Understanding Parallel Sysplex Performance Advanded Topics

Joan Kelley
IBM Corp
Poughkeepsie, NY

Trademarks and Disclaimers

The following are trademarks of the International Business Machines Corportation:

S/390	VTAM	MVS/ESA
CICS	RACF	Parallel Sysplex
IMS/TM	DFSMS	RMF
IMSDB	VSAM	

The following is a registered trademark:
DB2

Performance numbers were achieved in a controlled laboratory environment and may vary based on customer environments

CF Configuration Options

Many combinations

1. Standalone CF (ex. 2066 - 0CF, 2084 - 300)
   - Dedicated CPs - best choice for production
   - **Shared CPs**

2. Internal CF (ex. 2064 - 108)
   - Dedicated CPs (expensive - added into S/W license costs)
   - **Dedicated ICFs - good choice for production if..**
   - **CPs shared with MVS images**
   - **CPs shared with other CF images**

JdK

New format for ICFs - RMF OW42945

```
                                                 P A R T I T I O N   D A T A   R E P O R T

MVS PARTITION NAME                       Z3
NUMBER OF CONFIGURED PARTITIONS           5
NUMBER OF PHYSICAL PROCESSORS            10
                  CP                      8
                  ICF                     2
WAIT COMPLETION                          NO
DISPATCH INTERVAL                   DYNAMIC

------- PARTITION DATA --                 ---- AVERAGE PROCESSOR UTILIZATION PERCENTAGES --
                              PROCESSOR   LOGICAL PROCESSORS   ----- PHYSICAL PROCESSORS ----
  NAME     STATUS WGTS   CAP   NUM  TYPE  EFFECTIVE     TOTAL  LPAR MGMT   EFFECTIVE     TOTA
TPN           A     70    NO     8  CP        26.02     28.02       2.00       26.02    28.02
Z2            A     10    NO     2  CP        10.48     12.89       0.60        2.62     3.22
Z3            A     10    NO     2  CP        13.01     15.37       0.59        3.25     3.84
*PHYSICAL*                                                          1.67                 1.67
                                                                  ------      ------   ------
  TOTAL                                                             4.86       31.90    36.76

CF2           A    100    NO     2  ICF       98.51     98.59       0.09       98.51    98.59
CF3           A    280    NO     1  ICF        0.63      0.86       0.12        0.32     0.43
*PHYSICAL*                                                          0.96                 0.96
                                                                  ------      ------   ------
  TOTAL                                                             1.16       98.83    99.99
```

ICF activity reported separately

JdK

MVS using CF on same CEC

IF MVS is actively using structures in a CF on the same CEC, only certain structures should reside in this CF to avoid rebuild problems. See W98029 for full explanation.

- Resource Management structures are good candidates
  - IEFAUTOS
  - XCF
  - GRSSTAR
  - JES CKPT
  - Logger
  - ISTGENERIC
- Data Sharing structures -
  - Some may require time to rebuild - OW33615 improves
    - IMS Cache
    - DB2 Group Buffer pools
    - CICS TEMP STOR
  - Some will cause a sysplex wide subsystem outage
    - DB2 SCA
    - IRLM (IMS & DB2)
    - VSAM RLS lock
    - Logger (CICS & RRS)

JdK

Solution - CF duplexing

CF Duplexing establishes two copies of a given structure - changed data is written to both.

This provides:

1. An 'easy-to-implement' recovery mechanism for structures with no recovery
2. Faster recovery from CF failures
3. Failure isolation for internal CFs (ICFs)

JdK

More details

System creates two copies of a duplexed structure:
1. Primary (OLD)
2. Secondary (NEW)

Changed data is written to both the copies.

Requests to each structure copy are synchronized
1. User managed duplexing, by a held lock
2. System managed duplexing, by 2 CF-to-CF signals

Installation selects duplexing option for specific exploiters/structures in CFRM structure definition
1. DUPLEX (ENABLED) - duplexing automatically maintained
2. DUPLEX (ALLOWED) - duplexing may be manually started/stopped
3. DUPLEX (DISABLED) - duplexing not allowed - default

User (DB2) CF Duplexing

If DUPLEX specified, changed data is sent to secondary

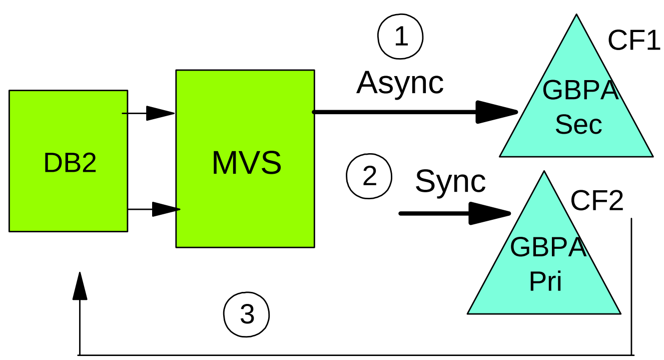

When primary completes, MVS forces completion of secondary and returns to DB2

On switch (Pri->Sec), all local cache buffers invalidated

Serialized by locks already held, so cost is low - 1-2%
Available in DB2 6.1, CFLEVEL 5, OS390 2.6.0

Primary and Secondary Structures on RMF

```
On CF1
 STRUCTURE NAME = DSNDB1G_GBP14     TYPE = CACHE  STATUS = ACTIVE PRIMARY
           # REQ    -------------- REQUESTS -------------
 SYSTEM    TOTAL             #     % OF  -SERV TIME(MIC)-
 NAME      AVG/SEC           REQ    ALL     AVG    STD_DEV

 Z0          454K   SYNC    440K   44.5    19.8     46.3
           252.5    ASYNC    14K    1.4   368.3    678.0
                    CHNGD      1    0.0  INCLUDED IN ASYNC
 -------------------------------------------------------------------------------
 TOTAL       988K   SYNC    965K   97.6    21.4     46.9      -- DATA ACCESS ---
           548.9    ASYNC    23K    2.4   414.7    701.2        READS        7688
                    CHNGD     20    0.0                         WRITES      11795
                                                                CASTOUTS      988
                                                                XI'S       524501
On CF2
 STRUCTURE NAME = DSNDB1G_GBP14     TYPE = CACHE  STATUS = ACTIVE SECONDARY
           # REQ    -------------- REQUESTS -------------
 SYSTEM    TOTAL             #     % OF  -SERV TIME(MIC)-
 NAME      AVG/SEC           REQ    ALL     AVG    STD_DEV

 Z0         7404    SYNC       0    0.0     0.0      0.0
            4.11    ASYNC   7404   68.3    95.0    179.7
                    CHNGD      0    0.0  INCLUDED IN ASYNC
--------------------------------------------------------------------------------
 TOTAL     10833    SYNC       0    0.0     0.0      0.0      -- DATA ACCESS ---
            6.02    ASYNC    11K    100   108.2    180.5        READS           0
                    CHNGD      0    0.0                         WRITES      11808
                                                                CASTOUTS        0
                                                                XI'S            0
```

JdK

System managed CF Duplexing

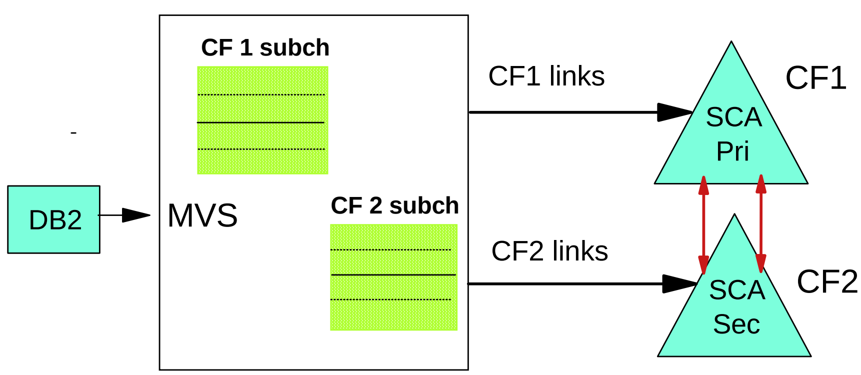

1. Simultaneous launch to maximize parallelism
2. Coordinated execution to preserve serialization
3. Both requests must complete successfully or both are 'backed out'

JdK

Requirements for SM Duplexing

Requires: CFLEVEL 11 for G5, G6
CFLEVEL 12 for z8xx, z9xx
(Note: Migration -> Structure size increase)
CF to CF links

H/W levels: z9xx
z800
G5/G6

z/OS 1.2 or higher with supporting APARs:
XES: Recovery - OA05328, MNPS - OA01513
RMF: Subch tuning - OW54802

Search RETAIN with keyword CFDUPLEX for latest fixes

JdK

Migrating to SM Duplexing

Install and migrate to product levels that support SM Duplexing
- IRLM 2.1, DB2 7, MQ 5.2, CICS TS 2.2...
- See White Paper for current list

Format the CFRM couple dataset using SMDUPLEX keyword
Modify CFRM policy for specific structures
- Control placement of structures via CF preference list
- Enable SM duplexing - DUPLEX(ENABLE)

