#SHAREorg

SHARE
Technology · Connections · Results

# Diagnosing Sysplex Problems

Mark A Brooks
IBM

Wednesday, February 6, 2013: 3:00 PM-4:00 PM
Session Number 13402

mabrook@us.ibm.com

SHARE in San Francisco 2013

# Trademarks

**The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.**

| | | | | |
|---|---|---|---|---|
| IBM® | MQSeries® | S/390® | z9® | IBM® (logo) |
| ibm.com® | MVS™ | Service Request Manager® | z10™ | AIX® BladeCenter® |
| CICS® | OS/390® | Sysplex Timer® | z/Architecture® | DataPower® |
| CICSPlex® | Parallel Sysplex® | System z® | zEnterprise™ | DS4000® |
| DB2® | Processor Resource/Systems Manager™ | System z9® | z/OS® | DS6000™ |
| eServer™ | PR/SM™ | System z10® | z/VM® | DS8000® |
| ESCON® | RACF® | System/390® | z/VSE® | POWER7® |
| FICON® | Redbooks® | Tivoli® | zSeries® | ProtecTIER® |
| GDPS® | Resource Measurement Facility™ | VTAM® | zEC12™ | Rational® |
| HyperSwap® | RETAIN® | WebSphere® | Flash Express® | System Storage® |
| IMS™ | GDPS® | | | System x® |
| IMS/ESA® | Geographically Dispersed Parallel Sysplex™ | | | XIV® |

**The following are trademarks or registered trademarks of other companies.**

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license there from.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
InfiniBand is a trademark and service mark of the InfiniBand Trade Association.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

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

**Notes**:

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

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

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

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

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

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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

# Downloads available from SHARE website

- Presented Slides
  - Material to be presented during the session
  - “See handout” is a reference to slides in the handout that will not be presented
- Handout Slides
  - All of the above plus additional slides that will not be presented
    - Has a more complete treatment of key concepts
    - Includes examples of relevant messages, display output, and reports to illustrate what you would be looking at when diagnosing problems
  - If following along with the handout during the presentation, you can use the slide numbers in the lower left corner to keep pace (there will be lots of skipped slides)

# Sysplex environment

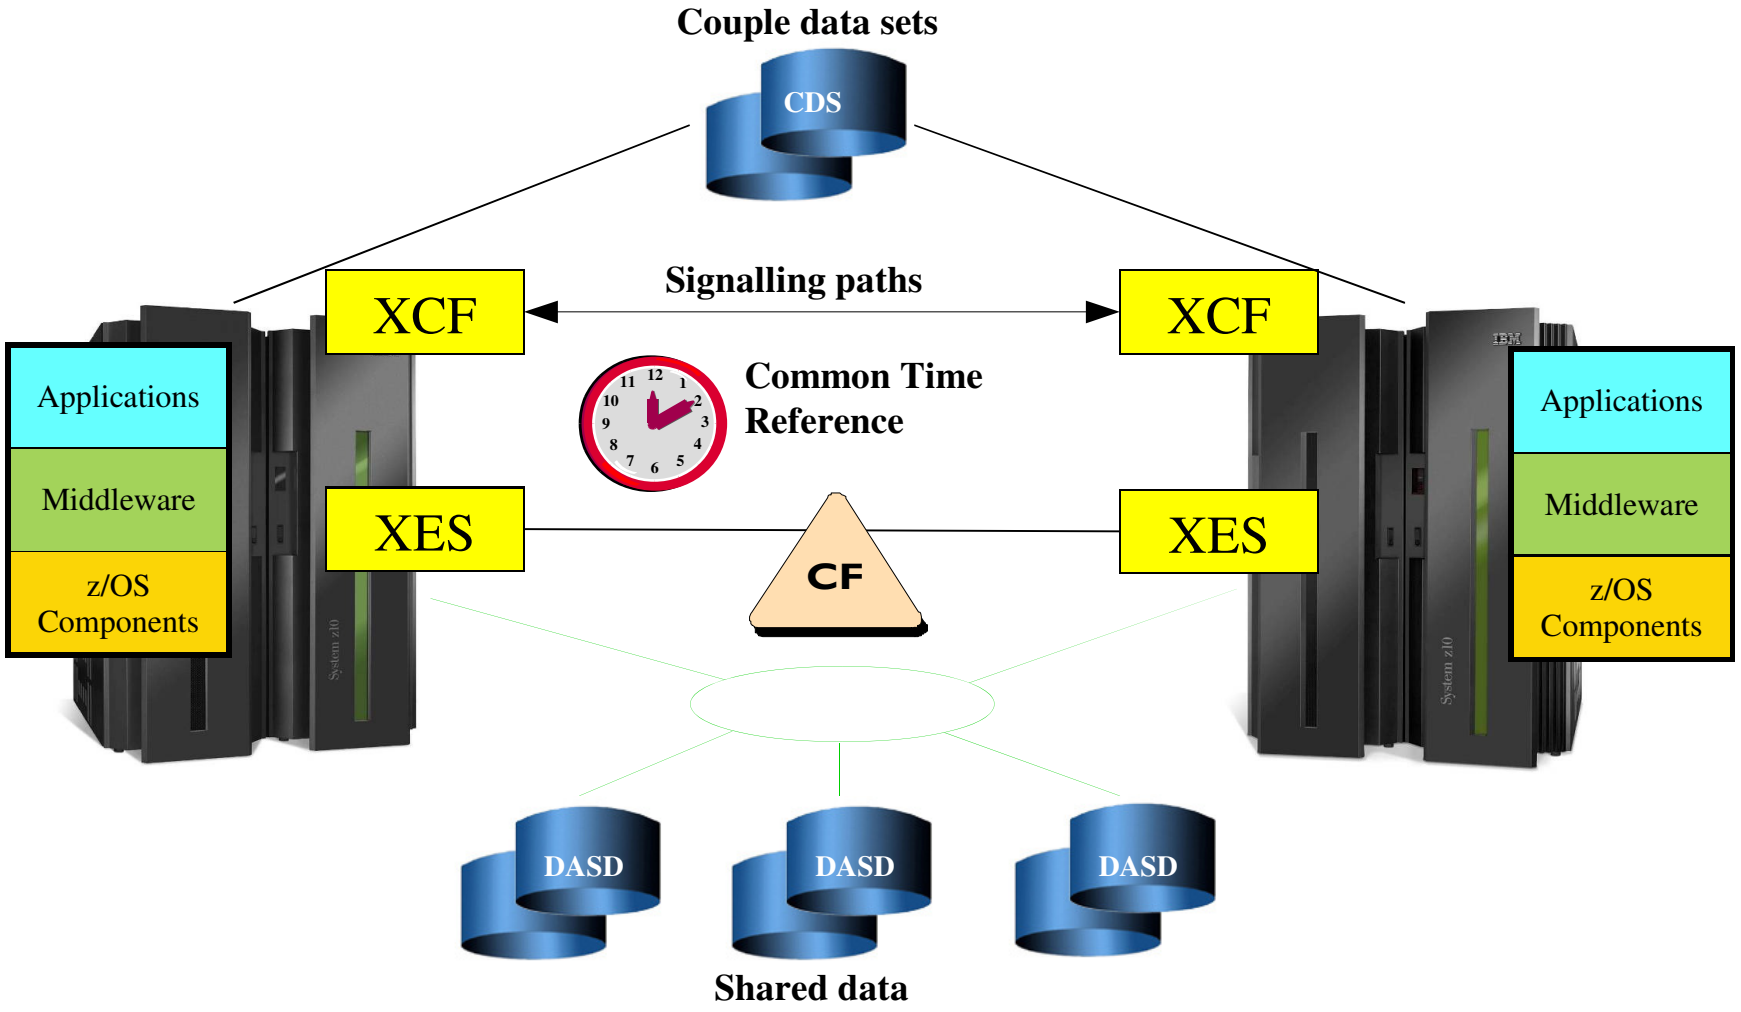

# Context

- In the sysplex, we have instances of z/OS, middleware, and applications running on individual systems
- These individual instances cooperate with one another to perform some function
- A given sysplex function may rely on the services of other sysplex functions, sometimes in rather non-obvious, even circular ways

# Sysplex Problems

- Sympathy sickness (hangs)
  - If an instance of a sysplex function is unresponsive, it may:
    - Hold serialization on shared resource
    - Stop sending replies
  - Other instances may hang as a result
- Performance issues (delays)
  - Instance is responsive, but “slow”
  - May impact response time of individual requests, which at high request rates, can lead to significant throughput problems and/or large queue effects that also give the appearance of hangs

# Signs that a sysplex might be sick

- IXC102A “reply down when system reset”
- IXC426D “sending signals but not updating status, how to proceed?”
- IXC431I “member stalled”
- IXC631I “member stalled and impacting other systems”
- IXC633I “member is impaired”
- IXL040E, IXL041E “structure connector not responding”
- IXL045E “having trouble getting connector SRBs to run”
- IXL044I “errors communicating with CF”
- ISG361A “waiting for list lock”
- ISG178E “ring disruption”
- D GRS,C
  - ISG343I output
- IOS071I “MIH timeout”

# Signs that a sysplex might be sick ...

- IXC426D “sending signals, not updating status.  How to proceed?”
- IXC427A “sending signals, not updating status.  SFM will act”
- IXC446I “sending signals, not updating status.  SFM will act”
- IXC256A “cannot remove CDS until these systems respond”

But the real question is: Why is my sysplex sick?

And our concern is that the symptoms are not always so obvious

# Getting to root cause of problem can be hard

- You don't have a “sysplex problem”, you have a “problem”
- The problem is more complex because of its sysplex context
  - Lots of “pieces” scattered throughout the sysplex
  - Any given symptom may have a root cause elsewhere in the sysplex
  - As problem persists, its impact tends to spread, which induces more symptoms, which makes the root cause harder to find
  - Lack expertise to understand relationships and dependencies
- But ultimately there is likely some one thing on some one system that is causing the problem
- How to find that thing?

# My thesis

- We do not have the expertise needed to do sysplex diagnosis
- We do not know:
  - Implementations, interactions, and dependencies
  - Causes, relationships, or relevance of various symptoms
  - How a given problem might impact the sysplex
  - All potential causes of a given symptom
- But we do know:
  - Sysplex application instances interact with one another
  - Most exploit sysplex services to do so
  - These services and their usage can be observed and analyzed
- So despite our imperfect knowledge, there is hope
  - We can make sure the sysplex infrastructure is sound
  - If applications can readily communicate and share data, we likely have a problem that will yield to traditional single system diagnosis
  - Hopefully our analysis will point us towards the vicinity of the culprit

# Objectives

- Describe a small corner of the universe of sysplex infrastructure problems and their symptoms
- Provide a methodology to reliably discover root causes
  - Or at least get closer to the real source of the problem
- As a side benefit, you should be better able to:
  - Prevent problems
  - Assess risks
  - Apply appropriate remedies to resolve problems
- And if you need help, you should be better able to:
  - Provide the data needed for diagnosis and resolution
  - Identify the trouble spot, and thereby start your service call with the right set of IBM component experts

# How Component Experts Diagnose Problems

Step 1: Investigate “my component”
Step 2: Look at everything in “my component’s” space
- Configuration
- Defect
- Workload
- Changes

Step 3: Share findings with customers
Step 4: Customer will try suggestions
Step 5: If sufficient relief attained then DONE. Otherwise, continue to Step 6.
Step 6: If it’s not “my component” consider what else it can be???
- Maybe the “correct” next component is selected

Step 7: Engage another component or transfer the PMR to the next component
Step 8: Go to Step 1

# Bridging the Gap

- You know your systems and your workload (I hope), but lack component expertise
- IBM has component expertise, but does not know your shop
- I hope this material helps to bridge the gap
  - Use knowledge of your shop
  - Identify relevant symptoms
  - Deduce potential causes
  - Engage the right component experts
- So that problem resolution is timely and effective

# What is normal? What changed?

- Understanding differences between past and current (problem) behavior is often helpful for diagnosis
  - Configuration
  - Workload
  - Utilization
  - Request rates
  - Maintenance, either software or hardware
  - Migration, either software or hardware

- Good change activity logs can be invaluable for helping to pin point what and when (why?) a potentially relevant change was made
  - Although sometimes the change that leads to a problem was long ago and far away

# Problem Taxonomy

- **Dead System**
- Sick System
- Sysplex Fabric
- Sysplex Componentry
  - Coupling Facility
  - Signalling Service
  - Couple Data Sets
  - External Time Reference
- Configuration / Capacity
- Software Issues

# Dead System

- A dead system can't participate in anything
  - Dead System = wait state and not removed from sysplex
  - The root cause of sympathy sickness often turns out to be a dead system that no one noticed !
- **So this should always be the first thing you check**
- Symptoms
  - IXC402D "system looks dead"
  - IXC102A "tell me when system is reset"
  - Other messages (see handout)

# Dead System Symptoms: XCF messages

- IXC101I “removing system from sysplex”
- IXC105I “removed system from sysplex”
- IXC102A “reset system and reply down”
- IXC108I “fencing system”
- IXC109I “fencing completed (or failed)”
- IXC256A “cannot remove CDS until these systems respond”
- IXC409D “lost signal connectivity, how to proceed?”
- IXC426D “sending signals but not updating status, how to proceed?”
- IXC427A “sending signals, not updating status.  SFM will act”
- IXC602I “SFM will take this action if system looks dead”
- IXC800I “ARM could not restart elements elsewhere”

# Dead system must be dealt with ASAP

- Failure to deal with a dead system in a timely manner can easily induce sympathy sickness that expands into a morass of extraneous and confusing symptoms because nearly every sysplex application is likely to be impacted sooner or later somewhere in the sysplex

- So you really want automatic removal
  - SFM policy with ISOLATETIME (see handout)
  - SFM with BCPii
- Or you have to rely on vigilant operators

# Fencing

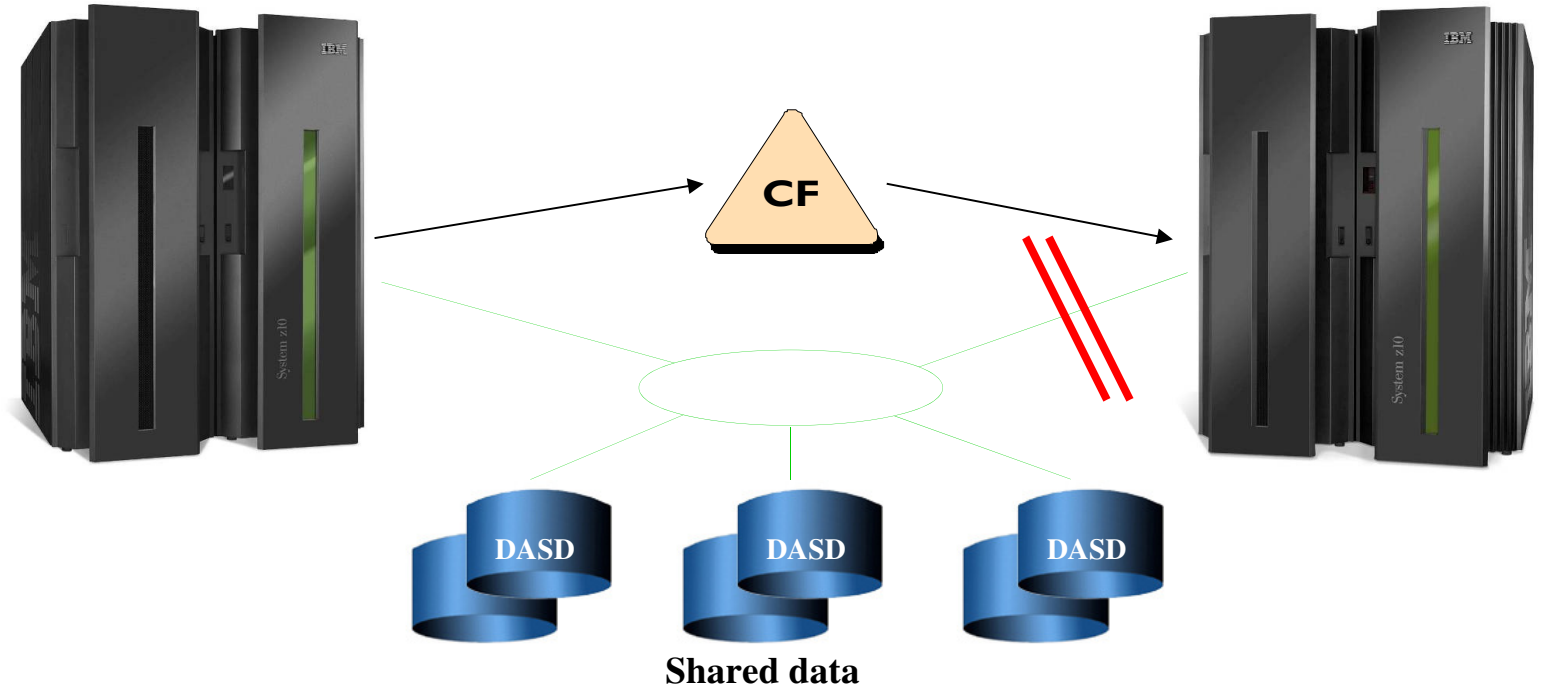


Fencing isolates a system so that it cannot access shared data, thus making it safe for the survivors to release serialization of the shared resources.

A command is sent via a CF to the target CEC. The target image will not be able to initiate any new I/O and ongoing I/O will be terminated.

# SFM with ISOLATETIME

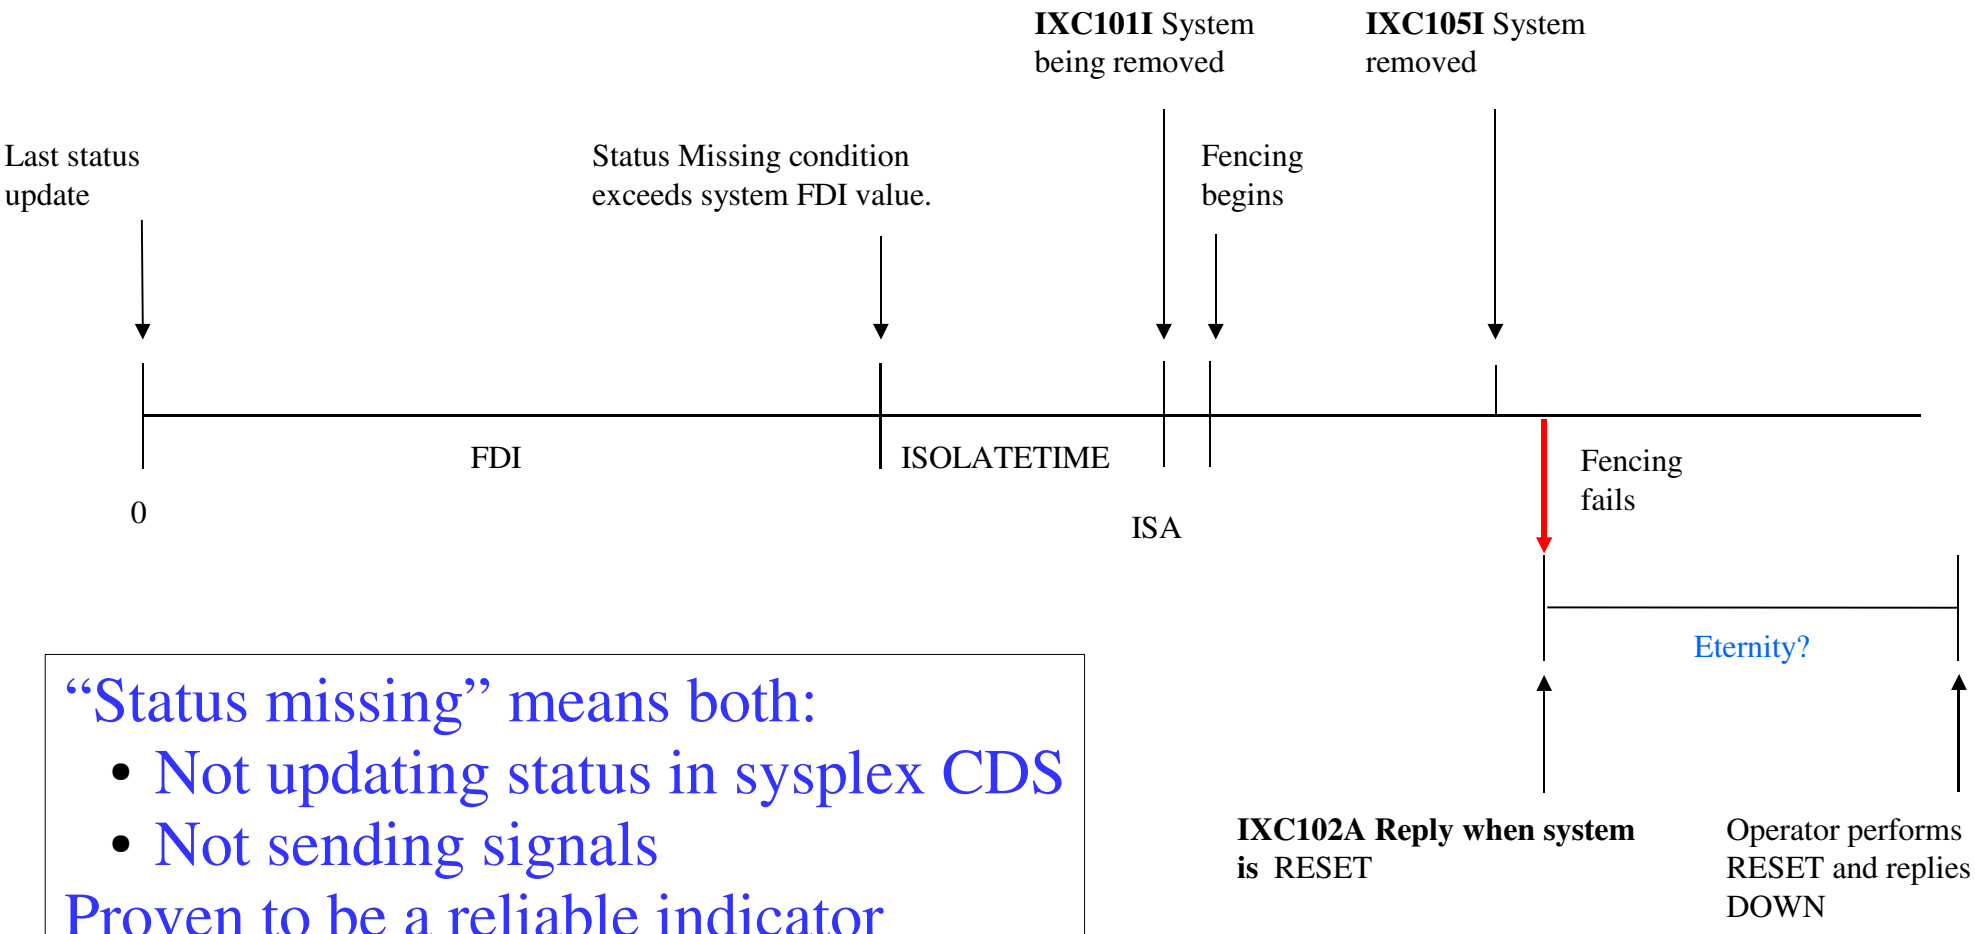


"Status missing" means both:

- Not updating status in sysplex CDS
- Not sending signals

Proven to be a reliable indicator

# SFM With BCPii

z/OS Images
(not VM guests)

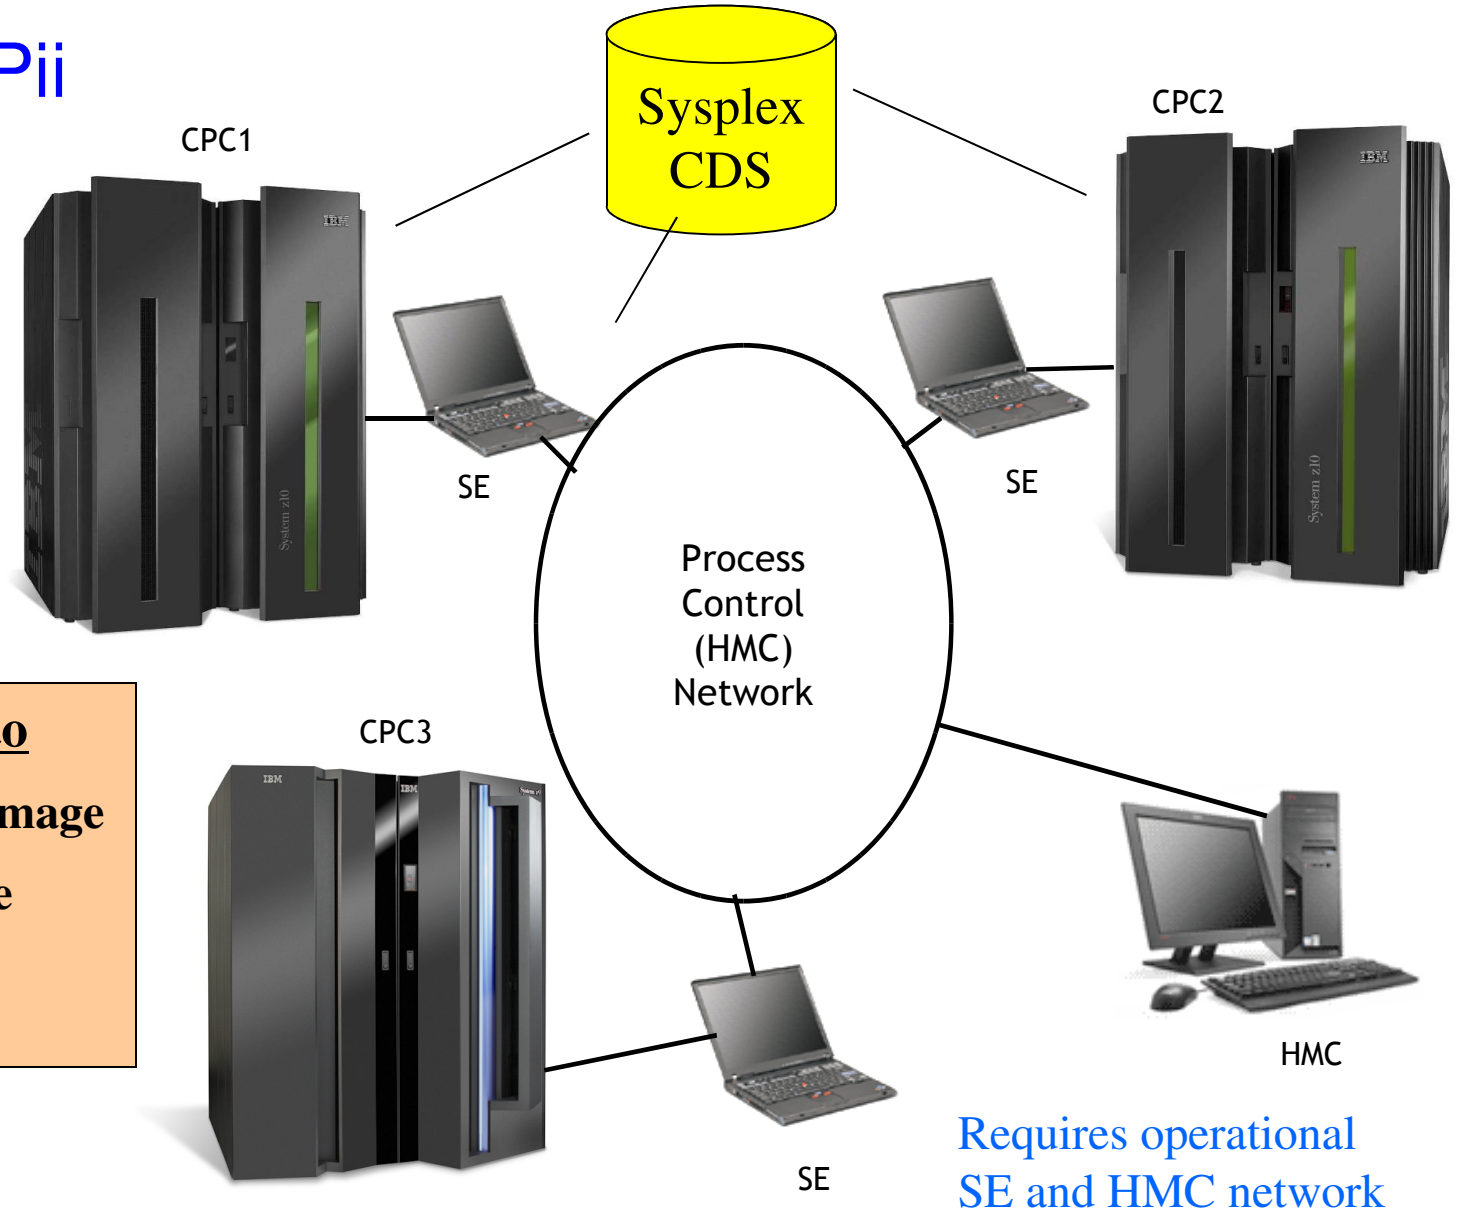


**XCF uses BCPii to**

- **Obtain identity of an image**
- **Query status of remote CPC and image**
- **Reset an image**

Requires operational SE and HMC network

# z/OS 1.11 SFM with BCPii

- Expedient removal of unresponsive or failed systems is essential to high availability in sysplex
- XCF exploits BCPii services to:
  - Detect failed systems
  - Reset systems
- Benefits:
  - Improved availability by reducing duration of sympathy sickness
  - Eliminate manual intervention in more cases
  - Potentially prevents human error that can cause data corruption

# SFM with BCPii

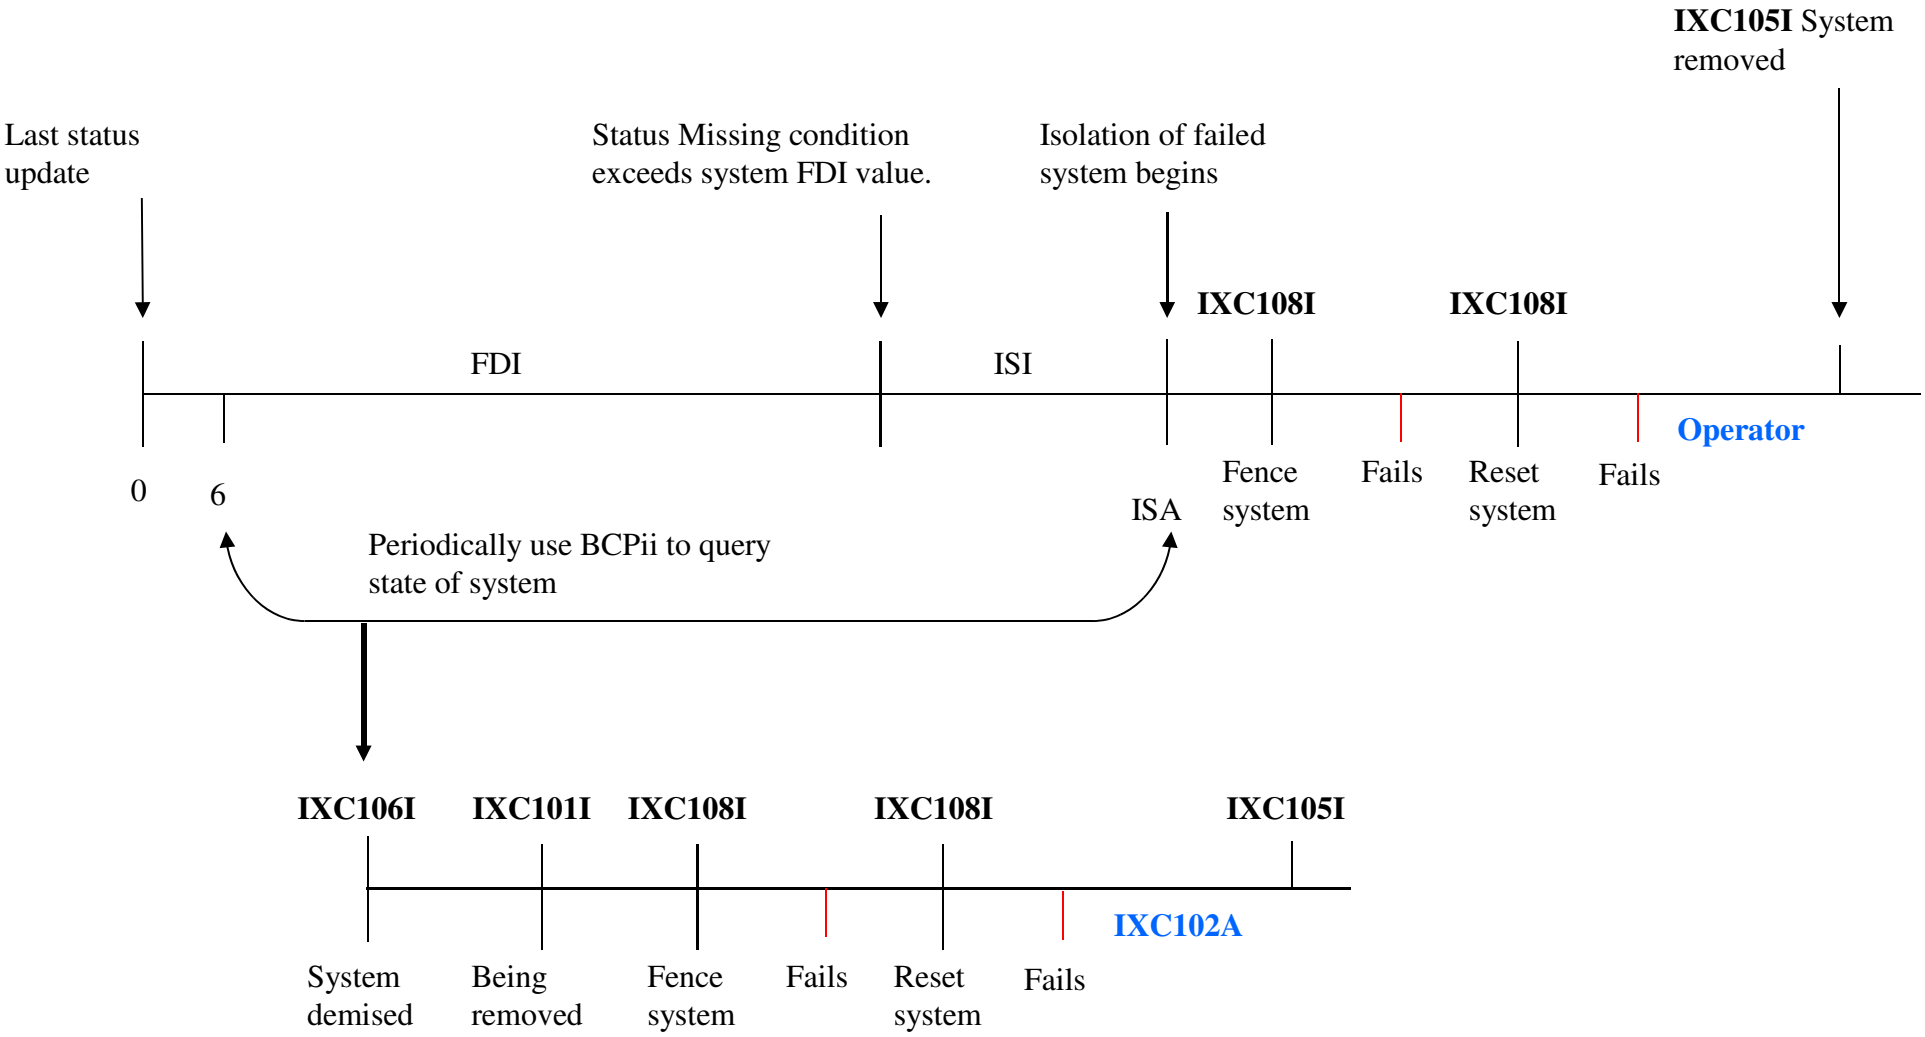


FDI = Failure Detection Interval

# z/OS 1.11 SFM with BCPii

- SFM will automatically exploit BCPii and as soon as the required configuration is established:
  - Pairs of systems running z/OS 1.11 or later
  - BCPii configured, installed, and available
  - XCF has security authorization to access BCPii defined FACILITY class resources
  - z10 GA2 with appropriate MCL's, or z196, or z114, or zEC12
  - **New version of sysplex CDS is primary in sysplex**
    - Toleration APAR OA26037 for z/OS 1.9 and 1.10
    - Does NOT allow systems to use new SSD function or protocols

If you have the appropriate environment, SET THIS UP !
You will likely eliminate this entire class of "dead system" issues

# Manual intervention may still be needed !

- Use of BCPii and SFM policies that permit the sysplex to automatically recover from dead systems will eliminate many of the cases where manual intervention is required
- XCF falls back on manual intervention if these fail
  - Need to be prepared for that possibility
    - Likely fewer opportunities for operators to maintain skills
  - So even if you exploit these technologies, checking for a dead system should still be first thing to consider

Suggestion: Page system programmer when IXC102A is issued

# Problem Taxonomy

- Dead System
- **Sick System**
- Sysplex Fabric
- Sysplex Componentry
  - Coupling Facility
  - Signalling Service
  - Couple Data Sets
  - External Time Reference
- Configuration / Capacity
- Software Issues

# Why Discuss Sick System ?

- Why discuss “single system diagnosis” in a “sysplex diagnosis” presentation?
- Many are fooled into thinking they have a sysplex problem when in fact they have a system problem
  - *So problem resolution is inefficient and takes longer than needed*
- A sick system can induce sympathy sickness
  - *So the root cause of a sysplex problem may well be one that is single system in scope*
  - *If the local components of the sysplex infrastructure are not running on a solid foundation, the sysplex will likely suffer*

# Sick System: Typical Root Causes

- Storage constraints
  - Not enough resource for system to run reasonably
  - Run away application
  - Defect
- CPU constraints
  - Not enough resource for system to run reasonably
  - SRBs looping
- Contention
  - ENQ, latches, local lock, spin locks
- DASD I/O issues
- Sick components

# Run Time Diagnostics (RTD)

- Reviews critical messages in the log
- Analyzes contention
  - GRS ENQ
  - GRS Latches
  - z/OS UNIX file system latches
- Examines address spaces with high CPU usage
- Looks for an address space that might be in a loop
- Evaluates local lock conditions
- Performs additional analysis based on what is found
  - For example, if XES reports a connector as unresponsive, RTD investigates the appropriate address space

*RTD can find many sick system issues. Give it a try.*

# Runtime Diagnostics

- Allows installation to quickly analyze a system experiencing “sick but not dead” symptoms
- Looks for evidence of “soft failures”
- Reduces the skill level needed when examining z/OS for “unknown” problems where the system seems “sick”
- Provides timely, comprehensive analysis at a critical time period with suggestions on how to proceed

- Runs as a started task in z/OS V1R12
  - S HZR
- Starts at IPL in z/OS V1R13
  - F HZR,ANALYZE command initiates report

- Message HZR0200I contains the report

# Sick System: Storage constraints

- System may not run well if storage constrained
- Shortages tend to induce paging
  - Resolving page faults introduces delay and system overhead
  - Swap out of address spaces implies programs are not running
- Typical Causes
  - You need more memory
  - Runaway application
  - Page pack performance (see “Sick System: DASD I/O issues”)
  - Defects

# Detecting Storage Constraints

- What changed?
- Who is consuming storage?
  - How much?
  - May need application knowledge to determine whether reasonable
  - But, is this typical for you?

