity to control image capacity:

- LPAR capping is related to active LPARs on CPC. If an LPAR is deactivated, capped LPAR gets a larger share of CPC, increasing image capacity (and therefore software bill)
- SCRT adds Initial and Reserved CPs to get image capacity if using number of logical CPs to limit image capacity. However Initial + Reserved is probably > online CPs
- Dedicated CPs do not make business sense. Results in wasted capacity or constrained LPARs while capacity is available elsewhere on CPC



ibm.com/redbooks

© 2001 IBM Corporation

Workload Charging

Defined capacities and IRD

- Can co-exist
- Not pre-reqs for each other
- WLM controls both, so it will not try to increase the weight of a soft-capped LPAR
- WLM may decide to vary off logical CPs while an LPAR is capped
- Depending on the workloads in the different LPARs, WLM may decide to move weight from a soft-capped LPAR to another LPAR in the LPAR Cluster that needs help



ibm.com/redbooks

© 2001 IBM Corporation

Workload Charging

Defined capacity value recommendations:

- If Development LPAR with discretionary workload, set defined capacity equal rolling 4-hour average
- If Production LPAR, set defined capacity = rolling 4-hour average plus 20% (as a starting point)
- Defined capacities are defined and reported in terms of announced MSUs (requires OW50998 available now)
- Ongoing monitoring:
 - ▶ Watch for % of time LPAR is soft-capped (in RMF PP report)
 - ▶ Online indicator in RMF Mon III that LPAR is soft-capped
 - ▶ Note that no message is issued when soft capping is invoked



ibm.com/redbooks

© 2001 IBM Corporation

Workload Charging

| Processor Model | SU/SEC | Calculated MSUs | Announced MSUs | Calculated minus Announced | Percent Difference |
|-----------------|------------|-----------------|----------------|----------------------------|--------------------|
| 2064-1C1 | 12112.0363 | 44 | 43 | -1 | -2.33% |
| 2064-1C2 | 11502.5162 | 83 | 83 | 0 | 0.00% |
| 2064-1C3 | 11142.0613 | 120 | 119 | -1 | -0.84% |
| 2064-1C4 | 10781.6712 | 155 | 153 | -2 | -1.31% |
| 2064-1C5 | 10540.1845 | 190 | 187 | -3 | -1.60% |
| 2064-1C6 | 10296.0103 | 222 | 217 | -5 | -2.30% |
| 2064-1C7 | 10056.5682 | 253 | 247 | -6 | -2.43% |
| 2064-1C8 | 9815.9509 | 283 | 276 | -7 | -2.54% |
| 2064-1C9 | 9575.1047 | 310 | 302 | -8 | -2.65% |
| 2064-110 | 9334.8891 | 336 | 327 | -9 | -2.75% |
| 2064-111 | 9211.2838 | 365 | 350 | -15 | -4.29% |
| 2064-112 | 8968.6099 | 387 | 372 | -15 | -4.03% |
| 2064-113 | 8724.1003 | 408 | 392 | -16 | -4.08% |
| 2064-114 | 8602.1505 | 434 | 410 | -24 | -5.85% |
| 2064-115 | 8359.4566 | 451 | 426 | -25 | -5.87% |
| 2064-116 | 8117.7067 | 468 | 441 | -27 | -6.12% |
| 2064-101 | 11585.8074 | 42 | 41 | -1 | -2.44% |
| 2064-102 | 10891.7631 | 78 | 78 | 0 | 0.00% |
| 2064-103 | 10430.2477 | 113 | 112 | -1 | -0.89% |
| 2064-104 | 10081.9156 | 145 | 143 | -2 | -1.40% |
| 2064-105 | 9732.3601 | 175 | 173 | -2 | -1.16% |
| 2064-106 | 9384.1642 | 203 | 199 | -4 | -2.01% |
| 2064-107 | 9153.3181 | 231 | 225 | -6 | -2.67% |
| 2064-108 | 8805.7237 | 254 | 245 | -9 | -3.67% |
| 2064-109 | 8456.6596 | 274 | 265 | -9 | -3.40% |



ibm.com/redbooks

© 2001 IBM Corporation

Workload Charging

- RMF Mon III reports LPAR defined capacity ¹, capacity used in interval ², rolling 4 hour average ³, max MSUs in last 4 hours ⁴, percent of time LPAR was soft capped ⁵, and optionally time left till soft cap

| RMF V1R2 CPC Capacity | | | | | | | | | |
|-----------------------|----------|--------------------|----------|----------------|----------------|------------|-----------|--------|--------|
| Samples: 59 | | System: RMF2 | | Date: 09/19/01 | Time: 23.14.00 | Range: 100 | sec | | |
| 2064 Model 1C7 | | --- 4 Hour MSU --- | | | | | | | |
| CPC capacity | 247 | Weight % of Max | | 25.5 | Average | | 32 | 3 | |
| Image capacity | 40 | WLM Capping % | | 5 | Max | | 4 | 62 | |
| Partition | --MSU -- | | Cap Proc | | -Logical | Util%- | -Physical | | Util%- |
| | Def | Act | Def | Num | Effect | Total | LPAR | Effect | Total |
| *CP | 1 | 2 | | | | | 1.1 | 16.0 | 17.1 |
| DOM1 | 100 | 82 | YES | 2.3 | 11.2 | 12.1 | 0.4 | 2.4 | 2.8 |
| RMF1 | 50 | 29 | NO | 8.2 | 9.3 | 9.8 | 0.1 | 9.2 | 9.3 |
| RMF2 | 40 | 25 | NO | 2.1 | 11.5 | 12.4 | 0.2 | 2.6 | 2.8 |
| RMF3 | 0 | 0 | YES | 2 | 8.8 | 9.3 | 0.3 | 1.8 | 2.1 |
| PHYSICAL | | | | | | | 0.1 | | 0.1 |



ibm.com/redbooks

© 2001 IBM Corporation

Workload Charging

WLC implementation summary:

- Install z900
- Download WLC and SCRT tools and start analyzing reports
- Ensure SMF Type 70 and 89 records are being recorded and kept
- Install z/OS
- Migrate all LPARs from OS/390 to z/OS
- Set up defined capacities based on tools and RMF reports
- Monitor on an ongoing basis



ibm.com/redbooks

© 2001 IBM Corporation

Workload Charging

WLC Reference information:

- Planning for Workload Licence Charges, SA22-7506
- S/390 Software Pricing Reference Guide, G326-0594
- SubCapacity Reporting Tool Users Guide, SG24-6522
- WSC FLASH on Defined Capacity, FLASH10099

All of these are available on the WLC Web site:

- http://www.ibm.com/zseries/wlc_lm

Required service relating to WLC:

- http://www.ibm.com/servers/eserver/zseries/wlc_lm/scrt_support.html



ibm.com/redbooks

© 2001 IBM Corporation

WLC - ISV products

One aspect that IBM cannot directly control is the price of non-IBM software.

Vendors that have announced support for WLC:

- ASG
- BMC
- Candle
- Compuware
- Isogon

However, to help customers compare notes, there is a forum that you can subscribe to where customers can compare experiences with non-IBM software. See


<http://www.can.ibm.com/isvcosts>



ibm.com/redbooks

© 2001 IBM Corporation

Feeding Time



ibm.com/redbooks

© 2001 IBM Corporation

z/OS 1.2 Sync/Async Heuristic Algorithm



© Copyright IBM Corp. 2000

CF Synch/Async heuristic algorithm

Prior to z/OS 1.2, synch to asynch conversion was based on hard-coded rules

- As a result, some synch requests could run for a long time, increasing the overhead of CF requests

In z/OS 1.2, decision about whether a request should be synch or asynch is based on dynamic information

This should reduce CPU overhead for requestors currently experiencing long synch response times

CF Synch/Asynch heuristic algorithm

Synchronous processing:

- Is efficient - requesting CP spins until response is received from CF, saving overhead of being undispached, swapping registers, being dispatched again, and so on
- Can burn MIPS, if CF takes a long time to respond

Asynchronous processing:

- Instead of holding the CP, XES releases the CP once the CF requests have been sent
- Consumes less CPU time than very long running synch requests
- Frees up CPU for other work while the request is processed
- Uses more instructions as dispatcher code is now involved and XES must be undispached then later re-dispatched



ibm.com/redbooks

© 2001 IBM Corporation

CF Synch/Asynch heuristic algorithm

In z/OS 1.2, XES selects the most efficient way (in terms of 'used' machine cycles) of handling each request.

This is based on:

- Current actual synch response times for each structure
 - For duplexed structures, XES keeps the time from the start of the first operation to the end of the last one
- Speed of the CPC that z/OS is running on (indicates number of instructions that could be executed while CP spins on synch request)



ibm.com/redbooks

© 2001 IBM Corporation

CF Synch/Asynch heuristic algorithm

Examples:

- IRLM request on z/900 to distant 9674-C04
 - Fast CPC, slow CF. CF responses are likely to be long, resulting in CP spinning for equivalent of (say) 50,000 instructions*. Asynch request might use 30,000 instructions*, so XES converts synch request to asynch
- IRLM request on 9672-R16 to 2064-100
 - Slow CPC, fast CF. Fast CF responses, resulting in CP spinning for equivalent of (say) 5,000 instructions*. In this case, it is more efficient to spin, so request is issued as synch
- For System-Managed Duplexed structures, requests to both CFs are handled in the same manner for each request - both synch or both asynch

* not actual values - these are just for illustration purposes



ibm.com/redbooks

© 2001 IBM Corporation

CF Synch/Asynch heuristic algorithm

RMF reporting:

- Requests that are converted to asynch are *not* shown as CHANGED in RMF. They are shown as asynch:
 - Changed still infers shared CP, slow CF, or lack of subchannel

| STRUCTURE NAME = PSMGAPPL01 TYPE = LIST STATUS = ACTIVE PRIMARY | | | | | | | | | | | |
|---|---------------|----------|------|------------|----------|---------|--------|------|------------------|---------|------|
| SYSTEM NAME | # REQ | REQUESTS | | TIME (MIC) | | | REASON | | DELATED REQUESTS | | |
| | TOTAL | # | % OF | -SERV | AVG | STD_DEV | # | % OF | AVG | STD_DEV | /ALL |
| NAME | AVG/SEC | REQ | ALL | AVG | STD_DEV | | REQ | REQ | /DEL | | |
| #@\$1 | 3047 10.16 | SYNC | 1899 | 26.4 | 94.6 | 123.4 | NO SCH | 0 | 0.0 | 0.0 | 0.0 |
| | | ASYN | 1148 | 16.0 | 501.5 | 631.8 | PR WT | 0 | 0.0 | 0.0 | 0.0 |
| | | CHNGD | 0 | 0.0 | INCLUDED | IN ASYN | PR CMP | 1715 | 56.3 | 7.9 | 7.7 |
| | | | | | | | DUMP | 0 | 0.0 | 0.0 | 0.0 |
| #@\$2 | 2232 7.44 | SYNC | 1106 | 15.4 | 107.8 | 142.3 | NO SCH | 0 | 0.0 | 0.0 | 0.0 |
| | | ASYN | 1126 | 15.7 | 488.8 | 524.5 | PR WT | 0 | 0.0 | 0.0 | 0.0 |
| | | CHNGD | 0 | 0.0 | INCLUDED | IN ASYN | PR CMP | 1259 | 56.4 | 10.0 | 5.0 |
| | | | | | | | DUMP | 0 | 0.0 | 0.0 | 0.0 |
| #@\$3 | 1907 6.36 | SYNC | 783 | 10.9 | 107.5 | 162.4 | NO SCH | 0 | 0.0 | 0.0 | 0.0 |
| | | ASYN | 1124 | 15.6 | 505.0 | 1171.4 | PR WT | 0 | 0.0 | 0.0 | 0.0 |
| | | CHNGD | 0 | 0.0 | INCLUDED | IN ASYN | PR CMP | 1212 | 63.6 | 10.8 | 0.5 |
| | | | | | | | DUMP | 0 | 0.0 | 0.0 | 0.0 |



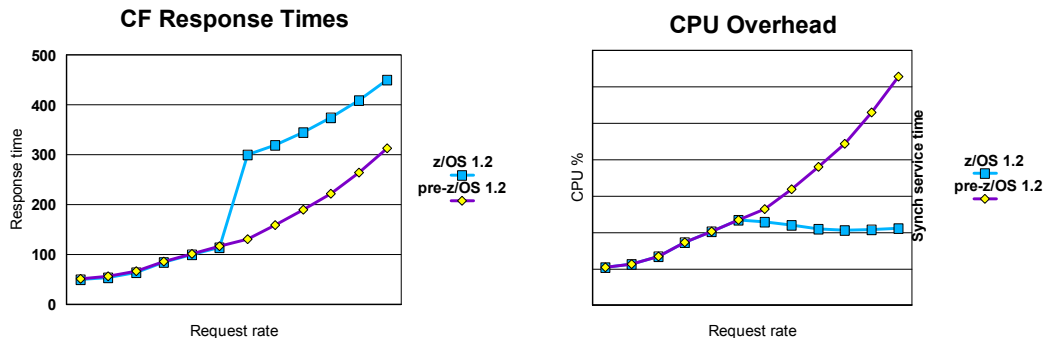
ibm.com/redbooks

© 2001 IBM Corporation

CF Synch/Asynch heuristic algorithm

Considerations:

- If you just install z/OS with no corresponding hardware changes, you may see increased CF response times as long-running synch requests get converted to asynch, however there should be a corresponding drop in overhead




ibm.com/redbooks

© 2001 IBM Corporation

CF Synch/Asynch heuristic algorithm

Considerations:

- If you upgrade just the CPC containing the operating system to a faster machine type, you may find that some CF response times *increase*, because the threshold is lower on a faster CPC, but overhead will decrease correspondingly



ibm.com/redbooks

© 2001 IBM Corporation

CF Synch/Asynch heuristic algorithm

Summary

- New algorithm is automatic, you cannot turn it on or off
- Should result in reduced overhead, especially when CFs are significantly slower or are at a distance from CPCs
- Response time for slower synch requests may increase significantly
- Overall system throughput should improve
 - Overall amount of time spent waiting for CF requests is very small
 - CPs now available for productive work more of the time



ibm.com/redbooks

© 2001 IBM Corporation



System-Managed Duplexing (Finally!)



Redbooks

International Technical Support Organization

© Copyright IBM Corp. 2000

System-Managed Duplexing



Topics:

- Intro
- How does it work
- How to set it up
- Performance considerations
- Who exploits it
- Operational considerations



ibm.com/redbooks

© 2001 IBM Corporation

System-Managed Duplexing



Introduced in z/OS 1.2

May make it possible to do data sharing with NO failure-isolated (standalone) CFs!

Extensive list of exploiters from Day 1

Provides faster recovery from CF or CF Link failures than existing mechanisms

Can significantly improve operability

Requires CFLEVEL 10 (available on z900 and 9672 G5/G6)

Requires links between the CFs



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Introduction



What does it do?

- Allocates and populates a duplicate copy of an existing structure, and establishes connections for all existing connectors
- Keeps both copies in synch, transparently to the connector
- In case of a failure, transparently falls back to simplex mode, keeping the instance which was unaffected by the failure
- Uses existing, familiar, commands and interfaces
 - Definition is identical to User Managed Duplexing
 - Operator commands are identical
 - ▶ Response indicates if structure uses User Managed or System Managed duplexing



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing



How does it improve recovery times? Prior to SM Duplexing, a structure could be recovered in one of three ways:

- Copying the data from the existing structure instance to a new one (fast (maybe, depending on the size of the structure)) - only an option for Link failure situations - source would be gone in case of a CF failure
- Recreating the data from in-storage data in the connected systems (slower)
- Reading the required data from DASD - this often involved restarting the affected subsystems (slow to *really* slow!)