To duplex LOG structures,
Format the LOGR couple dataset using SMDUPLEX keyword
Modify LOGR policy for specific logstreams
- Enable SM duplexing - LOGGERDUPLEX(COND)

Note: Earlier support used STG_DUPLEX and DUPLEXMODE parameters to define Dasd-only and staging datasets - See Setting up Sysplex

(At this point, you can still enable or disable duplexing, but returning to a downlevel CDS would be disruptive)

JdK

CF to CF links

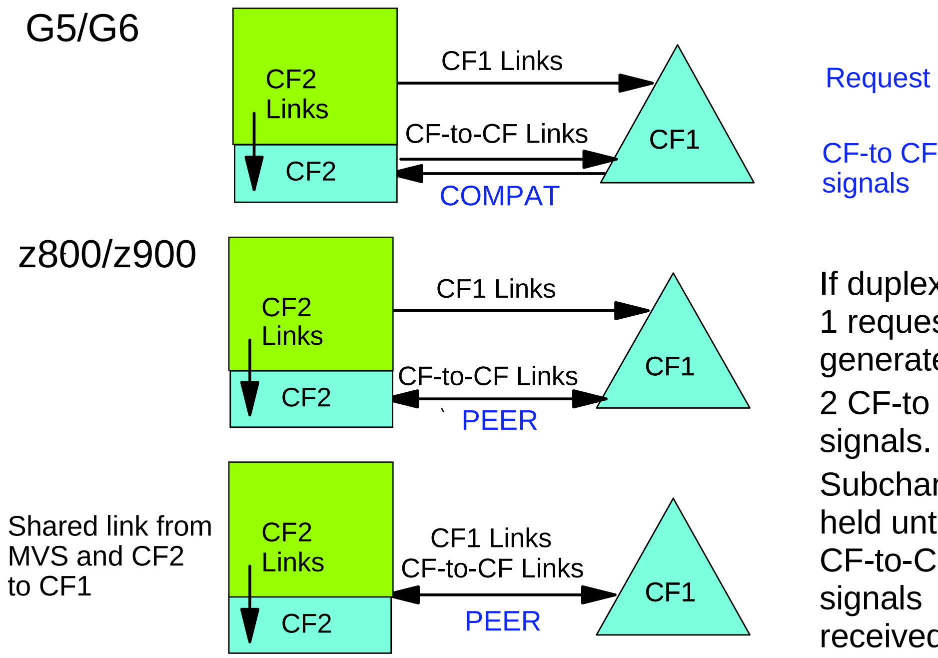

If duplexed, 1 request generates 2 CF-to CF signals.
Subchannel held until CF-to-CF signals received.

JdK

Duplexed Structures on RMF

1. RMF Mon I - CF report

COUPLING FACILITY USAGE SUMMARY

TYPE	STRUCTURE NAME	STATUS CHG	ALLOC SIZE	% OF CF STORAGE	# REQ	% OF ALL REQ	AVG REQ/ SEC
CF2	IRLMLOCK1	ACTIVE SEC	63M	1.0	1870K	7.1	2077.5
	.....						
CF3	IRLMLOCK1	ACTIVE PRIM	63M	1.0	1878K	15.3	2086.5

2. RMF Mon III - CF report

CF: ALL

Structure Name	Type	ST	System	--- Sync --- Rate	Avg Serv	--------- Async -------- Rate	Avg Serv	Chng %	Del %
FPMSGQ_STR	LIST		*ALL	1.4					
IRLMLOCK1	LOCK		*ALL	2077.5					
	LOCK		*ALL	2086.5					

Pop-up screen to find CF

RMF Coupling Facility

Lock Structure : IRLMLOCK1
Coupling Facility : CF2
System : *ALL

JdK

Display of duplexed structure

```
IXC360I  12.30.50  DISPLAY XCF 142
STRNAME: IGWLOCK00
 STATUS: REASON SPECIFIED WITH REBUILD START:
           POLICY-INITIATED
         DUPLEXING REBUILD
           METHOD       : SYSTEM-MANAGED
         ...
         REBUILD PHASE: DUPLEX ESTABLISHED
         ...
  DUPLEX        : ENABLED
   PREFERENCE LIST: CF2      CF3

 DUPLEXING REBUILD NEW STRUCTURE
  ...
  CFNAME        : CF3
  ...

 DUPLEXING REBUILD OLD STRUCTURE
  ...
  CFNAME        : CF2
```

JdK

Response Time - Simplex Request

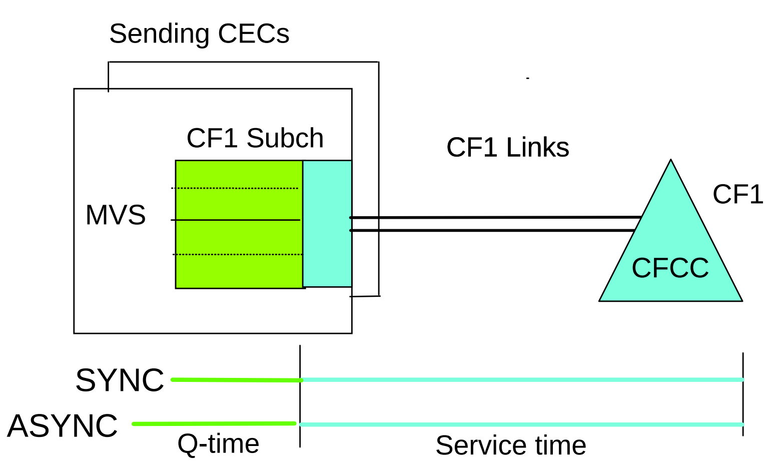

- MVS (XES) uses subchannels to manage link traffic
- 2 (COMPAT) or 7 (PEER) subchannels for each CF link
- Link types: ICS, ICB, IC or any combination. CEC favors fastest.

JdK

Response Time - Duplex Request

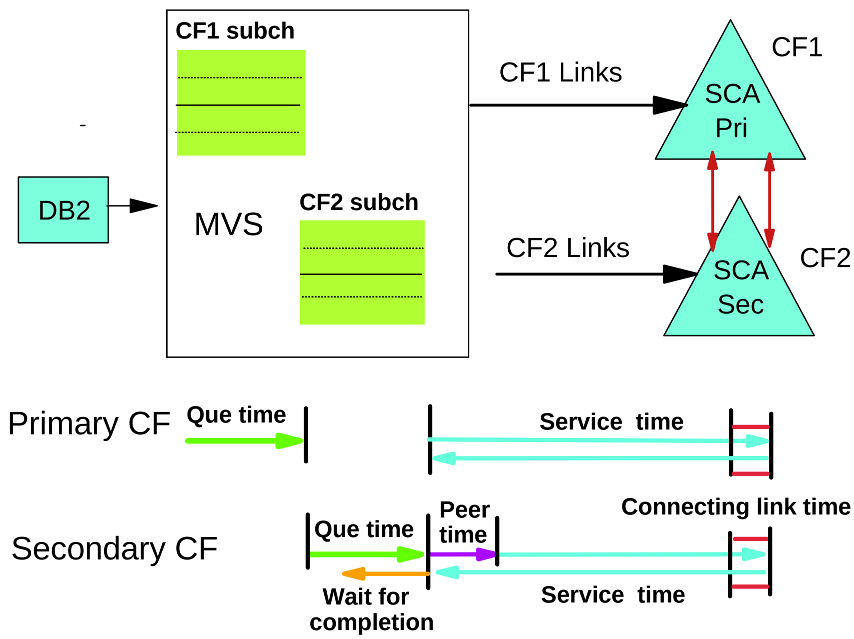

JdK

RMF- DB2 Lock, all requests duplexed

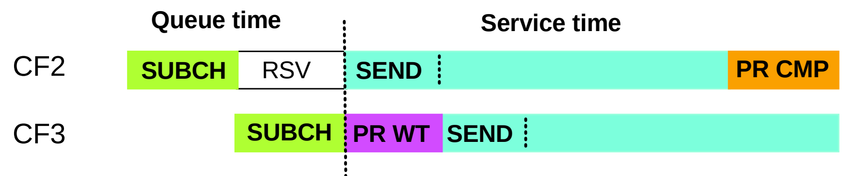


```
--------------------------------------------------------------------------------------
COUPLING FACILITY NAME = CF2
--------------------------------------------------------------------------------------
 STRUCTURE NAME = DSNDB1G_LOCK1     TYPE = LOCK   STATUS = ACTIVE PRIMARY
        # REQ    -------------- REQUESTS ------------    ------------- DELAYED REQUESTS -----------
SYSTEM  TOTAL                #     % OF  -SERV TIME(MIC)-   REASON     #    % OF   --- AVG TIME(MIC) ---
NAME    AVG/SEC             REQ    ALL     AVG   STD_DEV              REQ    REQ    /DEL  STD_DEV   /ALL

JA0        179K   SYNC    163K    91.5   108.5    137.6   NO SCH      0     0.0    0.0      0.0    0.0
         99.28    ASYNC    15K     8.5   483.9    474.0   PR WT       0     0.0    0.0      0.0    0.0
                  CHNGD     0      0.0  INCLUDED IN ASYNC PR CMP    47K    26.1   24.4     99.5    6.4

--------------------------------------------------------------------------------------
COUPLING FACILITY NAME = CF3
--------------------------------------------------------------------------------------
STRUCTURE NAME = DSNDB1G_LOCK1     TYPE = LOCK    STATUS = ACTIVE SECONDARY
        # REQ    -------------- REQUESTS ------------    ------------ DELAYED REQUESTS ------------
SYSTEM  TOTAL                #     % OF  -SERV TIME(MIC)-   REASON     #    % OF   --- AVG TIME(MIC)----
NAME    AVG/SEC             REQ    ALL     AVG   STD_DEV              REQ    REQ    /DEL  STD_DEV   /ALL

JA0        177K   SYNC    162K    91.4   103.8    147.8   NO SCH      0     0.0    0.0      0.0    0.0
         98.52    ASYNC    15K     8.6   473.6    464.9   PR WT    177K    100     5.4     15.7    5.4
                  CHNGD     0      0.0  INCLUDED IN ASYNC PR CMP   131K    73.7   15.1     56.2   11.1
```

JdK

RMF CF Structure Activity, cont.