# Detecting Storage Constraints

- Messages (see handout)
  - Auxiliary storage shortages
  - Pageable storage shortages
  - Frame shortages
- ABENDs
  - XCF 00C rsn xxxx006C implies frame shortage
- IPCS RSMDATA SUMMARY to see frame usage
- IPCS ANALYZE RESOURCE
  - Look for “ANY REAL FRAME”, presence implies storage constraint
- IPCS SYSTRACE ALL
  - Find “AVQ”. If low, implies RSM needs to replenish frames and requests for virtual storage may hang

# Sick system: Storage constraint messages

- IRA200E “auxiliary storage shortage”
- IRA201E “critical auxiliary storage shortage”
- IRA202I “auxiliary storage shortage relieved”
- IRA203I “consumer of auxiliary storage”
- IRA204E “consumer of auxiliary storage”
- IRA205I “consumed half of auxiliary storage”
- IRA206I “consumer of auxiliary storage”
- IRA210E “consumer set non-dispatchable”
- IRA211I “consumer set dispatchable”
- IRA220I “critical auxiliary storage shortage”
- IRA221D “show more or cancel consumer”
- IRA222I “consumer not canceled”
- IEE787A “pageadd command, how to proceed?”

# Sick system: Storage constraint messages

- IRA400E “pageable storage shortage”
- IRA401E “critical pageable storage shortage”
- IRA402I “pageable storage shortage relieved”
- IRA403E “swapped out address space”
- IRA404I “report storage consumed by address space”
- IRA405I “report high % fixed frames”
- IRA410E “set non-dispatchable”
- IRA411I “set dispatchable”
- IRA420I “consumers of fixed frames”
- IRA421D “show more or cancel consumer”
- IRA422I “consumer not canceled”

# Sick system: Storage constraint messages

- IRA100E “SQA shortage”
- IRA101E “critical SQA shortage”
- IRA102I “SQA shortage relieved”
- IRA103I “SQA expanded into CSA”
- IRA104I “SQA expansion into CSA relieved”
- IRA110E “high shared virtual storage shortage”
- IRA111E “critical high shared virtual storage shortage”
- IRA112I “high shared virtual storage shortage relieved”
- IRA120E “large frame shortage”
- IRA121E “critical large frame shortage”
- IRA122I “large frame shortage relieved”
- IRA130E “high common storage shortage”
- IRA131E “critical high common storage shortage”
- IRA132I “high common storage shortage relieved”

# Sick System: CPU Constraints

- "Something" will not be running
- Usually the "something" is deemed to be less important, so it may seem that you are getting the desired result
  - Test system
  - Discretionary work
- But failure to run the less important "something" can induce sympathy sickness for the important work if it:
  - Holds serialization for which there is contention
  - Holds resources, or worse, accumulates them while not running
    - For example, XCF Signal buffers
  - Not participating in cooperative processes

# CPU Constraint Considerations

- Latent demand?
  - If LPAR busy less than MVS busy, the physical processor is being ripped away despite the fact that MVS has work to do
- Blocked work?
  - If ready work is only getting “trickles”, the box is pretty much saturated
- Running high importance interactive/transaction oriented workloads at more than 90-92% busy is asking for trouble

- Did something change?
  - Configuration (upgrade, number CPs, capping, CF, ...)
  - Workload
  - CPU utilization by some job/space

# Detecting CPU Constraints

- RMF partition data reports
  - CPC Capacity
  - Distribution of IN-READY Queue
- RMF Monitor III
  - CPC Report
- Will likely need a timely dump to determine if there is a loop
  - Run Time Diagnostics might be able to detect it

# Post Processor: CPU Activity Report

C P U  A C T I V I T Y

z/OS V1R11 SYSTEM ID PP1A DATE 05/27/2012 INTERVAL 15.00.001

CONVERTED TO z/OS V1R12 RMF TIME 19.00.00 CYCLE 1.000 SECONDS

CPU 2097 CPC CAPACITY 3120 SEQUENCE CODE 0000000000007F24E

MODEL 750 CHANGE REASON=N/A HIPERDISPATCH=YES

H/W MODEL E56

| ---CPU--- | | ---------- TIME % ---------- | | | | LOG PROC | | --I/O INTERRUPTS-- | |
|---|---|---|---|---|---|---|---|---|---|
| NUM | TYPE | ONLINE | LPAR BUSY | MVS BUSY | PARKED | SHARE % | | RATE | % VIA TPI |
| 0 | CP | 100.00 | 40.99 | 42.37 | 0.00 | 100.0 | HIGH | 233.4 | 28.92 |
| 1 | CP | 100.00 | 64.68 | 64.35 | 0.00 | 100.0 | HIGH | 1113 | 27.85 |
| 2 | CP | 100.00 | 15.31 | 15.28 | 0.00 | 100.0 | HIGH | 56.55 | 27.45 |
| 3 | CP | 100.00 | 18.61 | 18.57 | 0.00 | 100.0 | HIGH | 74.50 | 28.44 |
| 4 | CP | 100.00 | 42.72 | 42.62 | 0.00 | 100.0 | HIGH | 268.6 | 28.69 |
| 5 | CP | 100.00 | 51.47 | 51.19 | 0.00 | 100.0 | HIGH | 884.9 | 27.19 |
| 6 | CP | 100.00 | 13.54 | 13.51 | 0.00 | 100.0 | HIGH | 52.71 | 27.76 |
| 7 | CP | 100.00 | 16.32 | 16.29 | 0.00 | 100.0 | HIGH | 66.75 | 28.83 |
| 8 | CP | 100.00 | 54.93 | 54.82 | 0.00 | 100.0 | HIGH | 309.3 | 28.63 |
| 9 | CP | 100.00 | 31.54 | 31.50 | 0.00 | 100.0 | HIGH | 114.5 | 28.00 |
| A | CP | 100.00 | 18.08 | 18.05 | 0.00 | 100.0 | HIGH | 61.62 | 28.28 |
| B | CP | 100.00 | 24.52 | 24.49 | 0.00 | 100.0 | HIGH | 87.97 | 28.65 |
| C | CP | 100.00 | 51.44 | 51.35 | 0.00 | 100.0 | HIGH | 247.8 | 28.48 |

etc.

# Post Processor: CPU Activity Report

```
-----------------------DISTRIBUTION OF IN-READY QUEUE-----------------------
NUMBER OF                0    10   20   30   40   50   60   70   80   90   100
ADDR SPACES     (%)      |....|....|....|....|....|....|....|....|....|....|

<=  N          100.0     >>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>>
 =  N +  1       0.0
 =  N +  2       0.0
 =  N +  3       0.0
<=  N +  5       0.0
<=  N + 10       0.0
<=  N + 15       0.0
<=  N + 20       0.0
<=  N + 30       0.0
<=  N + 40       0.0
<=  N + 60       0.0
<=  N + 80       0.0
>   N + 80       0.0
```

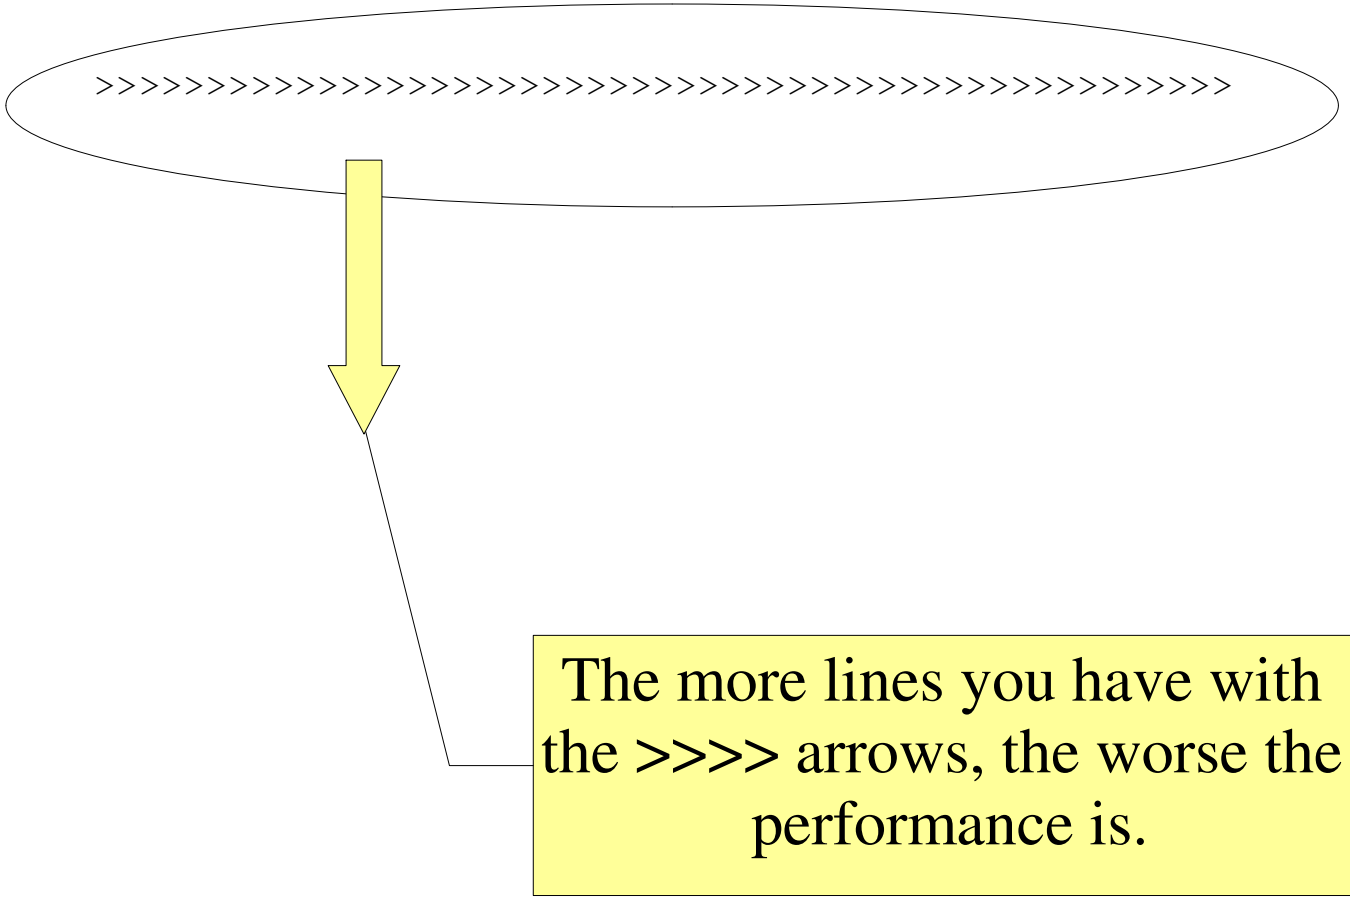


N = NUMBER OF PROCESSORS ONLINE UNPARKED (48.0 ON AVG)

# RMF Monitor III CPC Capacity

```
                          RMF V1R11  CPC Capacity                    Line 1 of 44

Samples: 100     System: SD0   Date: 08/07/12  Time: 10.36.40  Range: 100   Sec

Partition:   SD0         2817 Model 750
CPC Capacity:    4300    Weight % of Max: 10.0   4h Avg:    1   Group:   N/A
Image Capacity:  1376    WLM Capping %:    0.0   4h Max:   33   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                        88.0                        0.3     4.8    5.1
CSK            0     2  NO    8.0       0.3     0.3    0.0     0.0    0.1
SA0            0    39  NO    8.0       5.6     5.7    0.0     0.9    0.9
SB0            0    16  NO    8.0       2.2     2.3    0.0     0.4    0.4
SC0            0    43  NO    8.0       6.0     6.2    0.0     1.0    1.0
SD0            0     3  NO    8.0       0.4     0.4    0.0     0.1    0.1
SJ0            0     3  NO    8.0       0.4     0.4    0.0     0.1    0.1
etc.
```

# RMF Monitor III Processor Delay

```
                        RMF V1R11 Processor Delays                Line 1 of 138
Command ===>                                                     Scroll ===> HALF

Samples: 60      System: MVS1  Date: 11/28/09  Time: 09.10.00  Range: 60 Sec

             Service  CPU  DLY  USG EAppl  ---------- Holding Job(s) ----------
Jobname   CX Class    Type  %    %    %     %  Name      %  Name      %  Name

WSWS7     O  OMVS     CP    11   46  59.4   9 *ENCLAVE   7 DBS3DIST   7 WSP1S2F
WSP1S2FS  SO WASCR    CP     4    4  42.5   2 DBS3DIST   2 WSWS7      2 VTAM44
                      AAP    6    0  98.4   6 *ENCLAVE
WSP1S6FS  SO WASCR    CP     0    0   5.3
                      AAP    6    0   7.7   6 *ENCLAVE
DBS3DBM1  S  DB2HIGH  CP     2    6   0.8   2 XCFAS      2 DBS3DIST   2 WSP1S2F
WSP1S6F   SO WASCR    CP     0    2   1.9
                      AAP    2    2   0.7   2 *ENCLAVE
U078069   O  OMVS     CP     2    4   1.2   2 WSWS7      2 DBS3DIST   2 U078069
WSP1S4F   SO WASCR    CP     0    0   0.1
                      AAP    2    0   0.4   2 WSP1S6F
U078068   O  OMVS     CP     2    0   0.2   2 XCFAS      2 WSWS7      2 *ENCLAVE
DBS3DIST  SO DB2HIGH  CP     0   78 111.0
                      IIP    0    2  21.3
```

# Sick System: Contention

- Arises when the various ENQs, latches, or locks used for serialization are not immediately available to one or more work units because they are already held by another work unit
- The requesting work units then suffer delay
- I think there are two categories
  - Persistent – held too long (something is not working)
  - Competitive – lots of interest (workload)
- Generally need to stay focused on the holder, not the waiters

# Persistent Contention

- Work unit holds serialization for unreasonably long time
- Perhaps lots of work to do while serialized
  - Spikes?
  - Increase in volume?
- Perhaps work items are taking longer than normal
  - Is work unit running? enough?
  - Impacted by delays in some service it uses?
  - Hardware issues? (device recovery, retry, timeouts, ....)
  - Software issues? (error recovery, dumping, …)
- Perhaps a little of both
  - Workload changes create more work and this application comes up short on resources needed to run effectively

# Competitive Contention

- Work is running, but response/transaction times impacted
- In effect, the pieces don't play well together
- Tends to arise as the result of workload changes
- May be hard to detect and/or diagnose
  - The contention does not last long enough to see it, or
  - Constantly being released and acquired but always looks held when observed
    - Without data describing duration of the acquisition, this might seem to be persistent contention
  - Likely needs application specific understanding to figure out how the various pieces interact with each other
    - Are design changes needed?

# Sick System: Contention

- ENQ contention
  - Often arises as the result of submitting some new batch job that serializes same resource as existing work/job, or
  - Submitting multiple batch jobs that need to use the same data sets
- Latch contention
  - Often arises as the result of work load changes
  - USS, RRS, Logger are examples of users of latches
  - Latches are local but can induce sympathy sickness:
    - Get latch, send message, <delay?>, get response, release latch
    - <delay?> Get latch, formulate response, send response, release latch
  - Can give rise to a tangled web of sympathy sickness which can be quite difficult to diagnose (stay focused on the holder)

# Sick System: Contention ...

- Local Lock
  - Some applications use local lock for serialization
  - Many system services often need it as well
    - GETMAIN, ATTACH, …
  - Is holder running? enough?
  - Long queue effects?
- Spin Locks
  - If RSM lock, could be sign of real storage shortage
  - Hardware errors?
  - Long queue effects?
  - Software defect?

# Detecting contention

- Run Time Diagnostics
- D GRS,C "Most common command to display contention"
- D GRS,AN,BLOCKER "To find out the trouble makers"
- D GRS,AN,WAITER "To find the victims"
- D GRS,RES=(major_name,minor name) "To focus on a particular resource"
- RMF Serialization Delay Report
  - ENQ
  - Latches
  - Local lock (others as well)
- IPCS ANALYZE RESOURCE
- Messages
  - IEE601E "excessive spin loop"
  - IEE331A "excessive spin loop, how to proceed?"
  - IEE178I "spin loop recovery action"

# D GRS,C Sample Output

```
ISG343I 05.26.32 GRS STATUS 604
S=SYSTEM  SYSZTIOT
SYSNAME        JOBNAME          ASID     TCBADDR   EXC/SHR    STATUS
MVSS      JES2                  0031     00AC5D48    SHARE      OWN
MVSS      JES2                  0031     00AC5598    SHARE      OWN
MVSS      JES2                  0031     00ABD170 EXCLUSIVE    WAIT
MVSS      JES2                  0031     00ABBD58 EXCLUSIVE    WAIT
MVSS      JES2                  0031     00ABB058 EXCLUSIVE    WAIT
MVSS      JES2                  0031     00AC5AB8    SHARE      WAIT
MVSS      JES2                  0031     00ABD4C0    SHARE      WAIT
MVSS      JES2                  0031     00AC5828    SHARE      WAIT
MVSS      JES2                  0031     00ABB2B8    SHARE      WAIT
MVSS      JES2                  0031     00ABBAC8    SHARE      WAIT
NO REQUESTS PENDING FOR ISGLOCK STRUCTURE
NO LATCH CONTENTION EXISTS
```

# RMF Mon III ENQ Delays

```
                        RMF V1R11  ENQ Delays                          Line 1 of 2

Samples: 100     System: SD0   Date: 08/07/12  Time: 10.41.40  Range: 100   Sec

         DLY  -------------- Resource Waiting --------------  ---- Holding ----
Jobname   %    %  STAT Major/Minor Names (Scope)               %  Name/SYS STAT

DUMPSRV    1    1  EW  SYSZDAE         (SYSS)                   1 DUMPSRV   EO
                       DATASET                                   /SE0
```

# Identify owner of offending spin lock

- PSACLHS and PSACLHSE indicate which spin locks are held
  - Use IHAPSA mapping macro to decode the lock held strings
    - Comments identify z/OS component that uses the lock
  - Each CPU has its own PSA
- Need to get a dump and use IPCS
  - IP STATUS shows PSACLHS for each CPU
  - Use IP LIST PSAn to format PSA for CPU n
    - PSACLHS is at offset x2F8
    - PSACLHSE is at offset x4C4

  – Message IEE178I identifies offending CPU

If you identify relevant spin lock, you can better route the PMR.
If RSM lock, you might look for storage constraints.

# Contention considerations

- Contention is often induced by other problems
- Experience suggests:
  - Changing mix of batch jobs often induces ENQ contention
  - Spinning on RSM locks often implies real storage shortages
  - Workload changes often induce competitive contention
  - Lack of dispatch time can induce persistent contention
  - USS latch contention often arises due to issues on peer system
- So at this point in our methodology, if root cause is not obvious
  - *As it might be for batch jobs or real storage shortages*
- Take note of the contention, but continue diagnosis
  - *Who holds the resource? Why aren't they making progress?*

# Sick System: DASD I/O issues

- Performance, response times, throughput, even functionality may be impacted if applications encounter errors or delays while accessing data on DASD
- Experience suggests that potential for such delays exists with:
  - Synchronous mirroring
  - Slow DASD
  - Workload changes
    - Changes in request rates, device contention, ...
- Of particular concern for sysplex are impacts to
  - Couple Data Sets
  - Data sets used by Logger
  - Page packs

# Detecting DASD I/O Issues

- Messages (see handout)
  - IOS complaints regarding I/O errors and timeouts
  - XCF complaints about couple data sets
  - Logger complaints
- RMF DASD I/O reports. Response time issues?
  - Direct Access DASD Activity – post processor report
- RESERVES
  - IOS071I issued if delayed due to RESERVE
  - Run ENQ/DEQ Monitor to proactively identify any RESERVES (ISGAUDIT)
- IPCS ANALYZE RESOURCE
  - Indicate any outstanding paging I/O?

Best Practice: Eliminate all RESERVES. Run GRS STAR mode. Convert all RESERVES to global resources.

# Sick System: DASD I/O messages (IOS)

- IOS001E “some paths to device are inoperative”
- IOS002A “no paths to device”
- IOS050I “channel detected error”
- IOS051I “channel timeout”
- IOS052I “channel detected error, recovered and logged it”
- IOS071I “MIH timeout” (adjust MIH? HW error?)
- IOS079I “deleted queued request due to timeout”
- IOS075E “device has recurring MIH condition”
- IOS076E “MIH timeout” (various reasons, likely HW error)
- IOS077E “MIH timeout” (various reasons, likely HW error)
- IOS100I “boxed device”
- IOS101I “device boxed or forced offline”
- IOS102I “device boxed or forced offline”
- IOS107I “deferring boxing of device”
- IOS431I “someone holding reserve on device”

# Sick System: DASD I/O messages (XCF)

- IXC244E “cannot use this sysplex CDS”
- IXC246E “CDS experiencing I/O delays”
- IXC255I “cannot use this function CDS”
- IXC259I “I/O error on CDS”
- IXC267E “processing without alternate CDS”

# Sick System: DASD I/O messages (Logger)

- IXG114A “offload not making progress”
- IXG115A “fix offload problem.  Terminate task?”
- IXG271I “logger experiencing delays”
- IXG272E “logger task delayed, what to do?”
- IXG310I “offload not making progress”
- IXG311I “offload not making progress”
- IXG312E “offload delayed, what to do?”

# Post Processor: Direct Access Device Activity

D I R E C T  A C C E S S  D E V I C E  A C T I V I T Y



z/OS V1R11 SYSTEM ID NICF START 07/18/2012-18.46.00 INTERVAL 000.15.00

RPT VERSION V1R11 RMF END 07/18/2012-19.01.00 CYCLE 0.500 SECONDS

TOTAL SAMPLES = 1,800 IODF = 00 CR-DATE: 04/04/2012 CR-TIME: 13.06.35 ACT: ACTIVATE

| STORAGE GROUP | DEV NUM | DEVICE TYPE | NUMBER OF CYL | VOLUME SERIAL | PAV | LCU | DEVICE ACTIVITY RATE | AVG RESP TIME | AVG IOSQ TIME | AVG CMR DLY | AVG DB DLY | AVG PEND TIME | AVG DISC TIME | AVG CONN TIME | % DEV CONN | % DEV UTIL | % DEV RESV | AVG NUMBER ALLOC | % ANY ALLOC | % MT PEND |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| NF34 | 0B00 | 33909 | 32760 | MX0518 | 1.0H | 0007 | 0.519 | 2.60 | .000 | .016 | .000 | .136 | .862 | 1.60 | 0.08 | 0.13 | 0.0 | 18.3 | 100.0 | 0.0 |
| NF34 | 0B01 | 33909 | 32760 | MX0519 | 1.0H | 0007 | 0.051 | 1.70 | .000 | .011 | .000 | .134 | 1.29 | .273 | 0.00 | 0.01 | 0.0 | 27.6 | 100.0 | 0.0 |
| NF34 | 0B02 | 33909 | 32760 | MX0520 | 1.0H | 0007 | 0.044 | 1.27 | .000 | .003 | .000 | .122 | .906 | .246 | 0.00 | 0.01 | 0.0 | 17.3 | 100.0 | 0.0 |
| NF34 | 0B03 | 33909 | 32760 | MX0521 | 1.0H | 0007 | 0.031 | 4.00 | .000 | .023 | .000 | .146 | 3.61 | .251 | 0.00 | 0.01 | 0.0 | 7.3 | 100.0 | 0.0 |
| NF44 | 0B04 | 33909 | 32760 | NF4588 | 1.0H | 0007 | 0.001 | .256 | .000 | .000 | .000 | .128 | .000 | .128 | 0.00 | 0.00 | 0.0 | 1.0 | 100.0 | 0.0 |
| NF44 | 0B05 | 33909 | 32760 | NF4589 | 1.0H | 0007 | 0.001 | .256 | .000 | .000 | .000 | .128 | .000 | .128 | 0.00 | 0.00 | 0.0 | 0.0 | 100.0 | 0.0 |
| NF34 | 0B06 | 33909 | 32760 | MX0563 | 1.0H | 0007 | 1.734 | 10.4 | .000 | .019 | .000 | .147 | 3.83 | 6.44 | 1.12 | 1.78 | 0.0 | 11.8 | 100.0 | 0.0 |
| NF34 | 0B07 | 33909 | 32760 | MX0526 | 1.0H | 0007 | 1.262 | 10.7 | .000 | .021 | .000 | .149 | 3.77 | 6.77 | 0.85 | 1.33 | 0.0 | 10.3 | 100.0 | 0.0 |
| NFX7 | 0B08 | 33909 | 10017 | NFX825 | 1.0H | 0007 | 0.001 | .128 | .000 | .000 | .000 | .128 | .000 | .000 | 0.00 | 0.00 | 0.0 | 1.0 | 100.0 | 0.0 |
| NFX7 | 0B09 | 33909 | 10017 | NFX826 | 1.0H | 0007 | 0.001 | .128 | .000 | .000 | .000 | .128 | .000 | .000 | 0.00 | 0.00 | 0.0 | 2.0 | 100.0 | 0.0 |
| NFX7 | 0B0A | 33909 | 10017 | NFX827 | 1.0H | 0007 | 0.001 | .256 | .000 | .000 | .000 | .128 | .000 | .128 | 0.00 | 0.00 | 0.0 | 2.0 | 100.0 | 0.0 |
| NFX7 | 0B0B | 33909 | 10017 | NFX828 | 1.0H | 0007 | 0.014 | 2.43 | .000 | .020 | .000 | .148 | 2.08 | .207 | 0.00 | 0.00 | 0.0 | 6.0 | 100.0 | 0.0 |

# RMF Mon III Device Delays

```
                        RMF V1R11  Device Delays                        Line 1 of 1

Samples: 100     System: SD0   Date: 08/07/12  Time: 14.00.00  Range: 100   Sec

           Service  DLY USG CON  ------------ Main Delay Volume(s) ------------
Jobname  C Class     %   %   %    %  VOLSER   %  VOLSER   %  VOLSER   %  VOLSER

IXGLOGR  S SYSTEM     1  12  12    1 LGR10Q
```

# Sick System: Sick Components

- Experience suggests that the system will not be running well if the following components are having issues (no particular order)
  - Logger
  - RACF
  - JES3
  - JES2
  - RRS
  - Unix System Services
  - Consoles
  - GRS
  - SMF

*Many of these critical components exploit sysplex services. So bear in mind that they could be sick due to sysplex issues we have not yet covered.*

*My idealized methodology is trying focus on internal problems that impact the component independently of the sysplex infrastructure. In the real world we may not be able to achieve such isolation.*

- Certainly others, but these tend to be most prevalent with respect to having single system issues that impact the sysplex

# Detecting Sick Components

- Diagnostic Data
  - LOGREC entries
  - Dumps
- RMF Monitor III
  - Job Delay
  - Processor Delay
- Component messages (see handout)
- XCF Messages (see handout)
- XES Messages (see handout)

*List of key messages that Level 2 looks for to see if component is sick.*

*The XCF/XES messages point at components who may be sick. They generally do not imply that XCF/XES is sick.*

# Detecting Sick Components: Dumps

- IEA045I “SVC dump started”
- IEA145E “messages and commands lost, may need to reissue”
- IEA611I “finished a dump”
- IEA794I “captured SVC dump”
- IEA799I “unable to automatically allocate SVC dump data set”
- IEA911E “finished a dump”
- IEE711I “unable to take dump”

*If a component is creating dumps, it detected some sort of problem.*
*It suggests that the component is sick. Likely a good place to focus.*

# RMF Monitor III: Job Delay

```
                          RMF V1R11  Job Delays                        Line 1 of 1

Samples: 100      System: SD0   Date: 08/07/12  Time: 14.28.20  Range: 100   Sec

Job: XCFAS         Primary delay: Job is waiting to use the processor.

Probable causes: 1) Higher priority work is using the system.
                 2) Improperly tuned dispatching priorities.


------------------------- Jobs Holding the Processor --------------------------
Job:             XCFAS          Job:         CID1GA02          Job:         CID1GA03
Holding:            1%          Holding:            1%         Holding:            1%
PROC Using:         1%          PROC Using:         2%         PROC Using:         3%
DEV Using:          0%          DEV Using:          0%         DEV Using:          0%
---------------------------- Job Performance Summary ----------------------------
        Service       WFL -Using%- DLY IDL UKN ---- % Delayed for ---- Primary
CX ASID Class    P Cr  %   PRC DEV  %   %   %  PRC DEV STR SUB OPR ENQ Reason
S  0006 SYSTEM   1     50    1   0   1   0  99   1   0   0   0   0   0 XCFAS
```

# RMF Monitor III Delay Report

Samples: 100 System: SD0 Date: 08/07/12 Time: 13.05.00 Range: 100 Sec

| Name | CX | Service Class | Cr | WFL % | USG % | DLY % | IDL % | UKN % | % Delayed for PRC | DEV | STR | SUB | OPR | ENQ | Primary Reason |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| NFAGEN | T | TSO | | 0 | 0 | 17 | 15 | 2 | 0 | 0 | 0 | 17 | 0 | 0 | HSM |
| *MASTER* | S | SYSTEM | | 0 | 0 | 1 | 0 | 99 | 0 | 1 | 0 | 0 | 0 | 0 | HSMCDS |
| CID1GT01 | SO | CICSRGN | | 63 | 5 | 3 | 0 | 93 | 3 | 0 | 0 | 0 | 0 | 0 | CID1GA03 |
| IXGLOGR | S | SYSTEM | | 92 | 11 | 1 | 0 | 88 | 0 | 1 | 0 | 0 | 0 | 0 | LGR10M |
| SMSVSAM | S | SYSTEM | | 100 | 1 | 0 | 0 | 99 | 0 | 0 | 0 | 0 | 0 | 0 | |
| VTAM44ST | S | SYSSTC | | 100 | 1 | 0 | 0 | 99 | 0 | 0 | 0 | 0 | 0 | 0 | |
| DFHSM | S | SYSSTC | | 100 | 1 | 0 | 0 | 99 | 0 | 0 | 0 | 0 | 0 | 0 | |
| TCPCST | SO | SYSSTC | | 100 | 2 | 0 | | | | | | | | | |
| CID1GA01 | SO | CICSRGN | | 100 | 4 | 0 | | | | | | | | | |
| CID1GA02 | SO | CICSRGN | | 100 | 3 | 0 | | | | | | | | | |
| CID1GA03 | SO | CICSRGN | | 100 | 1 | 0 | | | | | | | | | |
| CID1GA04 | SO | CICSRGN | | 100 | 1 | 0 | 0 | 99 | 0 | 0 | 0 | 0 | 0 | 0 | |

User NFAGEN
delayed waiting for HSM
to recall datasets.

# Sick System Components: Logger

- IXG114A “offload not making progress”
- IXG115A “fix offload problem.  Terminate task?”
- IXG261E “logger CDS running out of space”
- IXG262A “logger CDS has run out of space”
- IXG271I “logger experiencing delays”
- IXG272E “logger task delayed, what to do?”
- IXG310I “offload not making progress”
- IXG311I “offload not making progress”
- IXG312E “offload delayed, what to do?”

# Sick System Components: RACF

- IRRX004A
- IRRX017I
- IRRC022I
- IRRC023I
- IRRC024I
- IRRC025I
- IRRC026I
- IRRC032I
- IRRC033I
- IRRI013I
- IRRI081I
- IRRN081I
- IRRO081I
- IRRQ081I

*Sorry, I did not get these interpreted yet*

# Sick System Components: RACF ....

- ICH501I “RACF not active”
- ICH505A “RACF initialization failed”
- ICH600A “How should RACF proceed?”
  - ICH586A “RACF DB will be corrupted if mixed usage”
  - ICH588A “RACF DB will be corrupted if mixed usage”
  - ICH589A “RACF DB will be corrupted if mixed usage”
  - ICH590A “RACF DB will be corrupted if being shared”
  - ICH591A “RACF DB will be corrupted if being shared”

# Sick System Components: JES3

- IAT1105 “file directories in use exceeds threshold”
- IAT7134 “console buffers in use exceeds threshold”
- IAT1016 “spool partition full, sysout processing suspended”
- IAT1017 “spool partition full, job selection suspended”
- IAT1018 “FYI, spool partition full”
- IAT8054 “queue being held”
- IAT6368 “checkpoint data set too small”
- IAT2008 “start cmd fails, likely due to storage or paging constraints”
- IAT6341 “running out of job numbers”
- IAT6395 “jobs are waiting”
- IAT6850 “too many WTO’s”

# Sick System Components: JES2

- $HASP9207 “checkpoint lock held for a long time”
- $HASP292 “waiting for response during checkpoint write”

# Sick System Components: RRS

- ATR247E “severe RRS error”
- ATR248E “RRS waiting for Logger to recover”
- ATR249E “RRS waiting for Logger to recover given logstream”
- ATR202D “gap in logstream, how should RRS proceed?”
- ATR210E “gap in RM DATA logstream, may need to cold start RRS members”
- ATR225D “cancel delayed, how should RRS proceed?”
- ATR226D “memterm delayed, how should RRS proceed?”
- ATR227D “cancel delayed, how should RRS proceed?”
- ATR228D “memterm delayed, how should RRS proceed?”
- ATR229D “cancel delayed, how should RRS proceed?”
- ATR230D “memterm delayed, how should RRS proceed?”
- ATR231D “cancel delayed, how should RRS proceed?”
- ATR232D “memterm delayed, how should RRS proceed?”
- ATR233D “cancel delayed, how should RRS proceed?”
- ATR234D “memterm delayed, how should RRS proceed?”

# Sick System Components: Unix System Services

- BPXB001E
- BPXF002I
- BPXF006I
- BPXF008I
- BPXF014D
- BPXF019I
- BPXF020I
- BPXF021I BPXF029E BPXF032D BPXF034I BPXF039I BPXF044I BPXF045A BPXF075I BPXF076I BPXF077S BPXF078W BPXF079S BPXF080I BPXF083I BPXF213E
- BPXF214E “unable to access BPXMCDS couple data set”
- BPXF215E “unable to access BPXMCDS couple data set”
- BPXF216E “file system partition cleanup delayed”
- BPXF217E “file system partition cleanup failed”
- BPXF218I BPXF221I BPXF222E BPXF226E BPXF230I BPXF242I BPXF243E BPXF244E BPXF245I BPXF247I BPXF249I BPXF252I BPXF253E BPXF254I BPXF255I BPXF256I BPXF257I BPXF259I BPXI004I BPXI005I BPXI016I BPXI017I BPXI018I BPXI019E BPXI026I BPXI031E BPXI027I BPXI028E BPXI029I BPXI030I BPXI032E BPXI033E BPXI035E BPXI036E AVAILABL BPXI039I BPXI040I BPXI043E BPXI055I BPXI056E BPXI058I BPXI068I BPXI060I BPXI061E BPXI062I BPXI064E BPXI065E BPXI066E BPXI067E BPXI068I BPXI075E BPXI076E BPXI077I BPXI078D BPXI082E BPXI083D BPXI084E BPXI085D BPXM048I BPXM032E BPXM050E BPXM055D BPXM057E BPXM120D BPXM056E BPXN002I BPXN003E BPXO043I BPXP003E BPXP004E BPXP006E BPXP022E BPXP022I BPXP001I BPXP007E BPXP008E

*Sorry, I did not get these interpreted yet*

# Sick System Components: Consoles

- CNZ2202E
- CNZ3003I
- CNZ3004E
- CNZ3005A
- CNZ3007I
- CNZ3009E
- CNZ3010I
- CNZ3014I
- CNZ4200I
- CNZ4201E
- CNZZ002E
- CNZZ007E
- CNZZ009E
- CNZZ014E
- CNZZ031E
- CNZZ033E

*Sorry, I did not get these interpreted yet*

# Sick System Components: Consoles

- IEA145E “messages and commands lost, may need to reissue”
- IEA230E “WTOR buffer shortage”
- IEA231A “WTOR buffer shortage”
- IEA359E “retained action message buffer shortage”
- IEA360A “retained action message buffer shortage”
- IEA367A “multiple console support inoperative, please acknowledge”
- IEA404A “WTO buffer shortage”
- IEA405E “WTO buffer shortage”
- IEA555E “unable to restart delayed WTOR processor”
- IEA556I “peer system’s console support inoperative”
- IEA557A “need operator to respond to reply ID 0 WTOR”
- IEA652A “discarding messages due to WTO buffer shortage”
- IEA654A “sympathy sickness due to WTO buffer shortage on peer system”

# Sick System Components: Consoles

- IEE141A “no master console, no alternates either”
- IEE624I “need to enable system console”
- IEE765E “syslog task failed”
- IEE767A “syslog buffer is full”
- IEE769E “system error in syslog”
- IEE775E “no storage for syslog buffer”
- IEE795I “syslog data loss”
- IEE806A “exceeded command limit”
- IEE824E “communication task failed”

# Sick System Components: GRS

- ISG361A “waiting for list lock”
- ISG362I “finally got the list lock”

- With “critical member” support in z/OS V1R12, the system may well be removed from the sysplex before it gets a chance to complain about the list lock

- White paper on Diagnosing GRS issues
  - *z/OS GRS: Performance Considerations*
  - www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101725

# Sick System Components: SMF

- IEE986E “SMF buffer space exceeds threshold”

# Detecting Sick Components: XCF messages

- IXC430E “system has stalled group members”
- IXC431I “group member is stalled”
- IXC440E “stalled group member impacting other systems”
- IXC615I “terminating group member”
- IXC631I “stalled group member impacting peer system”
- IXC640E “peer stalled group member impacting me”

*Generally these messages suggest that someone other than XCF is sick. XCF however, can be impacted by these problems, which in turn can induce sympathy sickness. Stay focused on the stalled group member and the system where the member resides.*

# Detecting Sick Components: XES messages

- IXL040E “structure connector not responding”
- IXL041E “structure connector not responding”
- IXL045E “having trouble getting connector SRBs to run”

*Generally these messages suggest that someone other than XES is sick. XES itself is unlikely to be impacted by these problems, but they can in turn induce sympathy sickness. Generally stay focused on the stalled connector and the system where it resides. However, some connectors will appear to be unresponsive because they are “waiting” for a peer connector to complete some task. In such cases, one really needs application specific knowledge to determine whether the unresponsive connector is the culprit, or a victim. You want to stay focused on the culprit.*

# Correlating messages and components

| Component | Prefix | Jobname | XCF Group | CDS | CF Structure |
|---|---|---|---|---|---|
| XCF | IXC | XCFAS | SYSXCF | Sysplex<br>ARM, SFM | IXC… |
| XES | IXL | (user's job) | IXCLO… | CFRM | |
| GRS | ISG | GRS | SYSGRS<br>SYSGRS2 | Sysplex | ISGLOCK<br>(if star mode) |
| RACF | ICH | | IRRXCF00 | | IRR |
| Consoles | CNZ | CONSOLES | SYSMCS<br>SYSMCS2 | | (via logger) |
| Logger | IXG | IXGLOGR | | LOGR | lots |
| JES | HASP or IAT | JES2<br>JESAUX<br>JESXCF | | | checkpoint |
| RRS | ATR | RRS (maybe) | SYSATR | | (via logger) |
| Unix System Services | BPX | | SYSBPX | BPXMCDS | |

# Methodology Consideration

- At this point in our sysplex diagnosis methodology, we may have determined that a particular component appears to be sick
  - *But the component might have dependencies or interactions with other systems in the sysplex*
  - *Without component knowledge, we may not be able to tell whether the sick component is actually suffering from sympathy sickness*
- So if root cause is not obviously a local issue, take note of the component and continue diagnosis

# Problem Taxonomy

- Dead System
- Sick System
- **Sysplex Fabric**
- Sysplex Componentry
  - Coupling Facility
  - Signalling Service
  - Couple Data Sets
  - External Time Reference
- Configuration / Capacity
- Software Issues

# Sysplex Fabric

- Consists of the various cables, links, channel paths, CHPIDs, and adapter cards that provide physical access to the sysplex componentry
- Along with the parmlib members and policies that govern logical access to the sysplex componentry
- Performance, response time, throughput, even functionality may be impacted if access to the sysplex componentry is impeded by
  - Error prone connections
  - Loss of connectivity
  - Outright loss or lack of access
- So now we make sure that the expected sysplex componentry exists and is accessible to each system in the sysplex

# Sysplex Fabric: Couple Data Sets

- Physical access
  - See: “Sick System: DASD I/O issues”
  - IODF
- Logical access
  - COUPLExx COUPLE statement for sysplex CDS
  - COUPLExx DATA statements for function CDS's
  - SETXCF COUPLE command
- Detecting Fabric Problems
  - D XCF,COUPLE,TYPE=ALL (see handout)
  - XCF Messages (see handout)
  - IOS messages (see handout)
  - Exploiter Messages (see handout)

Do the various CDS exist?
Have an alternate?
Accessible to all systems?