For duplexed structures, there is no recovery involved, so the recovery time is somewhere from a little faster to hugely faster



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Introduction



Starting duplexing

1. Duplexing is started using SETXCF START, RB, DUPLEX command, or automatically if DUPLEX(ENABLED) is specified in CFRM policy (can also be started using IXLREBLD instruction)
2. All existing connectors notified that structure will be temporarily unavailable
3. New structure instance is allocated in appropriate CF and all existing connectors connected to it
4. All or some of existing connected systems copy contents of existing structure over to new instance
5. Connectors notified that they can start using the structure again

Sounds familiar?? Same as beginning of a System-Managed Rebuild process, only "Old" structure instance does not get deleted



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Introduction



Maintaining duplex pair

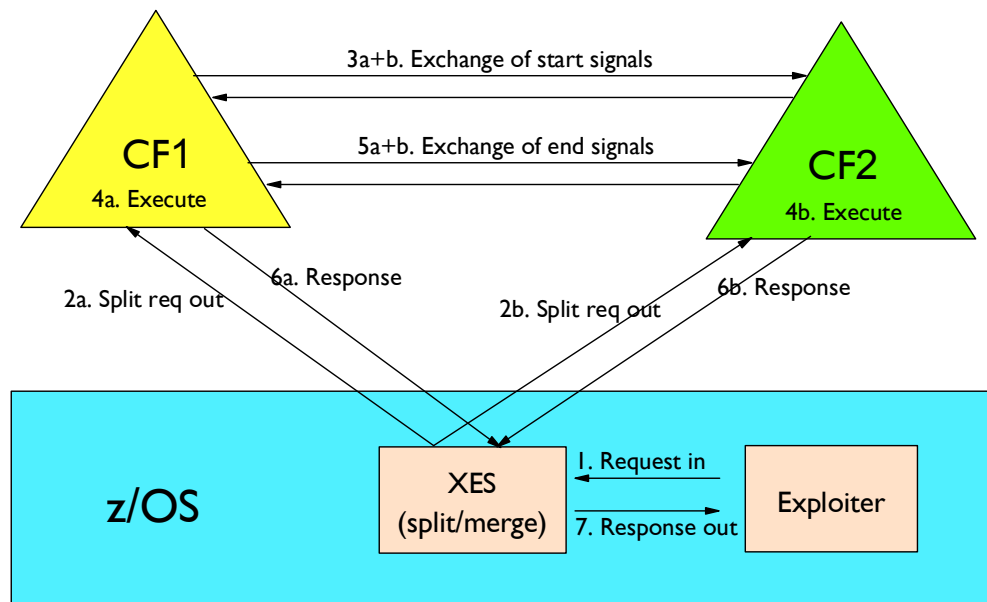
- (Most) Operations that update the structure contents are replicated to the secondary instance
- CFs communicate to synchronize start of processing the request - this serializes access to the data to be changed
- Request is processed
- CFs communicate to agree end of processing - this releases the serialization
- Both CFs send response back to requesting system
- Requesting system determines if request completed successfully on both CFs
- If yes, response is sent back to requestor
- If no, go into error processing



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Introduction



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Introduction



Duplexing is stopped:

- Manually via SETXCF operator command
- Via program interface (IXLREBLD)
- Automatically:
 - When duplex command reconciliation detects failure
 - When failure affects one of the structure instances
 - When CFRM policy change disallows duplexing



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Implementation



Hardware requirements:

- CPC containing CF must be 9672 G5/G6 or 2064 with CFLEVEL 10
- CPC containing operating systems must be 9672 G5/G6 or 2064, with latest microcode levels for optimum performance when sharing links between LPARs
- Need following 'spare' capacity:
 - Storage in CF containing secondary structure instance
 - CF links for new requests to secondary structure instances
 - CPU in both CFs, plus in operating system CPCs
- New links between the CFs
 - Each CF needs *at least* one sender plus one receiver to every other CF - two of each recommended for availability
 - If using ICFs, may be able to share CF Sender links with operating system LPARs - recommend ICB or ICB3 due to reduced latency




ibm.com/redbooks

© 2001 IBM Corporation


SM Duplexing - Implementation



Without Duplexing



With Duplexing (and dedicated CF-to-CF links)



ibm.com/redbooks

© 2001 IBM Corporation

-
- The diagram illustrates the CFLEVEL 10 architecture, showing two parallel processing paths. The left path includes components SYSA, SYSC, SYSD, and CF01. The right path includes components SYSB, TST1, SYSE, and CF02. Both paths are connected to an IRLM (Integrated Resource Link Manager) at the bottom. Arrows indicate data flow from SYSA to SYSE and from SYSC to TST1. A large arrow at the bottom points from the left path to the right path, labeled CFLEVEL 10.



A cartoon illustration of a man with a mustache, wearing a grey suit and a purple tie, looking into a large, ornate mirror. He is adjusting his tie. The background consists of a green triangle on the left and a blue triangle on the right.


The screenshot shows a terminal window titled "Session A - [32 x 80]". The menu bar includes "File", "Edit", "Appearance", "Communication", "Assist", "Window", and "Help". The main display shows a configuration menu with the following options:

- Goto
- Filter
- Backup
- Query
- Help

Below these options, there is a section titled "Available Partition Usage Types" with a table of partitions:

| Partition | Usage | Explanation |
|-----------|-------------------|---------------------------------------|
| 0 | Operating system | |
| 1 | Coupling facility | |
| 2 | CP/OS | Coupling facility or operating system |

The user is currently in "Configuration mode". The bottom of the screen shows a status bar with the text "Connected to remote server/host 9.12.6.54 using port 23".



A cartoon illustration of a man with a large nose and a grey suit, looking at his reflection in a tall, arched mirror. The background consists of a green triangle on the left and a blue triangle on the right.

-
- The diagram shows a hierarchical structure. At the top is a pink triangle labeled "CFLEVEL 10". Below it are three boxes: a yellow box on the left labeled "z/OS 1.1", a green box in the center labeled "z/OS 1.2", and another green box on the right labeled "z/OS 1.2". At the bottom is a cyan cylinder labeled "CFRM CDS" with "ITEM(SMDUPLEX)" below it. Lines connect "CFLEVEL 10" to each of the three boxes. Lines also connect each of the three boxes to the "CFRM CDS" cylinder. Two large red 'X' marks are placed over the lines connecting "CFLEVEL 10" to "z/OS 1.1" and "z/OS 1.1" to "CFRM CDS", indicating a configuration error or a step to be avoided.




```

Session A - [32 x 80]
File Edit Transfer Appearance Communication Assist Window Help
File Edit Edit_Settings Menu Utilities Compilers Test Help

EDIT ESA.SYS1.JCL(DEF CFRM) - 01.10 Columns 00001 00072
***** Top of Data ***** Scroll ==> CSA
Command ==>
000100 //DEFCFRM JOB (0,0),DEF_CFRM',CLASS=A,REGION=4M,
000200 MSGCLASS=X,NOTIFY=&&SYSUID,
000300 //STEP1 EXEC PGM=IXCL10SU
000400 //SYSPCAT DD DSN=MCAT.V@@$#M1,DISP=SHR
000500 //SYSPRINT DD *SOUT=*
000600 //SYSIN DD *
000700 DEFINEDS SYSDPLEX(@$#PLEX)
000800 MGR(SYS1,TEM(15))
000900 DSN(SYS1,XCF.CFRM01) VOLSER(@$#X2)
001000 CATALOG
001100 DATA TYPE(CFRM)
001200 ITEM NAME(POLICY) NUMBER(5)
001300 ITEM NAME(CF) NUMBER(8)
001400 ITEM NAME(STR) NUMBER(200)
001500 ITEM NAME(CONNECT) NUMBER(32)
001600 ITEM NAME(SMDPLEX) NUMBER(11)
001700 002301 ITEM NAME(SMDPLEX) NUMBER(1)
001800 DEFINEDS SYSDPLEX(@$#PLEX)
001900 MGR(SYS1,TEM(15))
002000 DSN(SYS1,XCF.CFRM02) VOLSER(@$#X1)
002100 CATALOG
002200 DATA TYPE(CFRM)
002300 ITEM NAME(POLICY) NUMBER(5)
002400 ITEM NAME(CF) NUMBER(8)
002500 ITEM NAME(STR) NUMBER(200)
002600 ITEM NAME(CONNECT) NUMBER(32)

```



SM Duplexing - Implementation



LOGR CDS

- New format - ITEM(SMDUPLEX) NUMBER(1) keyword
- Can PSWITCH to it non-disruptively, but *not* back
- Only necessary if you want to duplex Logger structures

```
Session A - [32 x 80]
File Edit Transfer Appearance Communication Assist Window Help
Edit Edit Edit_Settings Menu Utilities Compilers Test Help
Command ==>
***** Top of Data *****
000100 //DEFLOGR JOB (0,0) 'DEFINE LOGR CDS', CLASS=A, REGION=4M,
000200 //MSGCLASS=X, NOTIFY=&SYSUID
000300 //STEP1 EXEC PGM=IXCL10SU
000400 //SYSPRINT DD SYSOUT=*
000500 //SYSIN DD *
000600 DEFINEDS
000700   SYSPLX(##$#PLEX)
000800   MAXSYSTEM(4)
000900   DSN(SYS1.XCF.LOGR01) VOLSER(##$#X1)
001000   CATALOG
001100   DATA TYPE(LOGR)
001200     ITEM NAME(LSR) NUMBER(200)
001300     ITEM NAME(LSTR) NUMBER(120)
001400     ITEM NAME(DSEXTENT) NUMBER(10)
001500     ITEM NAME(SMDUPLEX) NUMBER(1)
001600   DEFINEDS
001700     SYSPLX(##$#PLEX)
001800     DSN(SYS1.XCF.LOGR02) VOLSER(##$#X2)
001900     CATALOG
002000     DATA TYPE(LOGR)
002100       ITEM NAME(LSR) NUMBER(200)
002200       ITEM NAME(LSTR) NUMBER(120)
002300       ITEM NAME(DSEXTENT) NUMBER(10)
002400 ***** Bottom of Data *****
```



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Implementation



Once new CDSs are defined and implemented, must define structures as supporting SM Duplexing:

- CFRM Policy updates:
 - Use existing DUPLEX(DISABLED/ALLOWED/ENABLED) keywords
- LOGR Policy updates:
 - DUPLEXing is specified at structure, not logstream, level
 - New DEFINE LOGSTREAM keyword (LOGGERDUPLEX) specifies whether Logger should duplex to DASD even if the logstream is in a duplexed Loader structure

```
Session A - [32 x 80]
File Edit Transfer Appearance Communication Assist Window Help
Edit Edit Edit_Settings Menu Utilities Compilers Test Help
Command ==>
***** Top of Data *****
000100 //LOGRUPD JOB (0,0) 'UPD LOGR DEFNS', CLASS=A, REGION=4M,
000200 //MSGCLASS=X, NOTIFY=&SYSUID
000300 //STEP1 EXEC PGM=IXCL10SU
000400 //SYSPRINT DD SYSOUT=*
000500 //SYSIN DD *
000600 DATA TYPE(LOGR) REPORT(YES)
000700 DEFINE STRUCTURE NAME(SYSTEM.OPERLOG) LOGSNUM(1)
000800   MAXBUFSIZE(65532) AVGBUFSIZE(535)
000900 DEFINE LOGSTREAM
001000   NAME(SYSTEM.OPERLOG) STRUCTNAME(SYSTEM.OPERLOG)
001100   LOGOFFLOAD(0) MODEL(NO) LS_SIZE(1024) LS_DATACLAS(LOGR24K)
001200   LOWOFFLOAD(0) HIGHOFFLOAD(80) STG_DUPLEX(YES)
001300   REIPD(2) AUTODELETE(YES)
001400   LOGSTREAM(1) DATACLAS(LOGR4K)
001500 ***** Bottom of Data *****
```



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Implementation

Considerations for defining duplexed structures:

- While you cannot back out to old format CFRM CDS, you can obviously stop duplexing a given structure at any time, non-disruptively
- Must specify at least two CFs in the PREFLIST
 - XES can only allocate in CFs you name in the PREFLIST
- CFs you specify must have CF-to-CF connectivity



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Performance



Considerations:

- First, have to put this in perspective - how much of a .5 sec (500,000 microsecond) CICS response time is spent waiting on the CF? Maybe 500 microseconds (0.1%)? Also, many large CICS customers still use DASD-only logging with response times around 1000 microseconds and still provide acceptable transaction response times
- Must evaluate savings related to structure recovery
- Performance impact of implementing SM Duplexing is more likely to be a capacity than a transaction response time issue



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Performance



Recovery response time considerations:

- Without duplexing, data must be moved from one CF to the other
- Elapsed time is a factor of the amount of data to be moved, the speed and number of links, the number of connectors, the response time for the CFRM CDS, the speed of the CFs, and the speed of the CPCs

With SM Duplexing:

- Elapsed time is a factor mainly of the number of connectors and the response time of the CFRM CDS. It is unaffected by the structure size
- In ITSO tests with 3 systems, recovery time was consistently about 3 seconds. Other tests with more systems and more structures are in the 12-15 second range. Traditional recovery took up to several minutes



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Performance



Resource considerations:

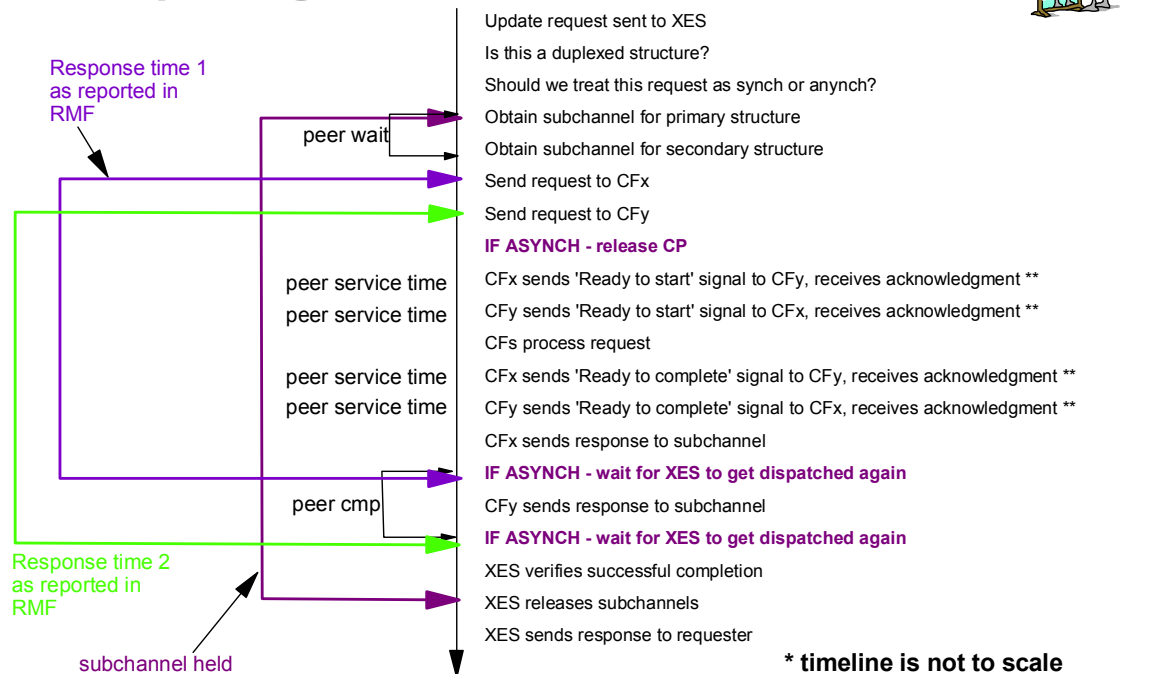
- More code in XES to handle duplexed requests
- Duplexed requests require more processing in *both* CFs - depends on percentage of requests that are duplexed. This will increase CF utilization
 - Rule of Thumb - Expect doubling of CF CPU% for primary structure
- Because duplexed requests take longer, some will be converted to asynch
 - Result of this is that subchannels are tied up for significantly longer
 - ▶ If simplex request takes 60 microseconds, one subchannel is busy for 60 microseconds.
 - ▶ If duplexed request takes 300 microseconds (because of synch to asynch conversion), two subchannels are tied up for 300 microseconds each - a tenfold increase



