

# Sysplex Failure Management (SFM) History and Proven Practice Settings

Mark A Brooks

IBM

mabrook@us.ibm.com

4:30 p.m. Wednesday March 14, 2012

Session 10850

# 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.com®	MVST™	Service Request Manager®	z10™
CICS®	OS/390®	Sysplex Timer®	z/Architecture®
CICSPIlex®	Parallel Sysplex®	System z®	zEnterprise™
DB2®	Processor Resource/Systems Manager™	System z9®	z/OS®
eServer™	PR/SM™	System z10®	z/VM®
ESCON®	RACF®	System/390®	z/VSE®
FICON®	Redbooks®	Tivoli®	zSeries®
HyperSwap®	Resource Measurement Facility™	VTAM®	
IMS™	RETAIN®	WebSphere®	
IMS/ESA®	GDPS®		
	Geographically Dispersed Parallel Sysplex™		

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

IBM, z/OS, Predictive Failure Analysis, DB2, Parallel Sysplex, Tivoli, RACF, System z, WebSphere, Language Environment, zSeries, CICS, System x, AIX, BladeCenter and PartnerWorld are registered trademarks of IBM Corporation in the United States, other countries, or both.  
DFSMSshm, z9, DFSMSrmm, DFSMSdftp, DFSMSdss, DFSMS, DFS, DFSORT, IMS, and RMF are trademarks of IBM Corporation in the United States, other countries, or both.  
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.  
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.

\* 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.

# History

- **MEMSTALLTIME**
  - **SSUMLIMIT**
  - **SFM with BCPii**
  - **System Default Action**
  - **XCF FDI Consistency**
  - **Critical Members**
  - **CFSTRHANGTIME**
- z/OS 1.8
  - z/OS 1.9
  - z/OS 1.11
  - z/OS 1.11
  - z/OS 1.11
  - z/OS 1.12
  - z/OS 1.12

SFM is the subcomponent within XCF that deals with the detection and resolution of sympathy sickness conditions that can arise when a system or sysplex application is unresponsive

# Terminology

- **FDI**
  - **ISA**
  - **ISI**
  - **SSUM**
  - INTERVAL
  - SSUM ACTION
  - SSUM INTERVAL
  - MONITOR-DETECTED STOP
- **Failure Detection Interval**
  - **Indeterminate Status Action**
  - **Indeterminate Status Interval**
  - **System Status Uppdate Missing**
  - = FDI
  - = ISA
  - = ISI
  - = SSUM
- *actionTIME(nostatus-interval)* - *action* = ISA, *nostatus-interval* = ISI  
–ISOLATETIME(n) or DEACTTIME(n) or RESETTIME(n)
- Sysplex partitioning
  - Remove system from sysplex
- SFM
  - Sysplex Failure Management