Using right COUPLExx?
What changed?

# Sysplex Fabric: Couple Data Sets
## D XCF,COUPLE

```
SYSPLEX COUPLE DATA SETS
PRIMARY    DSN: UTCXCF.SVPLEX1.COUPLE.PRI
           VOLSER: X1CPLP     DEVN: 1C6E
           FORMAT TOD         MAXSYSTEM MAXGROUP(PEAK) MAXMEMBER(PEAK)
           11/27/2011 13:21:2                                503   (56)
ALTERNATE  DSN: UTCXCF.SVPLEX
           VOLSER: X1CPLA
           FORMAT TOD                                        BER
           11/27/2011 13:21:26       32     120              503
ARM COUPLE DATA SETS
PRIMARY    DSN: UTCXCF.SVPLEX1.ARMR14.PRI
           VOLSER: X1CPLA     DEVN: B010
           FORMAT TOD         MAXSYSTEM
           04/11/2003 21:02:00       32
ALTERNATE  DSN: UTCXCF.SVPLEX1.ARMR14.ALT
           VOLSER: X1CPLP     DEVN: 1C6E
           FORMAT TOD         MAXSYSTEM
           04/11/2003 21:01:03       32
```

You run risk of a syplex outage if you don't have an alternate Sysplex CDS

# Sysplex Fabric: Couple Data Sets
## D XCF,COUPLE,TYPE=xxx ...

```
ARM COUPLE DATA SETS
PRIMARY    DSN: UTCXCF.SVPLEX1.ARMR14.PRI
           VOLSER: X1CPLA     DEVN: B010
           FORMAT TOD         MAXSYSTEM
           04/11/2003 21:02:00       32
           ADDITIONAL INFORMATION:
            FORMAT DATA
             VERSION 1, HBB7707 SYMBOL TABLE SUPPORT
             POLICY(8) MAXELEM(400) TOTELEM(600)
ALTERNATE  DSN: UTCXCF
           VOLSER: X1C
           FORMAT TOD
           04/11/2003
           ADDITIONAL
            FORMAT DA
             VERSION 1, HBB7707 SYMBOL TABLE SUPPORT
             POLICY(8) MAXELEM(400) TOTELEM(600)
ARM IN USE BY ALL SYSTEMS
```

Should have an alternate for availability

Should normally be accessible by all systems

# Sysplex Fabric: Couple Data Sets

- IXC244E “cannot use this sysplex CDS”
- IXC246E “CDS experiencing I/O delays”
- IXC253I “CDS removed from service”
- IXC255I “cannot use this function CDS”
- IXC256A “cannot remove CDS until these systems respond”
- IXC259I “I/O error on CDS”
- IXC267E “processing without alternate CDS”

# Sysplex Fabric: Couple Data Sets (IOS messages if relevant to CDS)

- See “Sick System: DASD I/O Messages (IOS)” for list of IOS messages that might be relevant to DASD I/O Issues
- You need to know where the CDS resides so that you can consider only those IOS messages related to the channel paths and devices that are relevant to the various CDS of interest

# Sysplex Fabric: Function Couple Data Sets

- ARM
  - IXC807I “rejected use of alternate CDS”
  - IXC808I “no access to CDS”
  - IXC809I “lost access to CDS”
  - IXC810I “unable to use CDS”

# Sysplex Fabric: Function Couple Data Sets

- BPXMCDS
  - BPXF214E “unable to access CDS”
  - BPXF215E “unable to access CDS”
  - BPXF226E “rejected use of CDS”
  - BPXF230I “rejected use of alternate CDS”
  - BPXI046I “unable to initialize new primary CDS”
  - BPXF050I “I/O errors accessing CDS”
  - BPXF058I “lost access to CDS”

# Sysplex Fabric: Function Couple Data Sets

- CFRM
  - IXC220W “lost access to CDS”
  - IXC520I “not using CDS”
- LOGR
  - IXG047I “unable to access CDS”
  - IXG054A “no CDS available”
- SFM
  - IXC610I “unable to use CDS”
- WLM
  - IWM047E “unable to access CDS”
  - IWM048E “no CDS”

# Sysplex Fabric: Coupling Facility

- Physical Access
  - CF
  - Links

  Does the CF exist?
  Accessible from all systems?
  Using the right CFRM policy?
- Logical Access
  - CFRM Policy
- Detecting Fabric Issues
  - D CF — physically connected?
  - D XCF,CF — logically connected?
  - D XCF,POLICY,TYPE=CFRM — right policy?
  - XCF Messages (see handout)
  - IXLERLOG LOGREC entries for link issues (see handout)
  - Check for flashing icons on the HMC — hardware issues?
  - HMC command to display CF — CF exist? Does it respond?

# Sysplex Fabric: Coupling Facility

- CF becomes inaccessible to z/OS if:
  - CF image is reset
  - CF aborts
  - CF suffers power outage
  - CEC on which CF resides goes down
  - Loses connectivity to the CF
- z/OS issues messages IXL157I and IXC517I if unable to access the CF. But two possible scenarios:
  - Could be due to loss of connectivity
  - Could be due to CF going down
- They both look the same to z/OS, but root cause is different

# Sysplex Fabric: Coupling Facility

- If CF does not respond to z/OS within 2 seconds, z/OS recycles the link under the assumption that there has been some sort of communication issue
- If z/OS recycles all the links at the same time, connectivity to the CF is lost
  - The CF may be up, but z/OS “disconnected” so to speak

# D CF,CFNAME=xxxx .....

If CF is not accessible, D CF may not show the CF at all. Alternatively, one or more sections of the output will not be available.

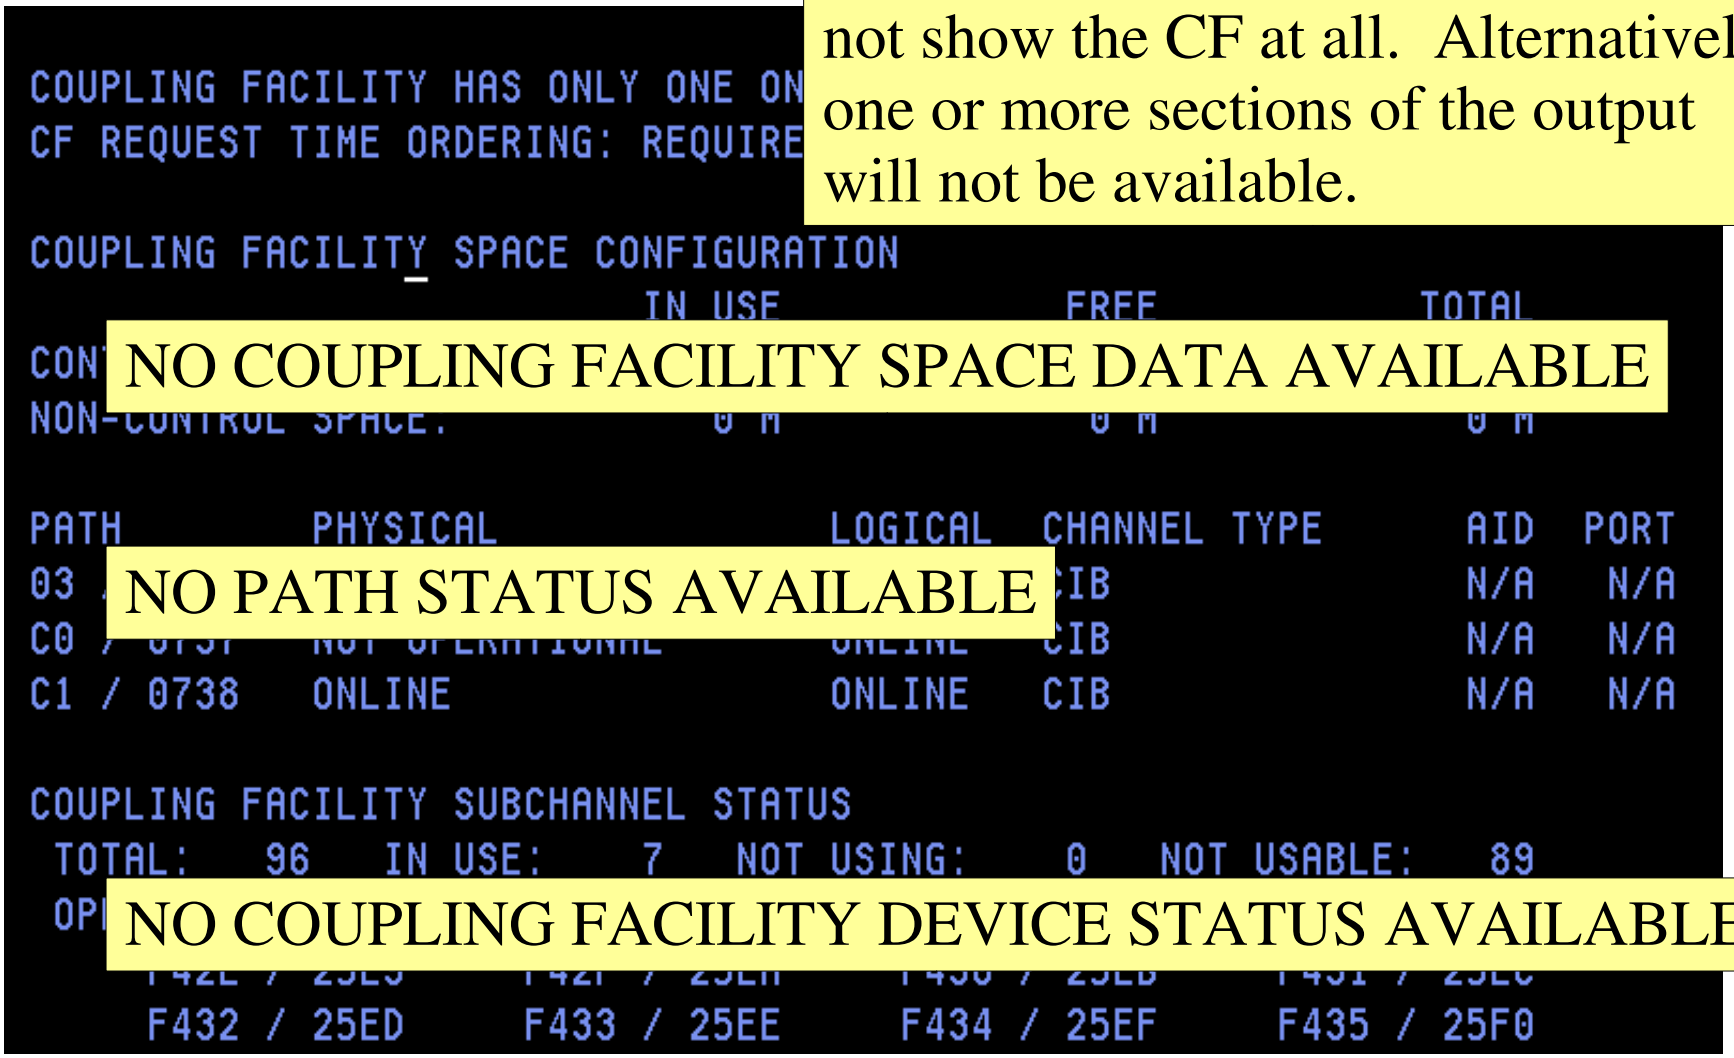

# D CF,CFNAME=xxxx .....

```
COUPLING FACILITY HAS ONLY ONE ONLINE SENDER PATH
CF REQUEST TIME ORDERING: REQUIRED AND NOT-ENABLED

COUPLING FACILITY SPACE CONFIGURATION
                          IN USE          FREE          TOTAL
CONTROL SPACE:             600
NON-CONTROL SPACE:           0

PATH         PHYSICAL
03 / 0727    NOT OPERATIONAL
C0 / 0737    NOT OPERATIONAL
C1 / 0738    ONLINE

COUPLING FACILITY SUBCHANNEL ST
 TOTAL:   96   IN USE:    7   N
 OPERATIONAL DEVICES / SUBCHANN
     F42E / 25E9     F42F / 25EA     F430 / 25EB     F431 / 25EC
     F432 / 25ED     F433 / 25EE     F434 / 25EF     F435 / 25F0
```

CF will not be accessible if "request time ordering" is required, but not enabled. Unless you are really down level on hardware, I would expect the function is installed. Thus "not enabled" likely implies that z/OS image and CF are either not in the same CTN, or are having other ETR related issues.

# D XCF,CF ....

```
D XCF,CF,CFNAME=SVT1CF1
IXC362I  22.06.42  DISPLAY XCF 411
CFNAME: SVT1CF1
  COUPLING FACILITY     :   002817.IBM.02.0000000F8E66
                            PARTITION: 09   CPCID: 00
  SITE                  :   SITE1
  POLICY DUMP SPACE SIZE:   6000 K
  ACTUAL DUMP SPACE SIZE:   6 M
  STORAGE INCREMENT SIZE:   1 M

  CONNECTED SYSTEMS:
  N64      N65      N66      N67

  MONITORING SYSTEM: N64

  STRUCTURES:
  DBSVPLX1_SCA           IRRXCF00_B001          IRRXCF00_B002
  ISTGENERIC(NEW)        IXCPLEX_PATH2          IXCPLEX_PATH4
  LOGGER_STR1            SYSZWLM_8E662817(NEW)
```

Shows CFs defined in the active CFRM policy. If CF not shown, you have a policy issue.

Right policy?
CF defined in the policy?
With the right Node ID?

# D XCF,POLICY

```
D XCF,POLICY,TYPE=ALL
IXC364I  22.40.25  DISPLAY XCF 927
TYPE: ARM
POLICY NOT STARTED
TYPE: BPXMCDS
      NOT SUPPORTED BY
TYPE: CFRM
      POLNAME:         IXCPOLX1
      STARTED:         07/19/2012 21:19:19
      LAST UPDATED:    07/19/2012 11:24:45
TYPE: LOGR
      NOT SUPPORTED BY DISPLAY XCF,POLICY
TYPE: SFM
      POLNAME:         IXCSFMX1
      STARTED:         04/17/2012 19:24:49
      LAST UPDATED:    04/17/2012 19:24:08

SYSPLEX FAILURE MANAGEMENT IS ACTIVE
TYPE: WLM
      NOT SUPPORTED BY DISPLAY XCF,POLICY
```

CFRM Policy started?
Right name?
Right instance?

# Sysplex Fabric: Coupling Facility

- IXC501A “use this CF or not?”
- IXC517I “system using this CF” was the problem before this?
- IXC518I “system not using this CF” but should it be?
- IXC519E “coupling facility is damaged”
- IXC512I “CFRM policy change pending”
- IXL051I “CF dump was taken for hardware support to review”

CF went down if dump was disruptive

- IXL044I “experiencing repeated IFCCs on path to CF”
- IXL157I “path to CF now operational”
- IXL158I “path to CF is not operational”

- IXL159E “detected error with notification vectors”

Is CF connected to same ETR as z/OS images?

- IXL160E “not using CF due to lack of request time ordering”
- IXL162E “not using CF due to lack of request time ordering”

# IXC501A “Use This CF or Not?”

- Operations must be very, very careful with IXC501A
- System programmer must be very careful with CFRM policy
- z/OS is asking operator to confirm whether this sysplex is supposed to be using the indicated CF
- The message is issued when it appears that the CF might be in use by some other sysplex
- If the operator responds “yes”, this sysplex will take over ownership of the CF
- If the CF is actually in use by some other live sysplex, that other sysplex will lose access to the CF
  - That other sysplex could potentially suffer a sysplex wide outage as a result of losing its CF

# Sysplex Fabric: Coupling Facility Symptom Record in LOGREC

```
TYPE:  SYMPTOM RECORD       REPORT:  SOFTWARE EDIT REPORT          DAY YEAR
                                                    REPORT DATE:  216  12
SCP:   VS 2 REL 3                                   ERROR DATE:   171  12
                            MODEL:   2097                         HH MM SS.TH
                            SERIAL:  03F25E                 TIME: 00:49:38.06
SEARCH ARGUMENT ABSTRACT:
    PIDS/5752SCIXL RIDS/IXLERLOG RIDS/IEANUC01#L FLDS/RESPCODE
    VALU/H00 FLDS/STATUSCODE VALU/H000000 FLDS/IRBSCC VALU/H00

SYSTEM ENVIRONMENT:
    CPU MODEL:  2097
    CPU SERIAL: 03F25E
    SYSTEM:     PP1C
    RELEASE LEVEL OF SERVICE RC
    SYSTEM DATA AT ARCHITECTURE
    COMPONENT DATA AT ARCHITECT
    SYSTEM DATA:  00000000 0000
COMPONENT INFORMATION:
    COMPONENT ID:
    COMPONENT RELEASE LEVEL:
    DESCRIPTION OF FUNCTION:
```

Look for IXLERLOG
Scroll down to find issue
IFCC if CF link issues
Other issues possible.

```
THE SYMPTOM RECORD DOES NOT CONTAIN A SECONDARY SYMPTOM STRING.
FREE FORMAT COMPONENT INFORMATION:
    KEY = F000     LENGTH = 000184 (00B8)
    +000    C396A497    93899587    40868183    899389A3    COUPLING FACILIT
    +010    A840D5C4    70401000    0503F0F0    F2F0F9F7    Y ND: ....002097
    +020    C           C           C           F0F0F0F1    E26IBM8300000001
    +030    C           C           C           89958740    E6CD.. COUPLING
    +040    8           8           8           857A4040    FACILITY NAME:
    +050    E           E           E           40404040    MABCF1
    +060    D           D           D           84A49385    REPORTING MODULE
    +070    7           7           7           40404040    : IXLE1REC
    +080    C           C           C           A289A27A    ERROR DIAGNOSIS:
    +090    C9D5E3C5    D9C6C1C3    C540C3D6    D5E3D9D6    INTERFACE CONTRO
    +0A0    D340C3C8    C5C3D240    C9D5C4C9    C3C1E3C5    L CHECK INDICATE
    +0B0    C4404040    40404040                            D
```

Physical link errors may arise due to:

- Bad or pinched cables
- Cables not well seated
- Dirt

# HMC: Operations System Messages For CF

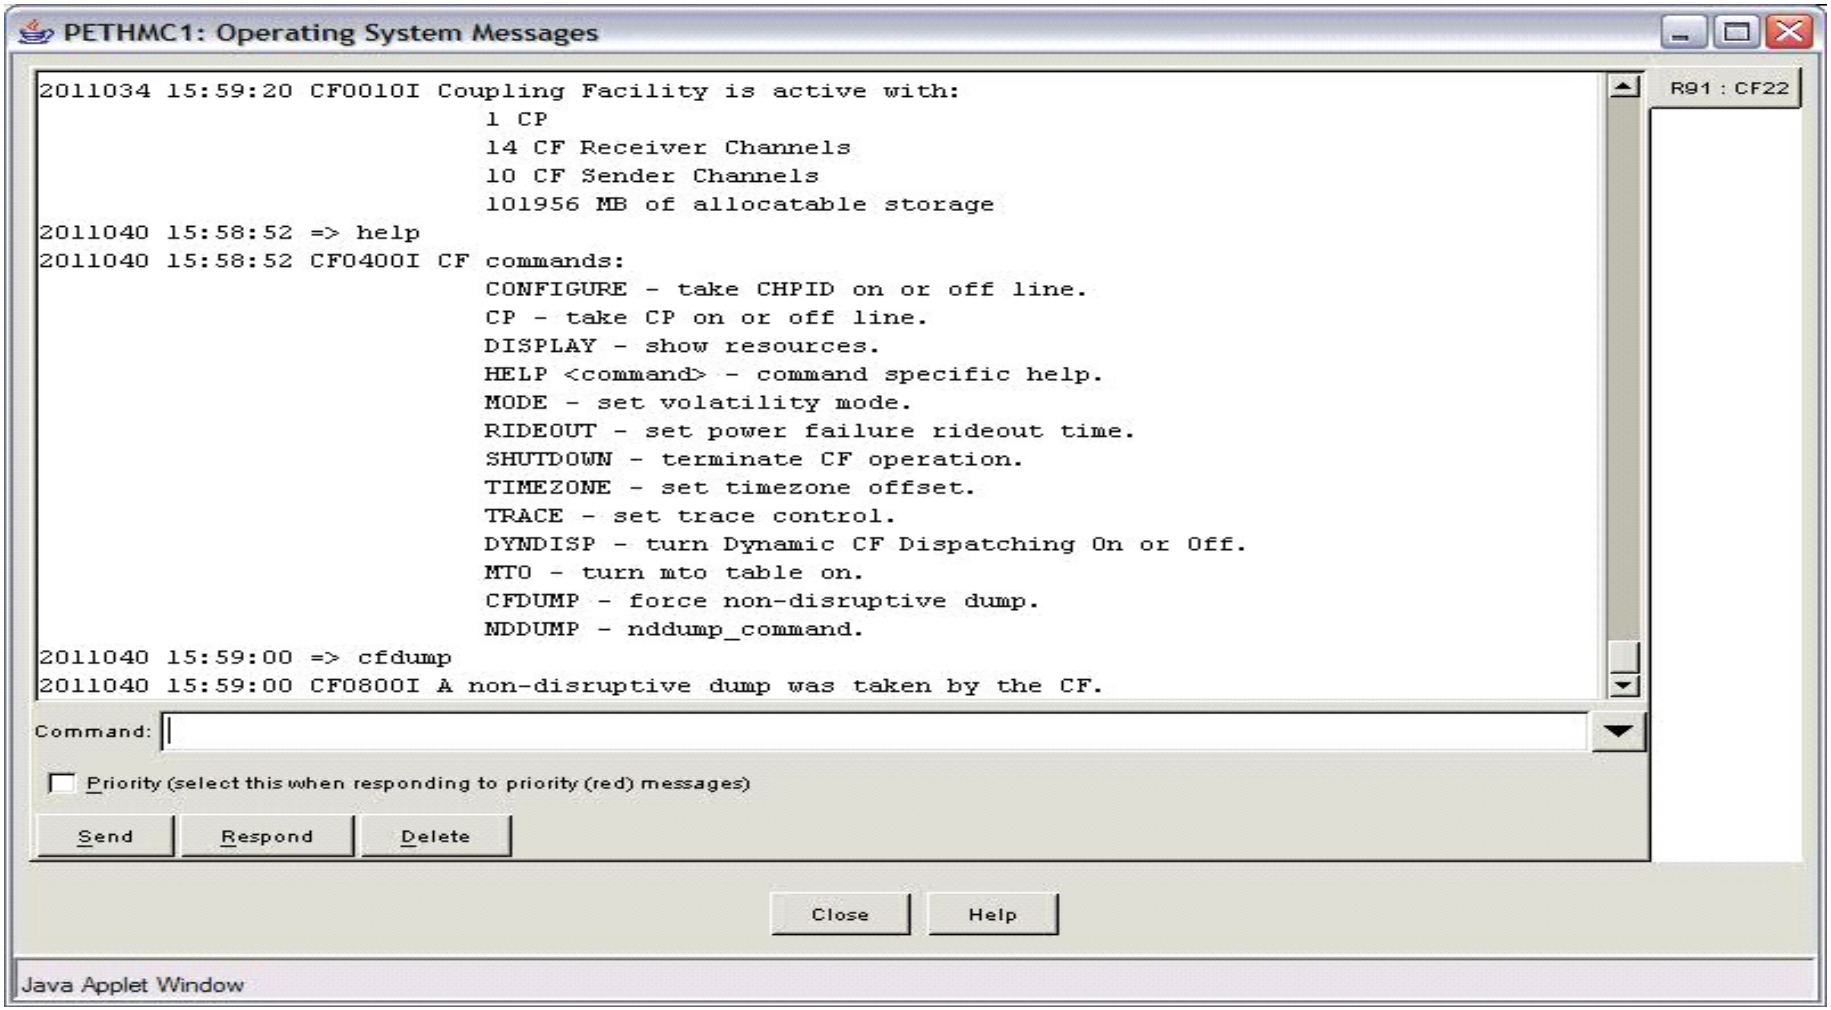

# Sysplex Fabric: CF Structures

- Physical Access
  - Do expected structures exist?
- Logical Access
  - CFRM Policy
  - IXLCONN
- Detecting Fabric Issues
  - D XCF,CF
  - D XCF,POLICY,TYPE=CFRM
  - RMF CF Activity Report
  - XCF/XES Messages (see handout)

Which structures do you normally run with?
Do they exist?

Running with expected CFRM policy?
Any failed structure allocations?

# Sysplex Fabric: CF Structures
## D XCF,CF ....

```
D XCF,CF,CFNAME=SVT1CF1
IXC362I  22.06.42  DISPLAY XCF 411
CFNAME: SVT1CF1
  COUPLING FACILITY     :   002817.IBM.02.0000000F8E66
                            PARTITION: 09   CPCID: 00
  SITE                  :   SITE1
  POLICY DUMP SPACE SIZE:   6000 K
  ACTUAL DUMP SPACE SIZE:
  STORAGE INCREMENT SIZE:

  CONNECTED SYSTEMS:
  N64      N65      N66      N67

  MONITORING SYSTEM: N64

  STRUCTURES:
  DBSVPLX1_SCA          IRRXCF00_B001          IRRXCF00_B002
  ISTGENERIC(NEW)       IXCPLEX_PATH2          IXCPLEX_PATH4
  LOGGER_STR1           SYSZWLM_8E662817(NEW)
```

Do your structures exist?
In the expected/desired CF?

# Sysplex Fabric: CF Structures
## D XCF,POLICY

```
D XCF,POLICY,TYPE=ALL
IXC364I  22.40.25  DISPLAY XCF 927
TYPE: ARM
POLICY NOT STARTED
TYPE: BPXMCDS
      NOT SUPPORTED BY
TYPE: CFRM
      POLNAME:          IXCPOLX1
      STARTED:          07/19/2012 21:19:19
      LAST UPDATED:     07/19/2012 11:24:45
TYPE: LOGR
      NOT SUPPORTED BY DISPLAY XCF,POLICY
TYPE: SFM
      POLNAME:          IXCSFMX1
      STARTED:          04/17/2012 19:24:49
      LAST UPDATED:     04/17/2012 19:24:08

SYSPLEX FAILURE MANAGEMENT IS ACTIVE
TYPE: WLM
      NOT SUPPORTED BY DISPLAY XCF,POLICY
```

CFRM Policy started?
Right name?
Right instance?

# Sysplex Fabric: CF Structures Messages

- IXL013I “application unable to connect to structure”
- IXL014I “application unable to connect to structure”

An application must connect in order to use a structure. The failure text might indicate that the structure could not be created at all.

- IXC453I “not enough signaling paths”
- IXC454I “unable to establish signaling connectivity”
- IXC455D “reply with interval or retry to re-initialize XCF”

You tend to see these messages on an IPLing system. Many root causes, but frequently implies that XCF signalling structures are not accessible.

- Applications may well issue their own messages to complain if their structures are not accessible

# Sysplex Fabric: Signalling Paths

Are CF structures used for signalling accessible?

Are CTC devices used for signalling online and operational?

- Physical Access
  - See “Sysplex Fabric: CF Structures”
  - CTC devices
- Logical Access
  - COUPLExx PATHOUT statements, or SETXCF START,PATHOUT
  - COUPLExx PATHIN statements, or SETXCF START,PATHIN
- Detecting Fabric Issues
  - D XCF,PO
  - D XCF,PI
  - Messages (see handout)
  - Apply “Sysplex Fabric: CF Structures” to signal structures
  - Check for IOS messages related to signal path CTC devices

# Sysplex Fabric: Signalling Paths
## D XCF,PO – summary view

```
D XCF,PO
IXC355I  22.32.57  DISPLAY XCF 867
PATHOUT TO SYSNAME:   ???????? - PATHS NOT CONNECTED TO OTHER SYSTEMS
 DEVICE (LOCAL/REMOTE): C140/???? C141/???? C142/???? C150/????
                        C151/???? C152/???? D140/???? D141/????
                        D142/???? D150/???? D151/???? D152/????
PATHOUT TO SYSNAME:   N64
 DEVICE (LOCAL/REMOTE): C110/C144 C111/C145 C112/C146 D110/D144
                        D111/D145 D112/D146
 STRNAME:               IXCPLEX_PATH1     IXCPLEX_PATH2
                        IXCPLEX_PATH3     IXCPLEX_PATH4
PATHOUT TO SYSNAME:   N65
 DEVICE (LOCAL/REMOTE): C120/C144
                        D121/D145
 STRNAME:               IXCPLEX_PA
                        IXCPLEX_PA
PATHOUT TO SYSNAME:   N66
 DEVICE (LOCAL/REMOTE): C130/C144
```

Shows perspective of local system

Paths not connected but should be?
Missing expected paths?
Paths to every other system?

"connected" does not imply "working"

# Sysplex Fabric: Signalling Paths
## D XCF,PO – detail

```
D XCF,PO,DEVICE=C110
IXC356I  22.34.29  DISPLAY XCF 886
LOCAL DEVICE     REMOTE   PATHOUT     REMOTE                  TRANSPORT
PATHOUT          SYSTEM   STATUS      PATHIN  RETRY  MAXMSG CLASS
C110             N64      WORKING     C144      100    5000 DEFAULT

LOCAL   REMOTE  REMOTE   PATHOUT     TRANSFR BUFFER  MSGBUF SIGNL MXFER
PATHOUT PATHIN  SYSTEM   STATUS
C110    C144    N64      WORKING
```

```
D XCF,PO,STRNAME=IXCPLEX_PATH1
IXC356I  22.39.14  DISPLAY XCF 916
STRNAME          REMOTE   PATHOUT
PATHOUT          SYSTEM   STATUS
IXCPLEX_PATH1             WORKING
                 N64      WORKING
                 N65      WORKING
                 N66      WORKING

STRNAME          REMOTE   PATHOUT     TRANSFR BUFFER  MSGBUF SIGNL MXFER
PATHOUT   LIST   SYSTEM   STATUS      PENDING LENGTH  IN USE NUMBR TIME
IXCPLEX_PATH1
            11   N64      WORKING           0   8124     110   239  1667
            18   N65      WORKING           0   8124     110   299  1012
            13   N66      WORKING           0   8124      56  4496  2703
```

Shows perspective of local system

"working" likely good, but not conclusive.

(If doubts, use detail D XCF,PI to verify flow from the perspective of target system)

Nonzero "pending" suggests delay or spike.

# Sysplex Fabric: Signalling Paths
## D XCF,PI - details

To check flow, issue detail pathin on the receiving side of signal path. Successive displays should show changes (unless path unused).

```
D XCF,PI,DEVICE=C114
IXC356I  22.37.52  DISPLAY XCF 910
LOCAL DEVICE    REMOTE   PATHIN      REMOTE                    LAST  MXFER
PATHIN          SYSTEM   STATUS      PATHOUT RETRY  MAXMSG RECVD TIME
C114            N64      WORKING     C140      100    2000  2351   248

LOCAL   REMOTE  REMOTE   PATHIN      DELIVRY BUFFER  MSGBUF SIGNL
PATHIN  PATHOUT SYSTEM   STATUS      PENDING LENGTH  IN USE NUMBR NOBUF
C114    C140    N64      WORKING           4    956       8  2351     0
```

Pending delivery of 4 is typical for CTC. 0 typical for list path. Bigger suggests msg exit delays.

Signal numbers on pathout side and pathin side should be close. Suggests transfer delay if not.

# Sysplex Fabric: Signalling Paths

- IXC458I “stopped signalling path”
- IXC459I “stopped signalling path unconditionally”
- IXC467I “restarting or stopping or rebuilding signalling path”

- IXC453I “not enough signaling paths”
- IXC454I “unable to establish signalling connectivity”

- Other messages
  - IOS for CTC devices used for signal paths
  - XCF/XES messages for CF or structures used for signal paths

# Sysplex Fabric: External Time Reference

- Physical Access
  - Coordinated Time Network
  - Timing links
- Logical Access
  - CTN ID
- Detecting Fabric Issues
  - D XCF,SYSPLEX,ALL
  - D ETR,DATA
  - Messages (see handout)
  - HMC

If z/OS image loses access to ETR, the system is in a wait-state.

So my "fabric detection" is either
- Proactive prevention
- Post mortem analysis

If lose ETR, "live" investigation is via the HMC and other systems that still have ETR access

Timer links operational?
Is the CTN ID correct?
PTS operational? BTS? Arbiter?
Is CTS the one you want?