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing timeline



* timeline is not to scale
** CF to CF signals may overlap



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Performance



| Structure/Product | Structure Type | Estimated % Duplexed Requests |
|-------------------------------|----------------|-------------------------------|
| JES2 | List | 80-100% |
| MQSeries | List | >95% |
| DB2 SCA | List | 50% |
| DB2 IRLM | Lock | >95% |
| IMS IRLM | Lock | >95% |
| IMS CQS | List | >95% |
| IMS VSO | Cache | 10->95% |
| MVS System Logger | List | up to 100% |
| VTAM GR | List | 80+% |
| VTAM MNPS | List | 80+% |
| Workload Manager | List | |
| CICS CF Data Tables | List | ? |
| CICS Named Counter Server | List | ? |
| CICS Shared Temporary Storage | List | ? |



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Monitoring



- RMF Mon III and PP show status (Pri/Sec), count, and rate

| COUPLING FACILITY ACTIVITY | | | | | | | | | | |
|---|------------------|----------------------|------|---------------------------|-----------------|--------------|--------------|--------------|-------------------------|-----------------------|
| z/OS V1R2 | | SYSPLEX #@\$#PLEX | | START 09/12/2001-17.05.00 | | INTERVAL 000 | | | | |
| | | RPT VERSION V1R2 RMF | | END 09/12/2001-17.10.00 | | CYCLE 01.000 | | | | |
| ----- | | | | | | | | | | |
| COUPLING FACILITY NAME = FACIL03 | | | | | | | | | | |
| TOTAL SAMPLES (AVG) = 299 (MAX) = 300 (MIN) = 298 | | | | | | | | | | |
| ----- | | | | | | | | | | |
| COUPLING FACILITY USAGE SUMMARY | | | | | | | | | | |
| ----- | | | | | | | | | | |
| STRUCTURE SUMMARY | | | | | | | | | | |
| ----- | | | | | | | | | | |
| TYPE | STRUCTURE NAME | STATUS | CHG | ALLOC SIZE | % OF CF STORAGE | # REQ | % OF ALL REQ | AVG REQ/ SEC | LST/DIR ENTRIES TOT/CUR | DATA ELEMENTS TOT/CUR |
| LIST | CIC_DFHSHUNT_001 | ACTIVE | | 3M | 0.5 | 0 | 0.0 | 0.00 | 1175 | 3525 |
| | | | | | | | | | 9 | 64 |
| LIST | I#\$#EMHQ | ACTIVE | | 4M | 0.8 | 0 | 0.0 | 0.00 | 3211 | 3209 |
| | | | | | | | | | 6 | 5 |
| | PSMGAPPL01 | ACTIVE | X | 5M | 1.0 | 3474 | 13.5 | 11.58 | 1097 | 6581 |
| | | | SEC | | | | | | 995 | 2000 |
| | PSMGCSQ_ADMIN | ACTIVE | | 10M | 2.0 | 39 | 0.2 | 0.13 | 9996 | 20K |
| | | | PRIM | | | | | | 5 | 80 |



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Monitoring



- RMF PP Structure Activity Report
 - One for each instance (Primary and Secondary)
 - PR WT always about 100% for secondary - monitor delay times
 - PR CMP shows disparity in response times of CFs

| COUPLING FACILITY STRUCTURE ACTIVITY | | | | | | | | | | | |
|--------------------------------------|---------|----------|-------------|--------------------|------------------|---------------------------|--------|-------|----------------|------|------|
| ----- | | | | | | | | | | | |
| STRUCTURE NAME = PSMGAPPL01 | | | TYPE = LIST | | | STATUS = ACTIVE SECONDARY | | | | | |
| ----- | | | | | | | | | | | |
| SYSTEM | # REQ | REQUESTS | | | DELATED REQUESTS | | | ----- | | | |
| NAME | TOTAL | # | % OF | -SERV TIME (MIC) - | REASON | # | % OF | ---- | AVG TIME (MIC) | ---- | |
| | AVG/SEC | REQ | ALL | AVG | STD_DEV | REQ | REQ | /DEL | STD_DEV | /ALL | |
| ----- | | | | | | | | | | | |
| #@\$1 | 1868 | SYNC | 284 | 8.2 | 103.1 | 184.7 | NO SCH | 0 | 0.0 | 0.0 | 0.0 |
| | 6.23 | ASYNC | 1584 | 45.6 | 461.0 | 475.7 | PR WT | 1887 | 101 | 4.6 | 0.7 |
| | | CHNGD | 0 | 0.0 | INCLUDED | IN ASYNC | PR CMP | 245 | 13.1 | 7.8 | 5.5 |
| | | | | | | | DUMP | 0 | 0.0 | 0.0 | 0.0 |
| #@\$2 | 744 | SYNC | 96 | 2.8 | 162.3 | 284.3 | NO SCH | 0 | 0.0 | 0.0 | 0.0 |
| | 2.48 | ASYNC | 648 | 18.7 | 481.3 | 414.6 | PR WT | 744 | 100 | 5.3 | 0.8 |
| | | CHNGD | 0 | 0.0 | INCLUDED | IN ASYNC | PR CMP | 53 | 7.1 | 8.4 | 4.0 |
| | | | | | | | DUMP | 0 | 0.0 | 0.0 | 0.0 |
| #@\$3 | 862 | SYNC | 194 | 5.6 | 111.9 | 168.2 | NO SCH | 0 | 0.0 | 0.0 | 0.0 |
| | 2.87 | ASYNC | 668 | 19.2 | 454.7 | 112.6 | PR WT | 868 | 101 | 5.2 | 0.9 |
| | | CHNGD | 0 | 0.0 | INCLUDED | IN ASYNC | PR CMP | 148 | 17.2 | 11.5 | 39.2 |
| | | | | | | | | | | | 2.0 |



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Monitoring



- RMF PP CF to CF Activity Report
 - One for each CF
 - Shows requests sent *from* that CF
 - Shows all requests as SYNC; in fact some are ASYNC and some are SYNC
 - Number of requests roughly 2.0 - 2.5 times number of requests to duplexed structures
 - Aim to keep service time as low as possible
 - Service time depends on number and type of links, extent of link sharing, CF CPU utilization

| CF TO CF ACTIVITY | | | | | | | | | | | |
|-------------------|-------|---------|----------------|-----|----------|----------------|-----------------------|-----|-------------|-------------|-----|
| PEER CF | # REQ | | -- CF LINKS -- | | REQUESTS | | | | DELATED REQ | | |
| | TOTAL | AVG/SEC | TYPE | USE | # REQ | SERVICE AVG | TIME (MIC) STD_DEV | | # REQ | % OF REQ | DEL |
| FACIL04 | 8914 | 29.7 | CBR | 1 | SYNC | 8914 | 7.5 | 0.0 | SYNC | 0 | 0.0 |



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Capacity planning




Need spare cycles in CPC

- Overhead will depend on:
 - Number of requests to duplexed structure
 - Percent of requests that are updates
 - Structure type

Need capacity on links:

- To secondary CF: all writes to duplexed primary structure will be replicated to secondary
- Between CFs - bandwidth requirement not expected to be large



Need storage in CF (should be there as white space)

Need spare cycles in CFs

- Recommend that CF CPU utilization not exceed 25-30%
- Need additional capacity in BOTH CFs



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Exploiters



Products supporting SM Duplexing at GA

- JES2
- MQ Series
- DB2 SCA
- IRLM (IMS, DB2)
- IMS Shared Message Queue
- IMS Shared VSO databases
- System logger (CICS, IMS, RRS, OPERLOG, LOGREC)
- VTAM
- WLM
- Smartbatch



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Exploiters



| Structure/Product | Required release or APAR |
|-------------------------------|------------------------------------|
| JES2 | z/OS 1.2 base |
| MQSeries | MQSeries V5.2 base |
| DB2 SCA | DB2 V7.1 base |
| DB2 IRLM | PQ52341 and PQ48996 |
| IMS IRLM | PQ45407 and PQ48823 |
| IMS CQS | IMS V7.1 plus APAR PQ47642 |
| IMS VSO | IMS V7.1 plus APAR PQ50661 |
| MVS System Logger | z/OS 1.2 base |
| VTAM GR | z/OS 1.2 base |
| VTAM MNPS | z/OS 1.2 base |
| Workload Manager | z/OS 1.2 base |
| BatchPipes | |
| CICS Shared Temporary Storage | CICS TS 2.2 plus z/OS APAR OW39892 |
| CICS CF Data Tables | CICS TS 2.2 plus z/OS APAR OW39892 |
| CICS Named Counter Server | CICS TS 2.2 plus z/OS APAR OW39892 |

* need OW41617 - XCF enabling APAR for SM Duplexing



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Exploiters



Products that have announced support in a future release:

- DFSMSHsm Common Recall Queue
- DFSMS VSAM Record Level Sharing Lock structure

Many products do not require failure independence (GRS Star, for example) and therefore are unlikely to ever add support for SM Duplexing - it simply is not needed for these



ibm.com/redbooks


© 2001 IBM Corporation

SM Duplexing - Exploiters



Configuration used for tests:

- 9672-Z47
- 4 LPs - 2 z/OS 1.2, 2 CFs
- z/OS LPs with 1 Ded CP each
- CF LPs with 1 Ded ICF each
- ICs for links from OSs to 1 CF
- ICBs for links from OSs to other CF
- ICBs for CF-to-CF links



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Tests



DISCLAIMER!

- Tests were *not* run in a controlled performance test environment
- Test jobs are *not* representative of production work
- Environment was more constrained than normal production configuration
- Tests were run on pre-GA hardware and pre-GA software
- Your results may differ significantly from these numbers
- Tests only had two systems and limited number of structures, so recovery times are probably a bit better than would be expected normally (less CFRM CDS contention)



ibm.com/redbooks


© 2001 IBM Corporation

SM Duplexing - JES2




Prior to duplexing:

- Without duplexing, any error affecting the JES2 Checkpoint structure results in JES2 forwarding to NEWCKPTx, probably on DASD, delaying JES2 operations.



With duplexing:

- If the checkpoint structure is duplexed, the failure is transparent (except for some XES messages). Checkpoint reconfiguration dialog is not invoked, jobs are not delayed, and operators are not involved.



Loss of CF



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - JES2



JES2 duplexing test

- Submit 4,000 jobs across 2 systems
- Test recovery - very short recovery, no need to respecify new NEWCKPTx
- Measure elapsed time, CPU utilization, and CF response time changes

| Test | No Dup | Dup |
|-----------------------------|--------------|------------------------|
| Elapsed (CHKPT) | 00:44:16 | N/A |
| Elapsed | 00:39:37 | 00:40:10 |
| CPU utilization | 15.5% | 15.8% |
| CF Utilization | 4.1/1.3% | 4.6/1.2% |
| % of requests duplexed | N/A | 81% |
| CF response time | 23.6/397 | 57/463(P) 58/255(S) |
| CF Req/sec | 271 | 270(P) 220(S) |
| CF failure recovery | Op Intv Reqd | 00:00:03 |
| CF link failure recov | Op Intv Reqd | 00:00:03 |
| CF to CF link failure recov | N/A | 00:00:02 |



ibm.com/redbooks


© 2001 IBM Corporation

SM Duplexing - MQ Series




Prior to duplexing:

- MQ structure could be rebuilt after a connectivity failure, but if there is a CF failure, all messages in structure are lost. After a connectivity failure, the MQ that lost connectivity abends.



With duplexing:

- Because MQ 5.2 supports System-Managed Rebuild, support for SM Duplexing is already there. In case of connectivity or CF failure, structure reverts to simplex mode, all MQs are unaffected.



Loss of CF



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - MQSeries



MQ duplexing test

- Run 3 jobs on 2 systems, each job writing 15000 messages (all to the same queue) then reading messages off the queue
- Test recovery - Make sure you duplex the ADMIN structure as well as the APPL ones!
- Measure elapsed time, CPU utilization, and CF response time changes

| Test | No Dup | Dup |
|-----------------------------|----------------------|----------------------|
| Elapsed | | |
| CPU utilization | 61.7%(1) 56.9%(2) | 68.5%(1) 67.6%(2) |
| CF Utilization | 8.5% | 15.5%(P) 13.9%(S) |
| % of requests duplexed | N/A | 100% |
| CF response time | 41.4 | 85.8(P) 82.4(S) |
| CF Req/sec | 1880 | 2080(P) 2077(S) |
| CF failure recovery | All MQs abend | 00:00:05 |
| CF link failure recov | Affected MQ Abends | 00:00:07 |
| CF to CF link failure recov | N/A | 00:00:10 |



ibm.com/redbooks


© 2001 IBM Corporation

SM Duplexing - DB2 SCA




Prior to duplexing:

- The DB2 SCA contains information about database errors and the restart status of members of the data sharing group. Normally, the SCA is recovered using information from the virtual storage of all connected members. If there is a double failure affecting both the structure and one of the connected members, a DB2 group restart must be done to recover the missing information.



With duplexing:

- Duplexing the SCA means that you can lose a CF **and** a connected DB2 member and still not have to do a group restart.



Double failure



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - DB2 SCA



DB2 SCA duplexing test

- Submit 2 DB2 jobs on each system, doing variety of reads and writes
- Test recovery from double failure

| Test | No Dup | Dup |
|-----------------------------|---------------|----------|
| Elapsed | | |
| CPU utilization | | |
| CF Utilization | | |
| % of requests duplexed | N/A | 50% |
| CF response time | | |
| CF Req/sec | | |
| CF failure recovery | 00:00:05(*) | 00:00:03 |
| CF + OS failure | Group Restart | 00:00:03 |
| CF link failure recov | | 00:00:02 |
| CF to CF link failure recov | N/A | 00:00:04 |



ibm.com/redbooks


© 2001 IBM Corporation

SM Duplexing - IRLM




Prior to SM Duplexing:

- The IRLM lock structure contains information about locks held by various members of the data sharing group. Normally, if the structure is lost, the information in the structure is rebuilt from information held in the virtual storage of all the connected IRLM subsystems. In case of a double failure without duplexing, a group restart is required.



With SM Duplexing:

- IRLM survives a double failure, with minimal recovery time. No need to restart surviving subsystems



Double failure



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - IRLM



IRLM duplexing test

- Submit 2 DB2 jobs on each system, doing variety of reads and writes
- Test recovery
- Measure elapsed time, CPU utilization, and CF response time changes
- Test recovery from double failure

| Test | No Dup | Dup |
|-----------------------------|--------------------|--------------------|
| Elapsed | | |
| CPU utilization | 23.2(1) 31.9(2) | 24.6(1) 40.7(2) |
| CF Utilization | 1.5%(P) 2.9%(S) | 5.1%(P) 7.6%(S) |
| % of requests duplexed | N/A | 100% |
| CF response time | 10.7 | 51.8(P) 44.5(S) |
| CF Req/sec | 1804 | 1729(P) 1729(S) |
| CF failure recovery | 00:00:05(*) | 00:00:03 |
| CF + OS failure | Group Restart | 00:00:03 |
| CF link failure recov | | 00:00:02 |
| CF to CF link failure recov | N/A | 00:00:04 |



ibm.com/redbooks


© 2001 IBM Corporation

SM Duplexing - IMS SMQ




Prior to SM Duplexing:

- If the message queue structure is lost, checkpoint information must be read from DASD, and changes from Logger reapplied. Recovery takes about 1 min for 3 mins of transactions since last checkpoint. More frequent ckpts = less recovery time, but CQS can't access queue during checkpoint.



With SM Duplexing:

- No need to recover anything from DASD - just revert to simplex mode. Can increase interval between checkpoints, improving CQS availability



Loss of CF



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - SMQ



SMQ duplexing test

- Submit 37 jobs on each of 2 systems. Each job places 500 messages on IMS queue
- Test recovery
- Measure elapsed time, CPU utilization, and CF response time changes

| Test | No Dup | Dup |
|-----------------------------|----------|----------|
| Elapsed | | |
| CPU utilization | | |
| CF Utilization | | |
| % of requests duplexed | N/A | 100% |
| CF response time | | |
| CF Req/sec | | |
| CF failure recovery | 00:00:20 | 00:00:06 |
| CF link failure recov | 00:00:15 | 00:00:03 |
| CF to CF link failure recov | N/A | 00:00:03 |



ibm.com/redbooks


© 2001 IBM Corporation

SM Duplexing - IMS VSO




Prior to SM Duplexing:

- Data is written to the VSO structure and later written to DASD. To protect the data, IMS can keep two copies of the structure. If one structure fails, IMS can continue with the remaining structure, but IMS commands required to re-duplex.