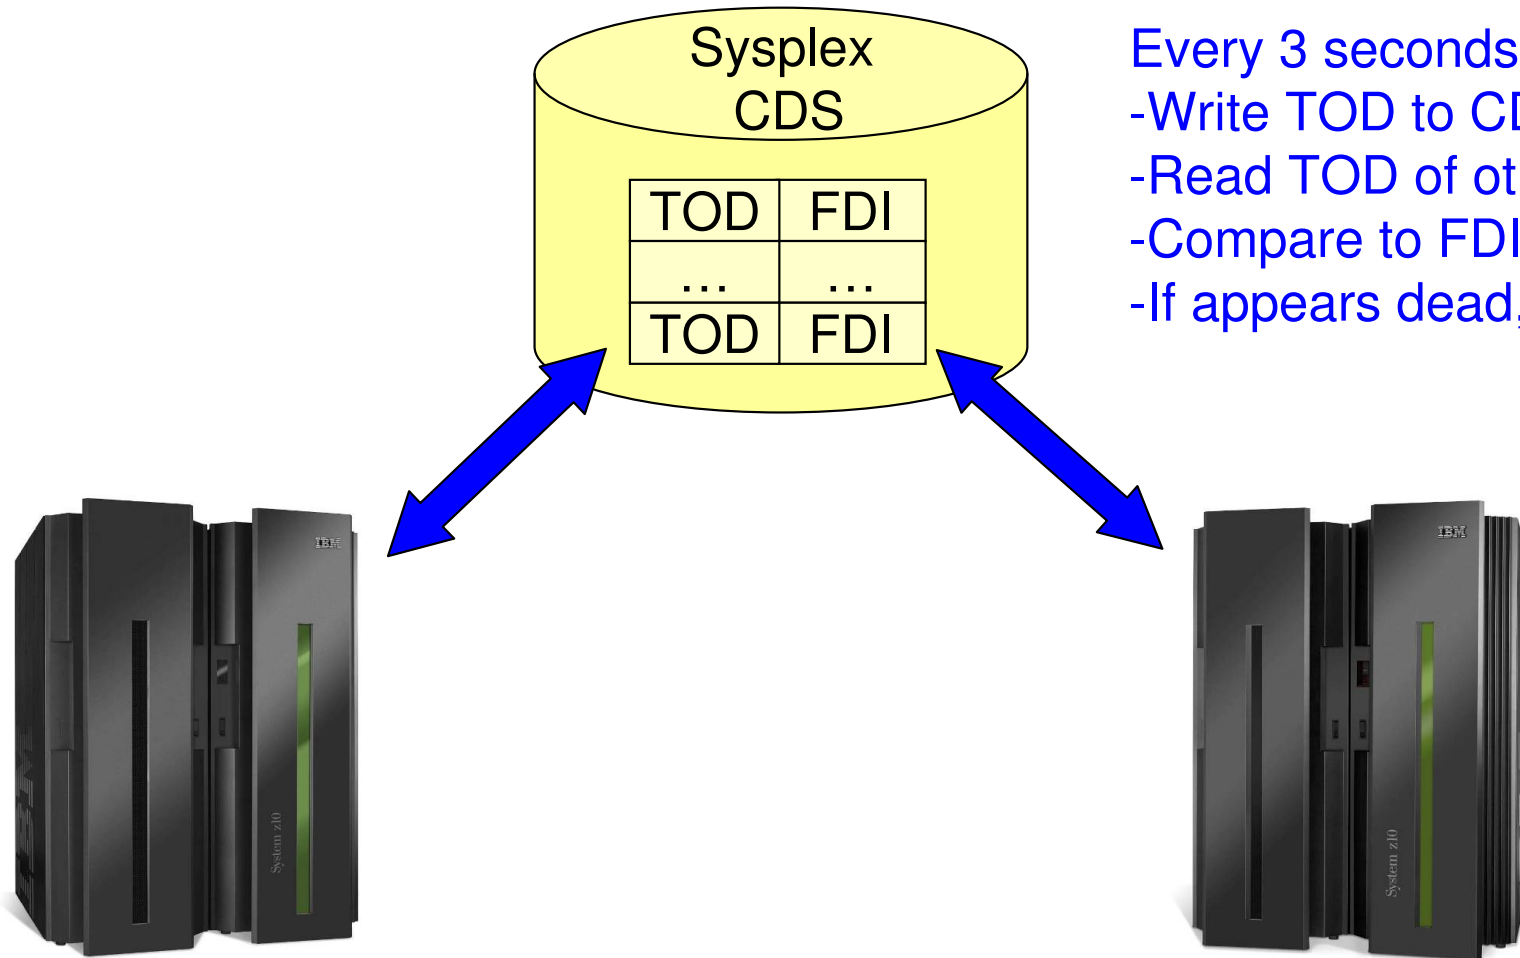
# Sympathy Sickness

- When a system becomes unresponsive
  - It may be serializing shared resources with RESERVEs, ENQ's, Locks
  - It may stop sending responses or otherwise fail to participate in various “group” protocols
- Other work may experience delays and hangs
  - Problem compounds as the sympathy sickness spreads
- Timely intervention is needed
  - One must correctly identify the culprit
  - Take corrective action

# Where we are going

- Look back in history at early days of sysplex, moving forward in time to explore the issues encountered and solutions provided along the way
  - Though not necessarily in actual historical order
- Initially focus on unresponsive systems
- Then broaden the scope to other forms of sympathy sickness

# System Status Monitoring




- Every 3 seconds:
- Write TOD to CDS
  - Read TOD of others
  - Compare to FDI
  - If appears dead, take action

FDI = Failure Detection Interval  
CDS = Couple Data Set



# Unresponsive System with Operator Prompt



Can we make eternity timely?