# Sysplex Fabric: External Time Reference

- IEA394A “lost access to ETR (STP)”
- IEA015A “lost access to ETR (sysplex timer)”

- IXC406I “not connected to same ETR”
- IXC438I “new or changed CTNID”
- IXL160E “required request time ordering not enabled”
- IXL162E “required request time ordering will never be enabled”

- IEA031I “STP alert was issued to HMC”
- IEA395I “switched to backup time server”

- And of course, the HMC

# Methodology Recap

- We have eliminated any local issues that might be root cause
  - Wait-states, constraints, device errors
  - Explainable contention
  - Lack of access to sysplex componentry
- We may be reserving judgment on local issues that might arise as the result of sympathy sickness
  - Unexplained contention
  - Sick components
- Or we might not have any potential culprits identified, but still believe that a problem exists

# Problem Taxonomy

- Dead System
- Sick System
- Sysplex Fabric
- Sysplex Componentry
  - **Coupling Facility**
  - Signalling Service
  - Couple Data Sets
  - External Time Reference
- Configuration / Capacity
- Software Issues

# Sysplex Componentry Considerations

- The sysplex components are rather intertwined and mutually dependent on each other. For example:
  - A CF structure might be used for signalling
  - But signals need to be sent to manage the structures in the CF
  - Management of the CF needs access to CDS
  - But signals need to be sent to manage the CDS
  - And the CDS defines who the signalling service can talk to
- I am largely ignoring this complexity and pretending that we can look at the components in isolation
  - But that could be a vast oversimplification if you happen to be rebuilding a signal structure while in the midst of performing a PSWITCH to an alternate CFRM CDS

# Methodology Recap

- At this point, we should have the following conditions:
  - All systems are operational
  - All sysplex componentry is accessible via error free connections
- We suspect there may be issues with sysplex componentry
  - Performance issues
  - Configuration or capacity issues
- So we need a detailed understanding of the how the sysplex componentry works so that we can determine
  - How such problems might arise
  - How they might be observed
  - How they might be resolved

# Coupling Facility Request Processing

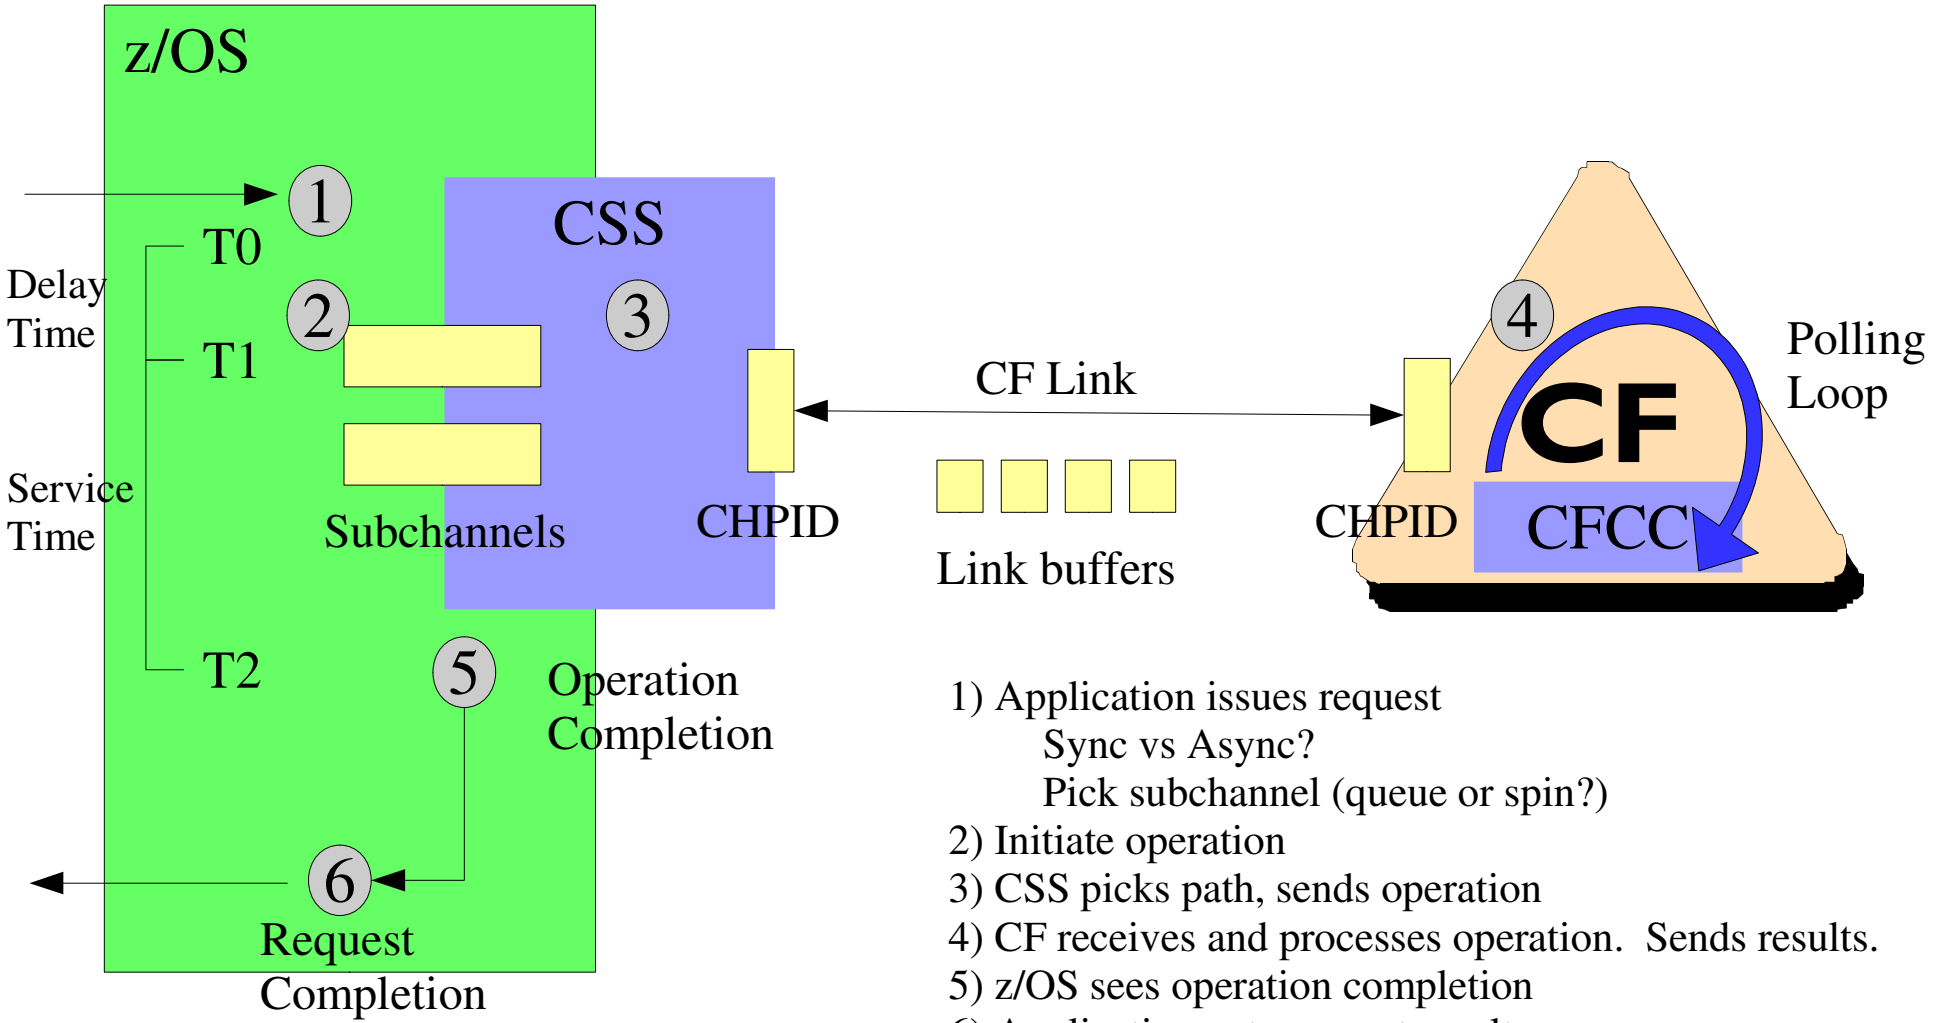


1) Application issues request
   - Sync vs Async?
   - Pick subchannel (queue or spin?)
2) Initiate operation
3) CSS picks path, sends operation
4) CF receives and processes operation. Sends results.
5) z/OS sees operation completion
6) Application gets request results

# CF Service Time Considerations

- Delay Time is time spent waiting for a subchannel
- Service Time is delta between sending operation and observing its completion
- Service time impacted by
  - Coupling Facility
    - Technology
    - Utilization
    - Contention
  - CF Links
    - Technology
    - Path busy conditions
    - Distance
  - Completion recognition

You get one number that encompasses all these factors.

One or more could be the issue.

# Coupling Facility Request Processing

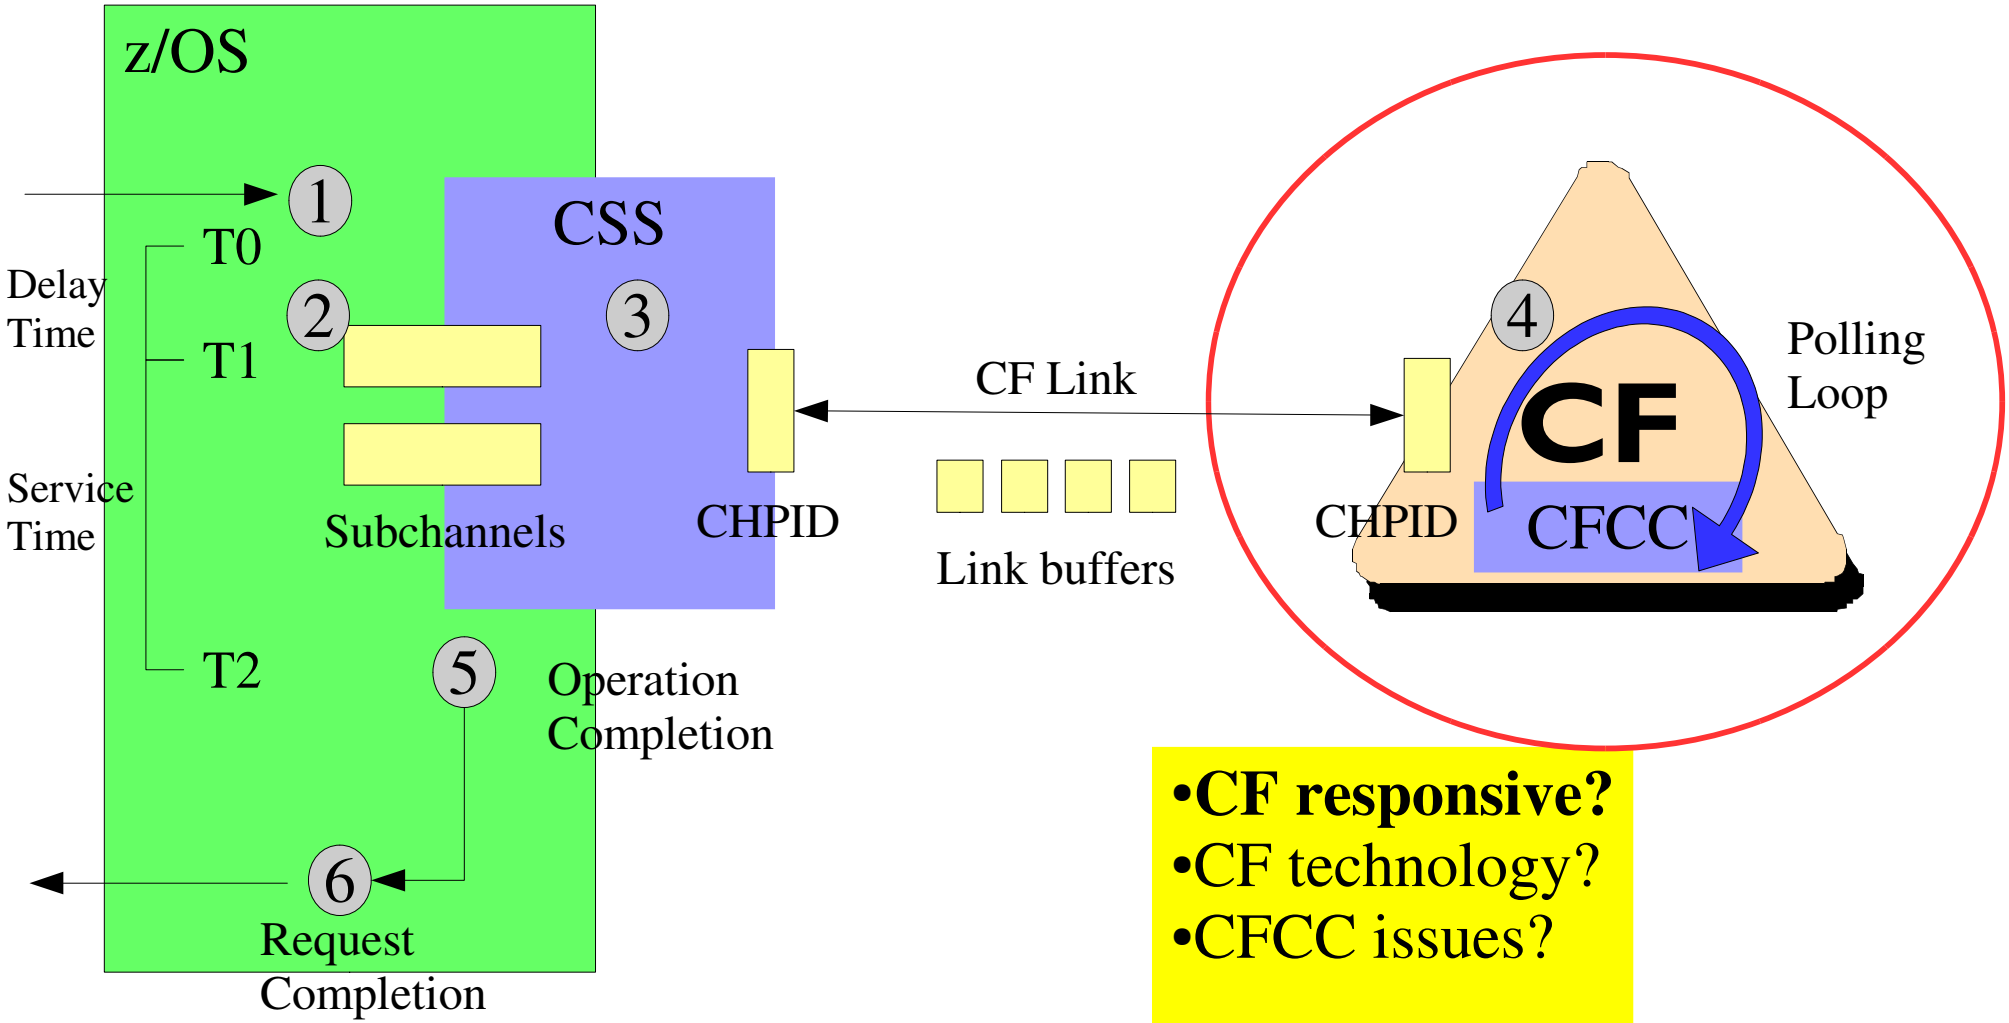


- **CF responsive?**
- CF technology?
- CFCC issues?

# CF Responsive?

- LPAR configuration must allow the CF to be sufficiently responsive so that the polling loop can run to receive commands in a timely manner
  - Dedicated CPs recommended
  - Shared CPs need sufficient weight
    - >50% for simplex or user managed duplexing
    - >95% for system managed duplexing
  - Dynamic Dispatch can give rise to erratic, elongated response times
  - Sharing CPs between z/OS and CF can impact response time
- CF utilization guidelines
  - <30% busy if single CP, otherwise <50%
  - Ensure sufficient capacity to handle structures that might be rebuilt into the CF as the result of failures or maintenance on a peer CF

# CF Technology Issues

- Our primary concern is the impact that CF service time has on the z/OS image that is accessing the CF
  - Faster CF request service time reduces overhead for z/OS image
- Type of CF processor determines how quickly an operation can be performed (but just one factor among many)
- Rule of thumb: CF processor should be no more than one generation behind the processor that hosts the z/OS images that use the CF

# CFCC Issues

- Stay current with maintenance

- Internal Contention
  - Very hard to detect, and seldom seen.  I mention it to be complete.
  - Contention issues within the CF generally arise from usage patterns for particular structures, and are isolated to those structures
  - Do not typically see “global” contention issues in the CF

# Detecting Coupling Facility Issues

- First resolve existence and accessibility issues
  - See “Sysplex Fabric: Coupling Facility”
- Performance
  - RMF CF Activity Report
  - RMF Partition Data Report
  - RMF Monitor III – CF Overview
  - RMF Monitor III – Sysplex CF Views

Getting enough physical CP?
Timely dispatch?
CF utilization within guidelines?

# RMF CF Activity Report

- CF Utilization within guidelines?
- Dynamic dispatching?
- Effective CPs?

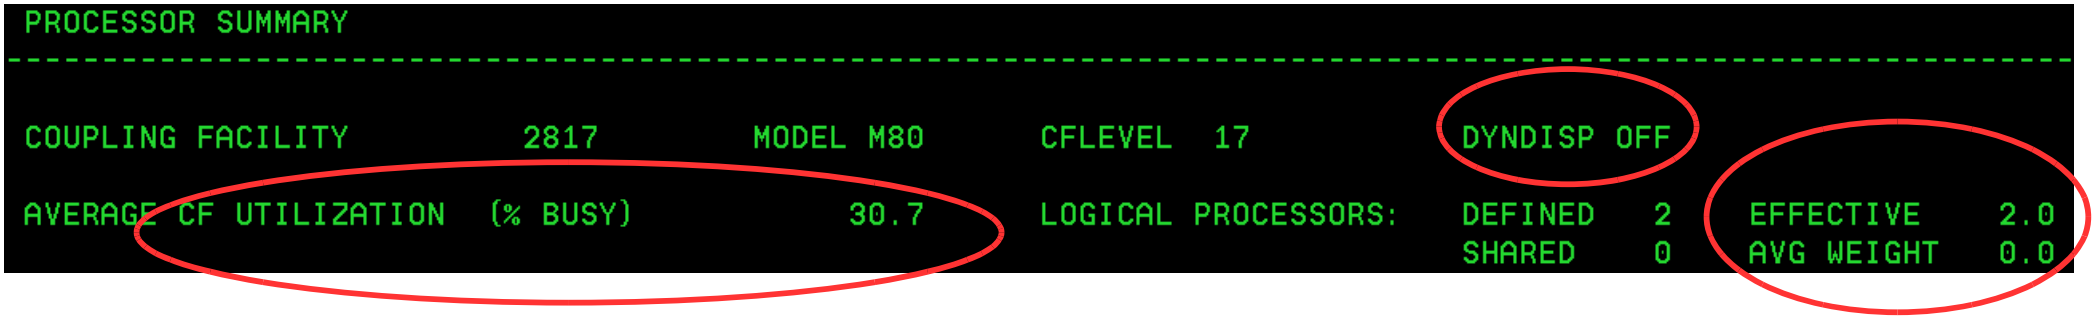


*May need to look at data over several intervals*
*What utilization is normal for you?*

# RMF Partition Data Report

```
          z/OS V1R13              SYSTEM ID SYS1            DATE 07/12/2012          INTERVAL 15.00.294
                                  RPT VERSION V1R13 RMF     TIME 20.15.00            CYCLE 1.000 SECONDS

--------- PARTITION DATA -----------------  -- LOGICAL PARTITION PROCESSOR DATA --   -- AVERAGE PROCESSOR UTILIZATION PERCENTAGES --
                 ----MSU----  -CAPPING--  PROCESSOR-  ----DISPATCH TIME DATA----   LOGICAL PROCESSORS  --- PHYSICAL PROCESSORS ---
NAME       S   WGT  DEF    ACT  DEF   WLM%  NUM   TYPE   EFFECTIVE       TOTAL      EFFECTIVE    TOTAL  LPAR MGMT  EFFECTIVE  TOTAL

CF1        A   DED                          1    ICF   00.15.00.269  00.15.00.275      100.0    100.0       0.00      50.00  50.00
CF2        A   DED                          1    ICF   00.15.00.170  00.15.00.176      99.99    99.99       0.00      49.99  49.99
*PHYSICAL*                                             00.00.00.048                                0.00                   0.00
                                                       ------------  ------------                 -----      -----       -----
  TOTAL                                                00.30.00.440  00.30.00.500                 0.00      99.99       100.0
```

Shared CP (suspect), or Dedicated CP (good) ?

*This is not CF Utilization. The logical CP looks 100% busy due to polling loop.*

Getting enough physical CP ?

# RMF Monitor III: CF Overview

```
                    RMF V1R11  CF Overview      - MABPLX           Line 1 of 3

Samples: 100     Systems: 9    Date: 08/07/12  Time: 14.00.00  Range: 100   Sec

CF Policy: CFRMPOL1    Activated at: 07/24/12 18.28.37

----- Coupling Facility ------  ------ Processor -------  Request  - Storage --
Name      Type  Model Lvl  Dyn  Util% Def Shr Wgt Effect  Rate     Size   Avail

CFRP      2817    M66  17  OFF   6.6    2   0       2.0      9441    20G    17G
CF1A      2094    S54  15  OFF   4.6    2   0       2.0      3123    20G    17G
CF1C      2817    M32  17  OFF   5.6    2   0       2.0      9095    20G    16G
```

# RMF Monitor III: CF System View

```
                    RMF V1R11  CF Systems        - MABPLXS           Line 1 of 27

Samples: 100      Systems: 9    Date: 08/07/12  Time: 14.00.00  Range: 100   Sec

CF Name  System    Subchannel   -- Paths --  -- Sync ---  ------- Async -------
                   Delay  Busy  Avail Delay    Rate  Avg    Rate  Avg  Chng  Del
                     %      %           %           Serv         Serv   %     %

CFRP     CSK         0.0    0.0    4     0.0     0.0     0    22.8   344  0.0   0.0
         SA0         0.0    2.1    4     0.0   112.6    11    2051   279 11.2   0.5
         SB0         0.0    0.8    4     0.1   104.6     7    1345   166  0.0   0.0
         SC0         0.0    2.6    4     0.0   120.4    11    3519   204  0.0   0.0
         SD0         0.0    0.2    4     0.0   108.8     6   664.1    81  0.0   0.0
         SE0         0.0    0.5    2     0.0    43.8    24   440.9   161  0.0   0.0
         SF0         0.0    0.4    2     0.0    38.7    25   381.0   157  0.0   0.0
         SG0         0.0    0.1    4     0.0    <0.1    72    <0.1   389  0.0   0.0
         SH0         0.0    0.2    2     0.0    <0.1    66    <0.1   412  0.0   0.9
CF1A     CSK         0.0    0.0    4     0.0    14.5    13    19.2   487  0.0   0.0
         SA0         0.0    0.3    4     0.0   140.6    24   446.3   164  0.5   0.2
         SB0         0.0    0.2    4     0.0    53.3    19   388.3   135  0.0   0.0
         SC0         0.0    0.2    4     0.0   156.5    23   475.0   116  0.1   0.1
         SD0         0.0    0.3    4     0.0    58.7    20   484.4   146  0.0   0.0
         SE0         0.0    0.6    2     0.0    33.9    32   348.4   224  0.0   0.0
         SF0         0.0    0.3    2     0.0     7.0    30   258.2   184  0.0   0.0
         SG0         0.0    0.1    4     0.0    12.4     6    44.4   348  0.0   0.0
         SH0         0.0    0.2    2     0.0    38.9    13    43.0   480  0.0   2.3
etc.
```

# Coupling Facility Request Processing

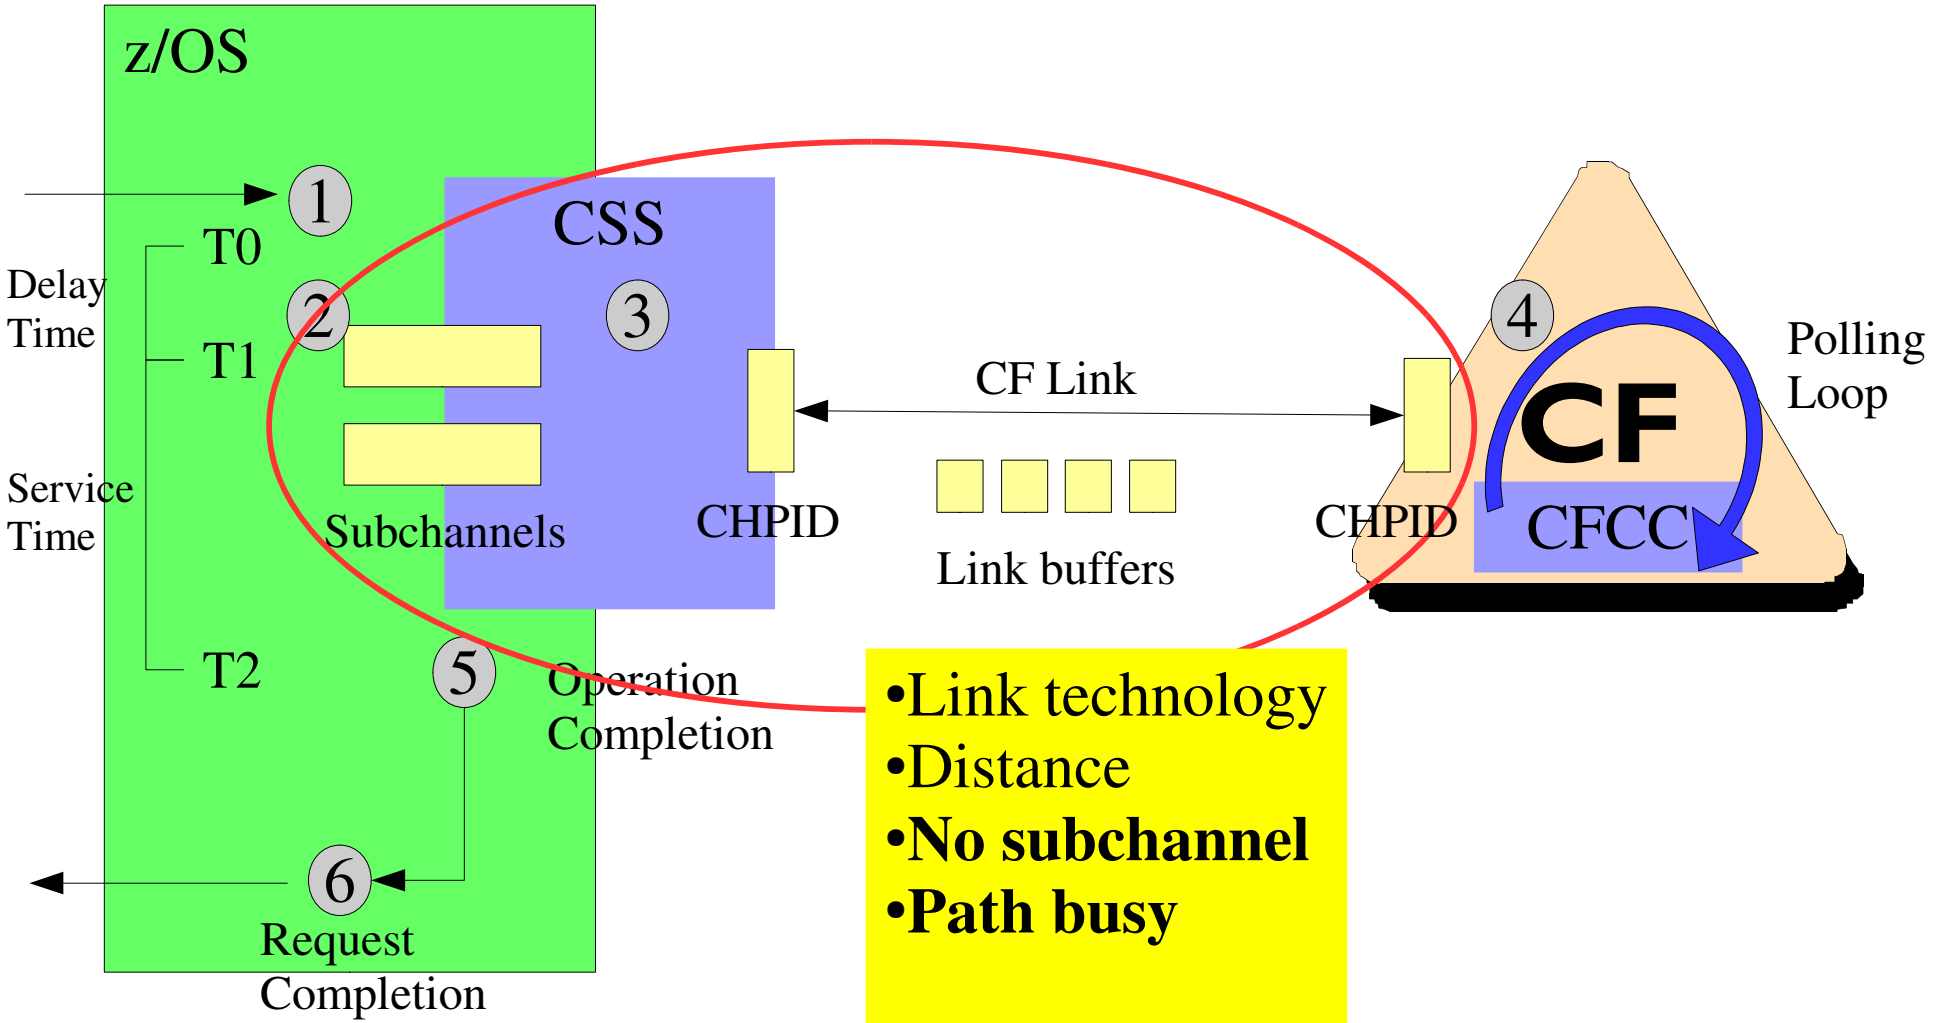

# CF Link Considerations

- Link technology and distance
  - Faster links improve transmit portion of service time
  - Distance increases service time by 10 mics/km
- No subchannel conditions
  - Bursts? Sustained load? Tuned due to path busy?
- Path busy conditions
  - Number of requests exceeds number of available link buffers
  - XES may tune number of subchannels to avoid this condition
  - Distance (link buffers in use for longer)
- Configured correct number of subchannels/CHPID?
  - 32 for HCA2-O LR or HCA3-O LR, otherwise 7

# Detecting CF link issues

- First eliminate all physical link errors
- First resolve or eliminate CF responsiveness issues
  - Unresponsive CF can induce link problems
  - Link buffers “linger”, which can induce path busy conditions
- RMF Report of CF Activity
  - Delayed requests implies “no subchannel”
  - Reports “path busy” conditions
- D CF,CFNAME
  - Shows configured links
  - How many subchannels available? Being used?

# RMF CF Activity: Subchannel Activity

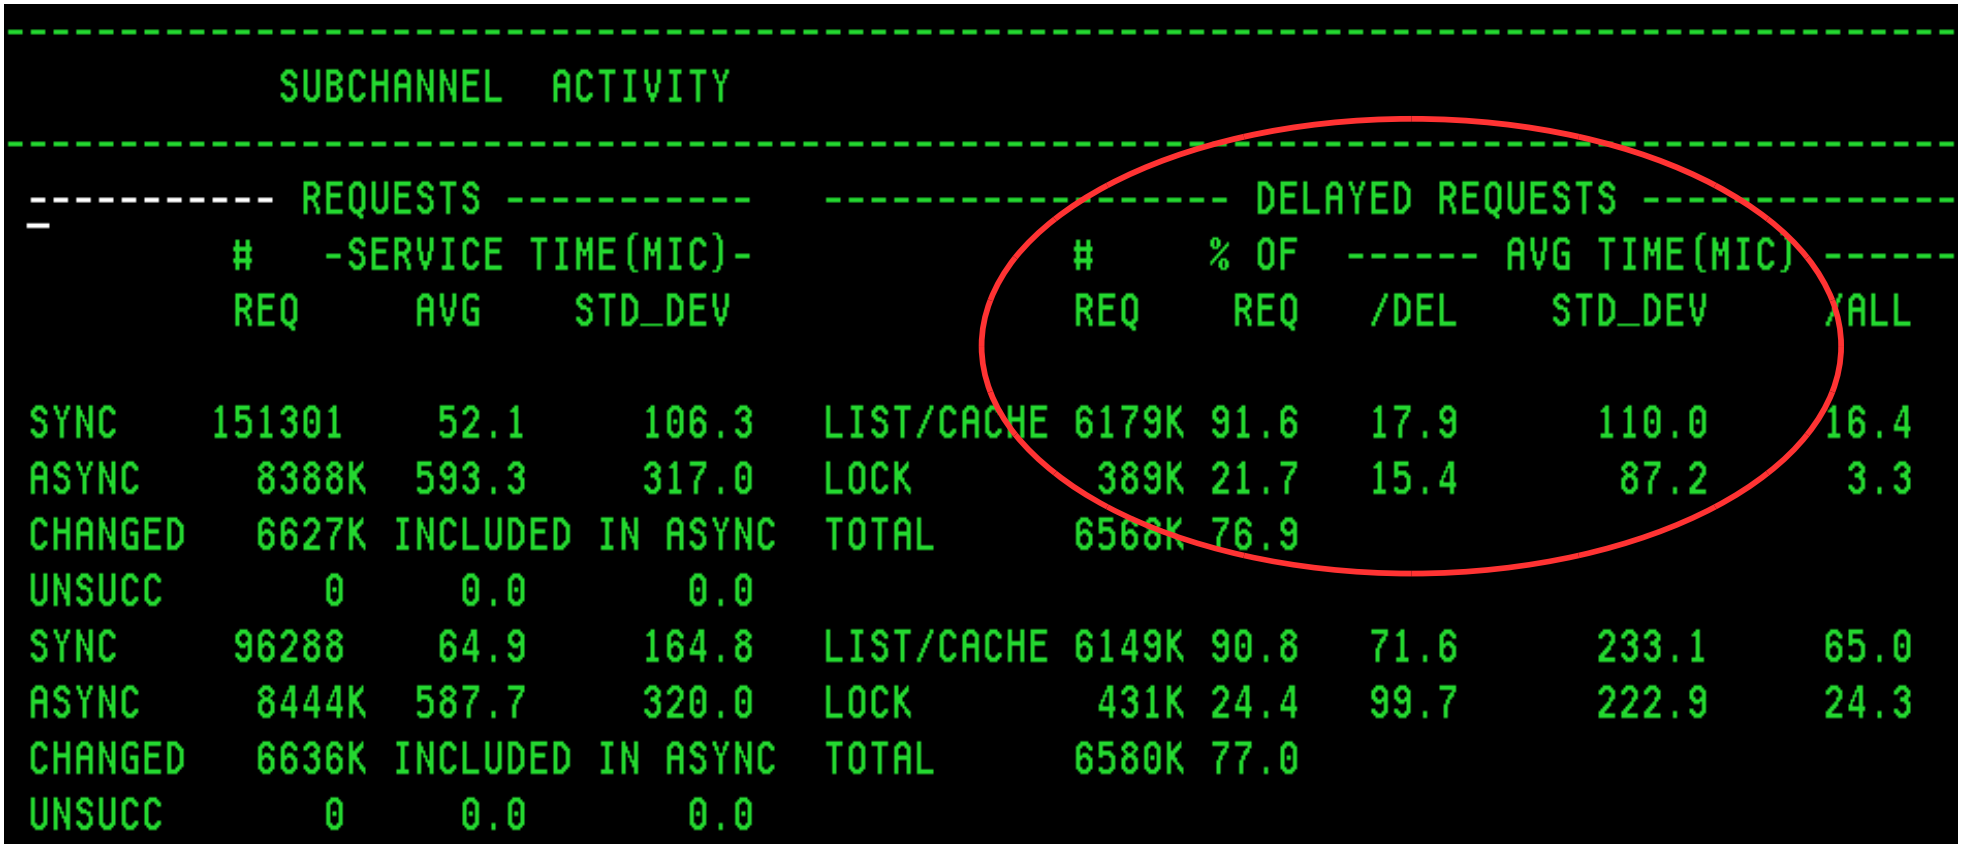

```
                    SUBCHANNEL  ACTIVITY

---------- REQUESTS ----------   ---------------- DELAYED REQUESTS --------------
_
         #   -SERVICE TIME(MIC)-              #     % OF  ------ AVG TIME(MIC) ------
         REQ     AVG    STD_DEV               REQ    REQ   /DEL    STD_DEV     /ALL

SYNC    151301    52.1     106.3   LIST/CACHE 6179K 91.6   17.9      110.0     16.4
ASYNC     8388K  593.3     317.0   LOCK        389K 21.7   15.4       87.2      3.3
CHANGED   6627K INCLUDED IN ASYNC   TOTAL      6568K 76.9
UNSUCC       0     0.0       0.0
SYNC     96288    64.9     164.8   LIST/CACHE 6149K 90.8   71.6      233.1     65.0
ASYNC     8444K  587.7     320.0   LOCK        431K 24.4   99.7      222.9     24.3
CHANGED   6636K INCLUDED IN ASYNC   TOTAL      6580K 77.0
UNSUCC       0     0.0       0.0
```

*Know your workload.*
*What is normal for you.*
*What changed?*

# RMF CF Activity: Subchannel Activity

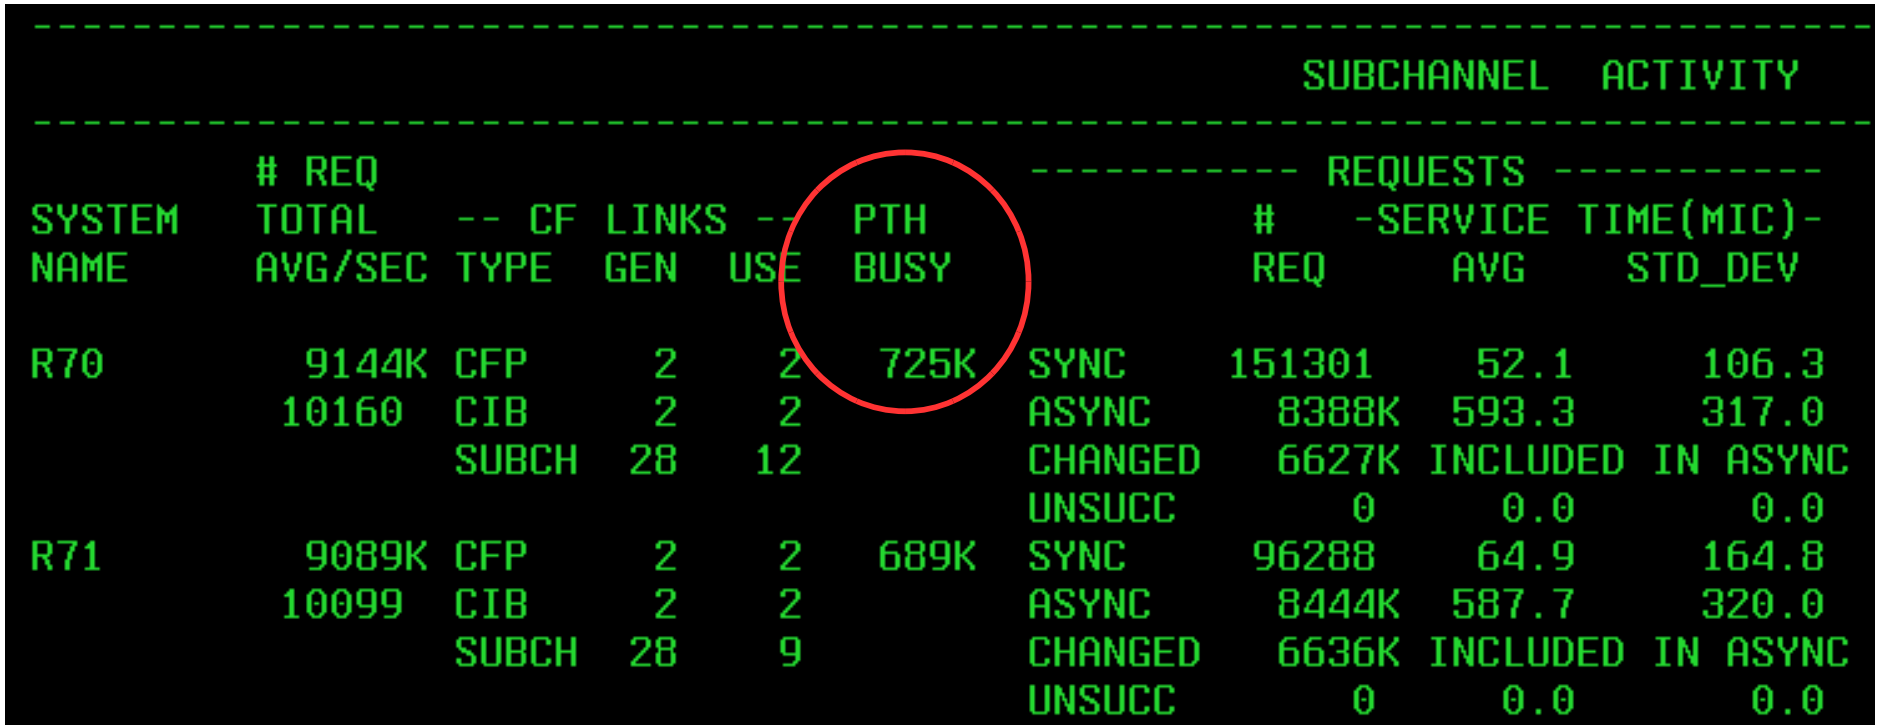

SUBCHANNEL ACTIVITY

| SYSTEM NAME | # REQ TOTAL AVG/SEC | -- CF LINKS -- TYPE | GEN | USE | PTH BUSY | REQUESTS | # REQ | -SERVICE TIME(MIC)- AVG | STD_DEV |
|---|---|---|---|---|---|---|---|---|---|
| R70 | 9144K | CFP | 2 | 2 | 725K | SYNC | 151301 | 52.1 | 106.3 |
| | 10160 | CIB | 2 | 2 | | ASYNC | 8388K | 593.3 | 317.0 |
| | | SUBCH | 28 | 12 | | CHANGED | 6627K | INCLUDED IN ASYNC | |
| | | | | | | UNSUCC | 0 | 0.0 | 0.0 |
| R71 | 9089K | CFP | 2 | 2 | 689K | SYNC | 96288 | 64.9 | 164.8 |
| | 10099 | CIB | 2 | 2 | | ASYNC | 8444K | 587.7 | 320.0 |
| | | SUBCH | 28 | 9 | | CHANGED | 6636K | INCLUDED IN ASYNC | |
| | | | | | | UNSUCC | 0 | 0.0 | 0.0 |

*This was a test intended to drive path busy conditions (it worked).*

*Know your workload.*
*What is normal for you.*
*What changed?*

# D CF,CFNAME=xxxx .....

```
PATH        PHYSICAL              LOGICAL  CHANNEL TYPE      AID  PORT
03 / 0727   NOT OPERATIONAL       ONLINE   CIB               N/A   N/A
C0 / 0737   NOT OPERATIONAL       ONLINE   CIB               N/A   N/A
C1 / 0738   ONLINE                ONLINE   CIB               N/A   N/A

COUPLING FACILITY SUBCHANNEL STATUS
 TOTAL:   96   IN USE:    7   NOT USING:    0   NOT USABLE:   89
 OPERATIONAL DEVICES / SUBCHANNELS:
                F42F /
                F433 /
```

Configured links with 32 SCH/CHPID but link only supports 7

Nonzero implies XES tuning to avoid path busy

Configuration error? Hardware problem?