With SM Duplexing:

- If one CF fails, XES automatically re-duplexes as soon as 2nd CF is available. No operator intervention required. No special procedures required for recovery of these structures



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - VSO



IMS VSO duplexing test

- Submit 1 job on each of 2 systems. Each job does a variety of reads, updates, inserts, and deletes.
- Remove definition of 2nd structure from IMS (specified as CFSTR2 on INIT.DBDS stmt)
- Test recovery
- Measure elapsed time, CPU utilization, and CF response time changes

| Test | No Dup | Dup |
|-----------------------------|--------|--------|
| Elapsed | | |
| CPU utilization | | |
| CF Utilization | | |
| % of requests duplexed | 0-100% | 0-100% |
| CF response time | | |
| CF Req/sec | | |
| CF failure recovery | | |
| CF link failure recov | | |
| CF to CF link failure recov | | |



ibm.com/redbooks


© 2001 IBM Corporation

SM Duplexing - CICS Log




Prior to SM Duplexing:

- Recommendation is that CICS DFHLOG should be duplexed to DASD if logstream is not failure-isolated. This has a performance penalty. If it is NOT duplexed, a single CPC failure could result in a CICS cold start.



With SM Duplexing:

- DFHLOG logstream no longer needs to be failure-isolated from CICS. If a CPC fails, there will be a copy of the data in the other CF. Instead of a cold start (or the cost of duplexing to DASD), there is practically no impact.



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - CICS Log



CICS Log duplexing test

- Run 100 concurrent transactions in 1 AOR in each of 2 systems
- Logstreams duplexed to staging data set because of lack of failure-isolated CF - with SM Duplexing, staging can be bypassed
- Test recovery
- Measure elapsed time, CPU utilization, and CF response time changes

| Test | Simp/no staging | Simp/staging | Dup/ no staging |
|-----------------------------|-----------------|--------------|---------------------|
| Elapsed | | | |
| CPU utilization | 13.1 14.0 | 14.5 11.8 | 15.3 11.2 |
| CF Utilization | 1.5(P) | 1.2(P) | 1.6(P) 1.1(S) |
| % of requests duplexed | N/A | N/A | 38.5% |
| CF response time | 61.2(P) | 67.9(P) | 75.7(P) 64.9(S) |
| CF Req/sec | 204.7 | 187.1 | 204.2(P) 78.4(S) |
| CF failure recovery | | | |
| CF link failure recov | | | |
| CF to CF link failure recov | | N/A | |



ibm.com/redbooks


© 2001 IBM Corporation

SM Duplexing - CICS Temporary Storage




Prior to SM Duplexing:

- Currently, CICS has NO recovery support for the TS structure. The data in the structure can be offloaded and reloaded, but this is disruptive. There is no way to non-disruptively move the structure, even for a planned outage. CICS 2.2 will provide rebuild support, but only for planned changes



With SM Duplexing:

- For planned or unplanned outages, structure simply reverts to simplex mode (assuming CICS 2.2 is installed)



CF failure



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Operations



Operational aspects:

- Changes to CF drain procedures
- Changes to CF recovery procedures
- Ongoing monitoring
- Changes to product recovery procedures
- Additional information on XCF commands



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Operations



Drain procedures:

- Without duplexing, draining a CF consisted of issuing "SETXCF START, RB, CFNM=cfname, LOC=OTHER" for all structures that supported rebuild, plus manual procedures for those that do not.
- With duplexing, you must first stop duplexing for any structures in the CF to be drained, with "SETXCF STOP, RB, DUPLEX, CFNM=cfname". This will deallocate any duplexed structures (OLD or NEW) that are in the named CF. Once this is complete, proceed as normal.



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Operations



Recovery prodedures

- When CF comes back online, XES will automatically re-duplex any structures defined with DUPLEX(ENABLED)
- Non-duplexed structures should be moved back into the recovered CF using SETXCF START, RB, POPCF=cfname
- After this completes, reduplex any DUPLEX(ALLOWED) structures with "SETXCF START, RB, DUPLEX, STRNM=whatever"



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Operations



Ongoing monitoring:

- When duplexing is stopped because of an error, the structure will end up in the CF with the best connectivity. When you start duplexing again, the NEW instance will go in the recovered CF. This might not be where you want it (remember, all reads are done to the OLD instance), so consider automation to check that OLD instance is in first CF in the PREFLIST
- If the CF-to-CF link fails, you may not be informed of that failure until it is recovered. Enhanced D CF command shows status of that link.
- If you make a change that requires a structure rebuild to complete (increasing the size of a lock structure, for example), you must stop duplexing that structure, do a rebuild in place, then start duplexing again (if DUPLEX(ALLOWED))



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Operations



Changes to product recovery procedures:

- JES2. If a CF (or a link to that CF) containing the JES2 checkpoint fails, it is no longer necessary to use the Reconfiguration Dialog, nor to specify a new NEWCKPTx. Recovery is the same as other products - using XCF commands.
- IMS VSO. If a CF (or a link to that CF) containing one of the VSO structures fails, it is no longer necessary to use /VUNLOAD and /STA commands to reduplex the structure. Structure will reduplex automatically (if DUPLEX(ENABLED)) or with SETXCF command.



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Operations



- Enhanced D CF command:

```
DISPLAY CF
SENDER PATH      PHYSICAL      LOGICAL      CHANNEL TYPE
      80          ONLINE        ONLINE        CFS
      C0          ONLINE        ONLINE        CFS
COUPLING FACILITY DEVICE      SUBCHANNEL  STATUS
CF00              0056        OPERATIONAL/IN USE
CF01              0057        OPERATIONAL/IN USE
CF02              0058        OPERATIONAL/IN USE
CF03              0059        OPERATIONAL/IN USE

REMOTELY CONNECTED COUPLING FACILITIES
CFNAME           COUPLING FACILITY
-----
LF02              SIMDEV.IBM.EN.ND0200000000
PARTITION: 0  CPCID: 00
CHPIDS ON TESTCF CONNECTED TO REMOTE FACILITY
RECEIVER:  CHPID  TYPE
           F1     CFR
SENDER:    CHPID  TYPE
           E1     CFS
LF01              SIMDEV.IBM.EN.ND0100000000
PARTITION: 0  CPCID: 00
CHPIDS ON TESTCF CONNECTED TO REMOTE FACILITY
RECEIVER:  CHPID  TYPE
           F0     CFR
SENDER:    CHPID  TYPE
           E0     CFS

NOT OPERATIONAL CHPIDS ON TESTCF
F2
```



ibm.com/redbooks

© 2001 IBM Corporation

SM Duplexing - Summary



SM Duplexing summary:

- Implementation is easy and externals are familiar
- Provides significant availability, ease of operation, and potentially cost benefits
- Performance cost must be evaluated
 - Use the new release of SPSSZR available later this year
- For ease of operation, specify structures as DUPLEX(ENABLED) once you complete successful testing

Documentation:

- z/OS 1.2 Setting Up a Sysplex
- System-Managed CF Structure Duplexing white paper (GM13-0103), available 10/01 - off Parallel Sysplex Home Page
- z/OS Version 1 Release 2 Implementation, SG24-6235



ibm.com/redbooks

© 2001 IBM Corporation

Fix time



ibm.com/redbooks

© 2001 IBM Corporation



z/OS Auto Tape Switching Improvements



© Copyright IBM Corp. 2000

Auto Tape Switching

Automatic Tape Switching

- Replaces IEFAUTOS for sharing assignable tape devices within a single sysplex
 - 3420 not supported
- Compatible with IEFAUTOS
 - Tapes still defined as AUTOSWITCHABLE in HCD, and ONLINE to all systems
- Planned Delivery:
 - Delivered via PTF after GA product delivery of z/OS R2



ibm.com/redbooks

© 2001 IBM Corporation

Auto Tape Switching

Removes use of IEFAUTOS structure

Tape drive information is maintained in the ALLOCAS address space

- Information is shared via XCF messaging

Tape drive serialization is managed via GRS

- Shared tapes are SYSTEMS scope
- Dedicated tapes are SYSTEM scope

Parallel Sysplex not required:

- Works in Basic Sysplex with GRS Ring but performance is significantly better with GRS Star



ibm.com/redbooks

© 2001 IBM Corporation

Auto Tape Switching

Enhanced debug via D GRS,[CIANALYZE]:

- 3480: SYSZATS.DEVxxxx
- Self describing: SYSZATS.NEDxxxxxxxxxxxxxx

Performance improvement (if GRS Star)

- Far fewer, shorter CF accesses

**Positions Allocation for future exploitation of GRS
Multi-Sysplex Support**

- Share tape drives across multiple sysplexes



ibm.com/redbooks

© 2001 IBM Corporation

Auto Tape Switching

- Updates to D U,,AS command
 - ▶ All online shared tape devices listed, even if not online to the system the command was issued on

```
IEE343I 22.44.04 UNIT STATUS 594
AUTOSWITCHABLE DEVICES CONNECTED TO SYSTEM FAGEN1
UNIT  TYPE  STATUS      SYSTEM  JOBNAME  ASID  VOLSER  VOLSTATE
05A2  348S
05B0  349S  A           FAGEN3   HOLDTAP2  0066  /REMOV
      FAGEN3
05B1  349S                               /REMOV
AUTOSWITCHABLE DEVICES NOT CONNECTED TO SYSTEM FAGEN1
UNIT  TYPE  STATUS      SYSTEM  JOBNAME  ASID  VOLSER  VOLSTATE
      /REMOV

FAGEN3(05A0,348S)
```



ibm.com/redbooks

© 2001 IBM Corporation



msys for Operations



© Copyright IBM Corp. 2000

Topics to discuss

msys for Ops functions

- Background
- Functions delivered by msys for Ops

msys for Ops installation

- FMIDs and relationship to NetView and SA/390
- Coexistence with NetView and SA/390
- Coexistence with other automation products
- Installation wizard
- Customization

msys for Ops future directions



ibm.com/redbooks

© 2001 IBM Corporation

msys for Ops - Background

Developed jointly by SA/390 Development and Sysplex customer support group in Poughkeepsie

Certain scenarios with common reactions

Based on analysis of customer outages

No direct relationship to msys for Setup

Advantage (over providing your own automation) is that IBM develops and maintains the code, including updates to match new system messages and capabilities (duplexing, for example)


Two aspects - automation for common events, panels for improved usability



ibm.com/redbooks

© 2001 IBM Corporation

msys for Ops functions



1. Check Log Stream Data Set 'SHAREOPTIONS'

2. System Log Recovery
3. Recovery of Log Stream Data Set Directory Shortage
4. Alternate Couple Data Set Allocation
5. WTO/R Buffer Shortage Recovery
6. INGPLEX SYSTEM
7. INGPLEX CONSOLE
8. INGPLEX CF
9. INGPLEX CDS



ibm.com/redbooks

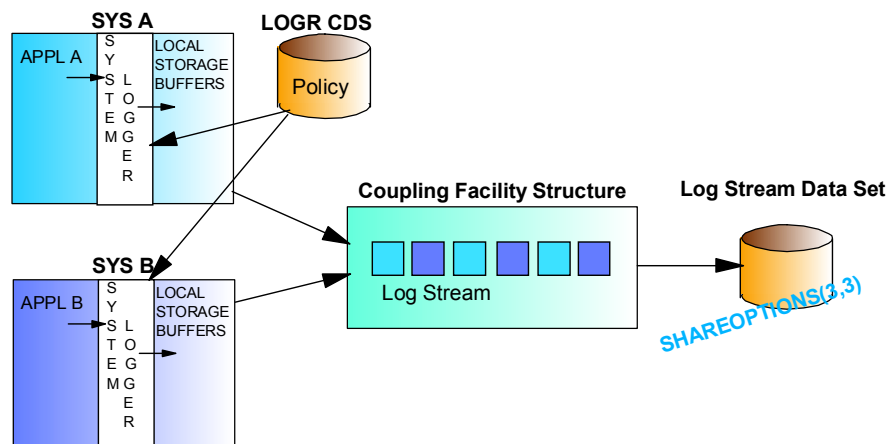
© 2001 IBM Corporation

Checking Logger shareoptions

■ SMS Requirement for System Logger : SHAREOPTIONS(3,3) !!!

■ If SHAREOPTIONS is not set to (3,3) then

- Allocating an Offload Data Set will fail
- Offload will not complete
- Other Systems may be unable to connect to structure



Literature: OS/390 MVS Setting Up a Sysplex, Chapter 9.4.1.4. / 9.4.5.7.



ibm.com/redbooks

© 2001 IBM Corporation

Checking Logger shareoptions

1. Find out the Log Stream Data Set names.
2. Check for SHAREOPTIONS(3,3).
3. If it's different: Issue AOF912I

Invoked :

1. Automatically scheduled every day at 11:10



ibm.com/redbooks

© 2001 IBM Corporation

Checking Logger shareoptions

In AOF CUST member (DSIPARM lib):

```
AUTO(  
  LOGGER  
)
```



ibm.com/redbooks

© 2001 IBM Corporation

msys for Ops functions

- ➔
1. Check Log Stream Data Set 'SHAREOPTIONS'
 2. System Log Recovery
 3. Recovery of Log Stream Data Set Directory Shortage
 4. Alternate Couple Data Set Allocation
 5. WTO/R Buffer Shortage Recovery
 6. INGPLEX SYSTEM
 7. INGPLEX CONSOLE
 8. INGPLEX CF
 9. INGPLEX CDS



ibm.com/redbooks


© 2001 IBM Corporation

System log recovery

IEE037D indicates that the SYSTEM LOG is NOT ACTIVE.

If the system log is critical an immediate log (re)start must be done by the operator.

If the operator misses that message, vital data may be lost.



ibm.com/redbooks

© 2001 IBM Corporation

System log recovery

Upon IEE043I or IEE533E followed by IEE037D:

WRITELOG START

Upon IEE769E (recursive abends) followed by IEE037D:

No automation

Upon IEE041I (syslog available again):

VARY SYSLOG,HARDCPY

Invocation:

The action is automatically taken upon receiving the above messages.



Literature: OS/390 MVS System Messages Vol.4



ibm.com/redbooks

© 2001 IBM Corporation

System log recovery

In AOFCUST member (DSIPARM lib):

**AUTO(
LOG
)**




ibm.com/redbooks

© 2001 IBM Corporation

msys for Ops functions

1. Check Log Stream Data Set 'SHAREOPTIONS'

2. System Log Recovery

 3. Recovery of Log Stream Data Set Directory Shortage

4. Alternate Couple Data Set Allocation

5. WTO/R Buffer Shortage Recovery

6. INGPLEX SYSTEM

7. INGPLEX CONSOLE

8. INGPLEX CF

9. INGPLEX CDS



ibm.com/redbooks

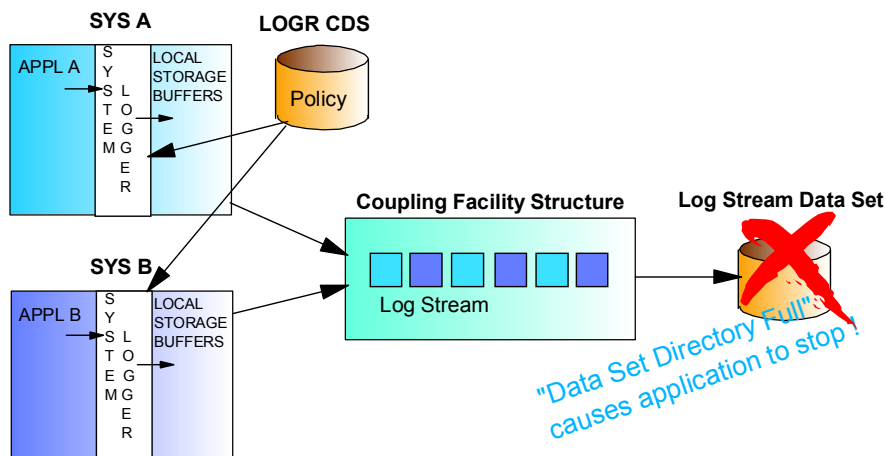
© 2001 IBM Corporation

Logstream directory shortage recovery

Resulting messages:

■ IXG257I DATA SET DIRECTORY FOR LOGSTREAM logstream IN STRUCTURE strname IS OVER 90% FULL.