Ex - MQ MSG - some requests duplexed

```
--------------------------------------------------------------------------------------------------
COUPLING FACILITY NAME = CF2   Internal z990
--------------------------------------------------------------------------------------------------
                                         COUPLING  FACILITY  STRUCTURE  ACTIVITY
--------------------------------------------------------------------------------------------------
 STRUCTURE NAME = MQGPMSGQ1         TYPE = LIST    STATUS = ACTIVE SECONDARY
         # REQ     -------------- REQUESTS -------------    ------------ DELAYED REQUESTS -------------
SYSTEM   TOTAL              #     % OF  -SERV TIME(MIC)-     REASON    #    % OF  --- AVG TIME(MIC) ----
NAME     AVG/SEC           REQ     ALL     AVG    STD_DEV            REQ     REQ    /DEL  STD_DEV   /ALL

JB0        198K    SYNC   5978     3.0   103.8     103.2     NO SCH    0     0.0     0.0      0.0    0.0
         110.0     ASYNC   192K   97.0   163.0     302.7     PR WT   198K    100     4.5      5.0    4.5
                   CHNGD     2     0.0  INCLUDED IN ASYNC    PR CMP  166K   83.9    44.6    139.8   37.4

--------------------------------------------------------------------------------------------------
COUPLING FACILITY NAME = CF3   Internal z900 Turbo
--------------------------------------------------------------------------------------------------
                                         COUPLING  FACILITY  STRUCTURE  ACTIVITY
--------------------------------------------------------------------------------------------------

 STRUCTURE NAME = MQGPMSGQ1         TYPE = LIST    STATUS = ACTIVE PRIMARY
          # REQ     -------------- REQUESTS -------------    ----------  DELAYED REQUESTS -------------
SYSTEM    TOTAL              #     % OF  -SERV TIME(MIC)-     REASON    #    % OF  --- AVG TIME(MIC)----
NAME      AVG/SEC           REQ     ALL     AVG    STD_DEV            REQ     REQ    /DEL  STD_DEV   /ALL

JB0        3319K    SYNC   3035K   91.4    27.6      28.0     NO SCH   26     0.0   643.9    560.8    0.0
           1844     ASYNC   284K    8.6   171.2     244.3     PR WT     0     0.0     0.0      0.0    0.0
                    CHNGD    26     0.0  INCLUDED IN ASYNC    PR CMP  32K     1.0    65.3    488.8    0.6
```

JdK

RMF CF Structure Activity - Tuning

Gated by slowest resource (CEC,Link, CF)

1. If PR WGT is high (>10) - improve links for OTHER structure
   - Upgrade link technology
   - Add additional link(s)
   - Shared Sender CPs
2. If PR CMP is high - improve CF 'speed' for OTHER structure
   - Configuration change
     - Dynamic Dispatch off
     - Dedicated CPs
   - Upgrade technology
   - Add additional CF CPs

JdK

RMF CF Structure Activity - Example

CF2 - Internal z990

```
STRUCTURE NAME = ISTMNPS           TYPE = LIST   STATUS = ACTIVE PRIMARY
          # REQ    -------------- REQUESTS -------------   -------------- DELAYED REQUESTS -------------
SYSTEM    TOTAL              #     % OF   -SERV TIME(MIC)-    REASON     #    % OF   ---- AVG TIME(MIC) -----
NAME      AVG/SEC           REQ    ALL      AVG    STD_DEV              REQ    REQ    /DEL      STD_DEV    /ALL

JA0        1277K    SYNC   5056     0.1    174.9    146.2     NO SCH 1753    0.1    49.1       127.6     0.1
          709.7     ASYNC  1271K   20.7    370.8    428.5     PR WT  1310K   103     4.5        12.0     4.6
                    CHNGD  1311     0.0   INCLUDED IN ASYNC   PR CMP 1082K  84.8    42.6       118.0    36.1
```

CF1 - S/A G5

```
STRUCTURE NAME = ISTMNPS           TYPE = LIST   STATUS = ACTIVE SECONDARY

JA0        1278K    SYNC   5056     0.1    221.7    177.6     NO SCH 2085    0.2    2142        2404     3.5
          709.7     ASYNC  1271K   20.7    396.7    399.6     PR WT     0    0.0     0.0         0.0     0.0
                    CHNGD  1312     0.0   INCLUDED IN ASYNC   PR CMP  194K  15.2   101.2       419.7    15.4
```

Two hours later - after recovery scenario

CF2 - Internal z990

```
STRUCTURE NAME = ISTMNPS           TYPE = LIST   STATUS = ACTIVE PRIMARY

JA0        1350K    SYNC   5393     0.1    141.9    566.2     NO SCH 1568    0.1   10080        9787    11.7
          750.0     ASYNC  1344K   17.8    298.2    889.6     PR WT  1376K   102     3.8         0.9     3.8
                    CHNGD   609     0.0   INCLUDED IN ASYNC   PR CMP  682K  50.5    80.0       680.7    40.4
```

CF3 - S/A z900 Turbo

```
STRUCTURE NAME = ISTMNPS           TYPE = LIST   STATUS = ACTIVE SECONDARY

JA0        1350K    SYNC   5393     0.1    177.5    573.8     NO SCH 1479    0.1   300.7        1374     0.3
          750.0     ASYNC  1344K   17.8    326.0   1007.1     PR WT     0    0.0     0.0         0.0     0.0
                    CHNGD   609     0.0   INCLUDED IN ASYNC   PR CMP  668K  49.5    34.1       456.1    16.9
```


RMF CF to CF Link - Tuning

Every duplexed request generates two exchanges of signals between CFs, so impact is doubled.

High service time?

If CFs are close, improve link technology

Improve CF utilization

- More CPs
- Improve technology

Overflow problem in reporting CF-to-CF service time

XES APAR OW55968

RMF APAR OW55586

CFCC CFLevel 11, 5.04; CFLevel 12, 12.08

RMF CF to CF - New report

```
CF1 - G5                                   CF TO CF ACTIVITY
--------------------------------------------------------------------------------
          # REQ                          ----------- REQUESTS -----------
PEER      TOTAL       -- CF LINKS --              #    -SERVICE TIME(MIC)-
CF        AVG/SEC     TYPE       USE             REQ       AVG     STD_DEV

CF3      581643       CBR          1     SYNC  581643      31.5       6.0
          323.1       CFR          5
```

```
CF3 - Z900 Turbo                                   CF TO CF ACTIVITY
--------------------------------------------------------------------------------------
          # REQ                          ----------- REQUESTS -----------    -----
PEER      TOTAL       -- CF LINKS --              #    -SERVICE TIME(MIC)-
CF        AVG/SEC     TYPE       USE             REQ       AVG     STD_DEV

CF1      580582       CFR          2     SYNC  580582      50.5      10.4      SYNC
          322.5
```

CF-to-CF time accounts for a sizable share of the service time.

JdK

RMF CF to CF - New Report

```
CF2 - z900 S/A                                  CF TO CF ACTIVITY
-------------------------------------------------------------------------------------------
          # REQ                      ----------- REQUESTS -----------   ---------- DELAYED REQ
PEER      TOTAL      -- CF LINKS --          #    -SERVICE TIME(MIC)-        #     % OF
CF        AVG/SEC    TYPE       USE        REQ       AVG     STD_DEV        REQ     REQ    /DEL

CF3        6215K     CBP          2    SYNC    6215K    14.8        5.7    SYNC      0     0.0     0.0
          3453.0     CFP          4
```

```
CF3 - z900 Turbo                                CF TO CF ACTIVITY
-------------------------------------------------------------------------------------------
          # REQ                      ----------- REQUESTS -----------   ---------- DELAYED REQ
PEER      TOTAL      -- CF LINKS --          #    -SERVICE TIME(MIC)-        #     % OF  -----
CF        AVG/SEC    TYPE       USE        REQ       AVG     STD_DEV        REQ     REQ    /DEL

CF2        6212K     CBP          2    SYNC    6212K    19.9       11.0    SYNC      0     0.0     0.0
          3451.1     CFP          4

CF3            0     CBP          2    SYNC        0     0.0        0.0    SYNC      0     0.0     0.0
             0.0     ICP          4
```

CF that looks like it's attached to itself is the result of shared sender/receiver CF links

JdK

Benefits of CF Duplexing

CF Duplexing benefits are all related to recovery:

1. Provides an 'easy-to-implement' recovery mechanism for structures with no recovery - see Appendix C for list

Subsystem	Structure	Structure Type	User Managed Rebuilt	System Managed/SM Duplex
CICA	Temp Storage	LIST	NO	YES
CICS	Shared Data Tables	LIST	NO	YES
CICS	Named Counter	LIST	NO	YES
IMS	VSO	CACHE	NO	YES
JES2	Checkpoint	LIST	NO	YES
MQSeries	Shared Queues	LIST	NO	YES
DFSMS	HSM Common Recall Queue	LIST	NO	YES
WLM	IRD	CACHE	NO	YES

JdK

Benefits of CF Duplexing, cont.