FDI = Failure Detection Interval

OPN = OPNOTIFY interval

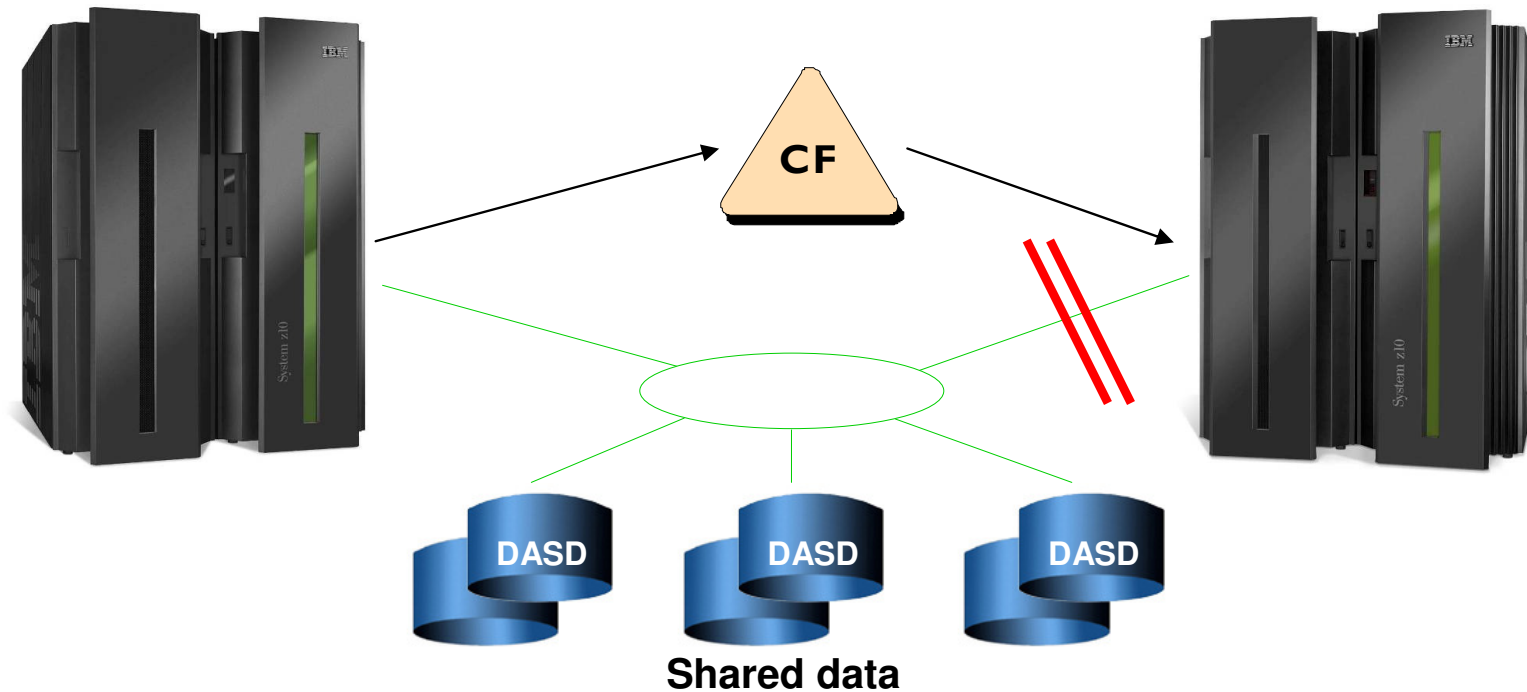


# Can We Shorten Eternity ?



- Sysplex Failure Management (SFM)
  - Use XCF Status Monitoring to detect unresponsive systems
  - Create a policy to specify whether XCF is to automatically remove unresponsive systems from the sysplex
- But what about the operator RESET of the system?
  - A critical step
  - Without it, data can be corrupted
    - So automating response of DOWN to IXC402D is likely bad