# Methodology Concern

- Subchannel and path busy conditions imply that there are not enough subchannels and link buffers to satisfy the workload
- Typically resolved by
  - Adding links, or CHPIDs (for IFB)
  - Upgrading link technology
- But …
  - Has workload grown?
  - Experiencing a spike? Just tolerate it?
  - Runaway application?

*Know your workload.*
*What is normal for you.*
*What changed?*

# Coupling Facility Request Processing

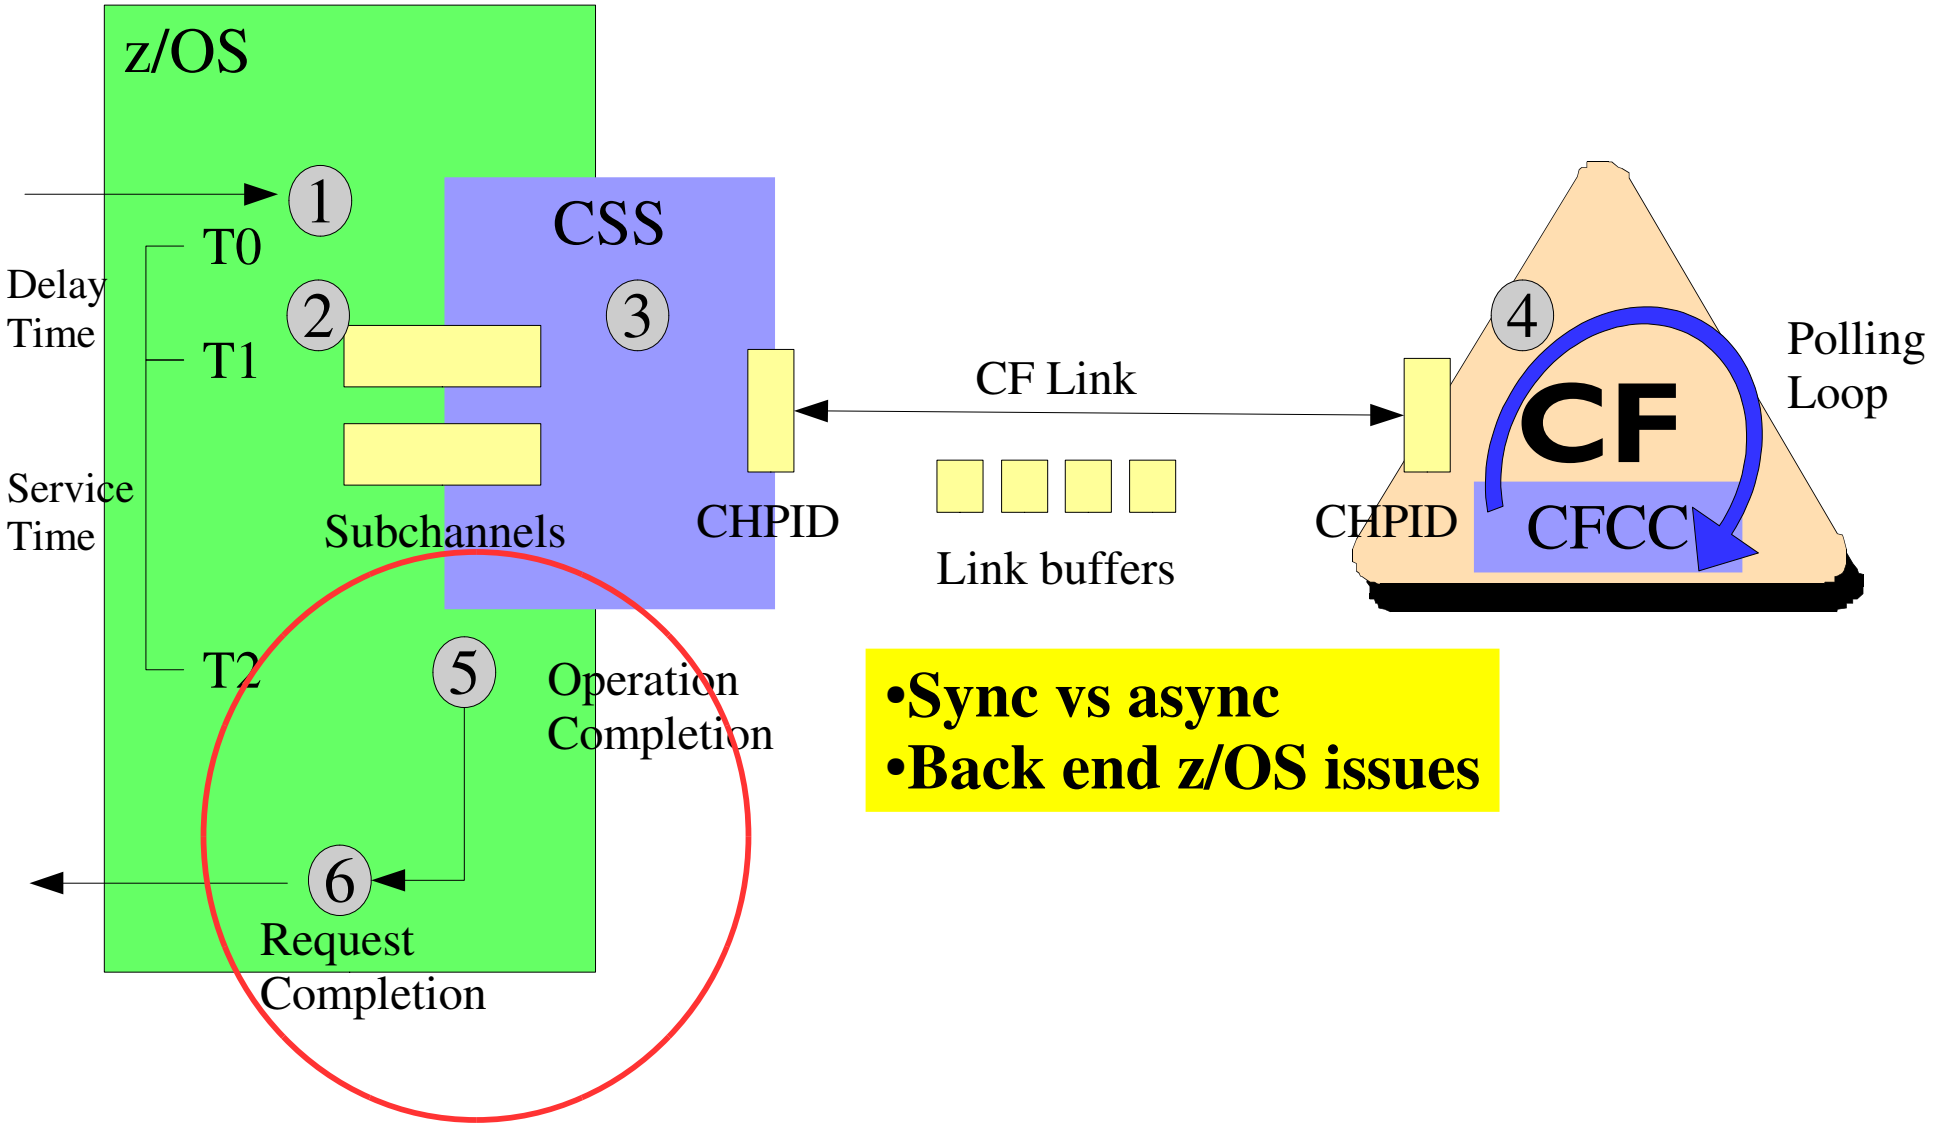


- **Sync vs async**
- **Back end z/OS issues**

# Asynchronous operation completion

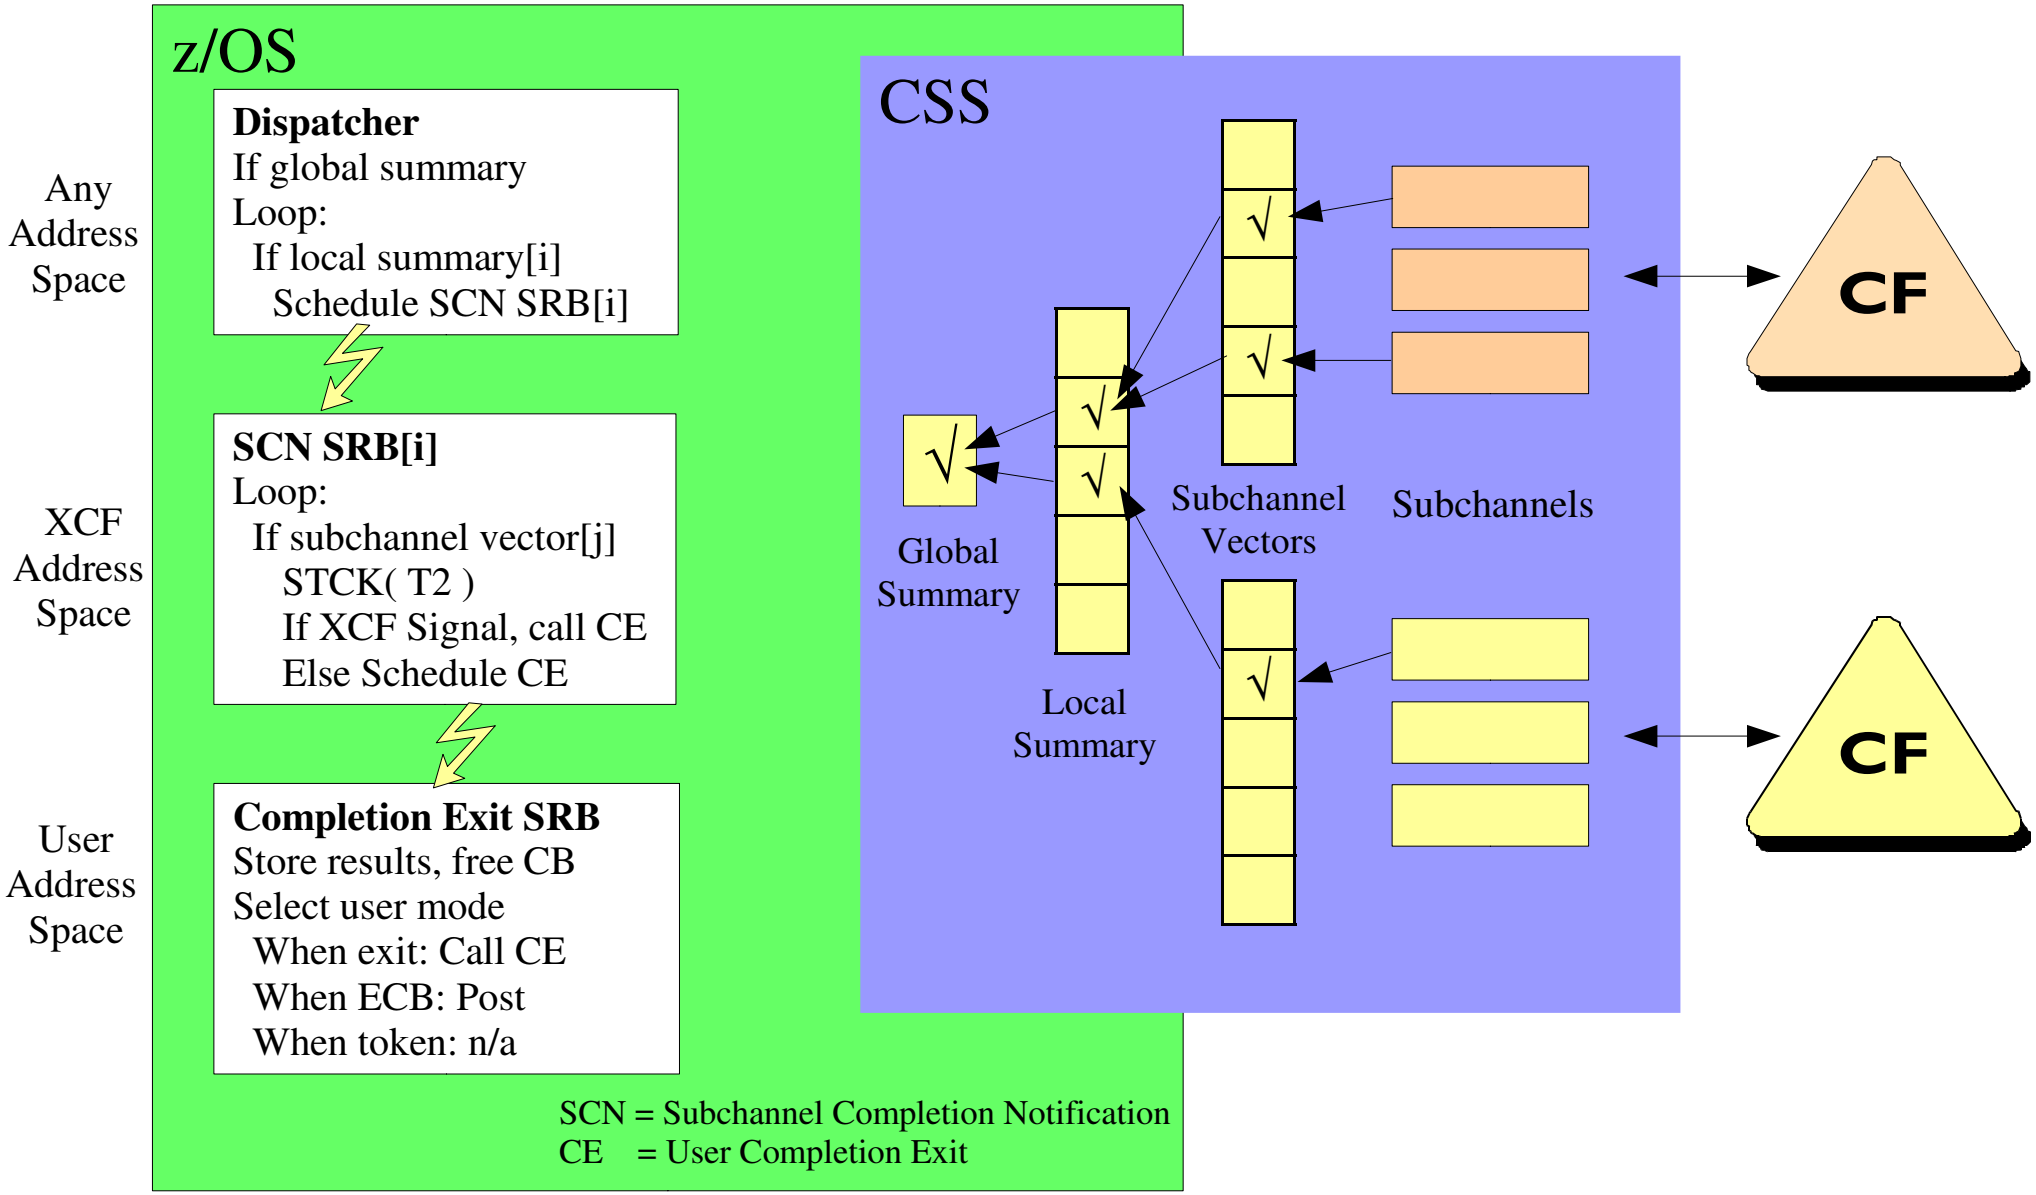

# Request Completion Issues

- Mismanagement of summary bits and subchannel vectors
  - (Extremely rare, only seen this once. I mention to be complete.)
  - Subchannel completion vector issues are likely “repaired” by synchronous requests, or an internal monitor
- Loss of physical processor
  - May elongate asynchronous CF service times
  - May elongate application response times
- Low loads
  - MVS may ask LPAR to take it out of its no work wait less often
  - If dispatcher runs less often, takes longer to notice completion
  - Which increases async service times

# Request Completion Issues ...

## z/OS Dispatching Issues

- Is application address space getting dispatched enough?
  - Won't show up in CF service time measurements
  - But application might appear to be sluggish, and
  - Increases dwell time of XES control blocks (common storage)
- XCF runs at high dispatch priority, so not usually an issue
  - Unless there is a more global issue
    - Storage constraints, spin loops, …
    - But we already eliminated those issues, right?

# CF Request Response Time Summary
## What does your “one number” tell you?

- Time z/OS spends waiting for subchannels

Delay time / Service time

- Time spent resolving path busy conditions
- Time spent transmitting request from z/OS to CF

Local

- Time spent waiting for physical dispatch of CF to receive request
- Time CF spends preparing, processing, and completing the request

Global

- Time spent transmitting results from CF to z/OS
- Time spent waiting for physical dispatch of z/OS to receive results

Local

Service time / Back End time

- Time spent in back end application completion processing

# Methodology

- Having eliminated issues for CF requests in general, there could be issues that are unique to specific structures
- Might use one or more of the following approaches to decide which ones to look at. You might review structures that:
  - Are known to be critical to the sysplex
    - XCF signalling paths, ISGLOCK, ...
  - Have service times or request rates out of line with past behavior
  - Are known to be used by applications under suspicion

*Bear in mind that the application could be using a service that exploits its own CF structure. You might miss the root cause if you only look at the structures you think are relevant to the suspicious application.*

- If you do identify something abnormal, you may need application specific expertise for deeper diagnosis

# CF Structure Specific Concerns

- Properly sized?
  - Always resize after CF/hardware upgrade
  - Resize as workload changes
  - Is CFRM policy change still pending for new size?
  - For lock structures: Contention? False contention?
- Hitting full thresholds?
  - Whether an issue depends on application
- Being altered?
- Being rebuilt?
- Newly allocated?
- Application design issues
  - Perhaps the implementation induces contention/queueing in CF

*See "z/OS Hot Topics" Issue 26 Lead article on CF Sizer*

*Sizing issues often arise long after CF upgrade for which structure was not resized*

*Contention implies delay, and induces signalling activity. Can fix false contention by increasing structure size appropriately.*

*Activity is quiesced (delayed) while in rebuild*

# CF Structure Specific Diagnosis
## Be Prepared

- Know your workload
- Periodically review structure usage
  - Request rates
  - Service times
  - Storage consumption (within the structure)
  - For lock structures, contention and false contention rates
- Relate changes in structure usage to workload variations
  - Track any “odd” behavior
- Maintain a list of which applications use which structures

# Diagnosing CF Structure Specific Issues

- What changed?
  - CFCC MCL upgrades
  - Processor upgrades
  - Software maintenance upgrades
  - Software migrations
  - Workload changes
- D XCF,STR
- RMF CF Activity Reports
  - Request rates and service times for each structure, each system
- RMF Monitor III – Sysplex CF Views

# D XCF,STR,STRNAME=xxx

```
D XCF,STR,STRNAME=IXCPLEX_PATH1
IXC360I  22.30.28  DISPLAY XCF 852
STRNAME: IXCPLEX_PATH1
 STATUS: ALLOCATED
 EVENT MANAGEMENT: POLICY-BASED
 TYPE: LIST
 POLICY INFORMATION:
  POLICY SIZE    : 278016 K
  POLICY INITSIZE: N/A
  POLICY MINSIZE : 0 K
  FULLTHRESHOLD  : 80
  ALLOWAUTOALT   : NO
  REBUILD PERCENT: N/A
  DUPLEX         : DISABLED
  ALLOWREALLOCATE: YES
  PREFERENCE LIST: SVT1     SVT1CF2  SVT1CF1
  ENFORCEORDER   : NO
  EXCLUSION LIST IS EMPTY
```

# D XCF,STR,STRNAME=xxxx .....

```
ACTIVE STRUCTURE
----------------
 ALLOCATION TIME: 07/19/2012 21:17:36
 CFNAME         : SVT1
 COUPLING FACILITY: 002818.IBM.02.0000000EC876
                    PARTITION: 01   CPCID: 00
 ACTUAL SIZE    : 272 M
 STORAGE INCREMENT SIZE: 1 M
 USAGE INFO      TOTAL     CHANGED    %
  ENTRIES:       62688           1    0
  ELEMENTS:      62659          16    0
 PHYSICAL VERSION: C9E48F21 65D57923
 LOGICAL  VERSION: C9E48F21 65D57923
 SYSTEM-MANAGED PROCESS LEVEL: NOT APPLICABLE
 DISPOSITION    : DELETE
 ACCESS TIME    : 0
 MAX CONNECTIONS: 32
 # CONNECTIONS  : 4

CONNECTION NAME  ID VERSION  SYSNAME  JOBNAME  ASID STATE
---------------- -- -------- -------- -------- ---- ------------
SIGPATH_01001B6C 01 000102DB N64      XCFAS    0006 ACTIVE
SIGPATH_02001B6F 02 00020213 N65      XCFAS    0006 ACTIVE
SIGPATH_03001B6E 03 000301F2 N66      XCFAS    0006 ACTIVE
SIGPATH_04001B6D 04 000401EE N67      XCFAS    0006 ACTIVE
```

# Post Processor: CF Activity Report – PART 1

```
------------------------------------------------------------------------------------------------------------------------
 COUPLING FACILITY NAME = CFMARK
 TOTAL SAMPLES(AVG) =   845  (MAX) =   861  (MIN) =   828
------------------------------------------------------------------------------------------------------------------------

                                          COUPLING  FACILITY  USAGE  SUMMARY
------------------------------------------------------------------------------------------------------------------------
 STRUCTURE SUMMARY
------------------------------------------------------------------------------------------------------------------------

                                                  % OF            % OF    % OF     AVG    LST/DIR  DATA      LOCK     DIR REC/
       STRUCTURE                       ALLOC      CF      #       ALL     CF      REQ/   ENTRIES  ELEMENTS  ENTRIES  DIR REC
 TYPE  NAME              STATUS CHG     SIZE      STOR    REQ     REQ     UTIL    SEC    TOT/CUR  TOT/CUR   TOT/CUR  XI'S

 LIST  MARK_LIST1        ACTIVE          20M       0.0   857662    1.2     0.4   952.96     24K      21K       N/A      N/A
                                                                                            323       19       N/A      N/A
       MARK_LIST2        ACTIVE         121M       0.1    18336    0.0     0.0    20.37    143K     285K       N/A      N/A
                                                                                            151      859       N/A      N/A
...lines omitted ...
 LOCK  MBROOKS_LOCK1     ACTIVE         768M       0.4   11206K   15.8     4.8    12452   1808K       0       134M      N/A
                                                                                             23K       0       463K      N/A
       MBROOKS_LOCK1     ACTIVE         500M       0.2    2453K    3.5     1.0   2726.0    702K       0       134M      N/A
                                                                                             53K       0        87K      N/A
...lines omitted ...
 CACHE BROOKS            ACTIVE         625M       0.3     720     0.0     0.0     0.80    441K     250K       N/A        0
                                                                                             11        4       N/A        0
       MBROOKS_GBP0      ACTIVE          49M       0.0      94     0.0     0.0     0.10     41K    8287       N/A        0
                                                                                              0        0       N/A        0
...lines omitted ...
                                       -------   ------  -------   -----   ------  -------
            STRUCTURE TOTALS             96G      48.3    70907K    100     100     78785
```

# Post Processor: CF Activity Report – PART 3

```
STRUCTURE NAME = ISGLOCK           TYPE = LOCK   STATUS = ACTIVE
          # REQ    -------------- REQUESTS -------------   -------------- DELAYED REQUESTS -------------
SYSTEM    TOTAL                 #     % OF  -SERV TIME(MIC)-    REASON    #    % OF  ---- AVG TIME(MIC) -----  EXTERNAL REQUEST
NAME      AVG/SEC               REQ    ALL     AVG    STD_DEV            REQ    REQ   /DEL    STD_DEV   /ALL   CONTENTIONS

BROOKS1     317K   SYNC    317K   18.1     9.9      2.3     NO SCH   59    0.0   19.7       44.8    0.0   REQ TOTAL      331K
          351.7    ASYNC     0     0.0     0.0      0.0     PR WT     0    0.0    0.0        0.0     0.0   REQ DEFERRED  7089
                   CHNGD     0     0.0  INCLUDED IN ASYNC   PR CMP    0    0.0    0.0        0.0     0.0   -CONT         7087
                                                                                                          -FALSE CONT   1334
...lines omitted...
-------------------------------------------------------------------------------------------------------------------------------
 TOTAL      1746K   SYNC   1584K   90.8    10.5      4.2     NO SCH   91    0.0   20.1       42.2    0.0   REQ TOTAL     1814K
            1940    ASYNC   161K    9.2   238.4    265.5     PR WT     0    0.0    0.0        0.0     0.0   REQ DEFERRED    46K
                    CHNGD     0     0.0                      PR CMP    0    0.0    0.0        0.0     0.0   -CONT           46K
                                                                                                          -FALSE CONT     12K
```

# RMF Monitor III: CF System View

```
                    RMF V1R11  CF Activity      - MABPLX      Line 1 of 154

Samples: 100     Systems: 9    Date: 08/07/12  Time: 14.00.00  Range: 100   Sec

CF: ALL          Type  ST System   CF    --- Sync ---    --------- Async --------
                                     Util   Rate   Avg     Rate   Avg   Chng   Del
Structure Name                        %             Serv           Serv   %      %

ADSW_DFHJ01      LIST  AP *ALL       0.0    0.0       0    0.8    305   1.3    0.0
                 LIST  AS *ALL       0.0    0.0       0    0.8    290   1.3    1.3
ADSW_DFHJ02      LIST  A  *ALL       0.0    0.8       6   <0.1     20  50.0   50.0
ADSW_DFHJ03      LIST  A  *ALL       0.0    3.5      16    0.7    991   1.5    1.5
ADSW_DFHJ04      LIST  A  *ALL       0.0    0.7      25    0.0      0   0.0    0.0
ADSW_DFHJ05      LIST  A  *ALL       0.0    0.7       7    0.0      0   0.0    0.0
ADSW_DFHJ06      LIST  A  *ALL       0.0   <0.1      18    0.0      0   0.0    0.0
ADSW_DFHJ07      LIST  A  *ALL       0.0    0.7      25    0.0      0   0.0    0.0
ADSW_DFHLGLOG1   LIST  AP *ALL       0.0    0.0       0    0.0      0   0.0    0.0
                 LIST  AS *ALL       0.0    0.0       0    0.0      0   0.0    0.0
APPCLOG          LIST  A  *ALL       0.0    0.0       0   <0.1    880   0.0    0.0
CI1_DFHLOG       LIST  A  *ALL      13.0  192.1      25  227.9    226   0.0    0.0
CI1_DFHLOG1      LIST  A  *ALL       4.2  236.0      22  197.3    208   0.0    0.0
CI1_DFHLOG2      LIST  A  *ALL       2.2  116.5      22  108.3    211   0.0    0.0
CI1_DFHSHUNT     LIST  A  *ALL       0.0    0.0       0    0.0      0   0.0    0.0
CI1_DFHSHUNT1    LIST  A  *ALL       0.0    0.0       0    0.0      0   0.0    0.0
CI1_DFHSHUNT2    LIST  A  *ALL       0.0    0.0       0    0.0      0   0.0    0.0
etc.
```

# CF Structure Specific Diagnosis
## XCF messages

- IXC512I “CFRM policy change pending”
- IXC521I “rebuild starting / rebuild complete”
- IXC522I “stopping rebuild of structure”
- IXC538I “duplexing rebuild not started (or stopped)”
- IXC552I “duplexed structures not failure isolated from each other”
- IXC553E “duplexed structures not failure isolated from each other”
- IXC573I “system managed rebuild failed”
- IXC585E “structure exceeds full threshold”
- IXC586I “structure below full threshold”
- IXC588I “altering structure”
- IXC589I “altering ended”
- IXC590I “auto alter completed, indicates if targets obtained or not”

# Problem Taxonomy

- Dead System
- Sick System
- Sysplex Fabric
- Sysplex Componentry
  - Coupling Facility
  - **Signalling Service**
  - Couple Data Sets
  - External Time Reference
- Configuration / Capacity
- Software Issues