2. Faster recovery from CF failures

- Much, much faster compared to log recovery (40x)
- Faster compared to rebuild (4x)

Total recovery time depends on the number of structures residing on the failing CF

Specifying DUPLEX(ENABLE) will increase failover time if there are more than 2 CFs in the sysplex

JdK

Benefits of Duplexing, cont.

3. Provides an alternative to staging datasets which might not be used because of performance.

Staging Datasets

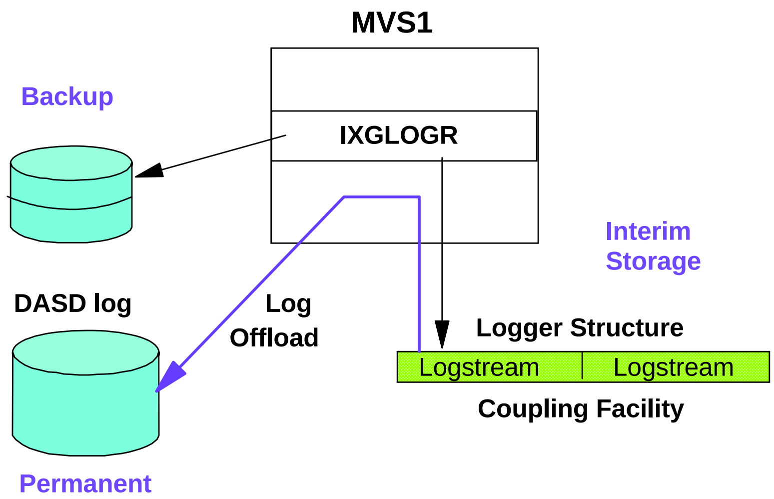

JdK

Benefits of Duplexing, cont.

4. Provides Failure Isolation for internal CFs

Without CF Duplexing, internal CFs are best suited to:

- CF for a single CEC sysplex
- CF which is not part of this sysplex
- Structures which don't need a local copy from the failing system to rebuild

► Using an ICF may allow you to upgrade CF technology and improve overhead

JdK

CF Overhead based on 9 CF ops /MIPS

Host CF	G4	G5	G6	z800	z900 1xx	z900 2xx	z890	z990 3xx
C04 - SM	11%	16%	19%	21%	22%	25%		
C05/HiPerLink	10%	14%	16%	18%	19%	22%	26%	30%
R06- HL	9%	12%	14%	16%	17%	19%	22%	26%
R06 - ICB	---	9%	10%	---	13%	14%	17%	20%
G5/G6 - IC	---	8%	8%	---	---	---		---
z800 - ISC	9%	11%	12%	11% Peer	12% Peer	13% Peer	15% Peer	18% Peer
z800 - ICB/IC	---	---	---	9% Peer	10% Peer	11% Peer	12% Peer	14% Peer
z900 - ISC	9%	11%	12%	10% Peer	11% Peer	12% Peer	14% Peer	16% Peer
z900 - ICB / IC	---	8%	9%	8% Peer	9% Peer	10% Peer	11% Peer	12% Peer
z890 - ISC	8%	8%	9%	9% Peer	10% Peer	11% Peer	13% Peer	15% Peer
z890 - ICB/IC	---	8%	8%	7% Peer	8% Peer	8% Peer	9% Peer	10% Peer
z990 - ISC	8%	8%	9%	9% Peer	10% Peer	11% Peer	13% Peer	14% Peer
z990 - ICB/IC	---	8%	9%	7% Peer	8% Peer	8% Peer	9% Peer	9% Peer

JdK

Cost of CF Duplexing

Cost of Duplexed pair of ops versus single op

Cost	Storage	CPU	CF link	CF CPU
User Managed	2x	2x	2x	2x
System Managed	2x	3x to 4x	4x to 5x	6x to 8x
Notes	No change - white space now "in-use"		May need less "white space"	May need less "white space"

- % of structure activity to be duplexed (thus pay above costs)
  - cache: 1% to 100% (typically 20%)
  - lock: 100%
  - list: near 100%

JdK

Ex. SMD Cost Estimates
Datasharing structures not duplexed

Host CPU capacity	simplex	impact	duplex
UM GBPs	5%	20%@2x	6%
SM lock	2%	0%@4x	2%
SM list	1%	100%@3x	3%
not duplexed	2%	NA	2%
TOTAL	10%		13%

Avg CF CPU busy	simplex	impact	duplex
UM GBPs	15%	20%@2x	18%
SM lock	5%	0%@5x	5%
SM list	4%	100%@4x	16%
not duplexed	6%	NA	6%
TOTAL	30%		44%

Avg CF link busy	simplex	impact	duplex
UM GBPs	5%	20%@2x	6%
SM lock	2%	0%@8x	2%
SM list	1%	100%@6x	6%
not duplexed	3%	NA	3%
TOTAL	11%		17%

JdK

Ex. SMD Cost Estimates
All structures duplexed

Host CPU capacity	simplex	impact	duplex
UM GBPs	5%	20%@2x	6%
SM lock	2%	100%@4x	8%
SM list	1%	100%@3x	3%
not duplexed	2%	NA	2%
TOTAL	10%		19%

Avg CF CPU busy	simplex	impact	duplex
UM GBPs	15%	0%@2x	18%
SM lock	5%	100%@5x	25%
SM list	4%	100%@4x	16%
not duplexed	6%	NA	6%
TOTAL	30%		65%

Avg CF link busy	simplex	impact	duplex
UM GBPs	5%	0%@2x	6%
SM lock	2%	100%@8x	16%
SM list	1%	100%@6x	6%
not duplexed	3%	NA	3%
TOTAL	11%		31%

JdK

Recommendations

- Duplexing is selectable by structure - each structure should be evaluated for value versus cost
  1. Start with structures that have no recovery mechanism
  2. Consider structures that require staging datasets or reading log datasets to recover
  3. If using internal CFs, duplex structures requiring local copy to rebuild.

Monitor resource utilization - same R.O.Ts as simplex

If more than two CFs, use DUPLEX(ALLOWED) to minimize failover time, especially for large number of structures

AutoAlter APAR OW50397 to allow 90% CF storage usage

Read System-Managed CF Structure Duplexing - whitepaper

JdK

Staged implementation

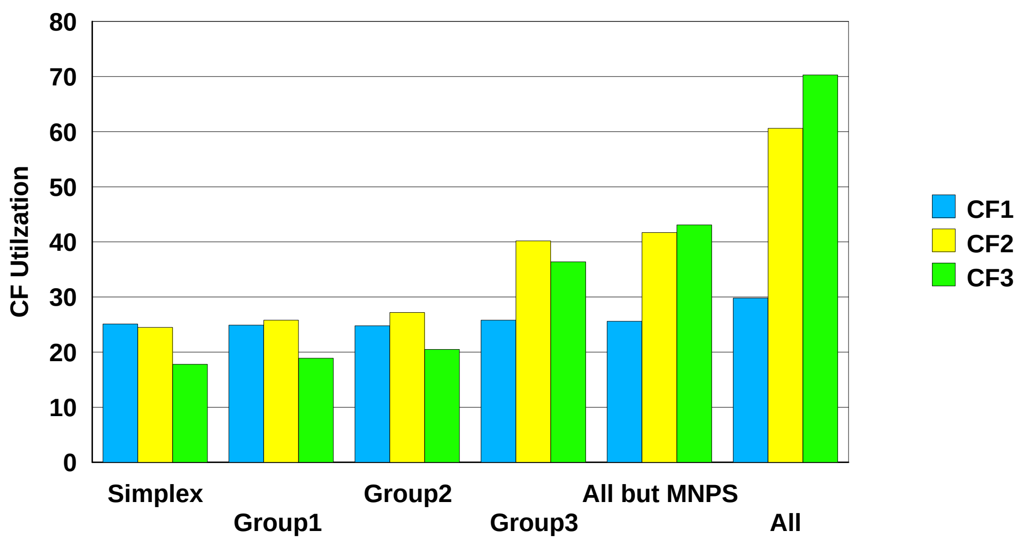

JdK

Sharing CF CP resource - examples

1. Want Production and Test CF on latest technology - but don't need capacity of 2 CF CPs
   - Test CF - performance not important
   - Production CF - optimal performance

   Goal: Minimal CP resource for test image with no impact to production image.