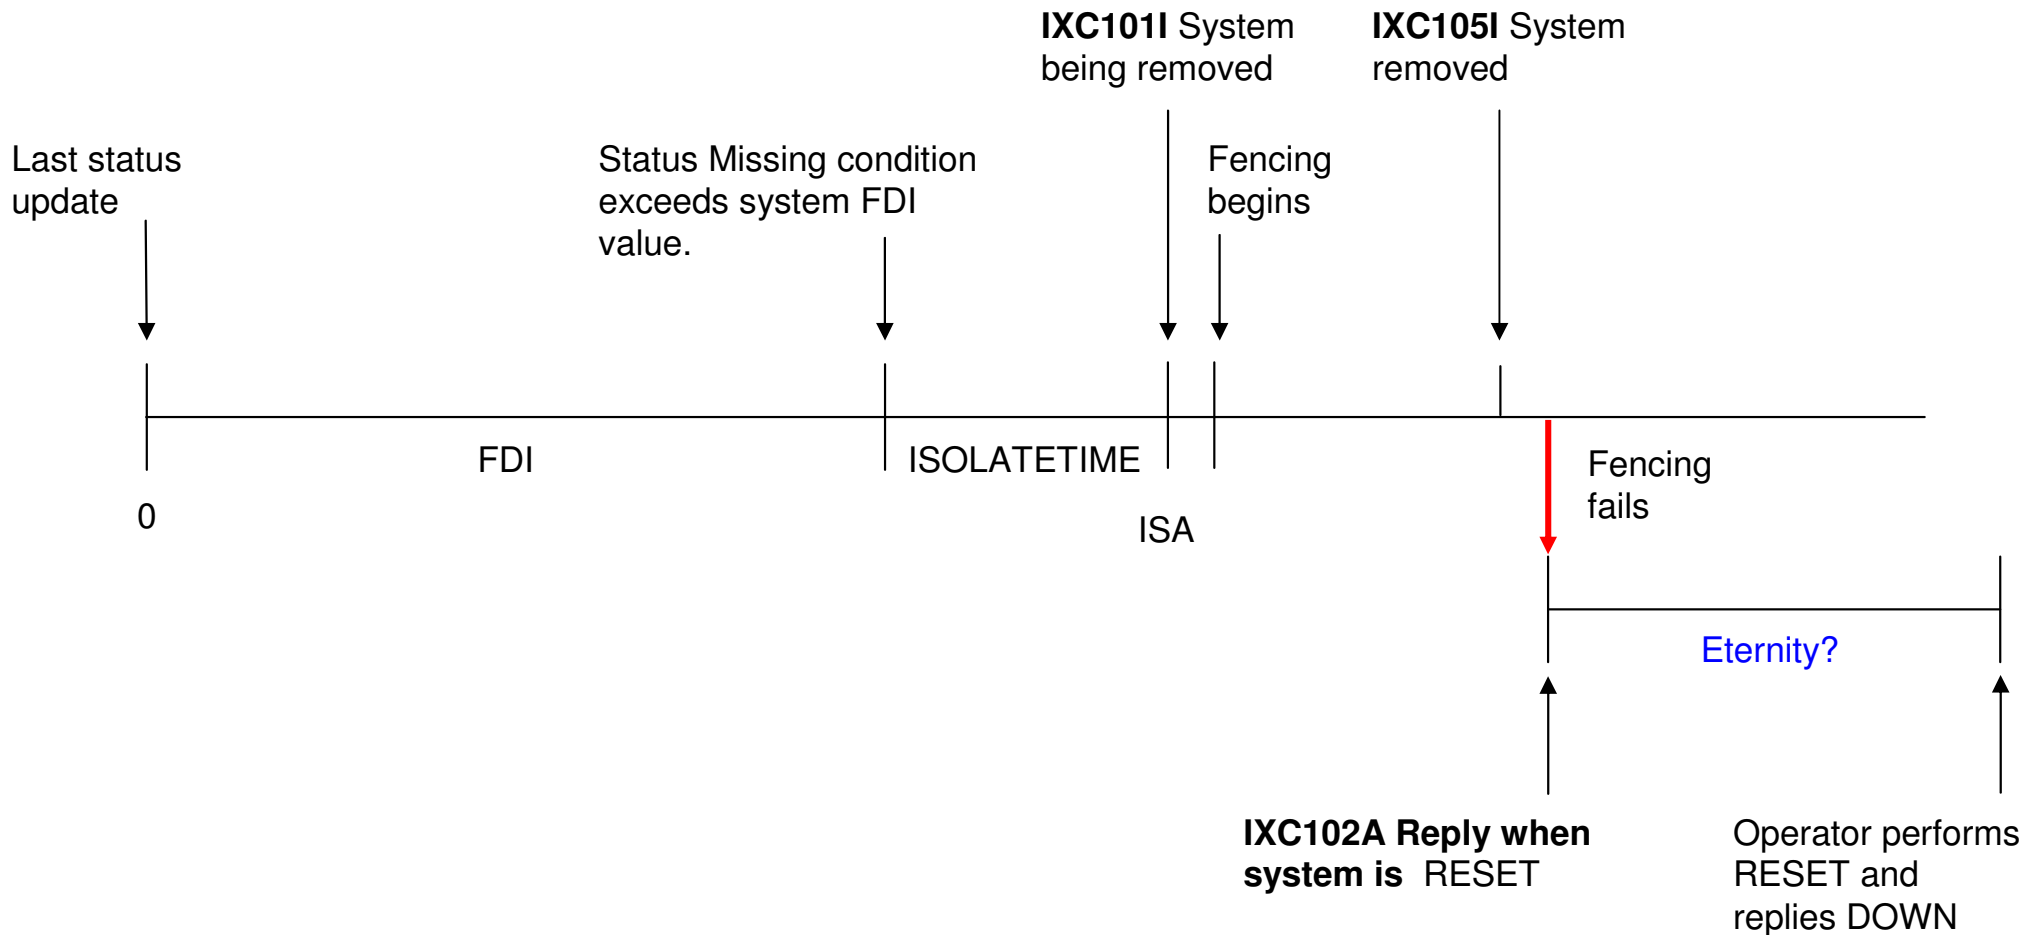
# Fencing



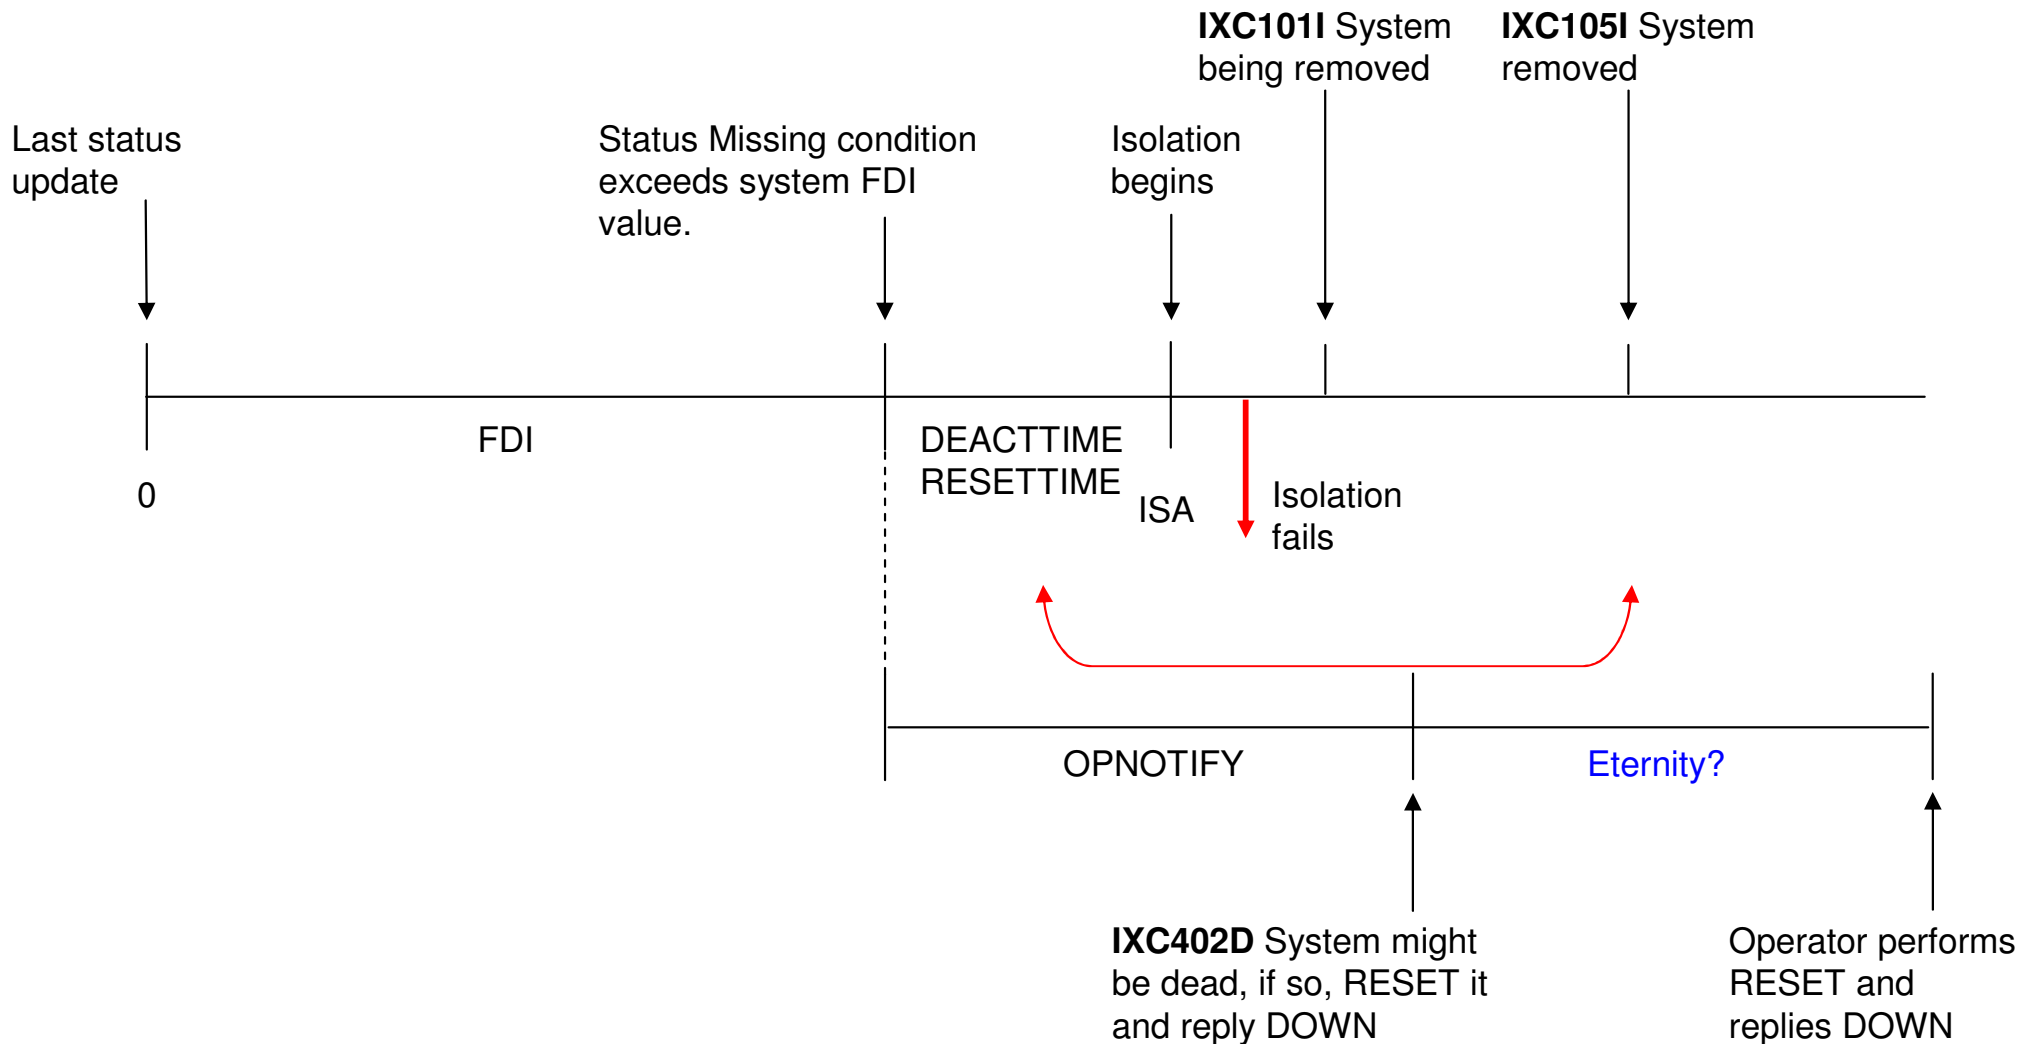
Fencing isolates a system so that it cannot access shared data, thus making it safe for the survivors to release serialization of the shared resources.