# Signalling Overview

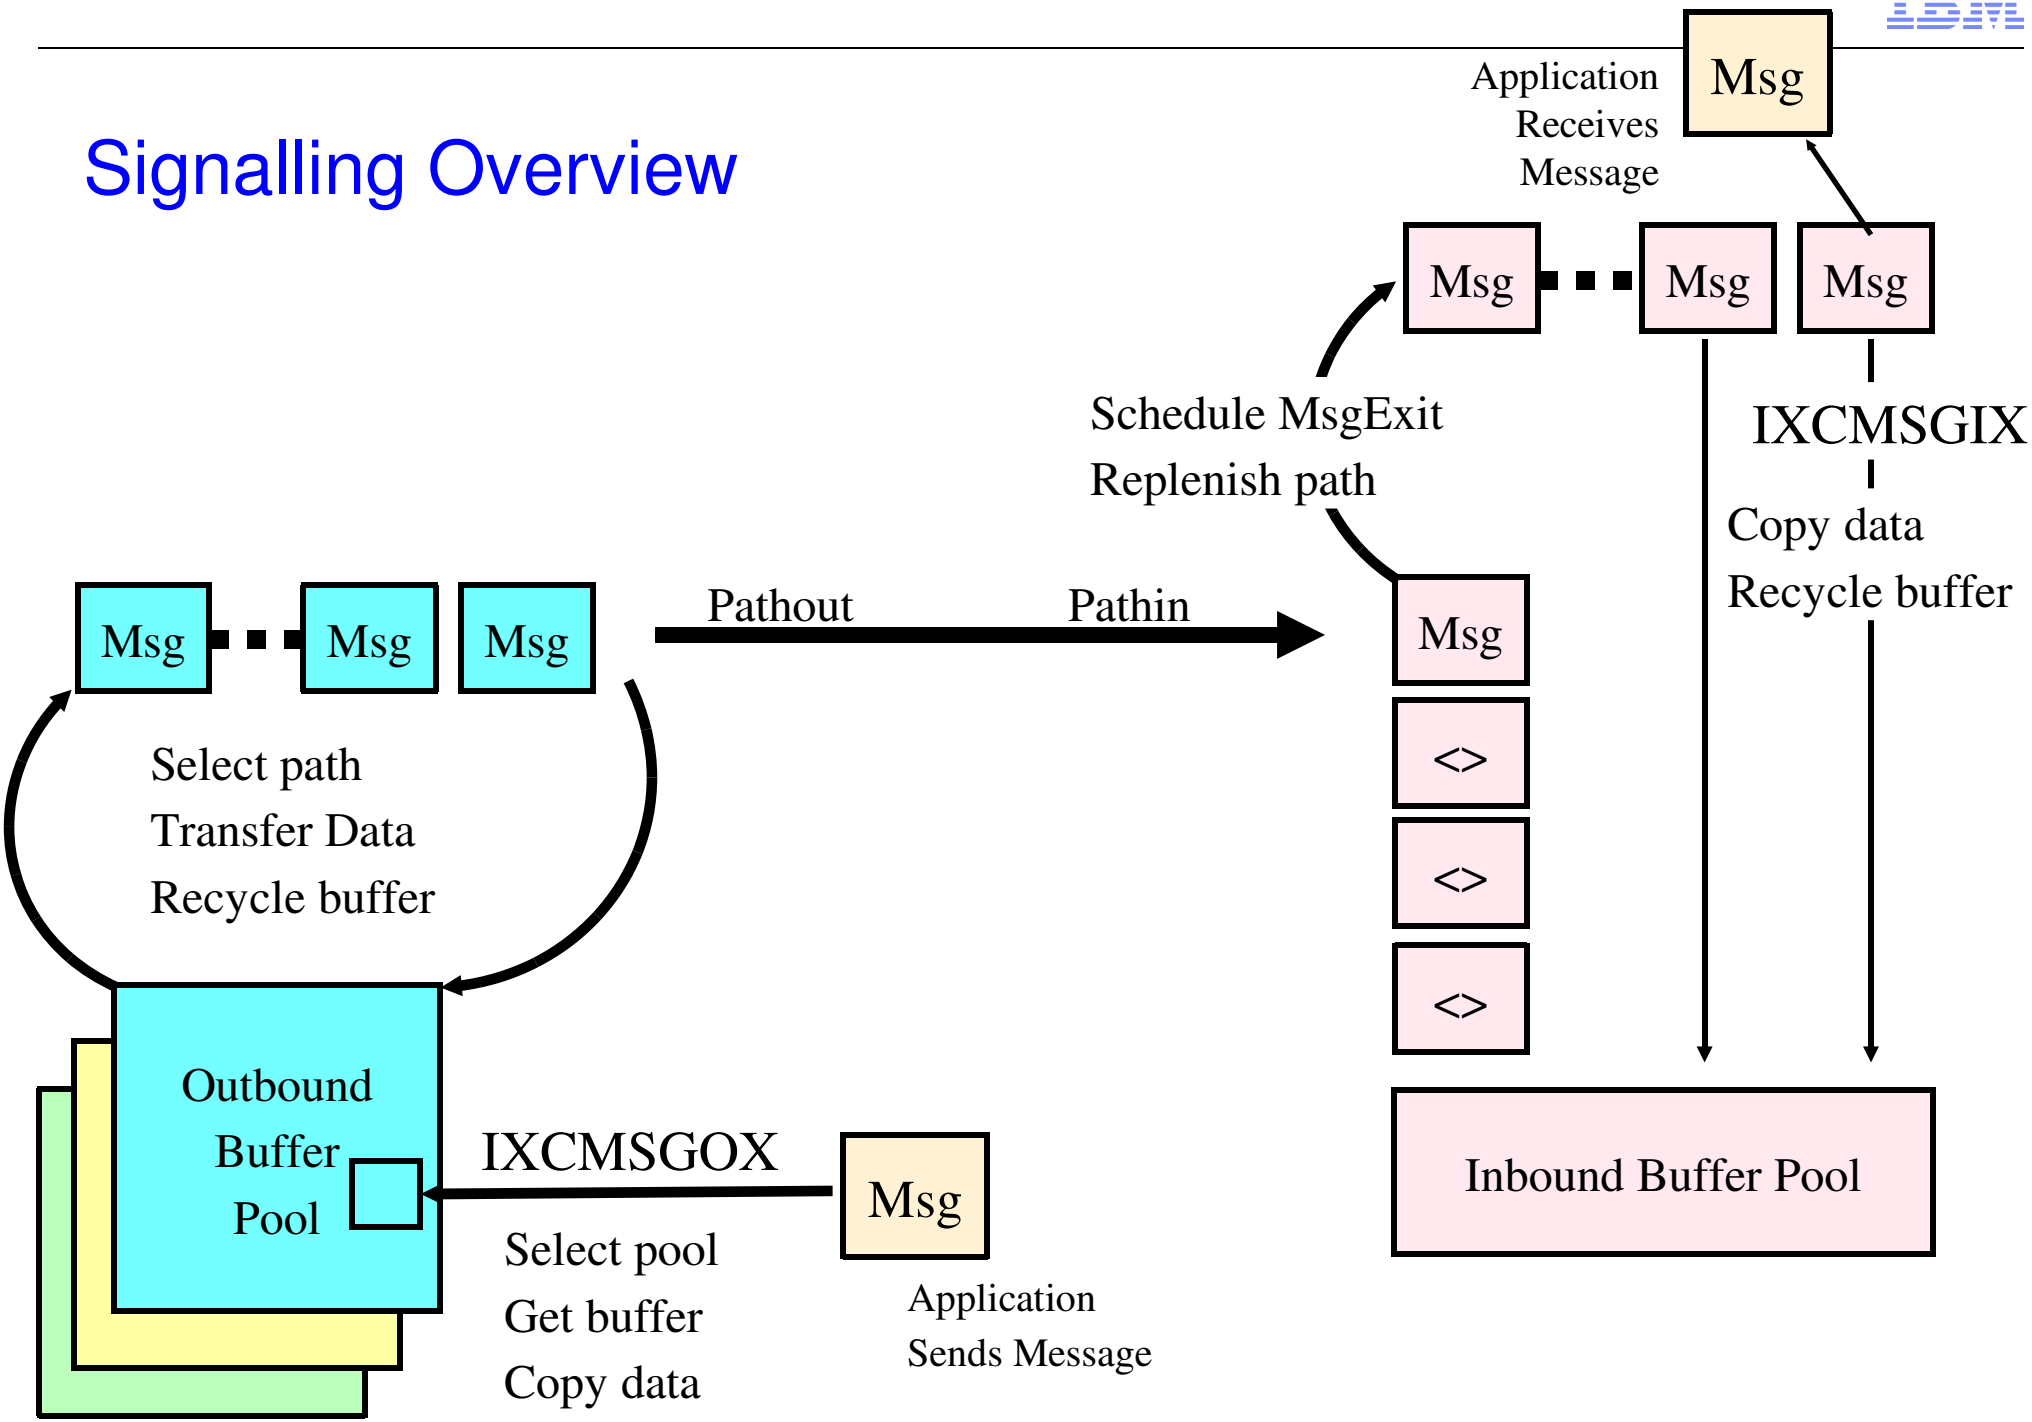

# Signalling Overview

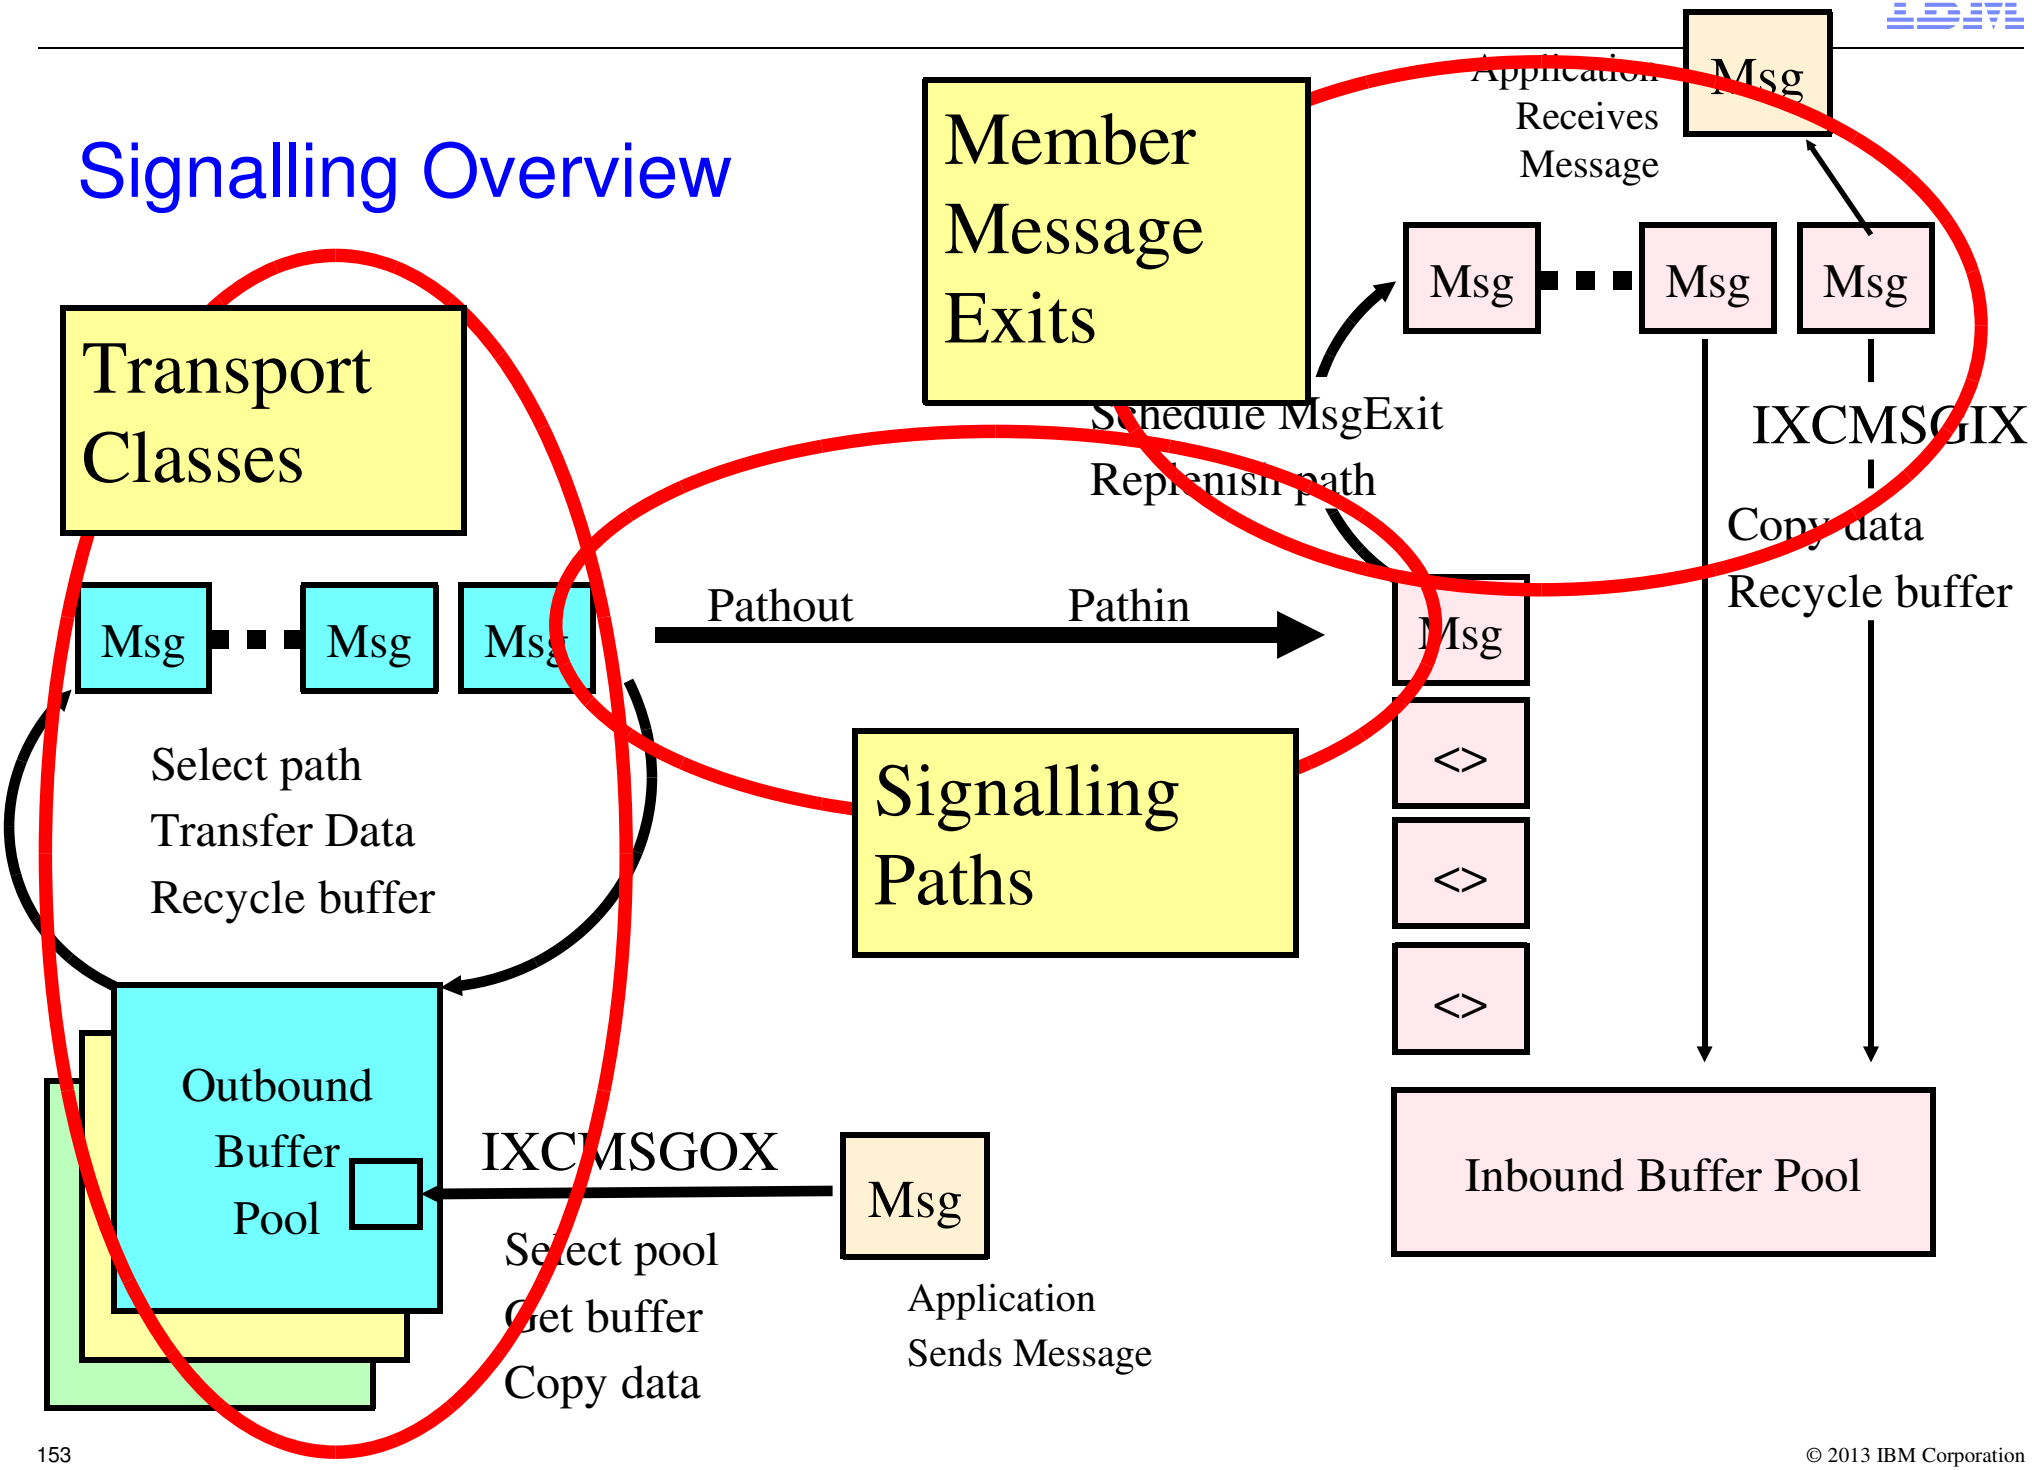

# Message Exit Concerns: Application Perspective

- Responsiveness
- Throughput

- Many applications pass messages to tasks for processing
  - Message delivery to the message exit may be timely
  - But the application can suffer if the tasks are lagging
  - XCF cannot detect such delays

  - Might the volume of message exit SRBs be starving the tasks?
  - Local lock or latch contention between the SRBs and the tasks inducing delays for one or the other?

# Message Exit Concerns: Sysplex Perspective

- Member delays/bursts can impede signal transfers, which can induce delays for other applications
- Long queue effects
  - Storage and buffer consumption
  - Processing time to run the queue

# Message Exit Delays

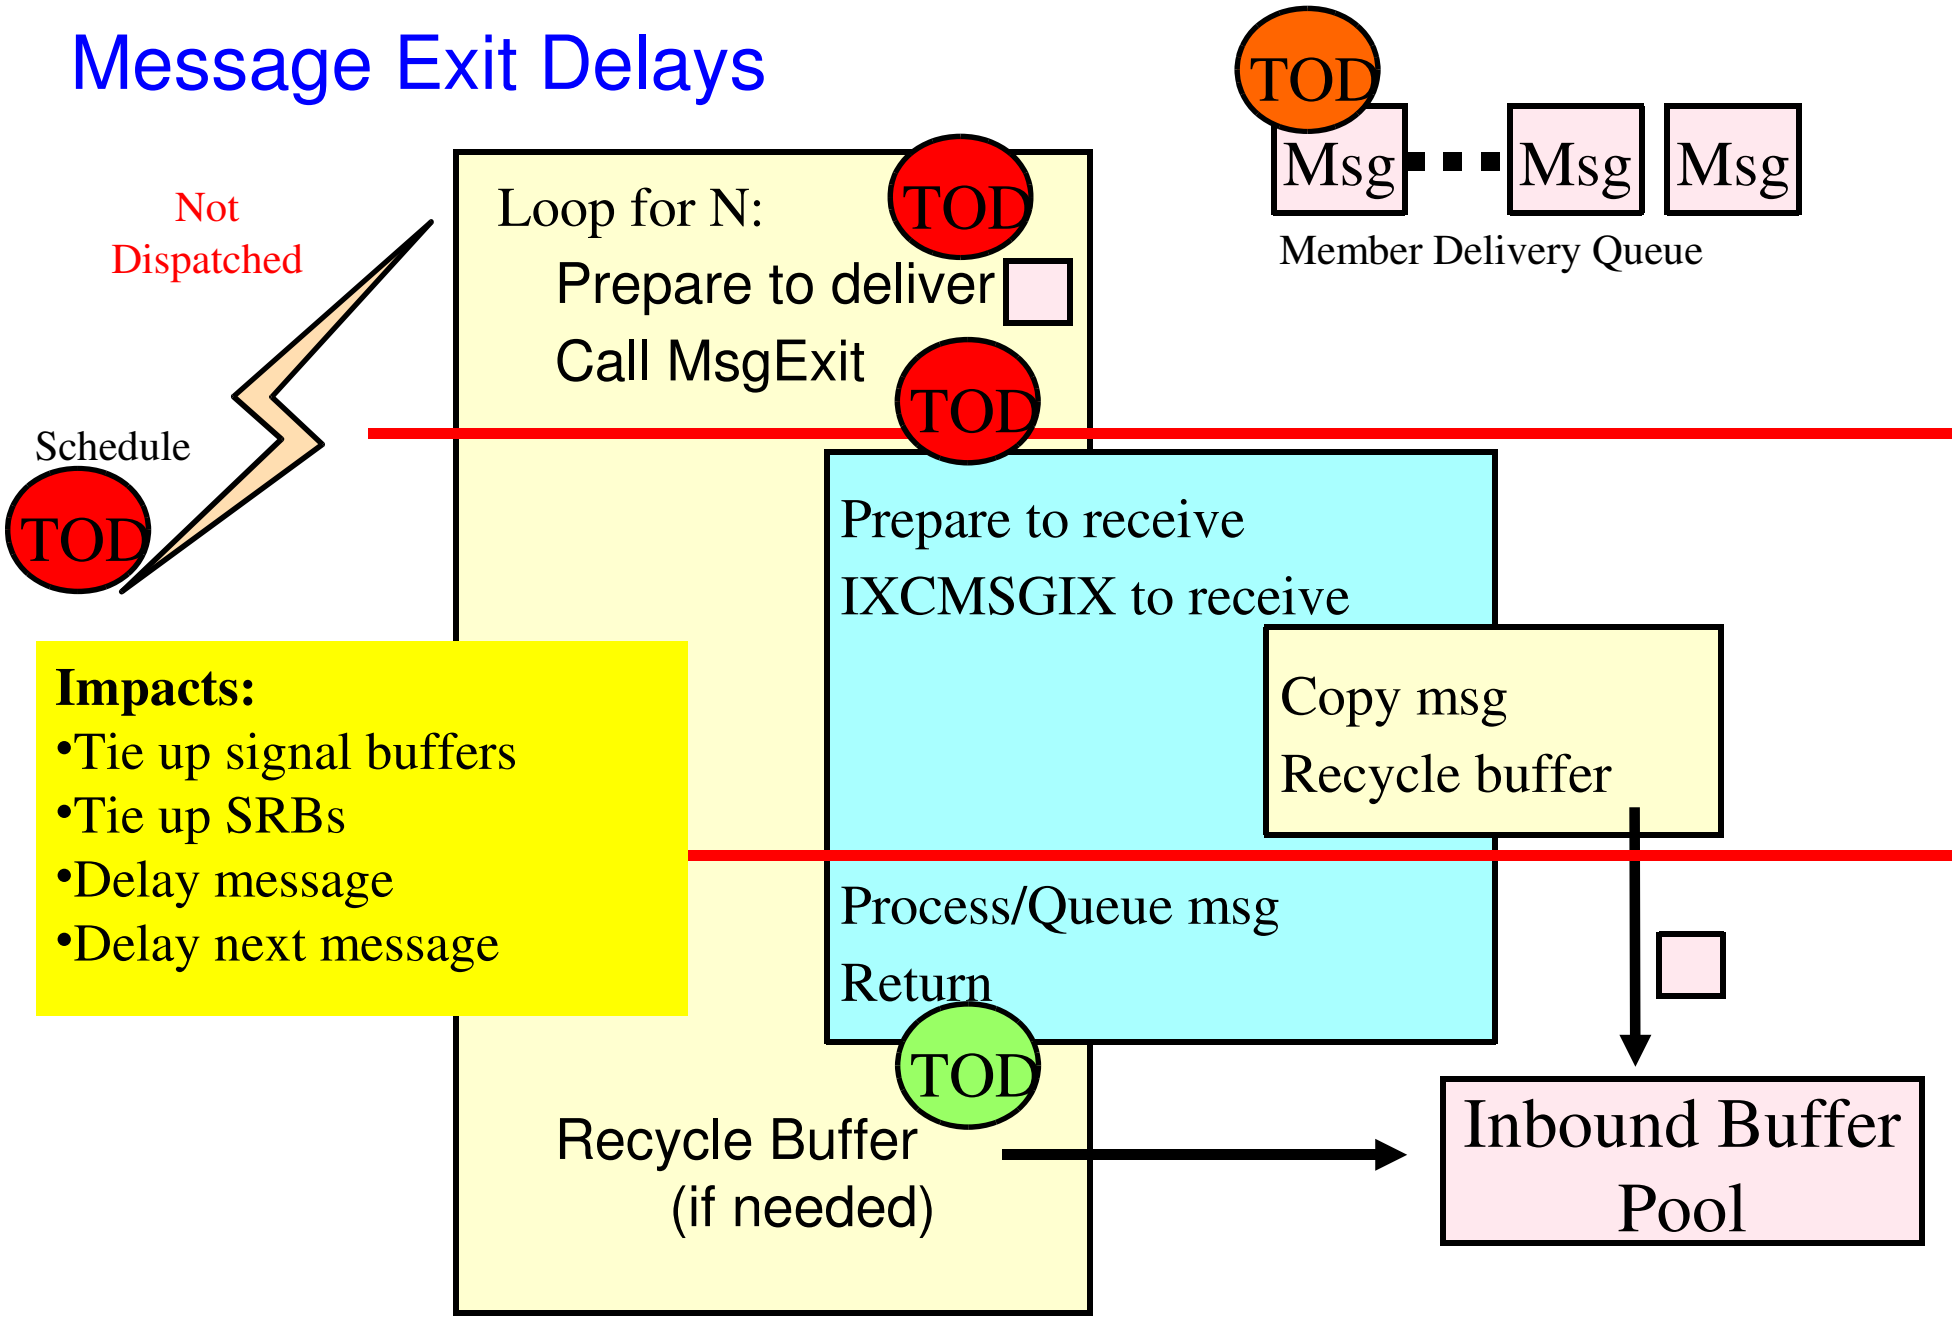


**Impacts:**

- Tie up signal buffers
- Tie up SRBs
- Delay message
- Delay next message

# Detecting Member Message Exit Issues

- XCF member stalled messages (see handout)
- D XCF,GROUP (see handout)
- IPCS COUPLE SIGNAL DETAIL report (see handout)

*These are not perfect in that they take a relatively long time to declare a stall condition.*

*We seem to be encountering more and more cases where short duration stalls impact the sysplex. Very hard to see.*

# Signalling Sympathy Sickness Indicators

| Impacted System | Culprit System |
|---|---|
| ▪ D XCF,G… shows stalls<br>▪ IXC467I Restart stalled I/O<br>Stalled Members | ▪ D XCF,G… shows stalls<br>▪ IXC431I member stalled<br>▪ ABEND 00C 020F0006<br>▪ IXC430E stalled members |
| ▪ IXC440E impacted<br>Sympathy Sickness | ▪ IXC631I member causing SS<br>▪ IXC640E if/when to act<br>▪ ABEND 00C 020F000C |
| If SFM allowed to take action → | ▪ ABEND 00C 020F000D<br>▪ IXC615I terminating<br>• ABEND 00C 00000160<br>• Wait State 0A2 rsn 160 |

# Timeline for stalled member messages

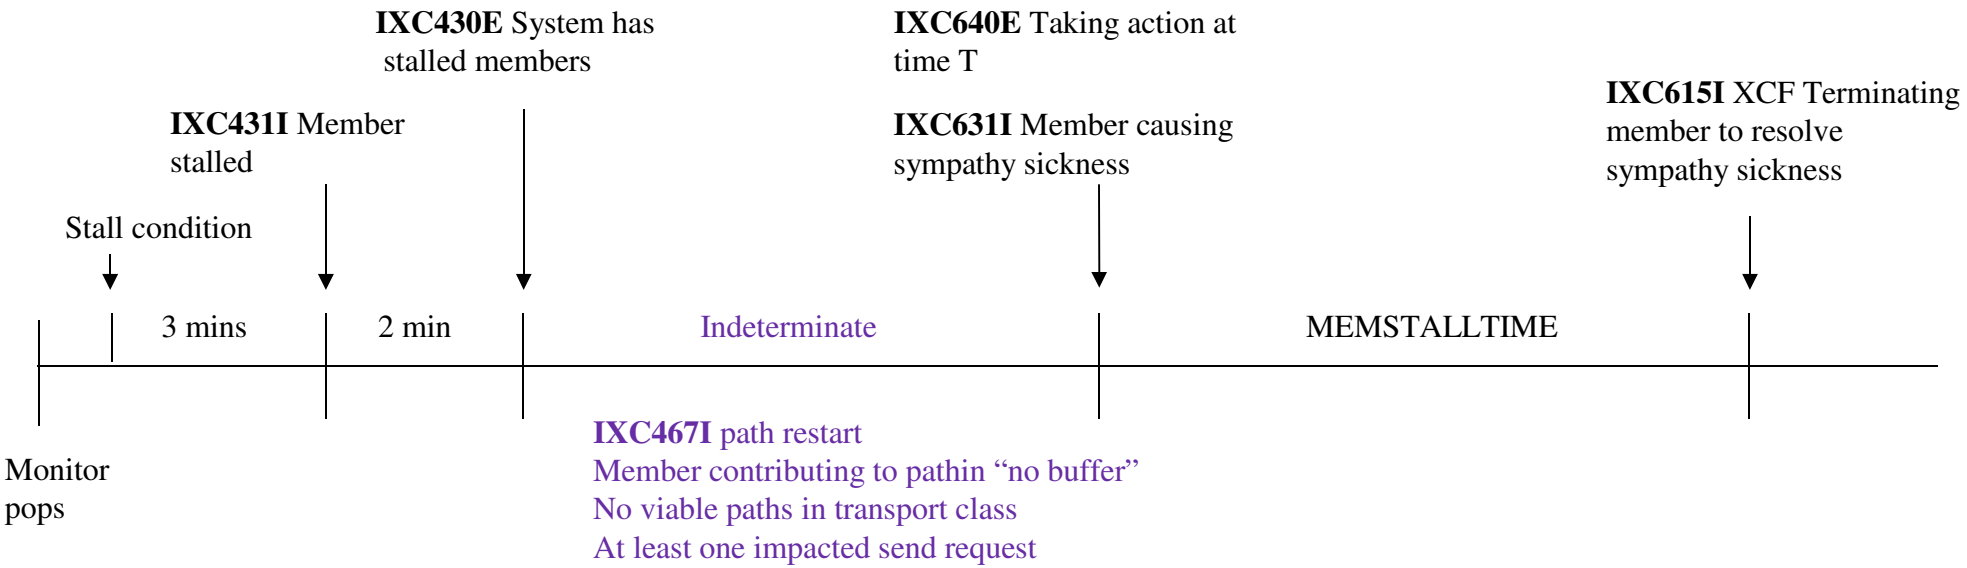


Stall condition = At least 1 exit stalled for 30 seconds or work item on head of queue for 30 seconds
Time T = "now" + MEMSTALLTIME

# Stalled Member Drilling

```
D XCF,G

IXC331I  09.43.27  DISPLAY XCF 297
   * INDICATES PROBLEM, ! INDICATES SEVERE PROBLEM
   GROUPS(SIZE):  *A0000002(2)      COFVLFNO(2)       CTTXNGRP(2)
                  EZBTCPCS(2)       ISTCFS01(2)       ISTXCF(2)
                  IXCLO02A(2)       IXCLO02B(2)       SYSCNZMG(2)
                  SYSDAE(4)         SYSENF(2)         SYSGRS(2)
                  SYSGRS2(1)        SYSIEFTS(2)       SYSIGW00(3)
                  SYSIGW01(3)       SYSIKJBC(2)       SYSIOSPX(2)
                  SYSIOS01(1)       SYSIOS02(1)       SYSJES(3)
                  SYSJ2$XD(3)       SYSMCS(7)         SYSMCS2(43)
                  SYSTTRC(2)        SYSWLM(3)         SYSXCF(2)
```

# Stalled Member Drilling …

```
D XCF,G,A0000002

IXC332I  09.43.27  DISPLAY XCF 300
 * INDICATES PROBLEM, ! INDICATES SEVERE PROBLEM
 GROUP A0000002:   *MEMBER1     MEMBER2
```

# Stalled Member Drilling …

```
D XCF,G,A0000002,MEMBER1

IXC333I  09.43.27  DISPLAY XCF 303
   INFORMATION FOR GROUP A0000002
   * INDICATES PROBLEM, ! INDICATES SEVERE PROBLEM
   MEMBER NAME:        SYSTEM:     JOB ID:    STATUS:
  *MEMBER1             SY1         XCAC3Z07   ACTIVE

 INFO FOR GROUP A0000002 MEMBER MEMBER1 ON SYSTEM SY1
 * INDICATES STALLS

 FUNCTION: TESTCASE XCJC3Z07
 MEMTOKEN: 01000064 001E0001      ASID: 0025       SYSID: 0100000B
    INFO: CURRENT             COLLECTED: 09/10/2010 09:47:03.876399

 ATTRIBUTES                      JOINED: 09/10/2010 09:42:53.741326
   JOIN TASK ASSOCIATION
   CRITICAL MEMBER
   LOCAL CLEANUP NOT NEEDED
   TERMLEVEL IS TASK
   MEMSTALL RESOLUTION IS JOIN TASK TERMINATION AFTER 603 SECONDS
   EXITS DEFINED: MESSAGE, GROUP
```

# Stalled Member Drilling …

```
SIGNALLING SERVICE
 MSGO ACCEPTED:          23  NOBUFFER:           0
 MSGO XFER CNT:           0   LCL CNT:          23  BUFF LEN:    956

                  SENDPND  RESPPND  COMPLTD  MOSAVED  MISAVED
 MESSAGE TABLE:         0        0        0        0        0

 MSGI RECEIVED:          23  PENDINGQ:           6
 MSGI XFER CNT:           0  XFERTIME:         N/A

                  IO BUFFERS       DREF   PAGEABLE   CRITICAL
 MSGI PENDINGQ:            0          0          5          0
 SYMPATHY SICK:            0

*ITEM 02053020: 09/10/2010 09:42:53.787346 ME SEQ:          19
*ITEM 02116050: 09/10/2010 09:46:27.096268 ME SEQ:          24
*EXIT 01F8CF00: 09/10/2010 09:42:53.784022 ME RUNNING
*EXIT 01F8D100: 09/10/2010 09:42:53.786051 ME RUNNING
*EXIT 01F8D300: 09/10/2010 09:43:26.842027 ME RUNNING
*EXIT 01F8D500: 09/10/2010 09:43:59.885911 ME RUNNING
 EXIT 01F8D700: 09/10/2010 09:46:28.024502 MV 00:00:00.000061
*EXIT 02003100: 09/10/2010 09:42:53.767349 ME RUNNING
*EXIT 02003500: 09/10/2010 09:42:53.769113 ME RUNNING
```

Delta
Pending
Preparing
Running

# IPCS COUPLE SIGNAL DETAIL as well

**Detail of signal exit SRBs for group:** SYSGRS member: SYSK

| SRB Addr | TOD When Service Called | FC | Duration | State |
|---|---|---|---|---|
| 08703600 | 06/17/2011 10:39:21.144506 | OM | *00:02:44.229274 | Running |
| 08704400 | 06/17/2011 10:39:21.147527 | OM | *00:02:44.226253 | Running |
| 089E8900 | 06/17/2011 10:39:21.148002 | NA | *00:02:44.225778 | Pending |
| 08705200 | 06/17/2011 10:39:21.148003 | NA | *00:02:44.225777 | Pending |

**Detail of work items queued for group:** SYSGRS member: SYSK

| ItemAddr | TOD When Queued | FC | Duration as of 10:42:05.373780 | ItemSeq# |
|---|---|---|---|---|
| 02799110 | 06/17/2011 10:39:21.144763 | OM | *00:02:44.229017 | 00000538 |
| 02BB7038 | 06/17/2011 10:39:21.145037 | OM | *00:02:44.228743 | 00000539 |
| 02BB7138 | 06/17/2011 10:39:21.145308 | OM | *00:02:44.228472 | 0000053A |
| 0299D0E8 | 06/17/2011 10:39:21.145768 | OM | *00:02:44.228012 | 0000053B |
| 0299D1E8 | 06/17/2011 10:39:21.146210 | OM | *00:02:44.227570 | 0000053C |

# Typical Causes of Message Exit Delays

- CPU Constraints
  - Higher priority work winning
  - Insufficient LPAR weight
- Storage Constraints
  - Page faults
  - Exhausted private
  - Fragmentation
- Contention
  - Waiting for ENQ, local lock, or latch
- Signal Volume
- Defects

Without a timely dump, virtually impossible to diagnose stalls, particularly those of short duration

Try Run Time Diagnostics

# Signalling Overview

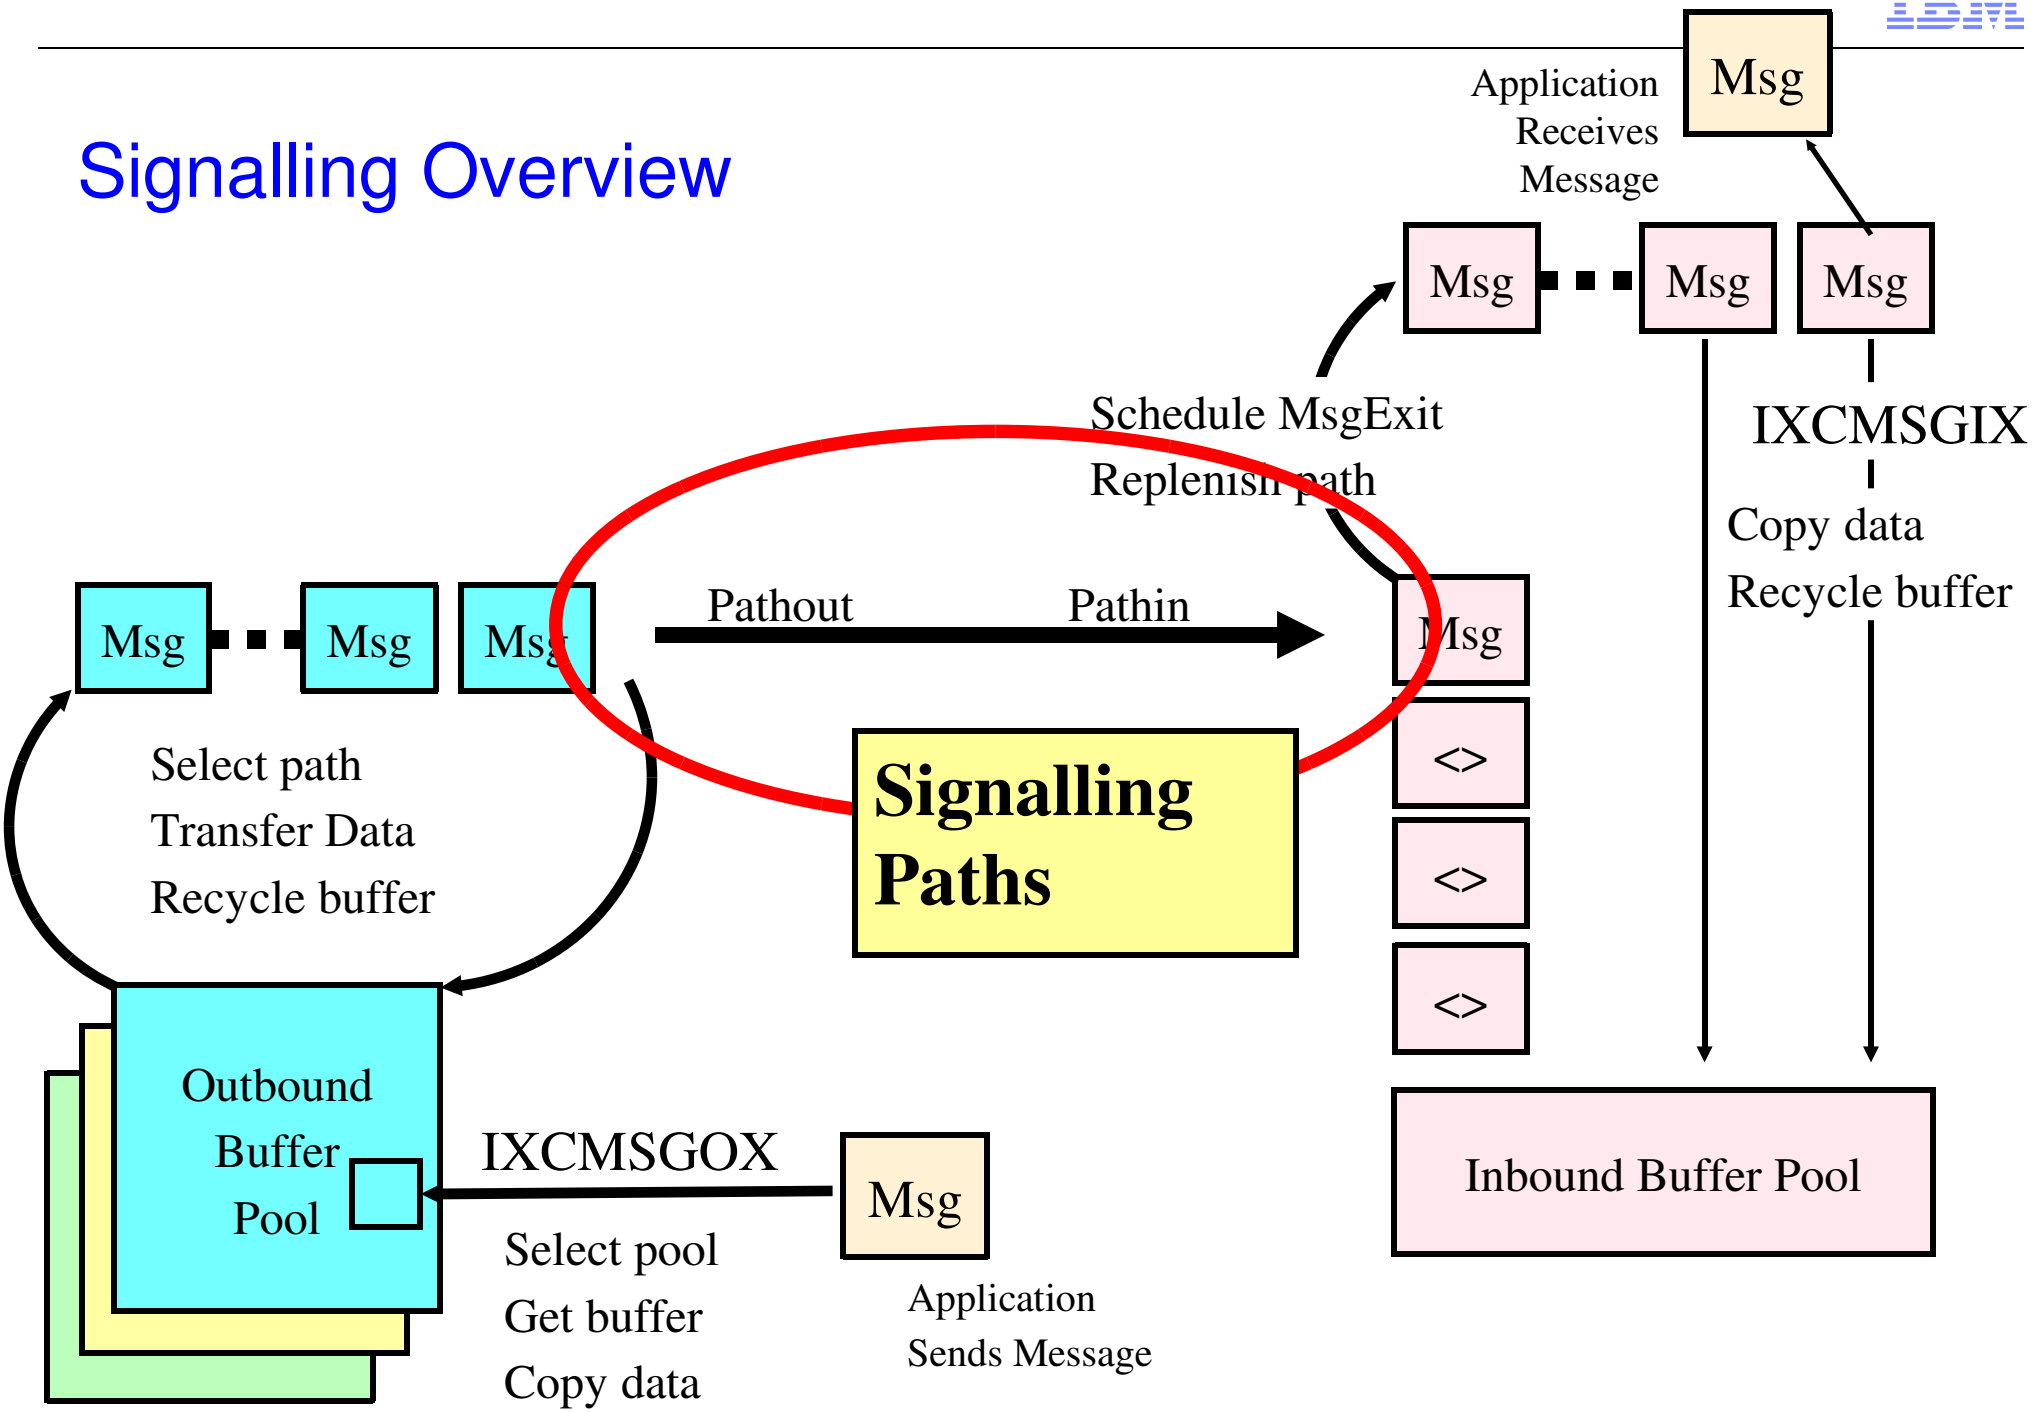

# Signal Path Concerns

- Target system operational?
- Path operational?
  - Restart? Stop? Rebuild?
- Inbound “no buffer” conditions
- Transfer time

# Signal Path Concerns: Inbound Buffers

- Stalled members consuming buffers?
- Lack of Transport Class Segregation on sending side?
  - *Every transport class must have a path for every target system*
  - *Signals may not be flowing via the expected paths*
- Transport Class Length relative to PATHIN MAXMSG?
  - *Number of buffers for given MAXMSG decreases as class length increases*
  - *If you define/modify class length, you may need to modify PATHIN MAXMSG to maintain number of buffers*
- PATHIN MAXMSG too small?
  - *May need more buffers if they tend to dwell in member message delivery*

The z/OS Health Checker looks for these problems

# Transport Class Specifications

- COUPLE
  - MAXMSG
- CLASSDEF(s)
  - CLASSLEN
  - GROUP
- PATHOUT
  - CLASS
  - How many paths?