2. Need additional CF capacity only during peak periods
   - One CF CF is normally adequate to handle load.
   - Ocassionally suffer performance lags during peaks.

   Goal: Automatically provide addition CF CP resource only when needed.

JdK

Implications of Sharing CF CP resource

A. Processor utilization
   - LIC codes runs in an 'active wait' polling for work, so LPAR sees the CF image as 100% busy and give the CF all the processor resources available, even at very low CF rates.
   - CP resource is apportioned by LPAR weight, **BUT**...LPAR gives each CF image control every 125 microsecs, so a CF image with low weight gets more resource than expected.

B. CF request response times

If a request to CF cannot be executed because it's timesliced out, the request waits:

1. If SYNC request, the sender waits
   - Service times go up -> more sender cycles used
   - Heuristic algorithm provides some relief by changing SYNC to ASYNC
2. Subchannel is held longer -> more channel utilization
3. Requests to System Managed CF Duplexing structures can timeout, "breaking duplexing"

JdK

How many CPs really assigned to the CF?

1. Standalone CF - No operating system so only RMF CF reports

Mon I - Post processor

```
                         COUPLING  FACILITY  USAGE  SUMMARY
---------------------------------------------------------------------------------
COUPLING FACILITY         9672        MODEL A04      CFLEVEL  13
AVG CF UTILIZATION (% BUSY)   22.7   LOGICAL PROCESSORS: DEFINED 4      EFFECTIVE    4.0
```

Mon III - Real time

```
                    RMF V1R5   CF Overview      - UTCPLXJ8          Line 1 of 3

Samples: 120     Systems: 14   Date: 08/10/04  Time: 10.11.00  Range: 120   Sec

---- Coupling Facility -----    ---- Processor -----   Request     -- Storage --
Name       Type   Model Level   Util% Defined Effect   Rate        Size     Avail

CF1        2086     A04    13   22.0       4     4.0     11623     6078M    4055M
CF2        2084     D32    13   32.2       3     3.0     27646     6078M    2836M
CF3        2064     212    13   26.4       4     4.0     13120     6078M    2095M
```


How many CPs really assigned to the CF?

2. Internal CFs - RMF CF report

```
                         COUPLING  FACILITY  USAGE  SUMMARY
---------------------------------------------------------------------------------
COUPLING FACILITY         2064        MODEL 106      CFLEVEL  12
AVG CF UTILIZATION (% BUSY)  8.3   LOGICAL PROCESSORS: DEFINED 1      EFFECTIVE    0.7
```

RMF Partitioned Activity Report

```
Number of Physical Processors            8
                          CP             6
                          ICF            2
...
                        ---MSU---    Cap  Processor-  Logical Processors -- Physical Processors --
Name   S   Wgt  Def      Act  Def   WLM%  Num   Type  Effective    Total  LPAR Mgt  Effect.  Total
 ---   -   ---  ----     ----  ---  -----  ---   ----  ---------    -----  --------- -------  -----
CF1A   A    10    0      N/A  NO     0.0    1   ICF      66.23    66.27      0.02    33.12  33.13
CF2A   A    10    0      N/A  NO     0.0    1   ICF      66.24    66.27      0.02    33.12  33.14
CF2B   A    10    0      N/A  NO     0.0    1   ICF      66.24    66.27      0.02    33.12  33.14
*PHYSICAL*                                                                   0.58            0.58
TOTAL                                                                        0.63    99.35  99.99
```

CF- Shared CPs on sending CEC

CF partition will use all the CP resource it can get

```
                                P A R T I T I O N   D A T A

MVS PARTITION NAME                S00
NUMBER OF CONFIGURED PARTITIONS          4
NUMBER OF PHYSICAL PROCESSORS            4
WAIT COMPLETION                         NO
DISPATCH INTERVAL                  DYNAMIC

---- PARTITION DATA ----             --- AVERAGE PROCESSOR UTILIZATION PERCENTAGE -----
                          # OF       -LOGICAL  PROCESSORS  ----- PHYSICAL PROCESSORS --
NAME  STATUS  WGHTS  CAP  LPs   ...   EFFECTIVE     TOTAL  LPAR MGMT    EFFECTIVE   TOTAL
S00      A      10    NO    2           82.43       82.77       0.09        20.61   20.69
S01      A      20    NO    3           39.37       39.57       0.15        29.53   29.68
S02      A      75    NO    4            4.46        4.70       0.24         4.46    4.70
CF01     A       3    NO    1           98.11       98.17       0.01        24.53   24.54
*PHYSICAL*                                                      0.60                 0.60
                                                                -----       ------   -----
  TOTAL                                                         1.00        58.51   59.52
```

- Contention will limit CF CP resource to wgt. defined
- If the CF is sharing CPs, do not CAP the partition and give it a respectable weight (at least 50%)
- Anything less than a CP will elongate service time

JdK

Shared CF CPs - SM duplexing

Recommendation - give CF partition enough weight so it will be dispatched least 95% of the time.

Ex. Want to add CF image (2CPs) to CEC with 14 shared CPs.
Total weight of MVS images on CEC is 1000
CF will contain SM Duplexed structures

$$\text{CF Wgt} = \frac{\text{Factor} * \#\text{ CF CPs}}{\text{Total \# Shared CPs}} * (\text{Total Wgt} + \text{CF Wgt})$$

$$\text{CF Wgt} = \frac{.80 * 2}{14} * (1000 + \text{CF Wgt}) = 126$$

CF CPs	Factor
1	.95
2	.80
3	.60
4	.50

Potential Problems - if less than 95%

1. Increased reponse time
2. Break duplexing
3. Application timeouts

JdK

Don't want to assign a whole CP to test CF?

Dynamic CF Dispatching - allows tradeoff between CF response time and CP Utilization

- **At low utilization, CFCC suspended for short periods**
  - **More CP resource for other partitions, but CF requests delayed**
- **As utilization increases, less CFCC suspension**
  - **Less CP resource for other partitions, but faster CF requests**

```
 --------- PARTITION DATA --------- --- AVERAGE PROCESSOR UTILIZATION PERCENTAGES
----
                           # OF    -LOGICAL  PROCESSORS  ----- PHYSICAL PROCESSORS
--
NAME  STATUS  WGHTS  CAP    LPs      EFFECTIVE    TOTAL   LPAR MGMT   EFFECTIVE
TOTAL
S18     A      50    NO      5   .    47.20       47.58      0.19       23.60
23.79
S19     A      50    NO      5        47.63       47.86      0.12       23.82
23.93
S1A     A      50    NO      5        47.67       47.92      0.12       23.84
23.96
S1B     A      50    NO      5        47.66       47.89      0.12       23.83
23.95
CF1     A      40    NO      2        17.77       18.63      0.17        3.55
3.73
*PHYSICAL*                                                   0.63
```

At low utilization, less CPU resource used but...

Dynamic CF Dispatching

But CF response time increases....

```
                          COUPLING  FACILITY  USAGE  SUMMARY
-------------------------------------------------------------------------------
AVG. CF UTIL. (%BUSY)   23.6%   LOGICAL PROCESSORS:  DEFINED  1   EFFECTIVE  0.0

                          COUPLING  FACILITY  STRUCTURE  ACTIVITY
 ------------------------------------------------------------------------------
STRUCTURE NAME = CFTWDB2_LOCK1      TYPE = LIST
           # REQ    -------------- REQUESTS ------------    ...
 SYSTEM    TOTAL              #      % OF  -SERV TIME(MIC)-
 NAME      AVG/SEC           REQ     ALL       AVG    STD_DEV

 J90         122    SYNC      54     3.6%  1219.6     1055.6
            2.03    ASYNC     68     4.5%  2004.2     2441.7
                    CHNGD      0     0.0%  INCLUDED IN ASYNC
```

As activity in the test CF partition increases, more CPU resource is used and CF response time improves.