■ IXG261E SHORTAGE OF DIRECTORY EXTENT RECORDS TOTAL n1 IN USE: n2 AVAILABLE: n3



Literature: OS/390 MVS Setting Up a Sysplex, Chapter 9.4.5.2



ibm.com/redbooks

© 2001 IBM Corporation

Logstream directory shortage recovery

1. Find out the number of directories
2. Extend the number of directories in LOGR policy (DSEXTENT).
3. Reformat and reallocate LOGR CDS's
4. Issue message when done



ibm.com/redbooks

© 2001 IBM Corporation

Logstream directory shortage recovery

1. ... automatically when receiving messages
IXG257I or IXG261E

IXG257I DATA SET DIRECTORY FOR LOGSTREAM logstream IN STRUCTURE strname IS
OVER 90% FULL.

IXG261E SHORTAGE OF DIRECTORY EXTENT RECORDS TOTAL decimalnumber1 IN USE:
decimalnumber2 AVAILABLE: decimalnumber3



ibm.com/redbooks

© 2001 IBM Corporation

Logstream directory shortage recovery

In AOFCUST member (DSIPARM lib):

```
AUTO(  
  LOGGER  
  CDS  
)
```


```
CDS(  
  HLQ hlq1.hlq2.hlq3  
  VOL (LOGR,vol1,vol2,vol3,vol4,vol5,vol6,vol7,vol8)  
)
```



ibm.com/redbooks

© 2001 IBM Corporation

msys for Ops functions

1. Check Log Stream Data Set 'SHAREOPTIONS'
2. System Log Recovery
3. Recovery of Log Stream Data Set Directory Shortage
-  **4. Alternate Couple Data Set Allocation**
5. WTO/R Buffer Shortage Recovery
6. INGPLEX SYSTEM
7. INGPLEX CONSOLE
8. INGPLEX CF
9. INGPLEX CDS




ibm.com/redbooks

© 2001 IBM Corporation

msys for Ops functions

Problem: Loss of Couple Data Set

- CDS's are critical resources in a Parallel Sysplex
 - SYSPLEX and CFRM CDS failures can cause sysplex outages
- Single-Point-Of-Failure Situation if no alternate CDS is allocated



ibm.com/redbooks

© 2001 IBM Corporation

Alternate CDS allocation

- msys for Operations finds out which alternate CDS is not allocated.
- It reads the candidate list of DASD's where an alternate CDS can be formatted and allocated.
 - The list of DASD's can be predefined by the user.
- It gets the first available DASD and formats/allocates the respective alternate CDS
 - Using the HLQ defined by the user.
- msys4Ops supports the following CDS types:
SYSPLEX, CFRM, ARM, LOGR, SFM



ibm.com/redbooks

© 2001 IBM Corporation

Alternate CDS allocation

1. ...at *msys for Operations* start up time
2. ...at runtime when receiving IXC253I₍₁₎



ibm.com/redbooks

© 2001 IBM Corporation

Alternate CDS allocation

In AOFCUST member (DSIPARM lib):

```
AUTO(  
  CDS  
)
```

```
CDS(  
  HLQ hlq1.hlq2.hlq3  
  VOL (SYSPLEX,vol1,vol2,vol3,vol4,vol5,vol6,vol7,vol8)  
  VOL (CFRM,vol1,vol2,vol3,vol4,vol5,vol6,vol7,vol8)  
  VOL (ARM,vol1,vol2,vol3,vol4,vol5,vol6,vol7,vol8)  
  VOL (LOGR,vol1,vol2,vol3,vol4,vol5,vol6,vol7,vol8)  
  VOL (SFM,vol1,vol2,vol3,vol4,vol5,vol6,vol7,vol8)  
)
```




ibm.com/redbooks

© 2001 IBM Corporation

msys for Ops functions

1. Check Log Stream Data Set 'SHAREOPTIONS'
2. System Log Recovery
3. Recovery of Log Stream Data Set Directory Shortage
4. Alternate Couple Data Set Allocation



5. WTO/R Buffer Shortage Recovery

6. INGPlex SYSTEM
7. INGPlex CONSOLE
8. INGPlex CF
9. INGPlex CDS



ibm.com/redbooks

© 2001 IBM Corporation

WTO Buffer shortage recovery

- When messages fill up the buffers, most jobs that issue a WTO or WTOR go into a wait until buffers are available.

Some reasons

- Unattended MCS consoles queueing up message buffers
- Program loops causing message flooding
- Repetitive job starts causing WTOR buffer shortages
- CONSOLxx PARMLIB member containing insufficient specifications



ibm.com/redbooks

© 2001 IBM Corporation

WTO Buffer shortage recovery

- Upon WTO,WTOR buffer shortage conditions (indicated by messages)
- Adjust buffer limits temporarily and permanent
- Perform buffer backlog analysis
- Adjust MCS console characteristics temporarily or permanent
- Support a user defined policy to KEEP or CANCEL jobs.



ibm.com/redbooks

© 2001 IBM Corporation

WTO Buffer shortage recovery

... when receiving

- IEE899I - Console Info (used for pre-analysis)
- IEA405E - WTO buffer shortage 80%
- IEA404A - Severe WTO buffer shortage 100%
- IEA406I - WTO buffer shortage relieved
- IEA230E - WTOR buffer shortage 80%
- IEA231A - Severe WTOR buffer shortage 100%
- IEA232I - WTOR buffer shortage relieved



ibm.com/redbooks

© 2001 IBM Corporation

WTO Buffer shortage recovery

In AOFCUST member (DSIPARM lib):

```
AUTO(  
  WTO  
)
```

```
WTOBUF(  
  jobname1  WTO  CANCEL  
  jobname2  WTOR  CANCEL  
  jobna*    *    CANCEL  
)
```

Note: Default then is " * * KEEP "



ibm.com/redbooks

© 2001 IBM Corporation

msys for Ops functions

1. Check Log Stream Data Set 'SHAREOPTIONS'
2. System Log Recovery
3. Recovery of Log Stream Data Set Directory Shortage
4. Alternate Couple Data Set Allocation
5. WTO/R Buffer Shortage Recovery
6. INGPLEX SYSTEM
7. INGPLEX CONSOLE
8. INGPLEX CF
9. INGPLEX CDS



ibm.com/redbooks

© 2001 IBM Corporation

Parallel sysplex operation

```
INGLX000      msys - Command Dialogs
Domain Id    = IPSFP      ----- INGPLEX -----
Operator Id  = HIR5
Date        = 03/08/01
Time        = 14:12:23

Sysplex . . . . . : KEY1PLEX

Select the desired command:

  1 Display systems (including ETR & signalling paths)
  2 Display consoles
  3 Control coupling facilities
  4 Control couple data sets

INGPLEX ...
SYStem
CONsole
CF
CDS

Command ==>
F1=Help      F2=End      F3=Return      F6=Roll
F12=Retrieve
```



ibm.com/redbooks

© 2001 IBM Corporation

Parallel sysplex operation

```
AOFKX100      msys - Command Dialogs
Domain ID    = IPSFP      ----- INGPLEX SYSTEM -----
Operator ID  = THIE
Date        = 03/19/01
Time        = 14:43:43

Sysplex . . . . . : KEY1PLEX
GRS Mode . . . . . : STAR

Display more info: C CPU E ETR I IPL O IOS S STOR/ESTOR
Signalling Path : D device T structure

Cmd  System  Status  Timestamp  INTERVAL  Action  SSUM  TIME  WEIGHT
---  -
KEY1  ACTIVE  14:43:41  86400     PROMPT  86400  1
KEY2  ACTIVE  14:43:41  86400     PROMPT  86400  1
KEY3  ACTIVE  14:43:39  86400     PROMPT  86400  1
KEY4  ACTIVE  14:43:42  86400     ISOLATE 0      1

Command ==>
F1=Help      F2=End      F3=Return  F6=Roll
F12=Retrieve  F9=Refresh
```



ibm.com/redbooks

© 2001 IBM Corporation

Parallel sysplex operation

- Using **INGPLEX SYS**, you can get information about

1. CPU (DISPLAY M=CPU)
2. ETR (DISPLAY ETR,DATA)
3. IPL (DISPLAY IPLINFO)
4. IOS (DISPLAY IOS,CONFIG(ALL))
5. STOR / ESTOR (DISPLAY M=STOR;DISPLAY M=ESTOR)
6. Device Paths (DISPLAY XCF,PATHIN,DEVICE=ALL)
7. Structure Paths (DISPLAY XCF,PATHIN,STRNAME=ALL,
DISPLAY XCF,PATHOUT,STRNAME=ALL)



ibm.com/redbooks

© 2001 IBM Corporation

Parallel sysplex operation

```
INGLX000      msys - Command Dialogs
Domain Id    = IPSFP      ----- INGPLEX -----
Operator Id  = HIR5
Date = 03/08/01
Time = 14:12:23

Sysplex . . . . . : KEY1PLEX

Select the desired command:

  1 Display systems (including ETR & signalling paths)
  2 Display consoles
  3 Control coupling facilities
  4 Control couple data sets

INGPLEX ...
System
CONsole
CF
CDS

Command ===>
F1=Help      F2=End      F3=Return      F6=Roll
              F12=Retrieve

ME d 30/015
```



ibm.com/redbooks

© 2001 IBM Corporation

Parallel sysplex operation

```

INGLX400                               SA 0S/390 - Command Dialogs                               Line 1 of 10
Domain Id = IPUFA                      ----- INGPLEX CONSOLE -----                               Date = 04/30/01
Operator Id = THIE                                                              Time = 17:44:57

Sysplex                               : AOCPLEX                               Master Console . . : MASTERA
Message Buffer Usage                   : 4 / 1500                               Reply Buffer Usage   : 4 / 26
Awaiting Replies                      : 4                               Eventual Action     : 0
Immediate Action                     : 0                               Awaiting Mounts     : 0
Critical Action                      : 0                               Operator Requests    : 0

-----
Cmds: D Details / R Requests
-----

  Console      Status      AUTH      NBUF      UD      Device      System      ALTGRP      MSCOPE
-----
MASTERA       MASTER      MASTER      0      Y      -none-      AOCA      --none--      *ALL
MASTERA1      INACTIVE  ALL      n/a      N      -none-      --none--      --none--      *ALL
MASTERB       ACTIVE    ALL      0      N      -none-      AOCB      --none--      *ALL
MASTERB1      INACTIVE  ALL      n/a      N      -none-      --none--      --none--      *ALL
MASTERC       ACTIVE    ALL      0      N      -none-      AOCC      --none--      *ALL
MASTERC1      INACTIVE  ALL      n/a      N      -none-      --none--      --none--      *ALL
MASTERD       ACTIVE    ALL      0      N      -none-      AOCD      --none--      *ALL
MASTERD1      INACTIVE  ALL      n/a      N      -none-      --none--      --none--      *ALL
TESTCONS      INACTIVE  INFO     n/a      N      -none-      --none--      --none--      AOCA
06            ACTIVE    INFO     n/a      N      -none-      MSCOPE=*A  --none--

Command ==>
F1=Help      F2=End      F3=Return    F4=Refresh    F6=Roll      F12=Retrieve

```



ibm.com/redbooks

© 2001 IBM Corporation

Parallel sysplex operation

- Using **INGPLEX CON**, you can get information about:
 1. Defined Consoles and their current status
(D C,L,L=*consname* ,
D R,L=*consname*,
D C)
 2. Details
(D C,CN=*consname*)
 3. Console Requests
(D R,L,T,CN=*consname*)



ibm.com/redbooks

© 2001 IBM Corporation

Parallel sysplex operation

```
INGLX000      msys - Command Dialogs
Domain Id    = IPSFP      ----- INGPLEX -----
Operator Id  = HIR5
Date = 03/08/01
Time = 14:12:23

Sysplex . . . . . : KEY1PLEX

Select the desired command:

  1 Display systems (including ETR & signalling paths)
  2 Display consoles
  3 Control coupling facilities
  4 Control couple data sets

INGPLEX ...
SYStem
CONsole
CF
CDS

Command ===
F1=Help      F2=End      F3=Return      F6=Roll
F12=Retrieve

3 = INGPLEX CF = INGCF

30/015
```



ibm.com/redbooks

© 2001 IBM Corporation

Parallel sysplex operation

```
INGLX900      msys - Command Dialogs      Line 1 of 2
Domain Id    = IPSFP      ----- INGCF -----
Operator Id  = THIE
Date = 03/19/01
Time = 18:05:44

Sysplex . . . . . : KEY1PLEX      SM process level . : 9
Cmds: D drain CF / E enable CF / P display sender paths / S display structures

  CF Name    Total Space    Free Space    Free%    Volatile    CF level
  -----
- CF01       378112 K       256768 K    67.91    YES         9
- CF02       378112 K       317184 K    83.89    YES         9

Command ==>
F1=Help      F2=End      F3=Return      F6=Roll
F9=Refresh   F12=Retrieve
```



ibm.com/redbooks

© 2001 IBM Corporation

Parallel sysplex operation

- Using **INGCF ENABLE**, you can:
 1. Activate the CF (not available at z/OS R2 GA)
 2. If the CF is not defined in the active policy a list of alternative policies may be displayed,
 1. Which have the CF defined and
 2. Which have defined all allocated structures and
 3. Which satisfy the needs of each structure (e.g.size)
 3. Set the sender paths to ONLINE for all systems which can connect to the CF
 4. Populate the CF with structures that should be allocated in the CF according to the preference list
- Advantages:
 1. Simplifies the process of bringing a CF into the sysplex
 2. Restricted access to some or all functions thru a SAF product



ibm.com/redbooks

© 2001 IBM Corporation

Parallel sysplex operation

- Using **INGCF DRAIN**, you can:
 1. Move the allocated structures from one CF to another
 2. If the structure is duplexed, stop duplexing on target CF
 3. Delete the allocated structures which cannot be moved
 4. Set the sender paths of all connected systems to OFFLINE
 5. Deactivate the CF (not available at z/OS R2 GA)
- Advantages:
 1. Simplifies the process of taking a CF out of a sysplex
 2. Restricted access to some or all functions thru a SAF product



ibm.com/redbooks

© 2001 IBM Corporation

Parallel sysplex operation

```
INGLK000      msys - Command Dialogs
Domain Id = IPSFP      ----- INGPLEX -----
Operator Id = HIR5      Date = 03/08/01
                        Time = 14:12:23

Sysplex . . . . . : KEY1PLEX

Select the desired command:      INGPLEX ...

  1 Display systems (including ETR & signalling paths)      SYStem
  2 Display consoles      CONsole
  3 Control coupling facilities      CF
  4 Control couple data sets      CDS

4 = INGPLEX CDS

Command ==>
F1=Help      F2=End      F3=Return      F6=Roll
                        F12=Retrieve

ME d      30/015
```



ibm.com/redbooks

© 2001 IBM Corporation

Parallel sysplex operation

```
INGKX300      Command Dialogs      Line 1 of 18
Domain ID = IPSFP      ----- INGPLEX CDS -----
Operator ID = HIR5      Sysplex = KEY1PLEX      Date = 03/08/01
                        Time = 15:49:06

System...: KEY4      Interval...: 86400      OPNotify: 86400
Maxmsg...: 999999      Cleanup...: 60      Retry...: 255
Classlen: 956      Max CFlevel: 9      COUPLExx: COUPLER1
SMREBLD...: 1      Max SMlevel: 9

Cmds: A allocate alternate CDS / C display CHPIDs
      D display CDS information / P switch alternate CDS to primary CDS

Type      MS      Volume      Dev      Couple Dataset Name
-----
- SYSPLEX
  PRIMARY...: 10      KEY1SP      260A      SYS1.KEY1.PXCFCDs
  ALTERNATE: 10      KEYLIB      2610      SYS1.KEY1.AXCFCDs
- ARM
  PRIMARY...: 8      KEYLIB      N/A      AOC.CDS.TEST.ARM.CDS01
  ALTERNATE: -- none --
- CFRM
  PRIMARY...: 32      KEY1SP      260A      SYS1.KEY1.SMREBLD.PXESCDs
  ALTERNATE: 32      KEYLIB      2610      SYS1.KEY1.SMREBLD.AXESCDs
- LOGR
  PRIMARY...: 32      KEY1SP      260A      SYS1.KEY1.PLOGCDs
  ALTERNATE: 32      KEYLIB      2610      SYS1.KEY1.ALOGCDs
- SFM
  PRIMARY...: 32      KEYLIB      N/A      AOC.CDS.TEST.SFM.CDS01

Command ==>
F1=Help      F2=End      F3=Return      F6=Roll
F8=Forward      F9=Refresh      F12=Retrieve

ME d      30/015
```



ibm.com/redbooks

© 2001 IBM Corporation

Parallel sysplex operation

- Using **INGPLEX CDS**, you can:
 1. Get information about the couple data sets at a glance
 2. Switch couple data sets
 3. Allocate alternate couple data sets including user-defined data sets
 4. Start policies
 5. View policy details
- Advantages:
 1. Managing couple data sets and policies by single keystroke functions
 2. Restricted access to the manipulation functions thru a SAF product



ibm.com/redbooks

© 2001 IBM Corporation

msys for Operations - installation

Installation:

- This uses NetView and SA/390, so are NetView and SA/390 skills needed to install it?
 - No, Wizard and installation checklist are provided that remove need for NetView or SA/390 skills
- Is the install time measured in days or weeks?!
 - Using the Wizard, should be possible to get up and running in half a day
 - Access Wizard at:
 - ▶ <http://www.ibm.com/servers/eserver/zseries/zos/wizards/>



ibm.com/redbooks

© 2001 IBM Corporation

msys for Operations - installation

- Wizard initial screen....

z/OS msys for Operations Configuration Assistant Task Menu - Netscape

IBM

z/OS msys for Operations Configuration Assistant : Interviews [Help](#)

We begin with a series of interviews in which you'll answer questions about the z/OS msys for Operations configuration that you are creating. When you have finished answering all of the interview questions, click **Build**. The z/OS msys for Operations Configuration Assistant will build a checklist of steps for you to follow, as well as customized jobs and other data sets for you to use.

☒ = Complete ☐ = Incorrect or Incomplete Data

| Interview Topics: Configure msys for Operations | Required | Status |
|--|----------|--------|
| Interview Topic #1: Specify your system environment | Yes | |
| Interview Topic #2: Specify your security product and authorize operators and commands | Yes | |
| Interview Topic #3: Specify your communications environment (VTAM and NetView) | Yes | |
| Interview Topic #4: Enable msys for Operations functions | Yes | |

[Your Privacy](#)



ibm.com/redbooks

© 2001 IBM Corporation

msys for Operations - installation

- 10 minutes later... Installation checklist....

| | Task |
|--------------------------|--|
| <input type="checkbox"/> | Copy procedure INGNVAP into a PROCLIB data set. |
| <input type="checkbox"/> | Copy procedures INGPOM and INGPXCU into a PROCLIB data set. |
| <input type="checkbox"/> | Update the active PROGxx member of SYS1.PARMLIB. |
| <input type="checkbox"/> | Add a Program Properties Table entry to the active SCHEDxx member of SYS1.PARMLIB. |
| <input type="checkbox"/> | Update the active Message Processing Facility List, MPFLSTxx in SYS1.PARMLIB. |
| <input type="checkbox"/> | Make the determined security definition changes. (SAF) |
| <input type="checkbox"/> | Make the determined security definition changes. (non-SAF) |
| <input type="checkbox"/> | Define the NVSS application major nodes to VTAM. |
| <input type="checkbox"/> | Activate the new VTAM APPLS. |
| <input type="checkbox"/> | Update the NVSS Style Sheet. |
| <input type="checkbox"/> | Define the system-unique VSAM datasets. |
| <input type="checkbox"/> | Customize the msys for Operations runtime functions. |
| <input type="checkbox"/> | Build the VTAM logon mode table. (optional) |
| <input type="checkbox"/> | Rate this wizard (optional). |



books

© 2001 IBM Corporation

msys for Operations - installation

Installation can be completed in half day, possibly without an IPL:

- 1 SCHEDxx definition, but this is usually included in samples
- Need AMODETAB in VTAMLIB
- 3 Libraries to add to APF list
- 1 Library to add to LNKLIST
- No LPALST or Subsystem definitions
- IRXANCHR must be updated - can be done dynamically
- Upload jobs to z/OS, eye-check, and run
- RACF definitions included in Wizard-generated jobs
- Setup 3 procs in PROCLIB



ibm.com/redbooks

© 2001 IBM Corporation

msys for Operations - installation

Customization:

- Customization information (system names, HLQs, volsers, and so on) is entered in Wizard
- All customization information is kept in single member (AOFCUST) which is created by the wizard based on information you provide
- Very easy to turn msys for Ops functions on and off - simply comment out keywords in AOFCUST



ibm.com/redbooks

© 2001 IBM Corporation

msys for Operations - installation

Relationship between msys for Operations and NetView and SA/390:

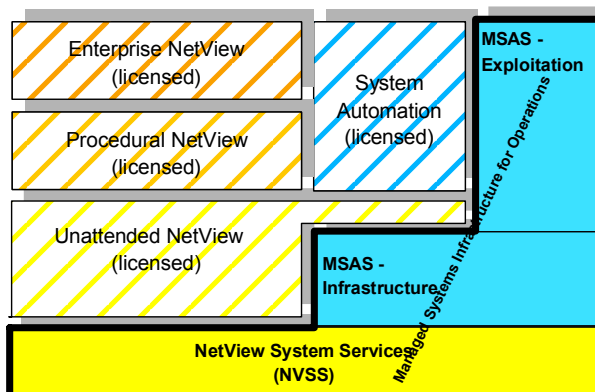
- In z/OS 1.2, base FMID from NetView has been included in z/OS. This is referred to as NVSS (NetView System Services) and only supports msys for Ops functions.
- Similarly the base FMID from SA/390 V2.1 has been included in z/OS 1.2. This also only supports msys for Ops.



ibm.com/redbooks

© 2001 IBM Corporation

msys for Ops - Installation



- ✓ New base element of z/OS
- ✓ Includes parts of NetView and System Automation
 - Automation Infrastructure
- ✓ Priced NetView and SA OS/390 build on top
- ✓ Not open for customer use, usage is locked
- ✓ Largely invisible



ibm.com/redbooks

© 2001 IBM Corporation

msys for Operations - installation

Coexistence with Full NetView and SA/390:

- FMIDs delivered with z/OS 1.2 are the base FMIDs for NetView and SA/390 V2.1
- Data sets delivered in z/OS 1.2 ServerPac are called NETVIEW.V1R4M0.whatever and ING.whatever (some of which are empty if you don't have full NetView or SA/390)
- If you order full NetView, ServerPac will come with additional FMIDs and same set of libraries will be used
- If you order full SA/390 V2, it will come with additional FMIDs and use the same set of libraries. msys for Ops function is included in SA/390 V2.1 (via SPE OW47903)
- If you are using SA/390 1.3, some functions must be disabled in either msys for Ops or SA/390. See manual for details



ibm.com/redbooks

© 2001 IBM Corporation

msys for Operations - installation

Coexistence with other automation

- msys for Ops will co-exist with other automation products
- All automation is message-driven
 - msys for Ops provides a list of the message it automates, so you can check that those messages are not automated elsewhere
 - Alternately, you can turn off some functions in msys for Ops if you wish to manage them using another product



ibm.com/redbooks

© 2001 IBM Corporation

msys for Ops - directions

Remember objective of msys for Ops:

- Improve system and application availability by timely reaction to potential problems
- Improve availability by making a sysplex easier to operate

Planned changes:

- Add local page data sets as needed
- Automate replies to IXC102A (Reply Down) messages
- Gather system data at IPL
- Add ability to configure CF links on and off and deactivate CFs
- Automatically rebuild structures in status of Policy Change Pending

Others - please send us your suggestions



ibm.com/redbooks

© 2001 IBM Corporation

msys for Operations - documentation

Documentation:

- Managed System Infrastructure for Operations Setting Up and Using, SC33-7968
- Installation checklist - on the Wizards Web site
- SPE description from SA/390 Web site:
 - <http://www.ibm.com/servers/eserver/zseries/software/sa/pdf/saptfmst.pdf>



ibm.com/redbooks

© 2001 IBM Corporation

GDPS Overview and latest enhancements



© Copyright IBM Corp. 2000

GDPS

What GDPS *is*:

- Disaster recovery solution for multi-site configurations
- Combination of services, software, specific support in IBM HW
- A business project, not an IT one
- A significant undertaking

What GDPS *is not*:

- A generalized automation solution
- Just some automation code and switching on remote copy
- A cheap solution:
 - HW (DASD, DWDM, CPUs)
 - Connectivity (significant inter-site bandwidth requirements)
 - The offering itself

GDPS

GDPS:

- Comes in various flavors, depending on the distance and data loss requirements
- Handles both planned and unplanned outages for DASD, systems, CPCs, sites
- Is a complete disaster recovery solution, spanning IMS, CICS/VSAM, DB2, batch - even Linux/390!



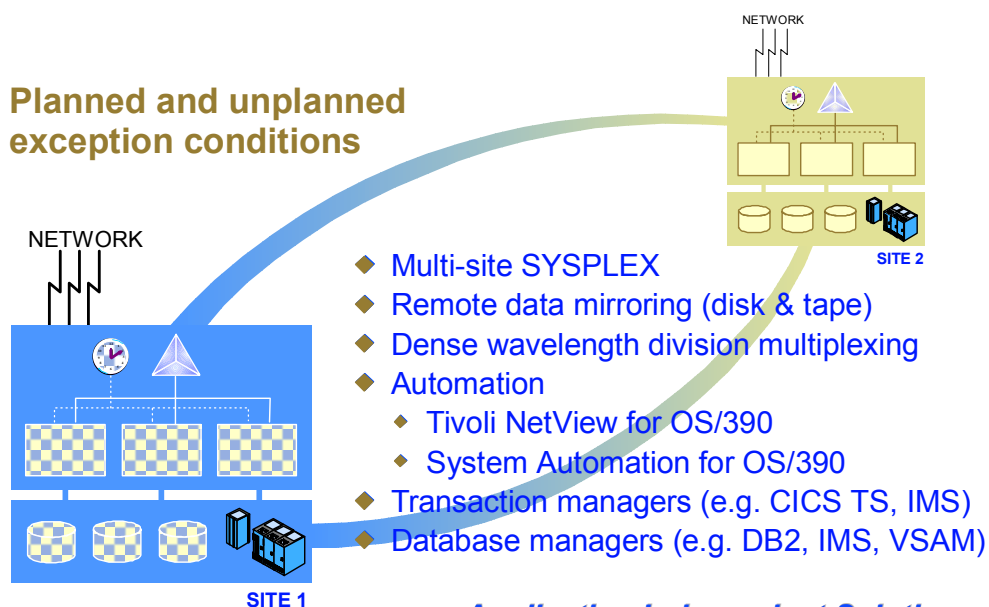
ibm.com/redbooks

© 2001 IBM Corporation

GDPS

GDPS/PPRC Offering

Planned and unplanned exception conditions



Application Independent Solution



ibm.com/redbooks

© 2001 IBM Corporation

GDPS

FREEZE Policy Options

FREEZE & GO

- ◆ Freeze secondary DASD configuration
- ◆ Allow applications to continue
- Optimize for remote restartability
- Least impact on application availability
- May lose data in case of real disaster

FREEZE & STOP

- ◆ Freeze secondary DASD configuration
- ◆ Stop all OS/390 images
- Optimize for remote restartability
- May impact application availability
- No data loss on primary site disaster

FREEZE & STOP Conditional

- ◆ Freeze secondary DASD configuration
- ◆ Determine reason for Suspend
 - ◆ If secondary HW problem then
FREEZE & GO
 - ◆ Other reason: **FREEZE & STOP**



ibm.com/redbooks

© 2001 IBM Corporation

GDPS

GDPS/PPRC Prereqs:

- Parallel Sysplex across 2 sites (maximum fiber length is 40KM)
- Replicate hardware across sites for redundancy
 - Processor(s), coupling facility and IBM 9037-2 sysplex timer
 - HMC automation infrastructure and associated processor support
 - Disk subsystems supporting PPRC Level 2 (Freeze)
 - Tape
- z/OS or OS/390 V2R6 or higher
- System Automation for OS/390 V1.3 or higher
- Tivoli NetView for OS/390 V1.2 or higher, or NetView 3.1
- Data required for restart must be disk resident and mirrored
 - Peer to Peer VTS support coming later in 2001




ibm.com/redbooks

© 2001 IBM Corporation

GDPS

Single Site Workload



One Sysplex with production systems in one site and monitoring system in other site
 Expendable workload in other site
 All primary data in one site; all secondary data in other site

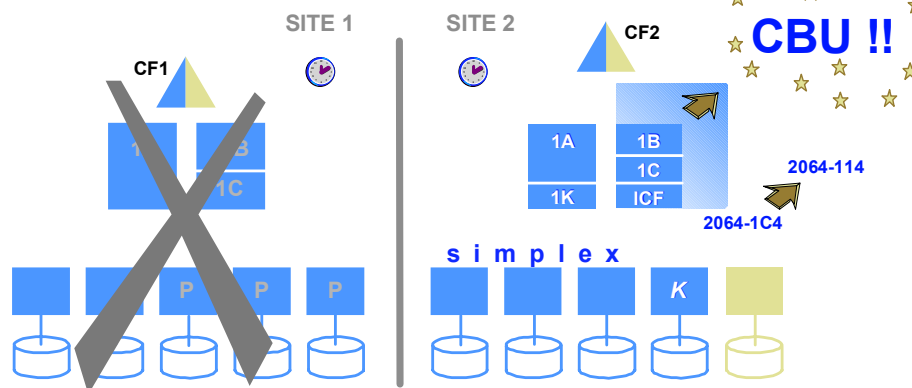


ibm.com/redbooks

© 2001 IBM Corporation

GDPS

Single Site Workload Site 1 Failure - Capacity Backup



Freeze and recover secondary disks
 Stop expendable systems
 Restart critical systems in site 2




ibm.com/redbooks

© 2001 IBM Corporation

GDPS

Multiple Site Workload



One SYSPLEX across both sites
Expendable workload in each site
All primary data in one site
All secondary data and controlling system in other site




ibm.com/redbooks

© 2001 IBM Corporation

GDPS

Multiple Site Workload Site 1 Failure / Maintenance



Freeze and recover secondary disks
Stop expendable systems
Restart critical systems in site 2

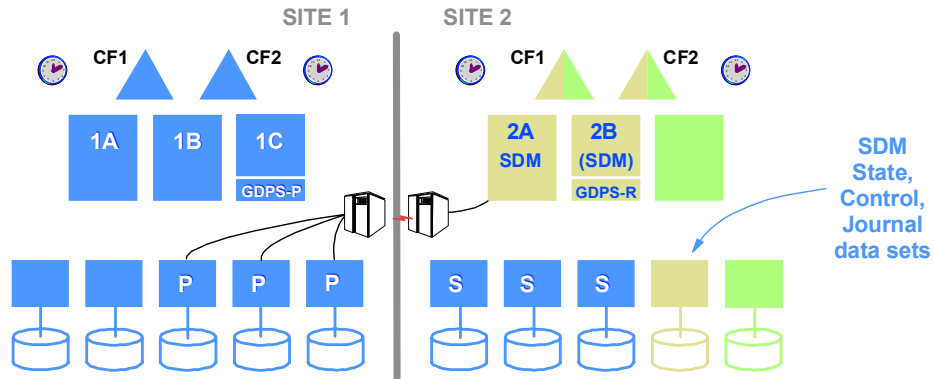


ibm.com/redbooks

© 2001 IBM Corporation

GDPS

GDPS / XRC Configuration



One SYSPLEX in site 1, all primary data in site 1
 All secondary data and controlling system in site 2
 (Coupled) System Data Mover and expendable workload in site 2




ibm.com/redbooks

© 2001 IBM Corporation

GDPS

GDPS/XRC - Primary Site Failure



Recover secondary disks and start GDPS-P in site 2
 GDPS-P stops expendable systems in site 2 and restarts Parallel Sysplex
 Restart site 1 Parallel Sysplex in site 2



ibm.com/redbooks

© 2001 IBM Corporation

GDPS

GDPS/XRC Prereqs:

- System Data Mover runs in a Parallel Sysplex
 - SDM Parallel Sysplex may be at any distance from primary site
- Parallel Sysplex preferred for production, but not required
- Replicate hardware in each site for redundancy
 - Processor(s), coupling facility and IBM 9037-2 sysplex timer
 - HMC automation infrastructure and associated processor support
 - Disk subsystems supporting XRC Level 1, 2 or 3 (level 3 - unplanned outage support - preferred)
 - Tape
- zOS V1R0 or OS/390 V2R6 or higher
- System Automation for OS/390 V1.3 or higher
- Tivoli NetView for OS/390 V1.2 or higher, or NetView 3.1
- Data required for restart must be disk resident and mirrored



ibm.com/redbooks

© 2001 IBM Corporation

GDPS - further information

Executive Overview

- www.s390.ibm.com/marketing/gf225114.html

Detailed Overview

- www.s390.ibm.com/marketing/gf225063.html

GDPS Service Offering

- <http://www.s390.ibm.com/nc/gdps.html>
- <http://www.as.ibm.com/asww/offerings/mww62b1.htm>
- <http://www.ibm.com/services/its/us/mus62b1.html>

E-mail for further information: gdps@us.ibm.com



ibm.com/redbooks

© 2001 IBM Corporation

Systems Programmer's Goody Bag



© Copyright IBM Corp. 2000

Sysprog Goody Bag


Useful commands

Parallel Sysplex Trainer Status

Support information

Useful Web sites

APARs of interest



Handy commands

Did you know you can...