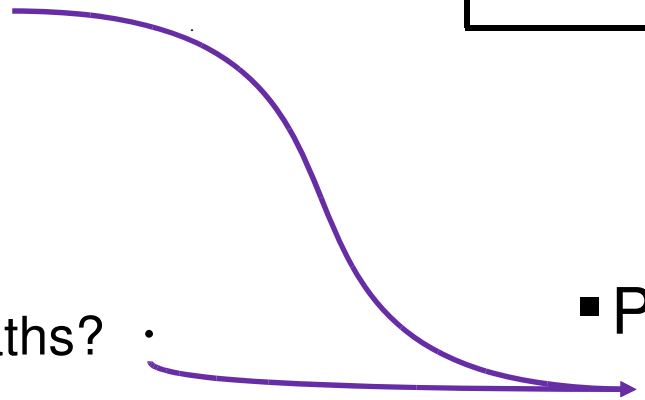

- PATHIN
  - MAXMSG

You must consider impact to target as well

**Sending System**

**Target System(s)**

# Transport Class Must Have Paths

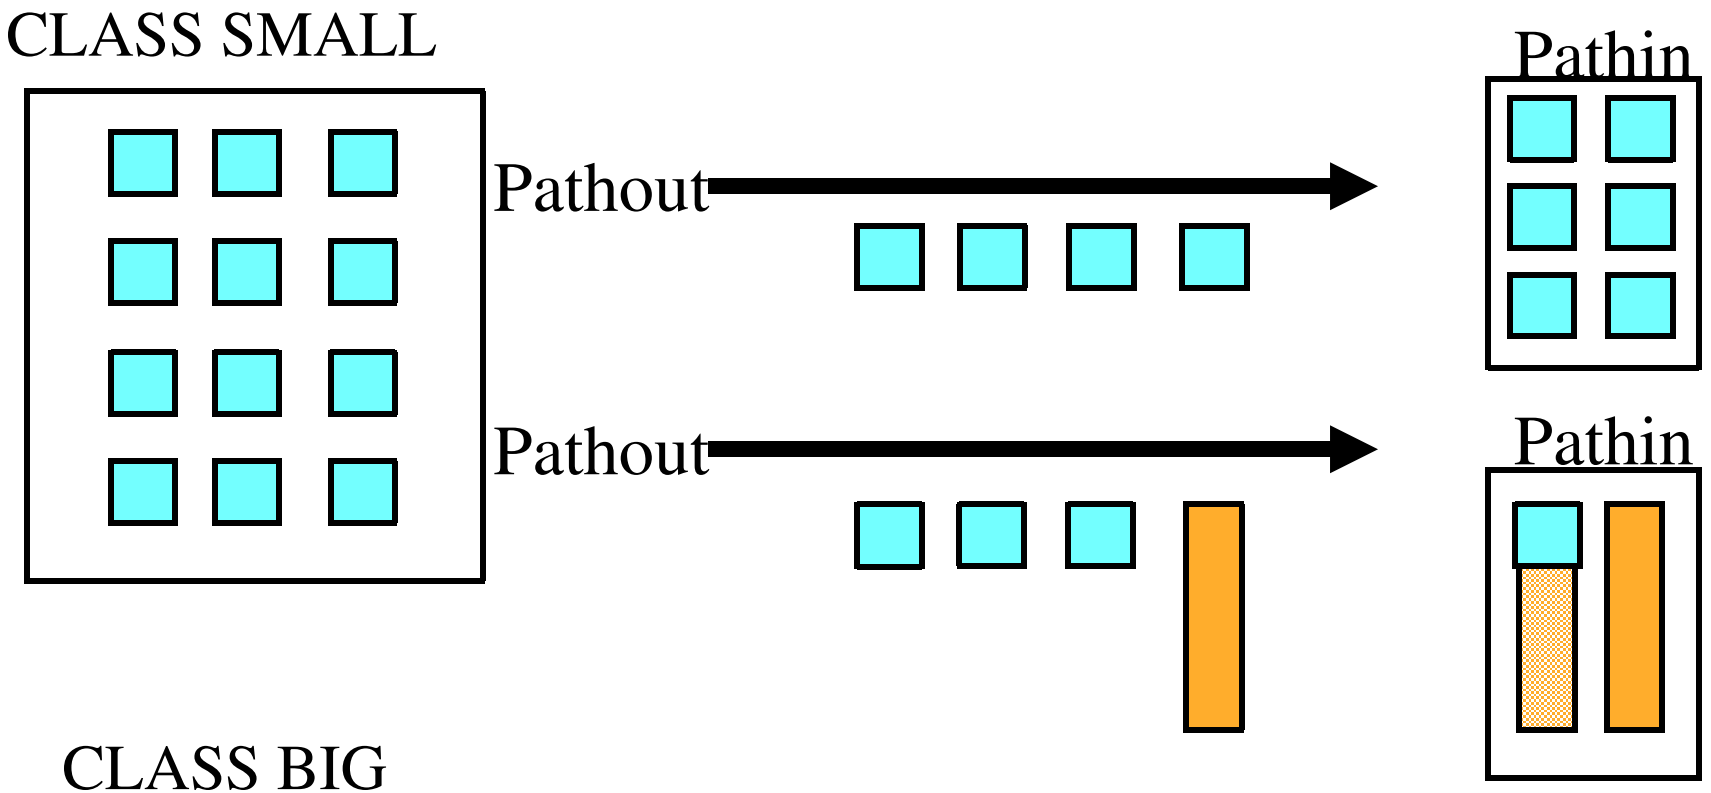


Big msg sent via paths in small class

Pathin buffer pool resized for larger msgs

So has fewer buffers for given MAXMSG

So pathin may get nobuff condition

So both big and small may get delays

# Detecting Signal Path Issues

- D XCF,PI,DEV=xxxx
- D XCF,PI,STRNAME=xxxx
- RMF Reports of XCF Activity
  - Path Statistics (inbound paths)

- D XCF,PO,DEV=xxxx
- D XCF,PO,STRNAME=xxxx
- D XCF,CLASSDEF
- RMF Reports of XCF Activity
  - Usage by System (outbound side)

*Primary concern is "no buffer" conditions on the inbound side of the path. Signals will not flow if the inbound side has no buffers with which to receive signals.*

# XCF Pathin Display

TClass? MAXMSG?

```
D XCF,PI,DEVICE=ALL

IXC356I  09.18.51  DISPLAY XCF
```

| LOCAL PATHIN | REMOTE PATHOUT | REMOTE SYSTEM | PATHIN STATUS | DELIVRY PENDING | BUFFER LENGTH | MSGBUF IN USE | SIGNL NUMBR | NOBUF |
|---|---|---|---|---|---|---|---|---|
| 2101 | 2401 | SYSA | WORKING | 4 | 956 | 10 | 76501 | 0 |
| 2201 | 2401 | SYSB | WORKING | 4 | 4028 | 30 | 95852 | 0 |
| 2301 | 2401 | SYSC | WORKING | 4 | 8124 | 52 | 24277 | 0 |
| 3101 | 3401 | SYSA | WORKING | 4 | 956 | 10 | 44552 | 0 |
| 3201 | 3401 | SYSB | WORKING | 4 | 956 | 10 | 82508 | 1 |
| 3301 | 3401 | SYSC | WORKING | 4 | 956 | 8 | 5595 | 0 |

These NOBUF counts are lifetime for the path and sending system instance

# RMF XCF Path Statistics (inbound)

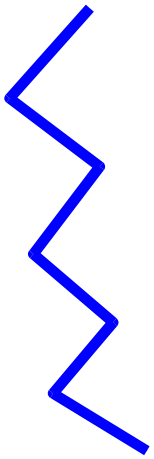

INBOUND TO SYSB

| FROM SYSTEM | T Y P | FROM/TO DEVICE, OR STRUCTURE | REQ IN | BUFFERS UNAVAIL | TRANSFER TIME |
|---|---|---|---|---|---|
| SYSA | S | IXCSTR1 | 848 | 0 | 1.369 |
| | S | IXCSTR2 | 6 | 0 | 0.585 |
| | S | IXCSTR3 | 167 | 0 | 1.042 |
| | S | IXCSTR4 | 219 | 0 | 1.671 |

These NOBUF counts are for the indicated reporting interval

# Inbound No Buffer Caveats

- CTC: wanted a buffer, but was signal in flight?
- List: wanted a buffer because signal is in flight
- Periodic retry could inflate “no buffer” count
  - Up to 200 millisecond delay, then every 10 milliseconds
  - So nonzero count indicates potential delay, but not necessarily “quantity”

# Signal Transfer Considerations

- Queue Time
- Transfer Time
- CTC vs CF Structure

*Signal transfers times are interesting in that they may indicate delays.*

*But signal transfers are generally always good unless you have other problems (stalled members, no buffer conditions, message flooding, CF service time issues, etc). If you resolve those problems, you won't have transfer time issues.*

*So the handout has plenty of details to explain the technical details. But generally there is nothing to see here. Move along..*

# CTC Path Signal Transfer

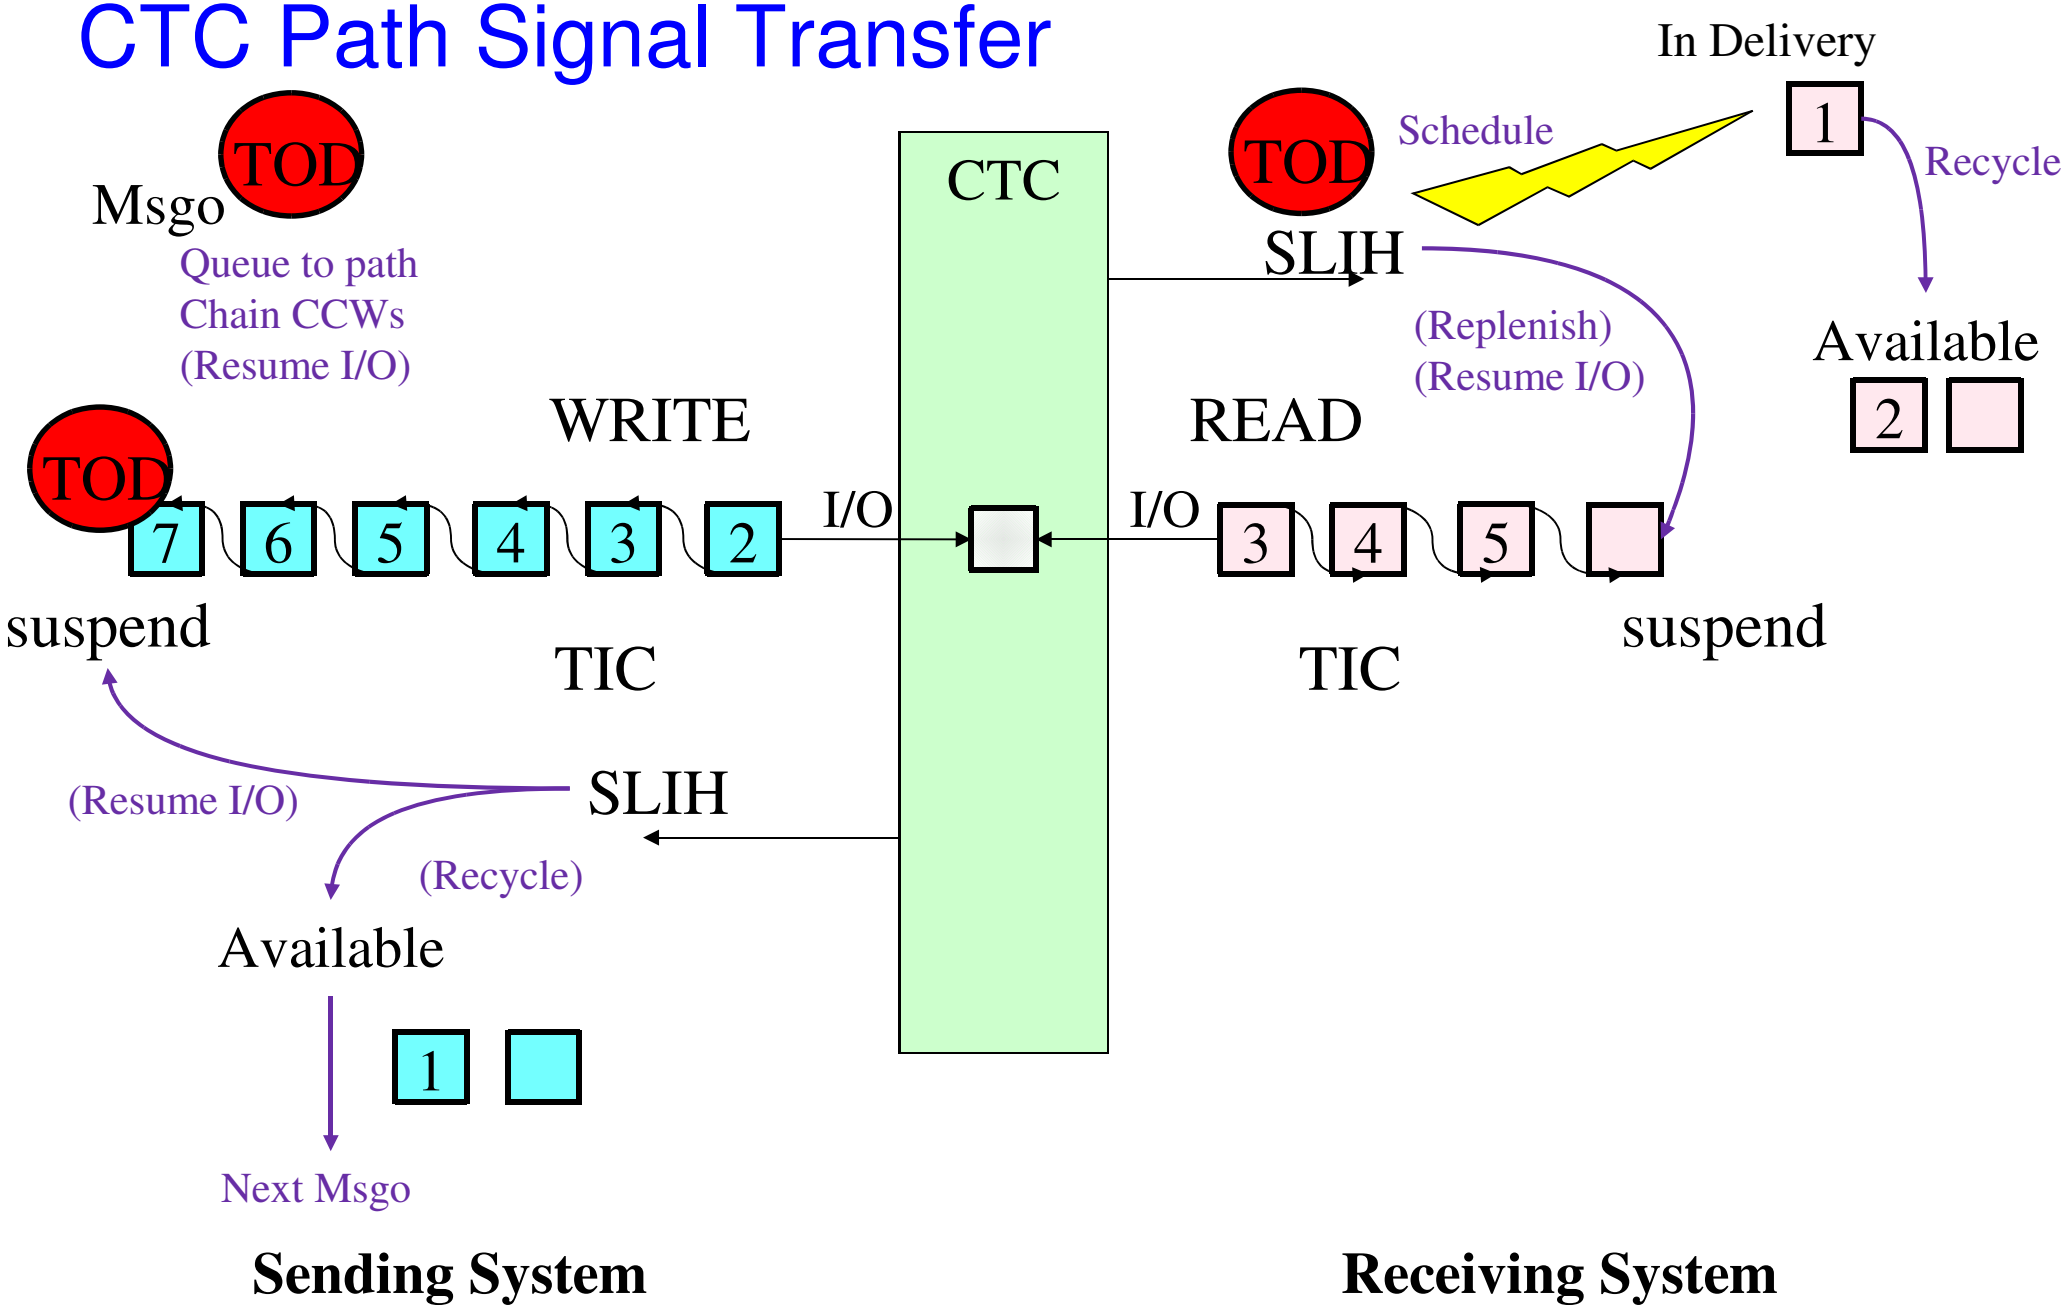

# CTC Path Timing

- Signals immediately chained for I/O, so would expect queue time to be quite small
- High queue time might be caused by:
  - Host issues between request and start I/O
  - Resend via alternate path after restart/stop of selected path
- High transfer times might be caused by:
  - Inbound buffer shortage
  - Host issues on target system
  - Device issues
  - Bursts, perhaps (long channel program chain)

# Signal Structure

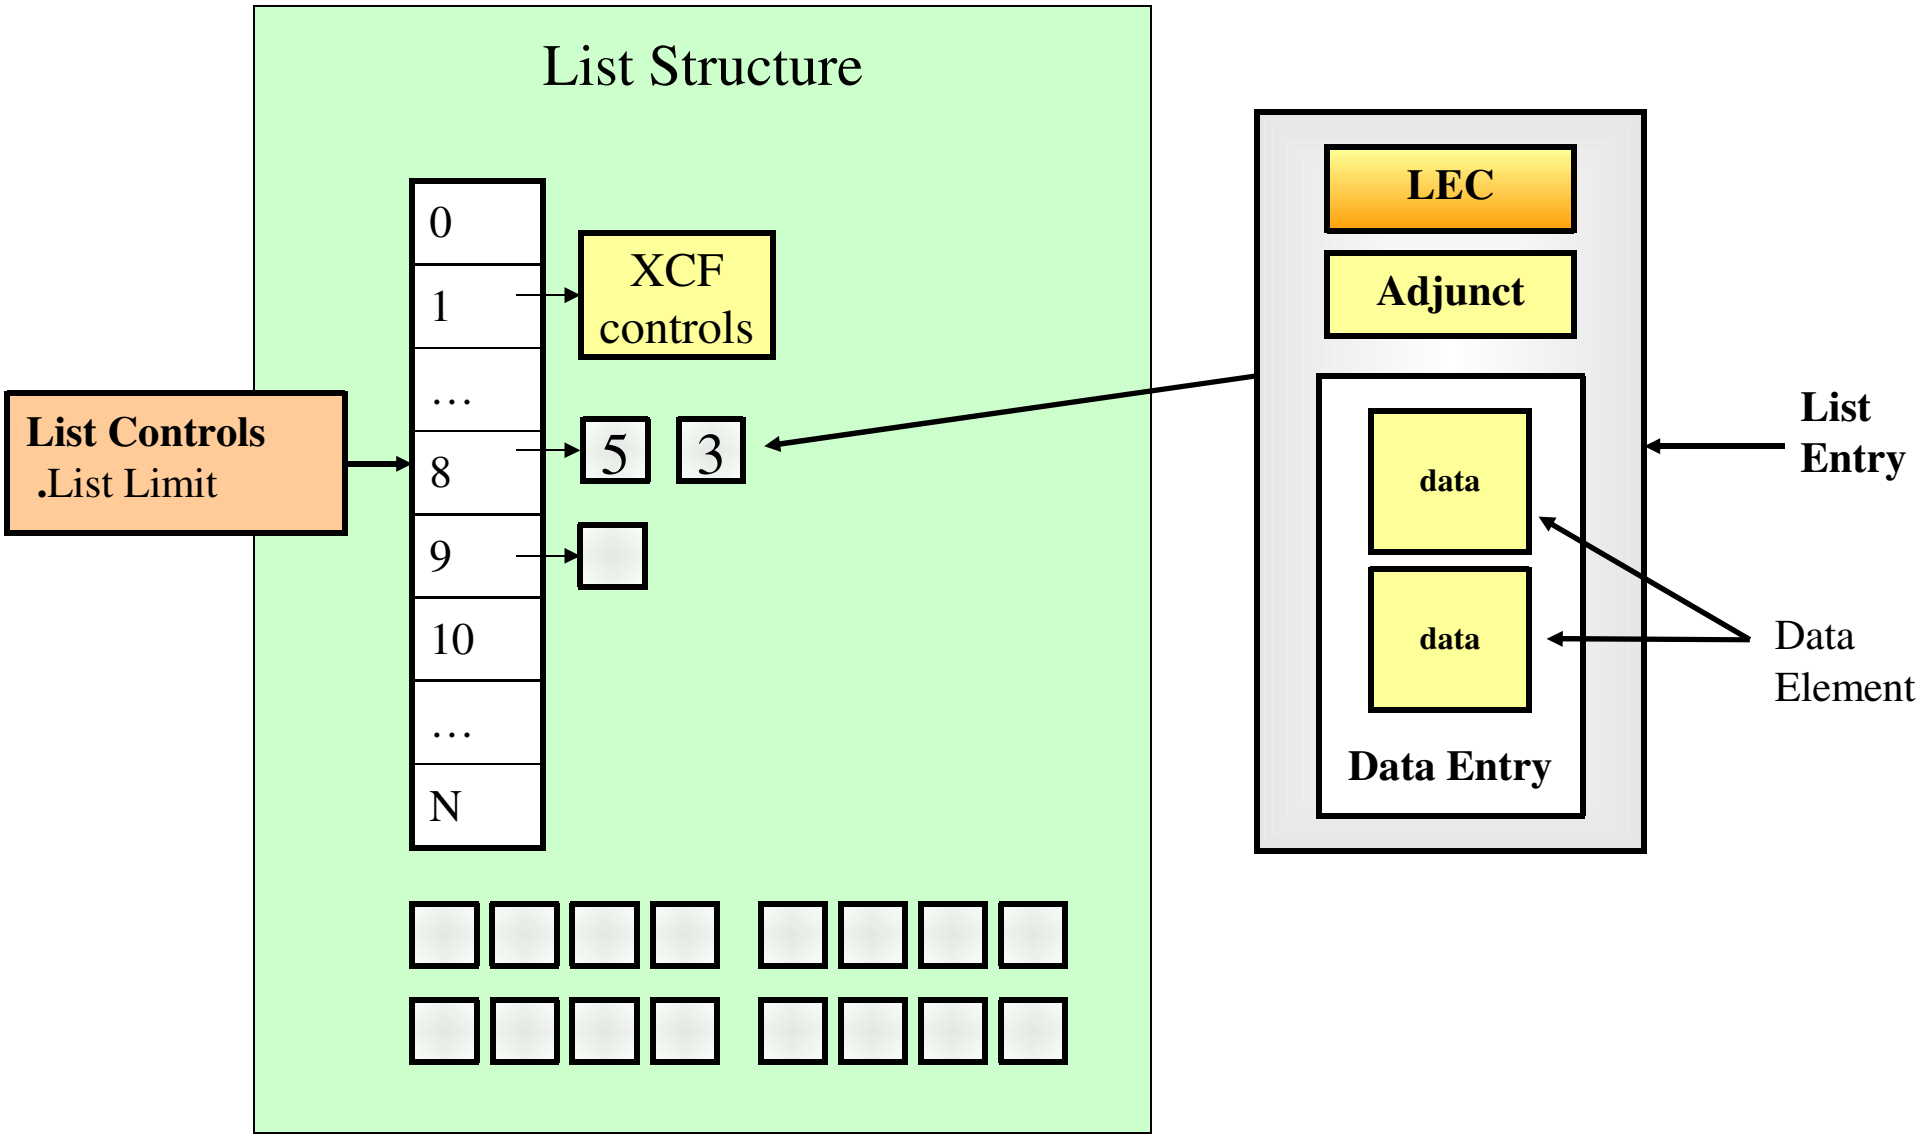

# List Path Signal Transfer

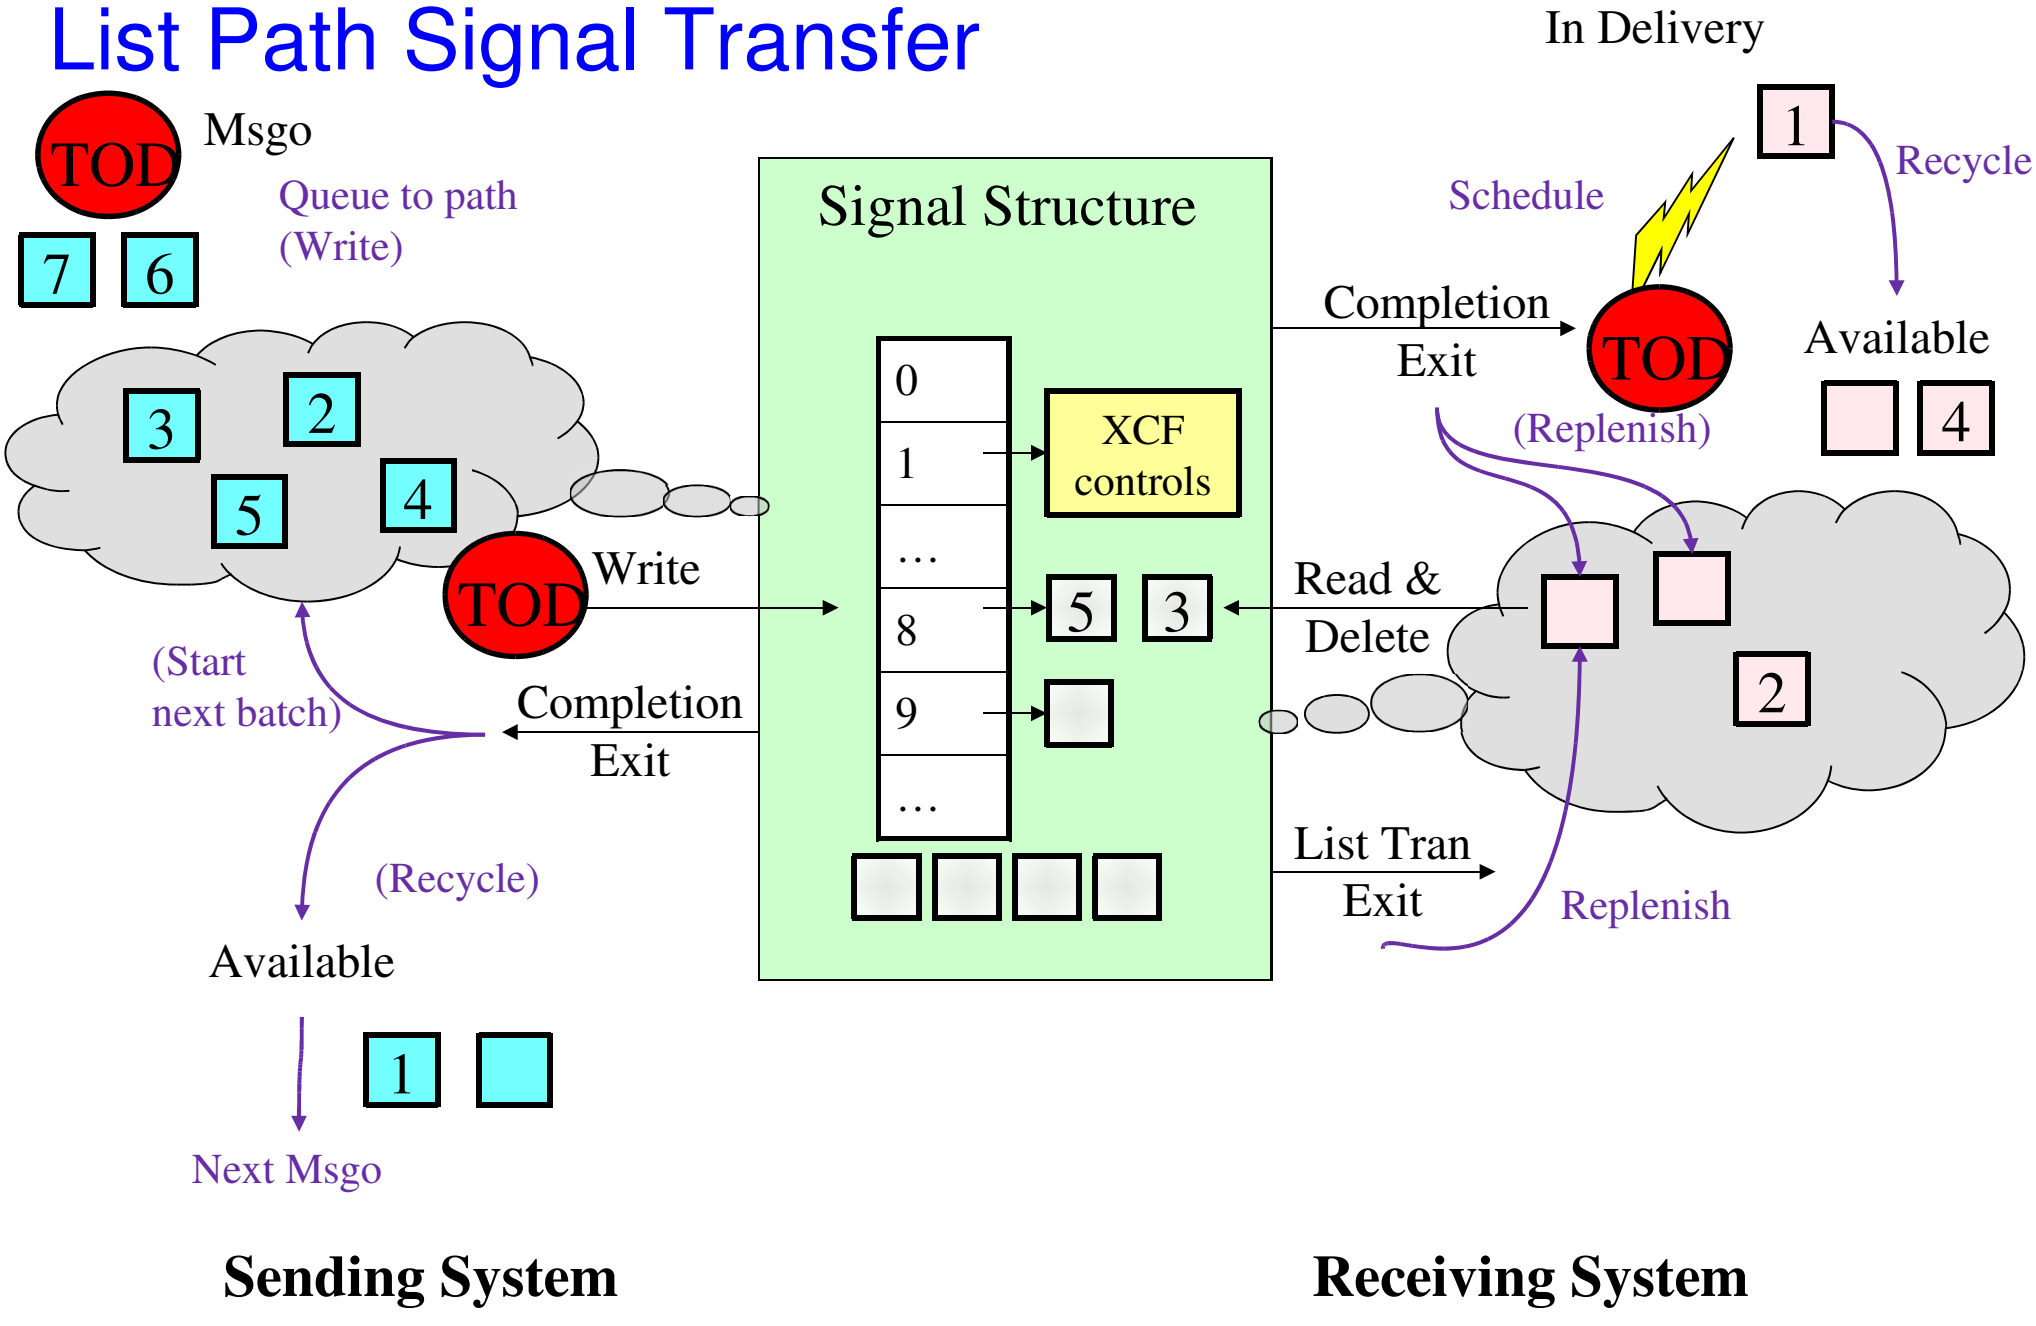

# List Path Timing

- Signal immediately written if in current batch, in which case expect small queue time (like CTC)
- High queue time might be caused by:
  - Host issues between request and start I/O
  - Host issues completing back end of current batch
  - List full conditions (CFSIZER ?)
- High transfer times might be caused by:
  - Inbound buffer shortage
  - Host issues on target system
  - Delayed empty to non-empty notifications
  - Issues getting to the CF (subchannel/path busy)
  - Issues in the CF (shared CP, contention, …)
  - Bursts, perhaps (long queue in CF)

# Queue Time

- "start I/O" minus "requested"
  - Includes time spent copying message to I/O buffer (page fault?)
  - Could lose CPU after accepted but before queued
  - Includes time spent initiating sends to peer targets for broadcast message
  - Could include time spent waiting for "no buffer" or "no path" conditions to clear
  - Could include time spent doing path restart (or stop) protocol to figure out that signal needed to be resent

    New "started I/O" time taken when signal is redriven

# Transfer Time

- "arrived" minus "start I/O"
  - For list path, could include list transition processing
    - Driving transition exit
    - Initiating request to read signal from CF
  - Includes z/OS time to get to back end routine
    - Interrupt processing for CTC path
    - Async command completion for list path
  - Includes XCF time to verify signal arrived intact

# Signal Timings

## Timings for Recent Signal Transfers

### CTC device

| TOD When MSGO Requested | Queue Time | Transfer Time | TOD When Arrived |
|---|---|---|---|
| 06/17/2011 10:35:40.031364 | 00:00:00.000001 | 00:00:00.000799 | 06/17/2011 10:35:40.032166 |
| 06/17/2011 10:35:44.229656 | 00:00:00.000000 | 00:00:00.000774 | 06/17/2011 10:35:44.230432 |
| 06/17/2011 10:35:49.477406 | 00:00:00.000000 | 00:00:00.000768 | 06/17/2011 10:35:49.478175 |
| 06/17/2011 10:35:51.576329 | 00:00:00.000002 | 00:00:00.000789 | 06/17/2011 10:35:51.577120 |
| 06/17/2011 10:35:57.873469 | 00:00:00.000000 | 00:00:00.000769 | 06/17/2011 10:35:57.874240 |

### List Path

| TOD When MSGO Requested | Queue Time | Transfer Time | TOD When Arrived |
|---|---|---|---|
| 06/17/2011 10:30:52.133303 | 00:00:00.000000 | 00:00:00.002652 | 06/17/2011 10:30:52.135956 |
| 06/17/2011 10:31:55.838311 | 00:00:00.000001 | 00:00:00.001375 | 06/17/2011 10:31:55.839687 |
| 06/17/2011 10:32:58.806705 | 00:00:00.000001 | 00:00:00.001941 | 06/17/2011 10:32:58.808647 |
| 06/17/2011 10:34:02.092010 | 00:00:00.000000 | 00:00:00.001946 | 06/17/2011 10:34:02.093958 |
| 06/17/2011 10:34:24.103905 | 00:00:00.000055 | 00:00:00.000956 | 06/17/2011 10:34:24.104917 |

Take dump. Run IPCS COUPLE SIGNAL DETAIL report

Recent timings reported for both paths and members

# Delivery Time

- "called msg exit" minus "arrived"
  - But we do not currently gather these
- Msg exit "returned" minus "called"
  - Available for residual SRBs
  - But no historical tracking

# Looking for Delay on Inbound Side

```
D XCF,PI,DEVICE=ALL
IXC356I  09.17.26  DISPLAY XCF 676

 LOCAL DEVICE     REMOTE   PATHIN      REMOTE                    LAST  MXFER
 PATHIN           SYSTEM   STATUS      PATHOUT RETRY   MAXMSG RECVD TIME
 2101             SYSA     WORKING     2301        20    2000  2076   379
 3101             SYSA     WORKING     3301        20    2000 10099   405

 LOCAL    REMOTE   REMOTE   PATHIN     DELIVRY BUFFER  MSGBUF SIGNL
 PATHIN   PATHOUT SYSTEM   STATUS     PENDING LENGTH  IN USE NUMBR NOBUF
 2101     2301     SYSA     WORKING          4    956       8  2076     22
 3101     3301     SYSA     WORKING          4    956       8 10099      0
```

STALL-IOPND

STALL-INOP

STALL-SS?

STALL-SS

# Looking for Delay on Outbound Side

```
D XCF,PO,DEVICE=ALL

IXC356I  09.17.11  DISPLAY XCF

LOCAL DEVICE     REMOTE   PATHOUT      REMOTE                    TRANSPORT
PATHOUT          SYSTEM   STATUS       PATHIN   RETRY   MAXMSG CLASS
2301             SYSC     WORKING      2101        20     2000 DEFAULT
3301             SYSC     WORKING      3101        20     2000 DEFAULT

LOCAL    REMOTE   REMOTE   PATHOUT     TRANSFR BUFFER  MSGBUF SIGNL MXFER
PATHOUT  PATHIN   SYSTEM   STATUS      PENDING LENGTH  IN USE NUMBR TIME
2301     2101     SYSC     WORKING           0    956      10  1642   357
3301     3101     SYSC     WORKING           0    956      10 10041   315
```

STALL-IOPND

STALL-INOP

STALL-SS?

STALL-SS

# Sanity Check Structure Size and Use

## From RMF Coupling Facility Usage Summary Report

| TYPE | STRUCTURE NAME | STATUS | CHG | ALLOC SIZE | LST/DIR ENTRIES TOT/CUR | DATA ELEMENTS TOT/CUR |
|---|---|---|---|---|---|---|
| LIST | DSNZPLEX_SCA | ACTIVE | | 33M | 38K | 76K |
| | | | | | 118 | 614 |
| | IXCSTR1 | ACTIVE | | 31M | 5255 | 5229 |
| | | | | | 1 | 18 |
| | IXCSTR3 | ACTIVE | | 31M | 5255 | 5229 |
| | | | | | 1 | 26 |

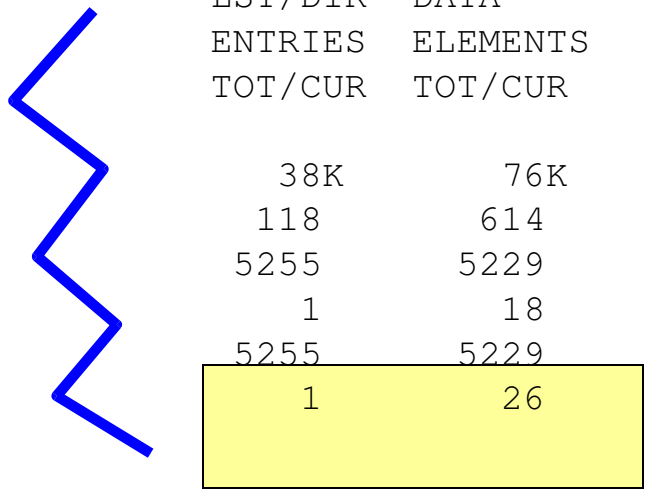

Expect mostly empty.
Potential for list full conditions?