Dynamic Dispatch can help...if done right

To prevent the test CF image from consuming CPU resource assigned to production CF

- ➤ Dyn Disp OFF for production CF
- ➤ Dyn Disp ON for test CF

DYN DISP ON

```
STRUCTURE NAME = DB2P0G_LOCK1    TYPE = LOCK
               # REQ     -------------- REQUESTS -------------   ------------- DELAYED
     NAME      AVG/SEC            REQ    ALL      AVG    STD_DEV            REQ    REQ
     SYSTEM    TOTAL              #      % OF    -SERV TIME(MIC)-   REASON   #     % OF
     A81G      66747    SYNC      67K    100%     554.3    667.6
               74.16    ASYNC      0     0.0%      0.0       0.0    NO SCH   0     0.0%
```

DYN DISP OFF

```
STRUCTURE NAME = DB2P0G_LOCK1    TYPE = LOCK
               # REQ     -------------- REQUESTS -------------   -------------- DELAYED
     SYSTEM    TOTAL              #      % OF    -SERV TIME(MIC)-   REASON   #     % OF
     NAME      AVG/SEC            REQ    ALL       AVG   STD_DEV            REQ    REQ
     A81G      62443    SYNC      62K    100%      22.6     38.6
               69.38    ASYNC      0     0.0%      0.0       0.0    NO SCH   0     0.0%
```

JdK

Production CF1 and test CF2 - 1 CP

Dynamic CF Disp is OFF for CF1, ON for CF2

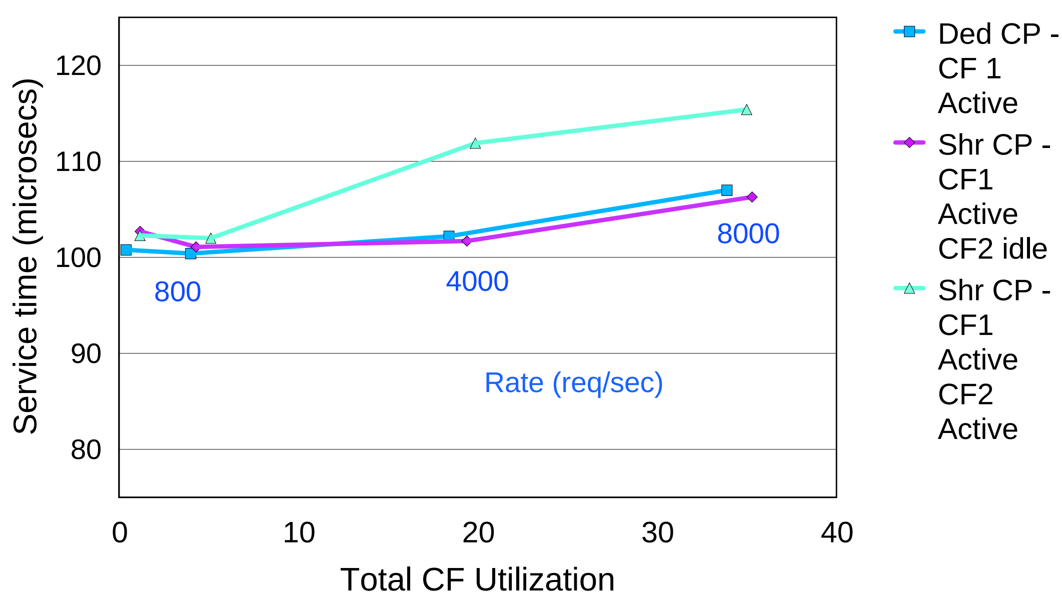

JdK

Production CF1 and test CF2 - 1 CP

Dynamic CF Disp is OFF for CF1, ON for CF2

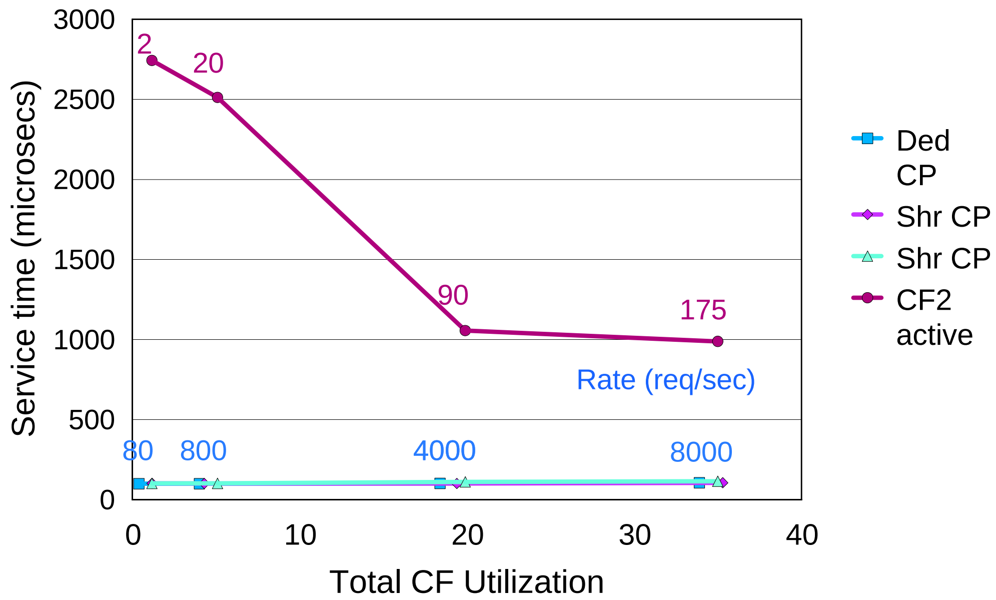

JdK

Dynamic Dispatch - CF utilization

CF utilization will be much higher on a CF with DYN DISP ON.
- **Dyn Disp puts CFCC code "to sleep" for periods of time.**
- **If requests are made to CF while it is "sleeping", when it wakes up, it will have have work to do, so it will look much busier than a CF with DYN DISP OFF.**

CF1 is processing 74.94 req/sec

DYN DISP ON

```
COUPLING FACILITY      2066        MODEL 0CF      CFLEVEL  12
AVG CF UTILIZATION  (% BUSY)  25.3   LOGICAL PROCESSORS:   DEFINED  1    EFFECTIVE   0.0
```

CF1 is processing 59.94 req/sec, almost the same amount

DYN DISP OFF

```
COUPLING FACILITY      2066        MODEL 0CF      CFLEVEL  12
AVG CF UTILIZATION  (% BUSY)   0.4   LOGICAL PROCESSORS:   DEFINED  1    EFFECTIVE   1.0
```

JdK

Dynamic ICF Expansion

Available of G3 and subsequent servers, CFLEVEL=5

Provides the advantages of ICFs and the ability to handle unexpected spikes in demand

CF partition with
- 1-n dedicated ICFs AND
- 1-n shared CPs OR 1-n shared ICFs

Recommendations for LPAR WGT
- For best response time, CF CPs should be dispatched as soon as they have work to do
  - ➤ At a minimum, CF CP should have weight of whole CP
  - ➤ CF CPs should be weighted 25 times higher than MVS CPs

JdK

Dynamic ICF Expansion

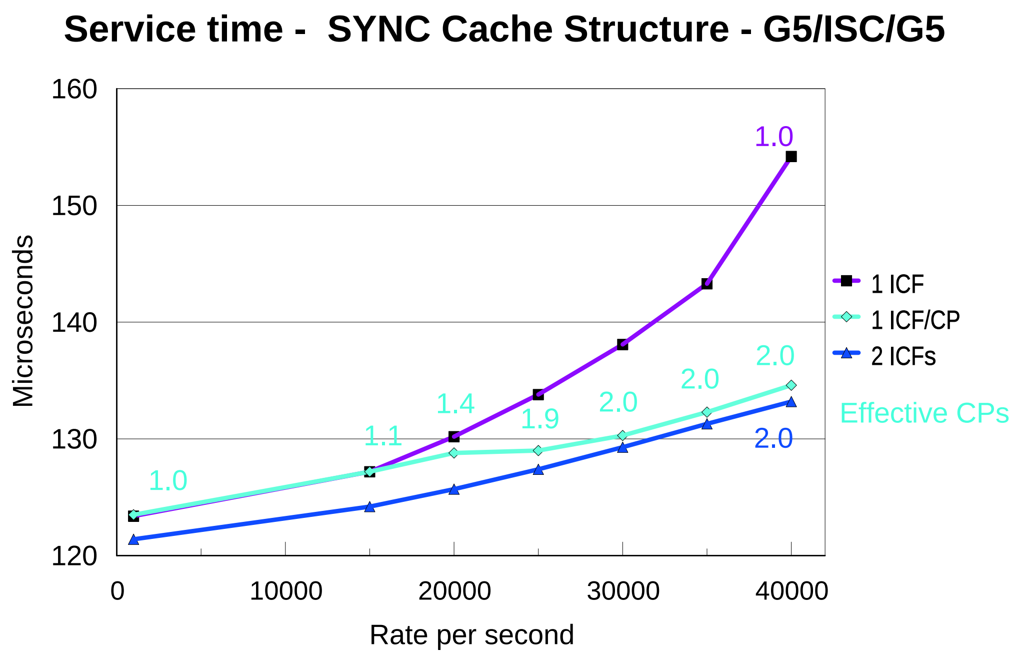

JdK

Dynamic CF Dispatching

In some configurations, Dyn Disp is forced ON

Coupling Facility Configuration	All Dedicated CFs or ICFs	Dedicated ICFs and Shared CPs	Dedicated ICFs and Shared ICFs	Shared ICF	Shared CP	Shared ICF
CF Model	-	-	-	z800 - 0CF z900 - 100	z800 z900	z800 z900
Dyn Disp Default Value	OFF	ON	ON	OFF	ON	OFF
Dyn Disp Forced?	No	Yes	Yes	No	No	No

In some configurations, Dyn Disp is defaulted ON

You can use D DYNDISP command on HMC to verify

JdK

Comparison of CP options

	CP definition	Dyn Disp	Effective CPs
CF1	2 Ded CPs	N/A	2.0
CF1	2 Shr CPs	OFF	2.0
CF1 CF2	2 Shr CPs 1 Shr CP	OFF ON	2.0 - 0.0 +
CF1 CF2	2 Shr CPs - wgt 75 1 Shr CP - wgt 25	OFF OFF	1.5 0.5
CF1 CF2	1 Ded/1 Shr CP 1 Shr CP	ON OFF	1.0 + 1.0 -
CF1 CF2	1 Ded 1 Ded	N/A N/A	1.0 1.0