- Delete any number of items in SDSF by typing //P beside the first one and // beside the last one?
- Display the time that a system last updated the sysplex couple data set using the D XCF,S,ALL command?
- Display information about LLA processing using the recently-documented (but long-available) D LLA command?
- Find out the first library in the LNKST to contain a named load module by using the following command:
`SETPROG LNKST,TEST,NAME=lnkstname,MODNAME=modulename`
- Find out the MVS level and the IPL date with the D IPLINFO command?
- Display and cancel long running or hung commands with new CMDS command in z/OS 1.2?



ibm.com/redbooks

© 2001 IBM Corporation

Handling LNKST expansion

If a LNKST library moves into a new extent, do the following:

- ▶ `SETPROG LNKST DEFINE NAME(whatever) COPYFROM(CURRENT)`
- ▶ `SETPROG LNKST ACTIVATE NAME(whatever)`
- ▶ `SETPROG LNKST UPDATE JOB(*)`

If you need to move a LNKST library

- ▶ `P LLA`
- ▶ `SETPROG LNKST DEFINE NAME(newname) COPYFROM(CURRENT)`
- ▶ `SETPROG LNKST DELETE NAME(newname) DSN(datasename)`
- ▶ `SETPROG LNKST ACTIVATE NAME(newname)`
- ▶ `SETPROG LNKST UPDATE JOB(*)`
- ▶ `S LLA,SUB=MSTR,LLA=00`

In both cases, beware of UPDATE JOB(*) command - see System Commands for information



ibm.com/redbooks

© 2001 IBM Corporation

Handy tools

Want to be able to logon to TSO on more than one system in a MAS? See RTA000089663 for sample code

Want to change a system symbol without doing an IPL? Use SYMUPDTE program, available from <ftp://www.redbooks.ibm.com/redbooks/SG245451/>

If you use this, make sure to download the latest level before any OS upgrades

Another way to find a module in LNKLIST - TSO ISRDDN command (use LI inside ISRDDN to display LNKLIST). To search for a module, use M modname LINKLIST

Another option is TSO ISRFIND. To see what ISRFIND can do, try it, or look in the source in ISP.ISPEXEC. ISRFIND is a diagnostic tool so is not officially supported



ibm.com/redbooks

© 2001 IBM Corporation

Bits and bytes

```
ISRFIND--ISPF/PDF LEVEL2 DIAGNOSTIC AID
```

```
APPLID   - ISR
ISPF LVL - ISPF 5.0
PDF LVL  - PDF 5.0
TIME     - 02:05
```

```
COMMAND ===>
```

```
DD Name      ==>      DD to search (blank for all)

Dataset info ==>      (blank) for no dataset info required
                  B for BASIC info (DSORG RECFM LRECL BLKSIZE
                  F for FULL info  (BASIC+ALLOC/USAGE+DIRECTORY)

Member Name  ==>      Member to search for (not required)
LOADMOD      ==>      Y if the member a LOADMOD.
                  (adds search of the LPA, LPALIST and LINKLIST)

Save Data    ==>      (blank) do not save ISRFIND output
                  S Save data in new dataset 'Userid.ISRFIND.SAVE'
                  A Append data to existing 'Userid.ISRFIND.SAVE'
```



ibm.com/redbooks

© 2001 IBM Corporation

Bits and bytes

- SNA Console Support in z/OS 1.1:
 - Full function console that can be used for anything except NIP
 - Requires VTAM to be up
 - No new STC
 - Requires changes to CONSOLxx and VTAMLST member

CONSOLxx

```
INIT      APPLID (SCSMCS&SYSCONE.)  GENERIC (SCSMCS$$)  
CONSOLE  DEVNUM (SMCS)  
          NAME (CON1)  
          ALTGRP (MASTER)  
          AREA (NONE)  
          AUTH (MASTER)  
          CMDSYS (*)  
          CON (N)  
          DEL (R)  
          LEVEL (ALL)  
          LOGON (REQUIRED)  
          MFORM (T,S,J,X)  
          MSCOPE (*ALL)  
          PFKTAB (PFKTAB1)  
          RNUM (28)  
          ROUTCODE (ALL)  
          RTIME (1/4)  
          SEG (28)  
          USE (FC)
```

APSMCSXX

```
SMCS      VBUILD TYPE=APPL  
SCSMCS&SYSCONE.  APPL
```



ibm.com/redbooks

© 2001 IBM Corporation

Bits and bytes

SMCS CONSOLE SELECTION

Enter the Console Name you want to access and press ENTER.

CONSOLE NAME ==> _____ (Required. This name must have been defined as an
SMCS console in CONSOLxx at IPL).

You are attempting to access:

SYSPLEX: ##\$#PLEX SYSTEM: ##\$1

Licensed Materials - Property of IBM
"Restricted Materials of IBM"
5694-A01 (C) Copyright IBM Corp. 2001
All rights reserved.

PF3/15=LOGOFF

```
IEE187I ENTER LOGON PARAMETERS  
LOGON      PASSWORD  
GROUP      SECLABEL  
IEE163I MODE= R
```




ibm.com/redbooks

© 2001 IBM Corporation

DFSMSHsm common recall queue

All Recall requests are placed onto a shared queue from which all hosts can process requests.

- Implemented using a CF List Structure.



ibm.com/redbooks

© 2001 IBM Corporation

DFSMSHsm common recall queue

Work Load Balancing

- Balances Recalls across entire HSMplex

Tape Mount Optimization

- Host that mounts a single tape will process ALL recall requests that require that tape

Priority Optimization

- Highest priority requests always processed first - ARCRPEXT exit used to assign priorities to requests (based on Userid and DSN)

Request Persistence

- Requests remain across failures

Flexible Configurations

- Recall Servers
- Not all hosts need to be connected to tape drives



ibm.com/redbooks

© 2001 IBM Corporation

Parallel Sysplex Training Environment

Parallel Sysplex Trainer:

- Developed in ITSO
- Provides load and go Parallel Sysplex with following exploiters:
 - DB2 V7
 - IMS V7
 - CICS 1.3
 - MQSeries 5.2
 - OS/390 R10
 - z/OS R1
- Can install native or under VM
- Installation takes about 3 hours
- Comes complete with Install Guide, Exercise Guide, and workloads to test all the exploiters



ibm.com/redbooks

© 2001 IBM Corporation

Parallel Sysplex Training Environment

Some of the things it delivers working examples of:

- VTAM GR use by CICS, TSO, SNA Consoles, IMS
- TCP/IP Dynamic VIPA takeover and Sysplex Distributor
- SNA Consoles support
- msys for Setup
- msys for Operations
- SDSF use of MQSeries
- CICS/RLS, DB2 data sharing, IMS data sharing, IMS Shared Messages Queues, MQSeries queues in CF, CICS TS in CF
- Extensive use of system symbols
- Use of Symbolic Alias Facility to control subsystem libraries
- Every possible CF exploiter except SmartBatch, NetView Access Services, and VTAM MNPS
- GRS Star, JES2 Checkpoint in CF, XCF structures, Tape Sharing, RACF sysplex data sharing, ECS, RRS, OPERLOG, WLM IRD



ibm.com/redbooks

© 2001 IBM Corporation

Parallel Sysplex Training Environment

Updated every 6 months, to coincide with latest operating system release

Comes with 2 OS releases and any combination of subsystems, based on what customer is licensed for

Latest sysplex-exploiting features included in each release

Marketed through IBM Learning Services

Available on subscription basis



ibm.com/redbooks

© 2001 IBM Corporation

Bits and Bytes

CFLEVEL/Operating System level support

| | CFLEVEL 1 | CFLEVEL 2 | CFLEVEL 3 | CFLEVEL 4 | CFLEVEL 5 | CFLEVEL 6 | CFLEVEL 7 | CFLEVEL 8 | CFLEVEL 9 | CFLEVEL 10 |
|-------------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|------------|
| OS/390 2.6 | YES | YES | YES | YES | YES | YES | YES | YES | YES | YES |
| OS/390 2.7 | YES | YES | YES | YES | YES | YES | YES | YES | YES | YES |
| OS/390 2.8 | YES | YES | YES | YES | YES | YES | YES | YES | YES | YES |
| OS/390 2.9 | YES | YES | YES | YES | YES | YES | YES | YES | YES | YES |
| OS/390 2.10 | YES | YES | YES | YES | YES | YES | YES | YES | YES | YES |
| z/OS 1.1 | YES | YES | YES | YES | YES | YES | YES | YES | YES | YES |
| z/OS 1.2 | YES | YES | YES | YES | YES | YES | YES | YES | YES | YES |

Operating System level support

| Operating System | Out-of-support Date |
|------------------|---------------------|
| MVS/SP 5.2 | March 2001 |
| OS/390 1.3 | March 2001 |
| OS/390 2.6 | March 2002 |
| OS/390 2.7 | March 2002 |
| OS/390 2.8 | September 2002 |
| OS/390 2.9 | March 2003 |
| OS/390 2.10 | September 2004 |
| z/OS 1.1 | - |
| z/OS 1.2 | - |



ibm.com/redbooks

© 2001 IBM Corporation

Bits and Bytes

CFLEVEL support

| | CFLEVEL 1 | CFLEVEL 2 | CFLEVEL 3 | CFLEVEL 4 | CFLEVEL 5 | CFLEVEL 6 | CFLEVEL 7 | CFLEVEL 8 | CFLEVEL 9 | CFLEVEL 10 |
|------------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|------------|
| 9121-511 | Yes | Yes | Yes | Yes | No | No | No | No | No | No |
| 9021-711 | Yes | Yes | Yes | Yes | Yes | No | No | No | No | No |
| 9672-Rn1 | Yes | Yes | Yes | Yes | No | No | No | No | No | No |
| 9672-Rn2/3 | Yes | Yes | Yes | Yes | Yes | No | No | No | No | No |
| 9672-Rn4 | N/A | N/A | Yes | Yes | Yes | Yes | Yes | Yes | No | No |
| 9672-Rn5 | N/A | N/A | N/A | Yes | Yes | Yes | Yes | Yes | No | No |
| 9672-Rn6 | N/A | N/A | N/A | N/A | N/A | Yes | Yes | Yes | Yes | Yes |
| 9672-Rn7 | N/A | N/A | N/A | N/A | N/A | N/A | N/A | Yes | Yes | Yes |
| 2064-1nn | N/A | N/A | N/A | N/A | N/A | N/A | N/A | N/A | Yes | Yes |

For more information, see:

<http://www.ibm.com/servers/eserver/zseries/pso/cftable.html>



ibm.com/redbooks

© 2001 IBM Corporation

Bits and Bytes

Architecture level sets

| | OS/390 2.6 | OS/390 2.7 | OS/390 2.8 | OS/390 2.9 | OS/390 2.10 | z/OS 1.1 | z/OS 1.2 | z/OS 1.3 | z/OS 1.4 |
|------------|---------------|---------------|---------------|---------------|----------------|----------|----------|----------|----------|
| 9121-511 | Yes | Yes | Yes | Yes | No | No | No | No | No |
| 9021-711 | Yes | Yes | Yes | Yes | No | No | No | No | No |
| 9672-Rn1 | Yes | Yes | Yes | Yes | No | No | No | No | No |
| 9672-Rn2/3 | Yes | Yes | Yes | Yes | Yes | No | No | No | No |
| 9672-Rn4 | Yes | Yes | Yes | Yes | Yes | No | No | No | No |
| 9672-Rn5 | Yes | Yes | Yes | Yes | Yes | No | No | No | No |
| 9672-Rn6 | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 9672-Rn7 | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| 2064-1nn | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |

See: <http://www.ibm.com/s390/os390/plug.html>



ibm.com/redbooks

© 2001 IBM Corporation

HIPER/PE Alert Service

Free service to notify you when a HIPER or PE APAR is opened for selected components

Subscribe to ResourceLink (<http://www.ibm.com/servers/resourcelink>) (must request a ResourceLink userid first).

In the future, you will also be able to subscribe to Red Alerts

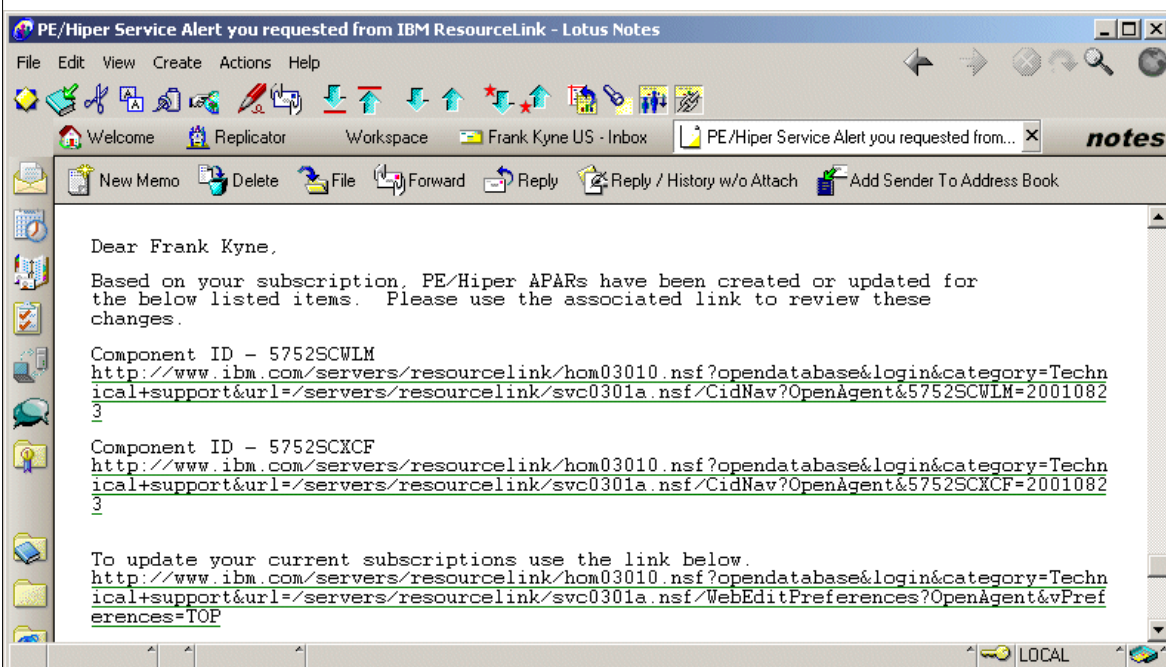
When a HIPER or PE is opened, you get emailed. Email contains hotlink to APAR in question



ibm.com/redbooks

© 2001 IBM Corporation

HIPER/PE Alert Service



ibm.com/redbooks

© 2001 IBM Corporation

Dynamic PAV Management

Ability for WLM to manage number of aliases on ESS devices, based on importance of work using a specific device

Pre-reqs:

- OS/390 2.7 + SPE or later
- WLM must be in Goal Mode (requires Monoplex or Multi-system sysplex)
- PAV feature code in the ESS

Customer feedback is very positive



ibm.com/redbooks

© 2001 IBM Corporation

Large Volume Support

Q: How many cylinders can you have in one z/OS volume?