# Signalling Overview

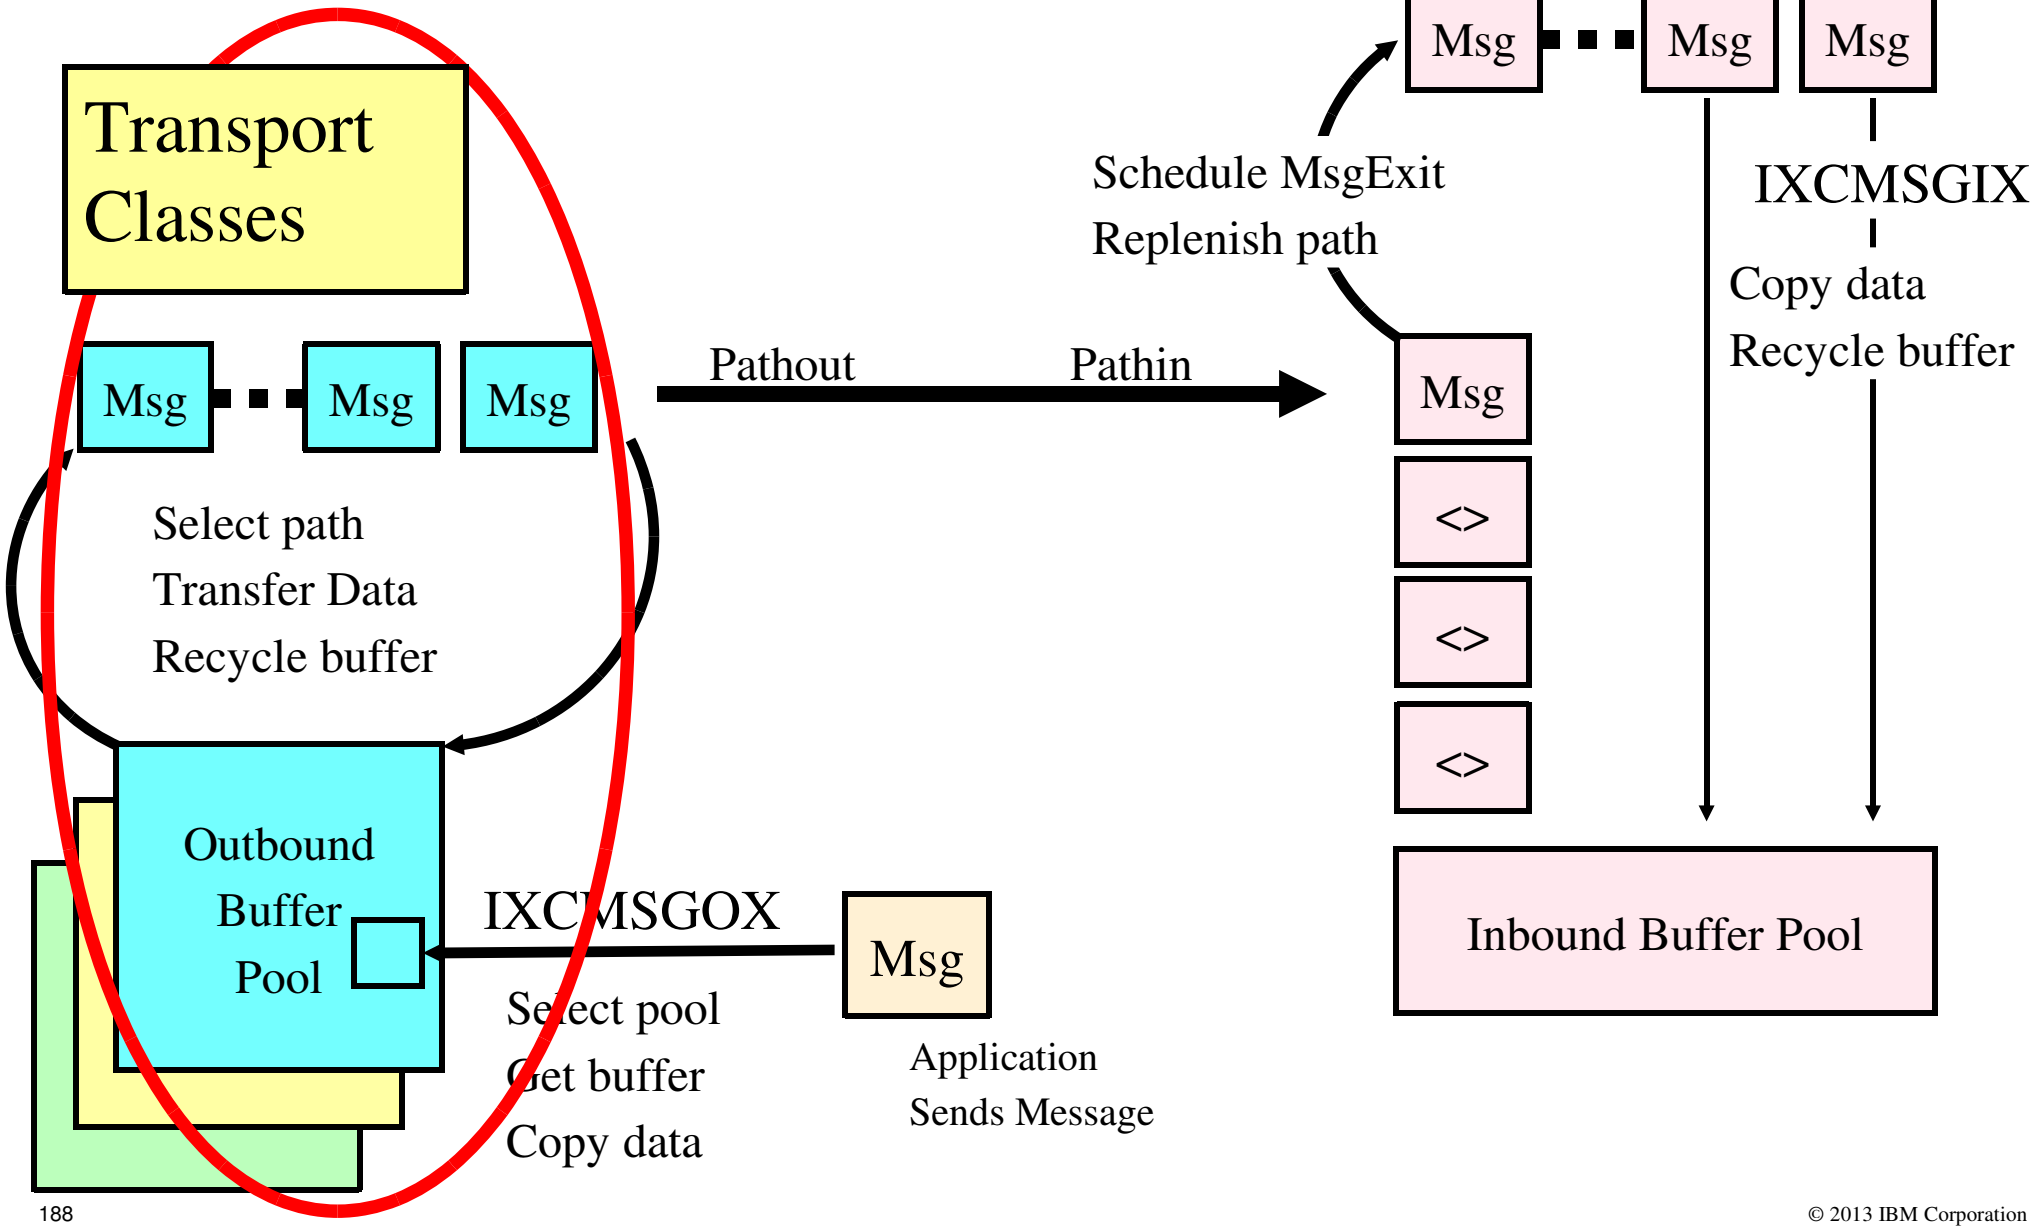

# Transport Class Concerns

- Appropriate segregation
  - Class Length
  - Group (?)
- Signalling Paths
  - **How many (must be > 0)**
  - Type
- Signal Buffers

*Make one class for 956 byte messages. Perhaps one, or two, maybe three classes for bigger message sizes.*

*Put at least two paths in each class for each target system.*

*Make sure the PATHIN MAXMSG values are reasonable for the chosen transport class length.*

# Segregation by Size

```
CLASSDEF CLASS(SML) CLASSLEN(956)
CLASSDEF CLASS(MED) CLASSLEN(8000)
CLASSDEF CLASS(BIG) CLASSLEN(32000)
```

```
SML =     0....956
MED =   957...8124
BIG = 8125..62464
```

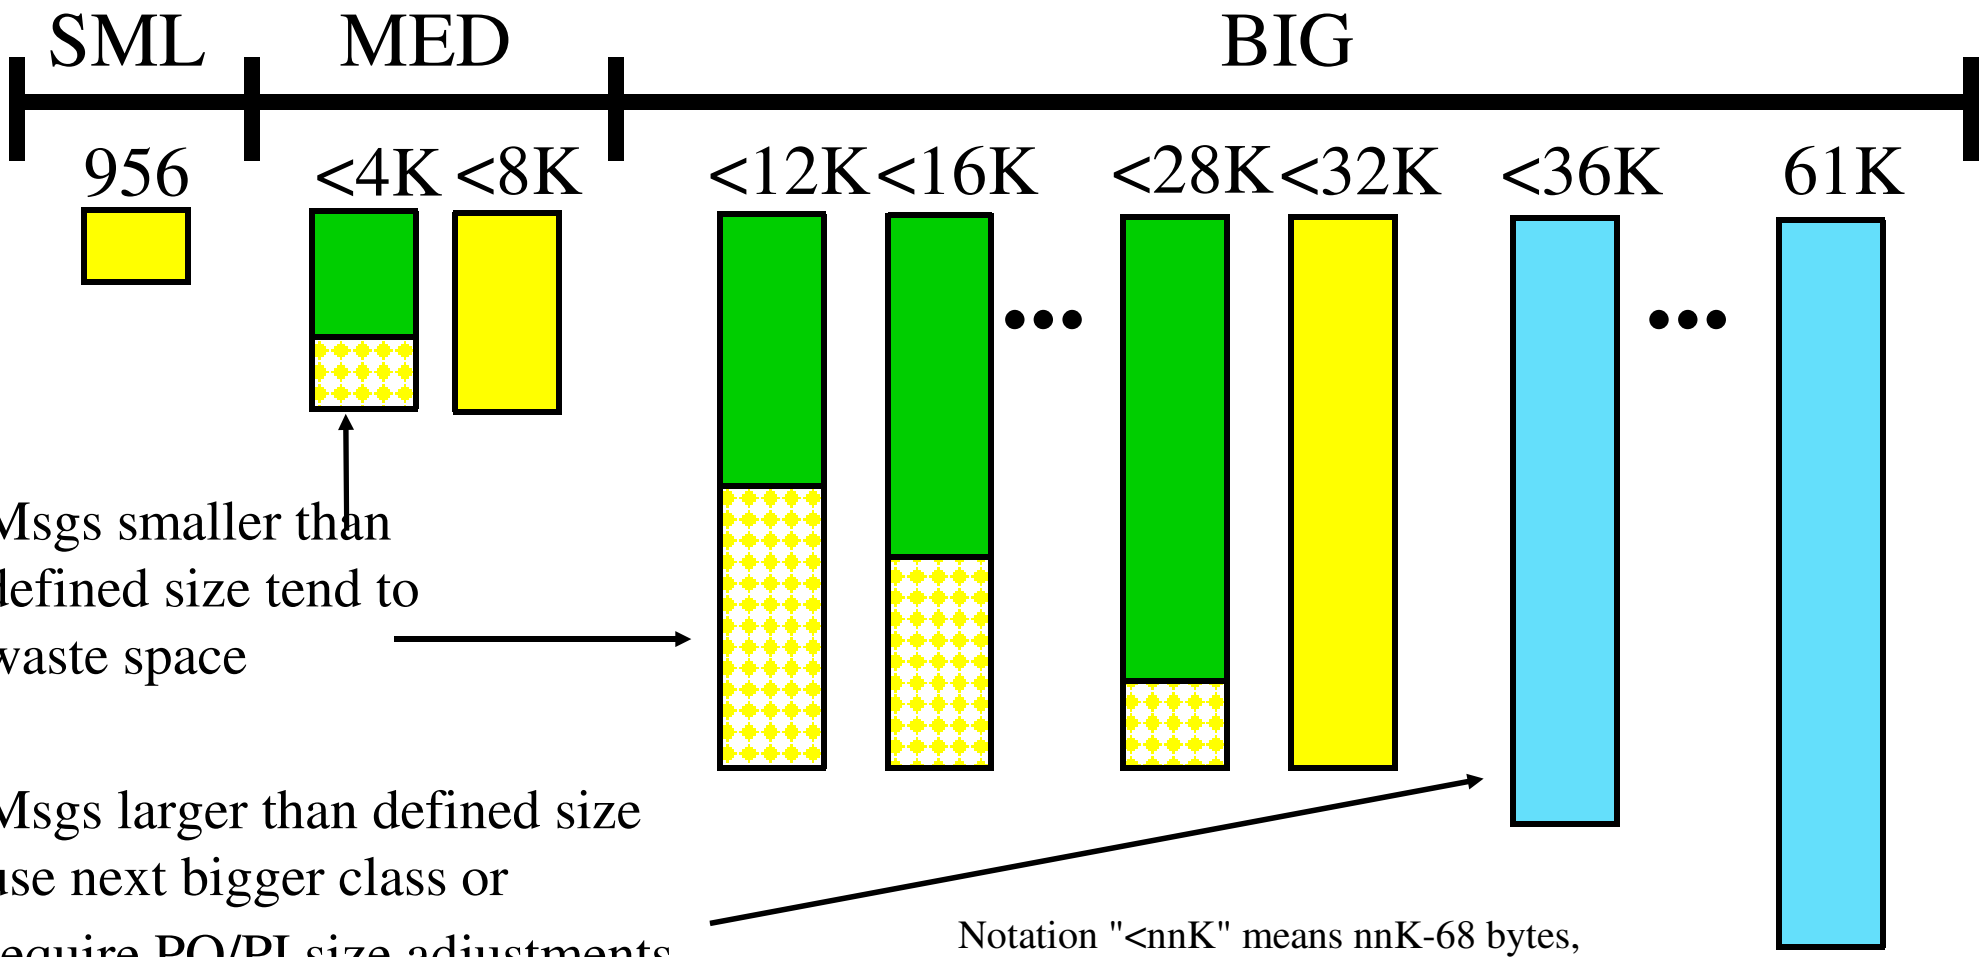

# Size Distributions by Class

```
SYS3  d xcf,classdef,class=all
SYS3  IXC344I  13.00.45  DISPLAY XCF
   TRANSPORT     CLASS         DEFAULT     ASSIGNED
   CLASS         LENGTH        MAXMSG      GROUPS
   DEFAULT         956           3000      UNDESIG

DEFAULT TRANSPORT CLASS USAGE FOR SYSTEM SYS1
SUM MAXMSG:      12000    IN USE:          18  NOBUFF:          0
  SEND CNT:      34523  BUFFLEN (FIT):   956
  SEND CNT:        256  BUFFLEN (BIG):  4028
  SEND CNT:          7  BUFFLEN (BIG):  8124
  SEND CNT:          1  BUFFLEN (BIG): 24508

DEFAULT TRANSPORT CLASS USAGE FOR SYSTEM SYS2
SUM MAXMSG:      12000    IN USE:          14  NOBUFF:          0
  SEND CNT:       9929  BUFFLEN (FIT):   956
  SEND CNT:        357  BUFFLEN (BIG):  4028
  SEND CNT:          6  BUFFLEN (BIG):  8124

DEFAULT TRANSPORT CLASS USAGE FOR SYSTEM SYS3
SUM MAXMSG:       3000    IN USE:           4  NOBUFF:          0
  SEND CNT:       3292  BUFFLEN (FIT):   956
  SEND CNT:        196  BUFFLEN (BIG):  4028
```

# Size Distributions by Member

```
D XCF,G,SYSMCS,ALL
IXC333I  15.45.54  DISPLAY XCF 303
   <snip>
 INFO FOR GROUP SYSMCS MEMBER SY1 ON SYSTEM SY1

 FUNCTION: Not Specified
 MEMTOKEN: 01000003 00070006      ASID: 0009       SYSID: 0100000B
     INFO: CURRENT           COLLECTED: 09/14/2010 15:45:54.275742

 ATTRIBUTES                     JOINED: 09/13/2010 17:12:33.022813
   JOIN TASK ASSOCIATION
   LASTING MEMBER
   SYSTEM CLEANUP PARTICIPANT
   LOCAL CLEANUP NOT NEEDED
   TERMLEVEL IS TASK
   MEMSTALL RESOLUTION IS NO ACTION
   EXITS DEFINED: MESSAGE, GROUP

SIGNALLING SERVICE
  MSGO ACCEPTED:       6820  NOBUFFER:          0
  MSGO XFER CNT:       3351   LCL CNT:       2532  BUFF LEN:   956
  MSGO XFER CNT:        399   LCL CNT:        419  BUFF LEN:  4028
  MSGO XFER CNT:        109   LCL CNT:         10  BUFF LEN:  8124
```

# RMF XCF Usage By System (outbound)

(5 minute interval)

REMOTE SYSTEMS

OUTBOUND FROM SYSA

| TO SYSTEM | TRANSPORT CLASS | BUFFER LENGTH | REQ OUT | BUFFER % SML | BUFFER % FIT | BUFFER % BIG | BUFFER % OVR | ALL PATHS UNAVAIL | REQ REJECT |
|---|---|---|---|---|---|---|---|---|---|
| SYSB | DEFAULT | 956 | 1,038 | 0 | 100 | 0 | 0 | 0 | 0 |
| | LARGE | 16,316 | 50 | 60 | 0 | 40 | 100 | 0 | 0 |
| | SMALL | 4,028 | 1 | 0 | 100 | 0 | 0 | 0 | 0 |
| SYSC | DEFAULT | 956 | 1,568 | 0 | 100 | 0 | 0 | 0 | 0 |
| | LARGE | 16,316 | 46 | 93 | 0 | 7 | 100 | 0 | 0 |
| | SMALL | 4,028 | 125 | 0 | 100 | 0 | 0 | 0 | 0 |
| SYSD | DEFAULT | 956 | 1,013 | 0 | 100 | 0 | 0 | 0 | 0 |
| | LARGE | 16,316 | 33 | 91 | 0 | 9 | 100 | 0 | 0 |
| | SMALL | 4,028 | 1 | 0 | 100 | 0 | 0 | 0 | 0 |
| TOTAL | | | 3,875 | | | | | | |

# Appropriate Number of Paths?

(5 minute interval)

OUTBOUND FROM SYSA

| TO SYSTEM | T Y P | FROM/TO DEVICE, OR STRUCTURE | TRANSPORT CLASS | REQ OUT | AVG Q LNGTH | AVAIL | BUSY | RETRY |
|---|---|---|---|---|---|---|---|---|
| SYSB | S | IXCSTR1 | DEFAULT | 843 | 0.00 | 843 | 0 | 0 |
| | S | IXCSTR2 | SMALL | 1 | 0.00 | 1 | 0 | 0 |
| | S | IXCSTR3 | LARGE | 50 | 0.00 | 47 | 3 | 0 |
| | S | IXCSTR4 | DEFAULT | 214 | 0.00 | 214 | 0 | 0 |
| SYSC | S | IXCSTR1 | DEFAULT | 357 | 0.00 | 357 | 0 | 0 |
| | S | IXCSTR2 | SMALL | 125 | 0.00 | 125 | 0 | 0 |
| | S | IXCSTR3 | LARGE | 46 | 0.00 | 46 | 0 | 0 |
| | S | IXCSTR4 | DEFAULT | 1,231 | 0.00 | 1,231 | 0 | 0 |
| SYSD | S | IXCSTR1 | DEFAULT | 408 | 0.00 | 408 | 0 | 0 |
| | S | IXCSTR2 | SMALL | 1 | 0.00 | 1 | 0 | 0 |
| | S | IXCSTR3 | LARGE | 33 | 0.00 | 33 | 0 | 0 |
| | S | IXCSTR4 | DEFAULT | 623 | 0.00 | 623 | 0 | 0 |
| TOTAL | | | | 3,932 | | | | |

*You do not need to drive BUSY to zero.*
*I tend not to worry about avail vs busy.*
*As Q Length approaches 1 or more, then assuming no other issues impacting transfers, you might add a path.*

# Detecting Transport Class Issues

- Run z/OS Health Checker

- D XCF,CLASSDEF
- D XCF,PI
- D XCF,GROUP
- COUPLE SIGNAL DETAIL
- RMF reports

# “XCF Lost My Signal”

- Preposterous !
  - XCF does not lose signals
  - But... signals can be delayed indefinitely
  - Usually a self-inflicted problem by exploiter
- Where might the signal be?
  - Pending transfer on the sending side
    - Signals flowing?
    - Any path in restart? rebuild?
  - Pending delivery to the message exit on the target system
    - Queued to the member delivery queue
    - Scheduled for delivery
  - (similarly for the response, except reverse roles of systems)
- Need timely dumps from both systems to investigate

# XCF Signal Floods

- We are seeing issues in the field where some application floods XCF with signals and others suffer as a result
  - Generally they can't keep up with their own flood
  - Or we get long queues of signals pending over the paths
  - Impact varies with who the “other” is
- Band Aids
  - Increase MAXMSG
  - Add signalling paths
- But understanding where the influx came from is critical
  - Know your workload
  - Use RMF Reports to see who is sending the signals
- Need timely dumps from both systems to investigate
  - *Why sending? Why so many? Why not keeping up?*

# Problem Taxonomy

- Dead System
- Sick System
- Sysplex Fabric
- Sysplex Componentry
  - Coupling Facility
  - Signalling Service
  - **Couple Data Sets**
  - External Time Reference
- Configuration / Capacity
- Software Issues

# Couple Data Set Concerns

- Connectivity
  - –Access (see “Sysplex Fabric: Couple Data Sets”)
  - –Implications for partial connectivity
- Capacity
  - –Be sure all formatted for same number of systems
  - –Otherwise needs to be formatted to meet needs of exploiter
    - Up to them to indicate when too small
- Performance
  - –Response time
  - –Contention
- Regression

# Couple Data Sets and Partial Connectivity

- Sysplex
  - System needs connectivity to IPL (except XCF-Local mode)
  - Systems that lose access wait-state
- ARM – system cannot use ARM
- BPXMCDS - functions may hang until access is restored
- CFRM - systems that lose access wait-state
- LOGR – lose logger functions that need CDS
- SFM - full sysplex-wide access required for CONNFAIL=YES
- WLM – lose functions that need CDS

# Sysplex Couple Data Set Capacity

- IXC202I “sysplex CDS is full (systems)”
- IXC700E “sysplex CDS is full (groups or members)”
- XCF_SYSPLEX_CDS_CAPACITY health check
- D XCF,COUPLE

```
SYSPLEX COUPLE DATA SETS
PRIMARY    DSN: UTCXCF.SVPLEX1.COUPLE.PRI
           VOLSER: X1CPLP     DEVN: 1C6E
           FORMAT TOD         MAXSYSTEM MAXGROUP(PEAK) MAXMEMBER(PEAK)
           11/27/2011 13:21:26       32     120  (110)      503   (56)
```

# Function Couple Data Sets and MAXSYSTEM

Function couple data sets should be formatted to support at least as many systems as supported by the sysplex couple data sets

```
D XCF,COUPLE
IXC357I  22.32.16  DISPLAY XCF 863

SYSPLEX COUPLE DATA SETS
PRIMARY    DSN: UTCXCF.SVPLEX1.COUPLE.PRI
           VOLSER: X1CPLP     DEVN: 1C6E
           FORMAT TOD         MAXSYSTEM MAXGROUP(PEAK) MAXMEMBER(PEAK)
           11/27/2011 13:21:26       32     120  (110)      503   (56)

BPXMCDS COUPLE DATA SETS
PRIMARY    DSN: UTCXCF.SVPLEX1.OMVSR14.PRI
           VOLSER: X1CPLP     DEVN: 1C6E
           FORMAT TOD         MAXSYSTEM
           04/11/2003 21:00:01        8

CFRM COUPLE DATA SETS
PRIMARY    DSN: UTCXCF.SVPLEX1.CFRMR18P.PRI
           VOLSER: X1CPLA     DEVN: B010
           FORMAT TOD         MAXSYSTEM
           12/01/2010 12:08:18       32
```

Oops

# Function Couple Data Set Capacity

- DISPLAY XCF,COUPLE,TYPE=ALL
  - Shows format parameters (limits) for all but WLM
- ARM
  - D XCF,ARMSTATUS for registered elements vs TOTELEM
  - Create of policy fails if MAXELEM exceeded
- BPXMCDS
  - MODIFY BPXOINIT,FILESYS=DISPLAY,GLOBAL
  - BPXI043E “approaching mount table limit”
  - AMTRULES – no way to tell limit reached?

# Function Couple Data Set Capacity

- CFRM
  - IXL013I, if “connection failure” = “max permitted connectors”
  - Create policy fails if exceed POLICY, CF, or STR limits

  - D XCF,STR will indicate need for larger CDS for reconciliation
  - D XCF,CF will indicate need for larger CDS for reconciliation
  - IXC502I “need bigger CDS for reconciliation”
  - IXC503I “need bigger CDS for reconciliation”
  - IXC514I “need bigger CDS for reconciliation”

# Function Couple Data Set Capacity

- LOGR
  - IXG010E “CDS too small”
  - IXG261E “CDS becoming full”
  - IXG262A “CDS is essentially full”
  - IXG270I “CDS becoming full”
- SFM
  - All capacity issues are detected when creating policies
- WLM
  - All capacity issues are detected when installing your service definitions

# Couple Data Set Performance

- Various CDS generally have relatively low request rates
  - Though may depend on application design/behavior
  - Sysplex monitors run amok?
    - Be careful how you run RMF and similar tools
- But good performance is needed to avoid
  - Removal of CDS
  - Application delays
- First resolve access issues if any
  - See “Sick System: DASD I/O Issues”
  - See “Sysplex Fabric: Couple Data Sets”

# Typical sources of CDS performance issues

- DASD
  - –Old and slow?
  - –Device caching enabled?
  - –Synchronous mirroring?
- Record Size
  - –CDS Formatted with too much white space?
- Contention
  - –Using MSGBASED processing for CFRM?
  - –Reserves?
  - –Other data sets with high request rates on same volume?
    - Increased request rates due to workload changes

# Detecting CDS Performance Issues

- Likely need baselines for comparisons
  - Changes to request rates
  - Changes to DASD response times
- RMF DASD I/O reports.
  - Direct Access DASD Activity – post processor report

# Couple Data Set Regression

- When IPLing the sysplex, your couple data sets should either be:
  - The ones most recently used by the sysplex (typical and best), OR
  - Freshly formatted, never before used (atypical, usually DR site)
- Couple data sets often contain status and configuration data
- Regressing to an older CDS is risky because the data in that CDS may not be consistent with:
  - The current configuration
  - Data recorded in other (non-regressed) couple data sets
  - Related application data recorded in other data sets

*You might get some messages prompting operator "Should we use this CDS?" when IPLing. Once the sysplex is up, I don't know of any reliable way to detect that regressed CDS was used.*

# Couple Data Set Regression ...

- Regression typically occurs when you use:
  - A backup copy of the CDS
  - A former primary/alternate CDS no longer actively in use
  - A CDS that was previously in use by some other sysplex
- NOTE: Bringing an old CDS into use via ACOUPLE is safe
- Resetting the TOD clock can also wreak havoc since time stamps recorded in the CDS may suddenly seem to be:
  - In the distant past
  - From the future

# Problem Taxonomy

- Dead System
- Sick System
- Sysplex Fabric
- Sysplex Componentry
  - Coupling Facility
  - Signalling Service
  - Couple Data Sets
  - **External Time Reference**
- Configuration / Capacity
- Software Issues

# External Time Reference

- System is in a wait-state if not time synchronized with sysplex
  - See “Sysplex Fabric: External Time Reference”
- Server Time Protocol (STP)
  - Must provide and maintain a resilient configuration as there is potential for sysplex-wide outage !
  - Understand Primary, Backup, and Arbiter roles
- Operational issues
  - Reassigning roles after failure (or planned outage) of any server that has a special role
  - Dealing with STP events
  - Dealing with loss of time synchronization

# External Time Reference

- D ETR,DATA shows STP info for local system
- Redundant timing links?
- Expected role?

```
D ETR,DATA
IEA386I 22.41.27 TIMING STATUS 932
SYNCHRONIZATION MODE = STP
  THIS SERVER IS A STRATUM 2
  CTN ID = POKSTP
  THE STRATUM 1 NODE ID = 002097.E40.IBM.00.0000000699DF
  NUMBER OF USABLE TIMING LINKS = 23
  THIS STP NETWORK HAS NO SERVER TO ACT AS ARBITER
```

# External Time Reference

- IEA031I “STP alert was issued to HMC”
- IEA382I “have single point of failure”
- IEA383I “have single point of failure”
- IEA388I “not connected to backup time server (or arbiter)”
- IEA389I “no server to act as backup timer (or arbiter)”
- IEA395I “switched to backup time server (or preferred)”

- IXC406I “not connected to same ETR”
- IEA394A “lost access to ETR (STP)”
- IEA015A “lost access to ETR (sysplex timer)”

# Operator Alerts for STP Events

- Operator alerts sent to z/OS console as well as HMC for STP related hardware & timing events:
  - Dial-out time service outside allowable tracking range
  - Dial-out access failure
  - NTP server failure
  - NTP servers not synchronized

- IEA031I STP ALERT RECEIVED.  ALERT CODE = *nn*

- Available in z/OS 1.11
- Aslo available at z/OS 1.9 and 1.10 with OA28323
- Supported by z10 and z9 servers with MCL driver 76

# External Time Reference

- Redbooks
  - *Server Time Protocol Implementation Guide*
    www.redbooks.ibm.com/abstracts/sg247281.html
  - *Server Time Protocol Planning Guide*
    www.redbooks.ibm.com/abstracts/sg247280.html
  - *Server Time Protocol Recovery Guide*
    www.redbooks.ibm.com/abstracts/sg247380.html

*Need to be very careful with CTN.*
*These books are excellent references.*
*Exposed to suffering outages during*
*migration to CTN, changing configurations,*
*or when operational errors are made.*

# Problem Taxonomy

- Dead System
- Sick System
- Sysplex Fabric
- Sysplex Componentry
  - Coupling Facility
  - Signalling Service
  - Couple Data Sets
  - External Time Reference
- **Configuration / Capacity**
- Software Issues

# Configuration and Capacity Concerns

- Resiliency
- Workload changes
  - Growth
  - Reconfiguration
- That might lead to more:
  - XCF signals
  - CF requests
  - CDS accesses

# Preventing Configuration Issues

- Adhere to best practices for availability
  - No single points of failure
  - Ensuring redundancy and fail-over capability often enough to mitigate configuration mistakes (perhaps until failures occur)
- Use IBM Health Checker for z/OS
  - My anecdotal experience leads me to believe that the number of multi-system outages has dropped significantly since the introduction of the health checker

White Paper: "Mission: Available"
www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101966

# Preventing Capacity Issues

- Good capacity planning
- Resize structures
- White space and head room for failover

# RMF Report of CF Activity

```
                                  COUPLING  FACILITY  USAGE  SUMMARY

STORAGE SUMMARY

                                        ALLOC      % OF CF      ------- DUMP SPACE -------
                                        SIZE       STORAGE    % IN USE   MAX % REQUESTED

TOTAL CF STORAGE USED BY STRUCTURES      1430M        9.5
TOTAL CF DUMP STORAGE                       6M        0.0          0.0                0.0
TOTAL CF STORAGE AVAILABLE              13573M       90.4
                                        ------
TOTAL CF STORAGE SIZE                   15009M

                                        ALLOC
                                        SIZE

TOTAL CONTROL STORAGE DEFINED           15009M        9.0
TOTAL DATA STORAGE DEFINED                  0K        0.0
```

- Bigger CF?
- Allow for failover? Maintenance?
- Structures properly sized?
- Force structures?

# Methodology Concern

- We might have what appears to be a sysplex configuration and capacity issue
- That is really an application design issue
- For example:
  - An application might choose to share data by sending copies of the data to all systems in the sysplex instead of sharing a single copy of the data in a CF that is accessible to all
- I proceed on the assumption that you are more likely to change the configuration than the application
  - But one can make a case for looking at application issues before considering configuration and capacity issues

# Typical causes of increased activity

- Signals
  - Merge of systems
  - Increase work
  - Application changes
  - GRS contention
  - Migration / Maintenance
- CF requests
  - Merge of systems
  - Increase in work
  - Change in type of work
  - Application changes
  - Processor upgrade or MCL application
- CDS access
  - Systems joining / leaving sysplex
  - Recovery
  - XCF groups joining / leaving

# Track impact on sysplex componentry as workload changes

- RMF Reports of XCF Activity
  - Member sending signals
  - Usage by System
  - Signal path usage
- RMF Reports of CF activity
  - CF as a whole
  - Structure by Structure
- RMF Reports of DASD activity
  - As related to Couple Data Sets

*Know your workload*

# Loss of capacity

- Expect failures.  Plan for them.
  - Provide spare capacity and head room for fail over
  - Reduce MTTR by automating restart of failed applications and systems

- See “z/OS MVS: Planning Operations” for information on using Auto-IPL to have failed system automatically IPLed back into the sysplex
  http://publib.boulder.ibm.com/infocenter/zos/v1r12/index.jsp?topic=%2Fcom.ibm.zos.r12.ieag300%2Fwsat.htm

# Problem Taxonomy

- Dead System
- Sick System
- Sysplex Fabric
- Sysplex Componentry
  - Coupling Facility
  - Signalling Service
  - Couple Data Sets
  - External Time Reference
- Configuration / Capacity
- **Software Issues**

# Methodology Recap

- At this point:
  - We have neither dead systems nor sick systems
  - Sysplex componentry is accessible, operational, and performing
  - Configuration is sufficient to support the workload
- If still have a problem, likely have an application specific issue

# Design decisions for sysplex applications

- Communication
  - XCF Signalling Service
  - CF List Structure
- Data sharing
  - Distributed
  - Centralized
- Serialization
  - ENQ
  - XES Locking Services
  - “Owner”
- Topology
  - Peer to peer
  - Master (fixed or dynamic)
  - Ring
  - Star
- Coordination
  - Signals
  - XES Sync points
  - ENQ, locks, latches

# Those design decisions have consequences

- Scalability
- Performance characteristics
- Failure modes

- All of which impact problem diagnosis
  - Hard to make progress if don't understand how its put together
  - Will likely need application specific expertise

# Preventing software issues

- Stay current with maintenance
- Stay current with releases

- IBM Best Practice: Apply z/OS maintenance RSU regularly and HIPERs more often
  - http://www-03.ibm.com/support/techdocs/atsmastr.nsf/PubAllNum/Flash10106
- IBM Best Practice: Subscribe to Red Alerts
  - http://www-947.ibm.com/support/entry/portal/Overview/Software/Software_support_(general)?pgel=wspace

# Summary

- Sysplex problems can be difficult to diagnose
  - Complex inter-dependencies
  - Symptoms may not be directly related to the root cause
- Prevention is highly recommended
  - **Enable SFM with BCPii** to resolve dead system issues
  - Exploit SFM policy to resolve sick but not dead issues
  - Use z/OS Health Checker to verify configuration
- Use the “Problem Taxonomy” to guide your diagnosis efforts. I believe it provides a disciplined approach that can help identify the likely area for root cause, even in the face of imperfect knowledge of the system

# For more information

- Setting Up a Sysplex
- Problem Determination Guide
- Initialization and Tuning Reference
- z/OS MVS Data Areas
- MVS Sysplex Services Guide
- MVS Sysplex Services Reference
- RMF Users Guide

#SHAREorg

SHARE
Technology · Connections · Results

# Questions?

# Diagnosing Sysplex Problems
**Session 13402**

Please fill out the online session evaluation at either:

- SHARE.org/SanFranciscoEval, or
- Aim your smartphone at this QR code below:

SHARE in San Francisco 2013

# Appendix

# Non-disruptive CF dump

SETXCF DUMPCF,

{CFNAME=cfname[,TYPE=NONDISRUPTIVE|DISRUPTIVE][,UNCOND=NO|YES]}

{STRNAME=strname}

CFNAME=cfname

Allows the operator to specify the CF to be dumped.

TYPE=DISRUPTIVE|NONDISRUPTIVE

Allows the operator to optionally request a disruptive CF dump.

Default: Nondisruptive

UNCOND=YES|NO

Allows the operator to bypass a confirmation if a disruptive CF dump is requested.

Default: No

STRNAME=strname

Allows the operator to request a CF dump by specifying a structure name. The CF(s) in which the structure resides will be dumped and the dump requested will be non-disruptive.

OA35342 – Introduced the z/OS operator command to collect a non-disruptive serialized CF dump. CF dumps are reviewed by IBM hardware support.

# Getting sysplex dumps

- If CF issue, activate SYSXES Ctrace
  - TRACE CT,16M,COMP=SYSXES
  - R XX,OPTIONS=(HWLAYER,REQUEST,CONFIG,CONNECT,RECOVERY,LOCKMGR),END
- If XCF issue, activate SYSXCF ctrace
  - TRACE CT,ON,COMP=SYSXCF
  - R xx,OPTIONS=(GRPNAME=(GEOPLEX),SIGNAL,GROUP,SFM),END
- Place dump command in IEADMCxx parmlib member:
- JOBNAME=(XCFAS),DSPNAME=('XCFAS'.*),
- SDATA=(ALLNUC,CSA,PSA,LPA,LSQA,NUC,RGN,SQA,SUM,SWA,TRT,XESDATA,COUPLE),
- REMOTE=(SYSLIST=*('XCFAS'),DSPNAME,SDATA)),END
- Initiate the dump on one system in the sysplex:
- DUMP COMM=(meaningful dump title),PARMLIB=xx