If Dyn Disp is ON, as rate to that CF increases, effective CPs increase (+) and decrease (-)

JdK

Shared IFLs and ICFs

IFLs and ICFs are managed in one pool,
MVC CPs in another

```
MVS PARTITION NAME                       D02
IMAGE CAPACITY                           302
NUMBER OF CONFIGURED PARTITIONS           10
NUMBER OF PHYSICAL PROCESSORS              8
                   CP                      5
                   ICF                     3
WAIT COMPLETION                           NO
DISPATCH INTERVAL                    DYNAMIC

 --------- PARTITION DATA -----------------  - LP DATA -   -- AVERAGE PROCESSOR UTILIZATION PERCENTAGES --
0                  ----MSU----  -CAPPING--  PROCESSOR-     LOGICAL PROCESSORS  --- PHYSICAL PROCESSORS ---
 NAME       S   WGT  DEF     ACT  DEF   WLM%  NUM   TYPE     EFFECTIVE     TOTAL  LPAR MGMT  EFFECTIVE  TOTAL
0A02        A    70    0      31  NO     0.0    2   CP         53.70     25.87       ****      21.48  10.35
 ...                                          .......                                                         ....
 TEST       A     1    0       0  NO     0.0    1   CP          0.00      0.00       0.00       0.00   0.00
 *PHYSICAL*                                                                          0.97              0.97
                                                                                   ------     ------ ------
   TOTAL                                                                             1.46      92.71  83.05

 DCF1       A    45    0      11  NO     0.0    1   ICF        18.28     18.28       0.00       6.09   6.09
 PCF1       A   333    0      59  NO     0.0    1   ICF        97.51     97.53       0.01      32.50  32.51
 XD1        A   621    0     111  NO     0.0    2   ICF        91.55     91.59       0.02      61.04  61.06
 XP1        A     1    0       0  NO     0.0    2   ICF         0.00      0.00       0.00       0.00   0.00
 *PHYSICAL*                                                                          0.18              0.18
                                                                                   ------     ------ ------
   TOTAL                                                                             0.21      99.63  99.84
```

JdK

Shared IFLs and ICFs, cont.

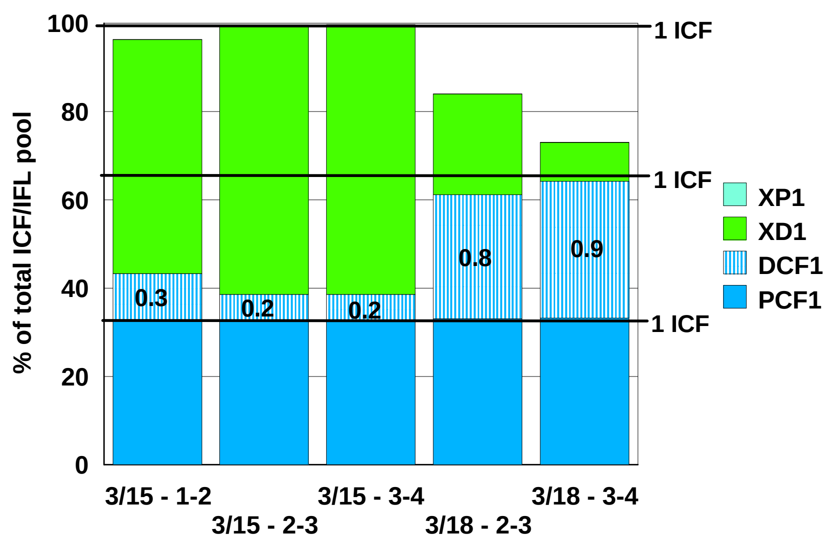

JdK

Shared MVS CPs - examples

1. Varying workloads
   - One MVS image runs OLTP which is "hot" during prime shift
   - Another MVS image runs the batch processing which is done almost exclusively at night.

   Goal: Maximum CP resource for each image during its peak activity

2. Migration testing
   - Test image with newest product
   - Full scale Production image

   Goal: Minimal CP resource for test image with no impact to production workload.

JdK

Allocating CP resources

Assign logical CPs and LPAR weight

➤ This is amount of CPU resource guaranteed to the image if there's contention.

If no contention, each image could get more.

- For example, test image may get production resource

What can go wrong?

1. Too many CPs defined on CEC

   R.O.T - Total # LPs ≤ 2 times the Total # CPs

   - Too many LPs causes additional overhead

2. Small image in a GRS Ring

   - Delays when XCF signal gets to small image and it has to wait for timeslice.

JdK

What's really running on this CEC?

RMF Mon III - Realtime view of EVERY image

RMF V1R5 CPC Capacity Line 1 o

Samples: 120 System: J90 Date: 08/09/04 Time: 15.36.00 Range: 120

Partition: J90 2064 Model 212
CPC Capacity: 445 Weight % of Max: 36.3 4h MSU Average: 91
Image Capacity: 297 WLM Capping %: **** 4h MSU Maximum: 123

Partition	--- MSU --- Def	Act	Cap Def	Proc Num	Logical Util % Effect	Total	- Physical Util % - LPAR	Effect	Total
*CP							2.3	94.7	97.0
JF0	0	114	NO	5.0	61.1	61.6	0.2	25.5	25.7
J90	0	105	NO	6.0	46.7	47.4	0.3	23.4	23.7
TPN	0	49	NO	5.0	26.1	26.5	0.2	10.9	11.1
Z0	0	135	NO	5.0	72.4	72.8	0.1	30.2	30.3
Z1	0	22	NO	3.0	19.1	19.6	0.1	4.8	4.9
PHYSICAL							1.4		1.4

JdK

What's really running on this CEC? - cont.

P A R T I T I O N D A T A R E P O R T

z/OS V1R5 SYSTEM ID SH0 DATE 09/01/2004 INTERVAL 29.59.931
RPT VERSION V1R5 RMF TIME 09.30.00 CYCLE 1.000 SECONDS

MVS PARTITION NAME SH0 NUMBER OF PHYSICAL PROCESSORS 16
IMAGE CAPACITY 31 CP 12
NUMBER OF CONFIGURED PARTITIONS 8 ICF 4
WAIT COMPLETION NO
DISPATCH INTERVAL DYNAMIC

-------- PARTITION DATA -----------------		----MSU----			-CAPPING--		-- LOGICAL PROCESSOR-		... - AVERAGE PROCESSOR UTILIZATION PERCENTAGE LOGICAL PROCESSORS ---		PHYSICAL PROCESSORS ---		
NAME	S	WGT	DEF	ACT	DEF	WLM%	NUM	TYPE	EFFECTIVE	TOTAL	LPAR MGMT	EFFECTIVE	TOTAL
SG0	A	396	0	94	NO	0.0	8.0	CP	37.67	37.76	0.06	25.11	25.17
SH0	A	DED	0	31		0.0	1	CP	99.96	99.99	0.00	8.33	8.33
SE0	A	45	0	19	NO	0.0	4.0	CP	15.31	15.57	0.09	5.10	5.19
SF0	A	99	0	55	NO	0.0	5.0	CP	35.22	35.61	0.16	14.68	14.85
SM0	A	171	0	57	NO	0.0	5.0	CP	36.49	37.08	0.25	15.20	15.45
SN0	A	189	0	59	NO	0.0	5.0	CP	37.29	37.91	0.26	15.54	15.80
PHYSICAL											0.91		0.92
											------	------	------
TOTAL											1.73	83.96	85.09

- Total CPs = 12, Total LPs = 28

  Why are some Logical Processors numbers expressed as decimals?

JdK

Intelligent Resource Manager (IRD)

IRD - distributes CP and I/O resources across a group (**LPAR cluster**) of MVS partitions which are on the same CEC and in the same sysplex.

Using sysplex goals

- WLM adjusts LPARs CP resource distribution
- WLM prioritizes I/O resource allocation.

Parallel sysplex - distributes workload across a sysplex which spans one or more CECs.

- WLM manages resources in a single image and gives recommendations to workload managers
- LPAR distributes CP resource within the CEC based on predefined Weights

JdK

Benefits of IRD

At very low cost, LPAR clustering improves systems management by managing:

1. CPU resources
   A. **Dynamic** distribution of capacity within LPAR cluster while protecting capacity outside LPAR cluster
   B. Improves efficiency
   C. Uses upgraded capacity immediately

2. IO resources
   A. Prioritizes work when I/O is constrained (CSSQ)
   B. Improves channel configuration efficiency (DCM)

Most useful when multiple images are consolidated on one CEC and/or workloads change dynamically.

JdK

Sharing CF Link resource - examples

1. Pool link resources for multiple MVS images on the same CEC to the same CF
   - Instead of 1 Ded CF link to each of 2 images
   - 2 CF links shared by each image

   Goal: Provide redunancy and additional link resource without impacting performance.