A: It depends!

- Prior to new Large Volume Support in z/OS 1.2, you could have up to 10019 cylinders
- With LVS SPE (12/01), z/OS now supports up to 32760 cylinders per 3390 logical volume

| Function | FMID | APAR | PTF | Function | FMID | APAR | PTF |
|----------------|------------|-------------------------------|-------------------------------|---------------|------------|--------------------|---------|
| Device Support | HDZ11F0 | OW46067 | UW81156 | SAM | HDZ11F0 | OW47161 | UW81156 |
| | HDZ11E0 | OW46067 | UW81155 | FAMS | HDZ11F0 | OW46841 | UW81156 |
| | HDZ11D0 | OW46067 | UW81154 | ISMF Navquest | HDZ11F0 | OW47210 OW50217 | UW81156 |
| CVAF | HDZ11F0 | OW50734 | UW99406 | US Panels | JDZ11FB | OW47210 OW50217 | UW81157 |
| DADSM/CVAF | HDZ11F0 | OW46068O W50268 | UW81156 UW81156 | KANJII Panels | JDZ11FC | OW47210 OW50217 | UW81158 |
| IEHLIST | HDZ11F0 | OW49291 | UW81156 | ISMF | | OW51073 | UW82690 |
| Device Support | HDZ11F0 | OW47497 | UW99406 | DFSORT | Release 14 | PQ47926 | |
| ICKDSF | Release 16 | PQ42534 | | HSM | HDZ11F0 | OW49147 OW49148 | |
| BCP | HBB7703 | OW48394 OW48197 OW48197 | UW79169 UW78960 UW78963 | DSS | HDZ11F0 | OW50405 | UW81900 |
| JES2 | | OW49317/73 | | ISPF | HIF5A02 | OW47857 | UW77465 |
| JES3 | HJS7703 | OW49477 | | | | | |

Bits and bytes

z900 and Compression:

- z900 delivers significant improvements in compression performance compared to 9672 G6 Turbo models.
- Compression typically 3-4 times more efficient
- Expansion 2-3 times more efficient
- ==> If you have data that you previously discounted compressing because of the performance impact, have another look once all CPCs in the plex are migrated to z900.



ibm.com/redbooks

© 2001 IBM Corporation

Bits and Bytes - "Wizards"

Interactive, Web-based tools to help you install and customize IBM products.

Available at: <http://www.ibm.com/servers/eserver/zseries/zos/wizards/>

Current wizards:

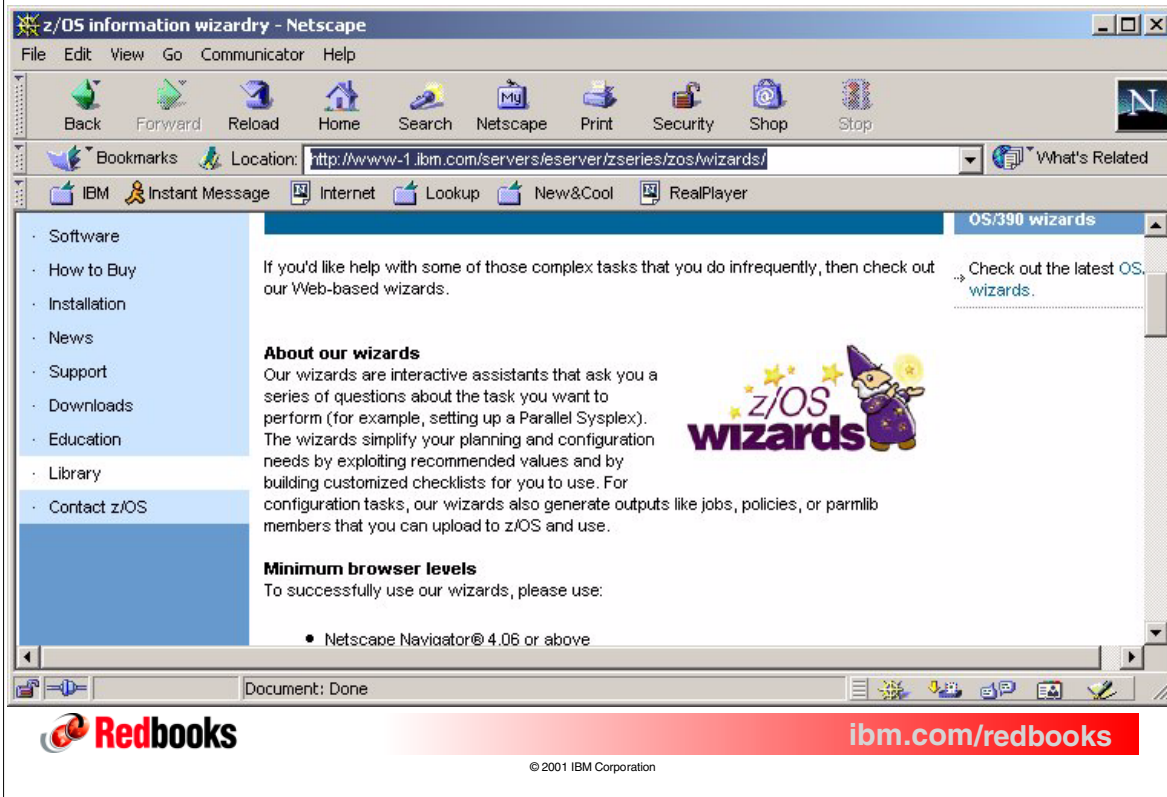
- z/OS and OS/390 Planning Wizard for e-business
- z/OS V1R1 Installation Planning Wizard
- z/OS V1R1 DFSMS Migration Planning Wizard
- z/OS V1R1 DFSMSrmm Customization Wizard
- z/OS V1R1 SDSF Customization Wizard
- z/OS Parallel Sysplex Customization Wizard
- z/OS Base Sysplex Customization Wizard
- z/OS V1R1 IP Configuration Wizard
- z/OS V1R1 UNIX Customization Wizard



ibm.com/redbooks

© 2001 IBM Corporation

Bits and Bytes - Wizards



ShopZSeries - Online ordering

Ability to order CBPDOs and ServerPacs directly through the Internet

Order at:

- <https://www.software.ibm.com/ShopzSeries>

Currently available in:

- Austria
- Belgium
- Denmark
- Finland
- France
- Germany
- Ireland
- Italy
- Netherlands
- Norway
- Spain
- Sweden
- Switzerland
- United Kingdom
- United States



ibm.com/redbooks

© 2001 IBM Corporation

ShopZSeries - Online ordering

How ShopZSeries works:

- Customer uploads "installed" inventory from target system image
- Customer selects desired upgrade offering (PDO or ServerPac) and desired "view" of product catalog
- Application presents analysis of "installed" inventory vs. current product catalog and IBM Licensed Inventory
- Upgrades are preselected for the customer; In case of ServerPac, the entire new configuration is pre-selected for the customer!
- Customer can add new products and perform technical requisite checking (and what-if's)
- Price & Terms 'n Conditions provided by IBM Order Center for "priced upgrades"
- Customer submits the order and can track order status through delivery



ibm.com/redbooks

© 2001 IBM Corporation

ShopZSeries - Online ordering

ShopzSeries - Order Details for kynef

Your software inventory has not been uploaded.

Order CBPDO 2001-09-27 15.46.54

Tip: [Made changes? Refresh your current order.](#)

Tip: [Add products to your order.](#)

Products in order

There are no products in this order.

CBPDO service options

Tip: All available service will be included for the products in your order.

Tip: Use the following options to include service for your existing products.

Tip: You can order service without ordering any products.

☒ Include service for all entitled MVS products starting at PUT level from my last order

For the base operating system, only include service for the following releases:

z/OS version 1.1
OS/390 version 2.10
OS/390 version 2.09
OS/390 version 2.08
OS/390 version 2.07
OS/390 version 2.06

Tip: Select all releases in your enterprise under customer number Ireland-992934. Hold down your Ctrl key to select multiple releases.

Select all



ibm.com/redbooks

© 2001 IBM Corporation

ShopZSeries - Online ordering

ePDO:

- Electronic deliver CBPDO (planned to be available 2002)
- Exploits RECEIVE FROMNETWORK (SMP/E V3R1)
- All products orderable except z/OS and OS/390 and their optional features
- Orders must be placed via ShopzSeries



ibm.com/redbooks

© 2001 IBM Corporation

Software Order Status

New capability to check status of already-placed orders

Current availability outside US varies

Available at:

- http://service.software.ibm.com/software_order_status



ibm.com/redbooks

© 2001 IBM Corporation

Software Order Status



End of Support dates

New function to get end of support date for IBM products

Requires JVM 1.3

Accessible at:

- <http://www.ibm.com/services/sl/products/java.html>

End of Support dates

Country: United States
 Operating System: OS/390 - IBM z/OS and OS/390
 Product Group: All
 Name, ID, End Date: [] [] Year [] Month []

800 products supported.

| OS | Group | Product Name | ID | VRM | End Date |
|--------|-------|---|----------|-------|------------|
| OS/390 | COM | eNetwork Host On-Demand for S/390 | 5655-A81 | 2.0.0 | |
| OS/390 | BASE | EREP MVS | 5658-260 | 3.5.0 | |
| OS/390 | DM | Extended Terminal Option (ETO) Support fo... | 5697-D68 | 1.1.0 | 2002-03-31 |
| OS/390 | AD/AE | Fault Analyzer for OS/390 | 5697-F19 | 1.1.0 | 2002-07-31 |
| OS/390 | AD/AE | File Manager for OS/390 | 5697-F20 | 1.1.0 | 2002-07-31 |
| OS/390 | BASE | First Failure Support Technology (FFST) for ... | 5695-044 | 1.2.0 | |
| OS/390 | TS | FlowMark/MVS | 5655-129 | 1.1.1 | 2001-09-30 |
| OS/390 | BASE | GDDM | 5695-167 | 3.2.0 | |
| OS/390 | AD/AE | GDDM-GKS | 5668-802 | 1.1.3 | |
| OS/390 | AD/AE | GDDM-IMD | 5668-801 | 2.1.3 | |
| OS/390 | AD/AE | GDDM-IVU | 5668-723 | 1.1.3 | |

Search Print

Redbooks ibm.com/redbooks © 2001 IBM Corporation

S/390 Support Site

Technical Information

Authorized Program Analysis Reports

Service bulletins

Hints and tips

Product-related publications

Other resources

Red Alerts

search PTFs/APARs

electronic PTF ordering

Report software defect

Review status of open PMR

Tools

And much more...

<https://service.software.ibm.com/s390/support>



ibm.com/redbooks

© 2001 IBM Corporation

LookAT

LookAT is an OS/390 & z/OS online help system specifically intended to provide users with message explanations when they need them

Enter a message ID and LookAT will return the message explanation reducing retrieval time significantly.

Searches APARs for DOC Hold information, as well as messages manuals

Where?

- <http://www.ibm.com/servers/s390/os390/bkserv/>
- <http://www.ibm.com/servers/eserver/zseries/zos/bkserv>



ibm.com/redbooks

© 2001 IBM Corporation

LookAT

Users can enter the LookAT command:

- At the TSO "READY" prompt
- In the ISPF Command panel (Option 6)
- From any ISPF command line
- From the UNIX System Services Shell

If you want to use it on MVS, the messages and codes collection bookshelf must exist on MVS.

You can download LookAT for TSO and VM from the following IBM Web site:

- <ftp://ftp.software.ibm.com/ps/products/ibmreader/tools/lookat/>



ibm.com/redbooks

© 2001 IBM Corporation

LookAT

LOOKAT Messages

1 Enter your Message ID:

2 Select an area to LookAt:

| z/OS™ | OS/390™ | z/VM™ | VSE/ESA™ |
|----------------------------|--|-------|----------------------------|
| <input type="radio"/> V1R1 | <input checked="" type="radio"/> V2R10 <input type="radio"/> V2R9 <input type="radio"/> V2R8 | | <input type="radio"/> V2R5 |

Or, for messages that are not enabled for LookAt, select an area to search using BookManager®:

| |
|--|
| <input type="radio"/> V2R7 <input type="radio"/> V2R6 <input type="radio"/> V2R5 <input type="radio"/> V2R4 <input type="radio"/> V1R3 <input type="radio"/> V1R2 <input type="radio"/> V1R1 |
|--|

APAR and ++HOLD Docs

| | |
|------------------------------------|------------------------------------|
| <input type="radio"/> All releases | <input type="radio"/> All releases |
|------------------------------------|------------------------------------|

3 Click **Go**


List Books

Help

News

Feedback

Survey

 **ibm.com/redbooks**

© 2001 IBM Corporation

Maintenance strategy update

New IBM Maintenance Strategy

- Referred to as Consolidated Service Test (CST)
 - New, additional intensive z/OS and major subsystems service test
 - Maintenance is tested first before recommendation
 - ▶ Provides a consistent, installable, maintenance recommendation - (the new RSU)
 - Service recommendations updated on a monthly and quarterly basis as testing continues.
- <http://www.ibm.com/servers/eserver/zseries/zos/servicetst>
- WSC Flash 10106 <http://www.ibm.com/support/techdocs>
- Watch out for upcoming update to service recommendations white paper by Barbara Bryant:
 - <http://www.ibm.com/servers/eserver/zseries/library/whitepapers/psos390maint.html>

SYMUPDTE program

Program to modify system symbols without an IPL

Documented in Appendix in Redbook SG24-5451

Available from:

- <http://www.redbooks.ibm.com>
- Go in on Redbooks Online, Additional Materials, then SG245451

IMPORTANT:

- Recent bug found resulting in storage overlays
- Fixed version available on ITSO Web Site from Sept 28.
- New version works with z/OS 1.2 and previous releases as well
- If you download this, please send an email to me so I can inform you of future problems or enhancements



ibm.com/redbooks

© 2001 IBM Corporation

Other useful Web sites

Cheryl Watson:

- www.watsonwalker.com

MVS-MAIN Listserver

- Subscribe to and access from www.google.com



ibm.com/redbooks

© 2001 IBM Corporation

APARs of Interest

- OW48672** **HIPER** WLM Stops swapping in address spaces 51 days after last IPL
- OW37621** Adds performance information to D XCF,Px,STRNM=xxx command
- PQ42127** Adds 64-bit real support to IMS
- PQ25914/36933** DB2 V6 support for data space buffers pools above 2GB real
- PQ44114** Provides ability to specify how many entries you want in the IRLM lock str
- OW47667** WLM APAR - removed disconnect time from I/O velocity calculation
- OW49807** RMF support for defined capacity measurements
- OW50998** Uses announced MSUs to report and specify Defined Capacities
- OW43778** Changes XES so that a rebuilt structure always has the same number of elements as the original
- OW49867** FICON Native support for XCF plus performance enhancements when used with CSS I/O Priority Queueing
- OW50382** Address WTO buffer shortage problems



ibm.com/redbooks

© 2001 IBM Corporation

Questions??



ibm.com/redbooks

© 2001 IBM Corporation