2. "Stretch" limited number of CF links when adding an internal CF
   - CF images require receiver links
   - System managed duplexing also requires CF sender links.

   Goal: Allow MVS and CF images to share CF sender links to "free up" additional CF links

Sharing CF links

Individual MVS images manage CF link traffic by queuing requests in S/W subchannels. But they have no knowledge of traffic generated by other sharing MVS images. This results in "path busy" conditions.

```
                                                       SUBCHANNEL  ACTIVITY
---------------------------------------------------------------------------
         # REQ                             ----------- REQUESTS -----------
SYSTEM   TOTAL   -- CF LINKS --  PTH                #   -SERVICE TIME(MIC)-
NAME     AVG/SEC TYPE  GEN  USE  BUSY               REQ      AVG     STD_DEV

JA0        4952K CBP     1    1  4570   SYNC       2470K   49.5       32.3
         5502.7  CFP     3    3         ASYNC      2488K  166.0      459.1
                 SUBCH  28   28         CHANGED    1680  INCLUDED IN ASYNC
                                        UNSUCC        0     0.0        0.0
JB0        4039K CBP     1    1  5573   SYNC       1353K   60.4       33.5
         4487.8  CFP     3    3         ASYNC      2675K  173.8      363.4
                 SUBCH  28   28         CHANGED    4649  INCLUDED IN ASYNC
                                        UNSUCC        0     0.0        0.0
```

z/OS 1.2 - Change to path busy

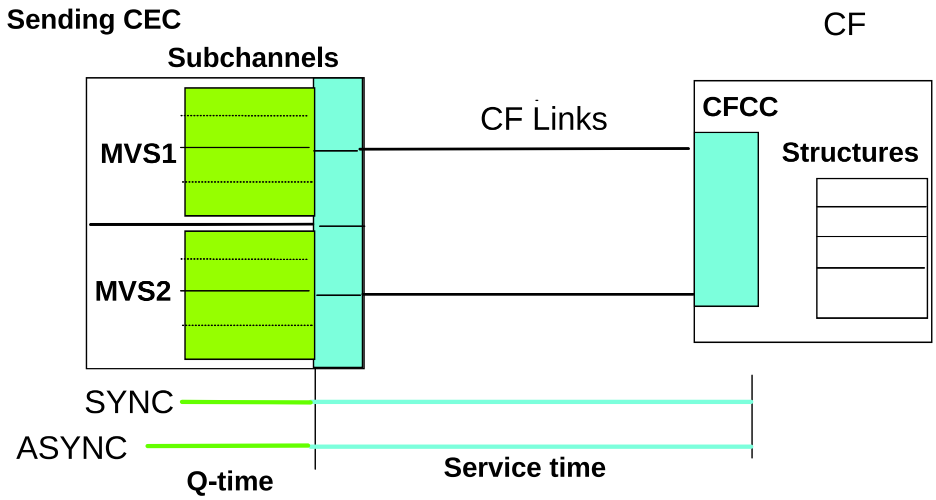

- <z/OS 1.2, SYNC reqs retried immediately / ASYNC requeued
  z/OS 1.2, all reqs retried immediately.
- Path busy retry time included in service time in all cases.
  - ►R.O.T. - Path Busy < 10% of Total Requests

JdK

Subchannel tuning - OW54796

Peer subchannels (2 -> 7 subchannels) aggravates path busy problem

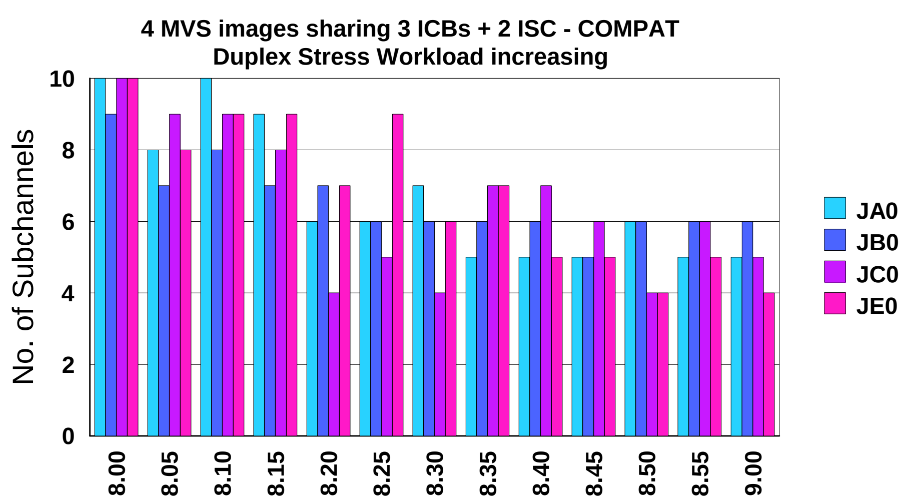

OW54796 Shifts point of queuing from path busy to subchannel busy

1. Allocates/deallocates subchannels (always at least 2)
   - keeps path busy in the 10-20% range
2. Improves efficiency in link constrained environments

JdK

Subchannel tuning - OW54796

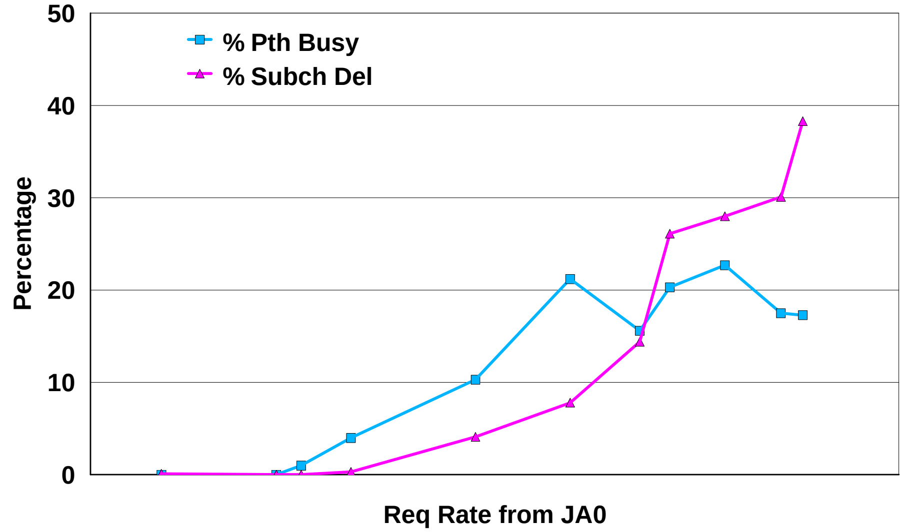

RMF APAR OW54820 supports subchannel tuning
- includes data from deallocated subchannels

JdK

Ex. Path busy, subchannel tuning

RMF Mon I

SYSTEM NAME	TOTAL AVG/SEC	-- CF LINKS -- TYPE	GEN	USE	PTH BUSY		# REQ	-SERVICE TIME(MIC)- AVG	STD_DEV
JB0	1564K	CBS	2	2	257K	SYNC	98555	55.7	80.3
	868.9	CFS	2	2		ASYNC	602754	166.2	154.6
		SUBCH	8	2		CHANGED	1701	INCLUDED IN ASYNC	

RMF Mon III

RMF V1R2 CF Systems - UTCPLXJ8

Samples: 120 Systems: 14 Date: 10/08/03 Time: 10.18.00 Range: 120 Sec

CF Name	System	Subch Delay %	-- Paths -- Avail	Delay %	-- Sync --- Rate	Avg Serv	------- Async ------- Rate	Avg Serv	Chng %	Del %
CF1	JA0	0.0	6	2.7	185.4	64	3145	200	0.0	0.2
	JB0	0.0	4	22.3	42.8	689	1007	187	0.1	0.3
	JC0	0.0	4	36.9						
	JE0	0.0	6	1.6						
	JF0	0.0	6	3.2						
	JG0	0.0	3	1.5						
	JH0	0.0	3	2.6						
	J80	0.0	2	0.0						
	J90	0.0	6	1.4						

RMF Coupling Facility - Subchs and Paths
Details for System : JA0
Coupling Facility : CF1

Subchannels Generated : 8
In Use : 2
Max : 8

Path IDs : 00 10 11 01
Types : CFS CBS CBS CFS

JdK

Additional Information

- Websites www.s390.ibm.com/servers/eserver/zseries
  - Parallel sysplex (CF sizer, CFLevel description)   .. /pso
    - System-Managed CF Structure Duplexing Implimentation Summary
    - System-Managed CF Structure Duplexing
  - RMF (tools, presentations, newsletters)   .. /rmf
- WSC FLASHs
  - Flash10159 New Heuristic Algorithm for CF Request Conversion
  - W99037 Performance Impacts of Using Shared ICF CPs
  - W9609 LPAR Performance Considerations for Parallel Sysplex Environments
- Publications
  - Setting up a Sysplex (SA22-7625-06)
  - z/Series 900 System Overview (SA22-1027-03b)
  - z/Series 990 System Overview (SA22-1032-00a)
  - Processor Resource/System Manager Planning Guide (SB10-7036-01)

JdK