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.

# Unresponsive System with SFM ISOLATETIME



# Unresponsive System with SFM DEACTTIME or RESETTIME



# z/OS 1.11 System Default Action



- SFM Policy defines how XCF is to deal with an unresponsive system
- Each system “publishes” in the sysplex couple data set the action that is to be applied by its peers
- A system publishes “default action” if:
  - The policy does not specify an action for it, or
  - There is no SFM policy active
- “Default action” is:
  - PROMPT prior to z/OS 1.11
  - ISOLATETIME(0) as of z/OS 1.11

# z/OS 1.11 System Default Action ...



- The resulting “default action” depends on who is monitoring who:
  - z/OS 1.11 will isolate a peer z/OS 1.11
  - z/OS 1.11 will PROMPT for lower level peer
  - Lower level system will PROMPT for z/OS 1.11
- D XCF,C shows what the system *wants*
  - *But it may not get that in a mixed sysplex*
- Note: z/OS 1.11 always tries fencing whenever it is needed
  - Lower level releases performed fencing only when an SFM policy was active

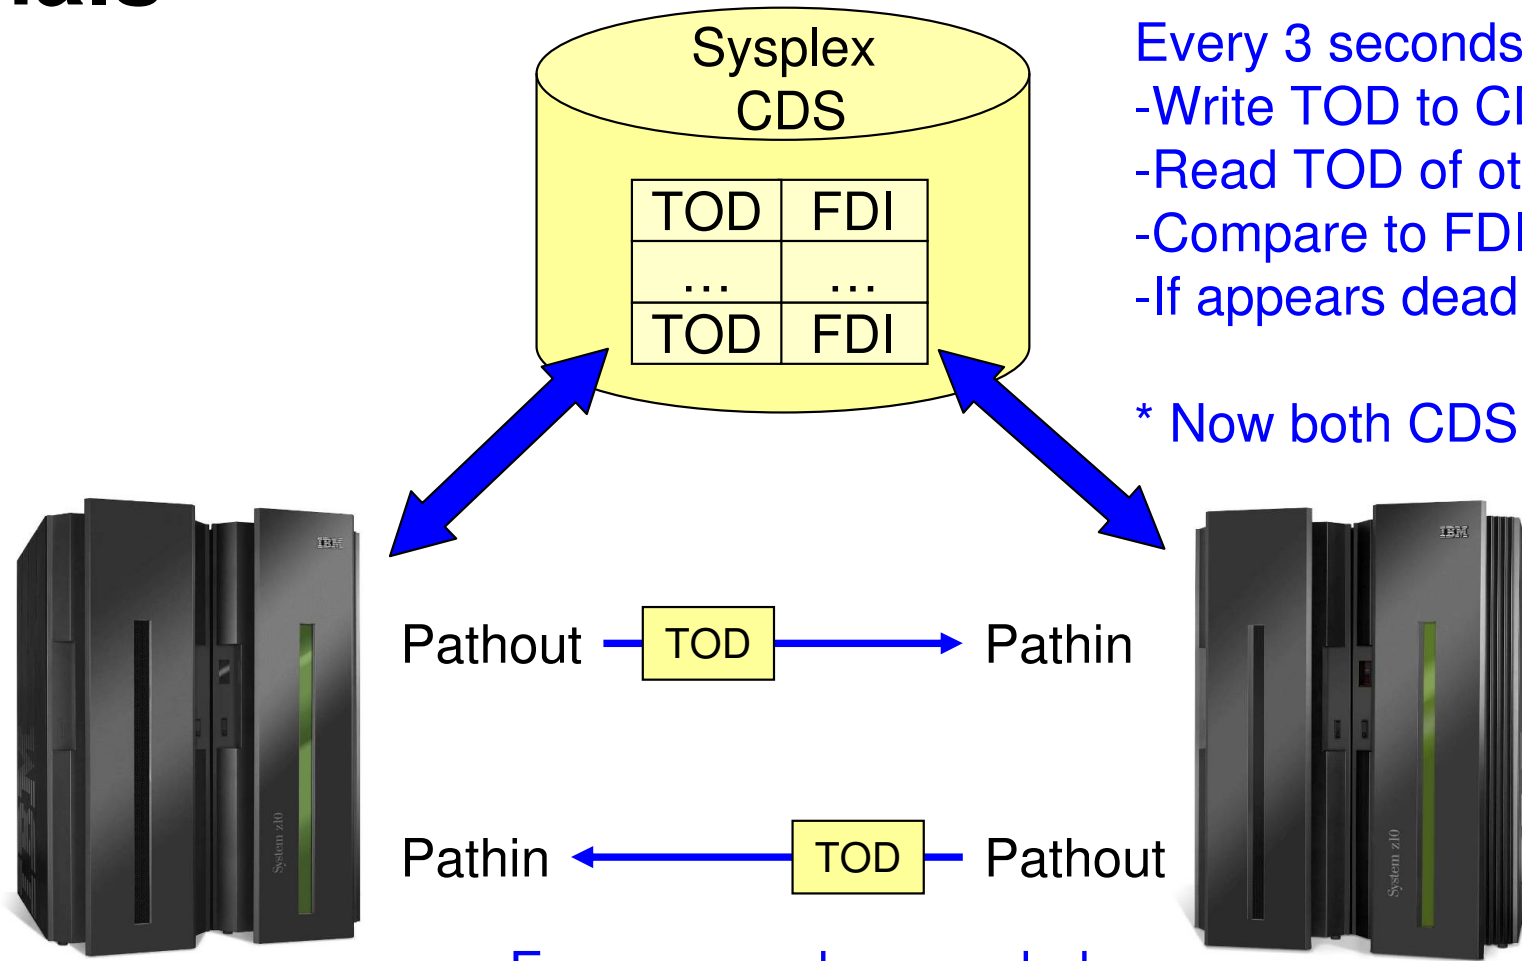
# Historically, there were problems



- Systems were being “needlessly” removed by SFM
- Failure to update status in the sysplex CDS was not a sufficiently reliable indicator of system failure
- So the system monitor was enhanced to watch for XCF signal traffic as well



# System Status Monitoring Plus Signals

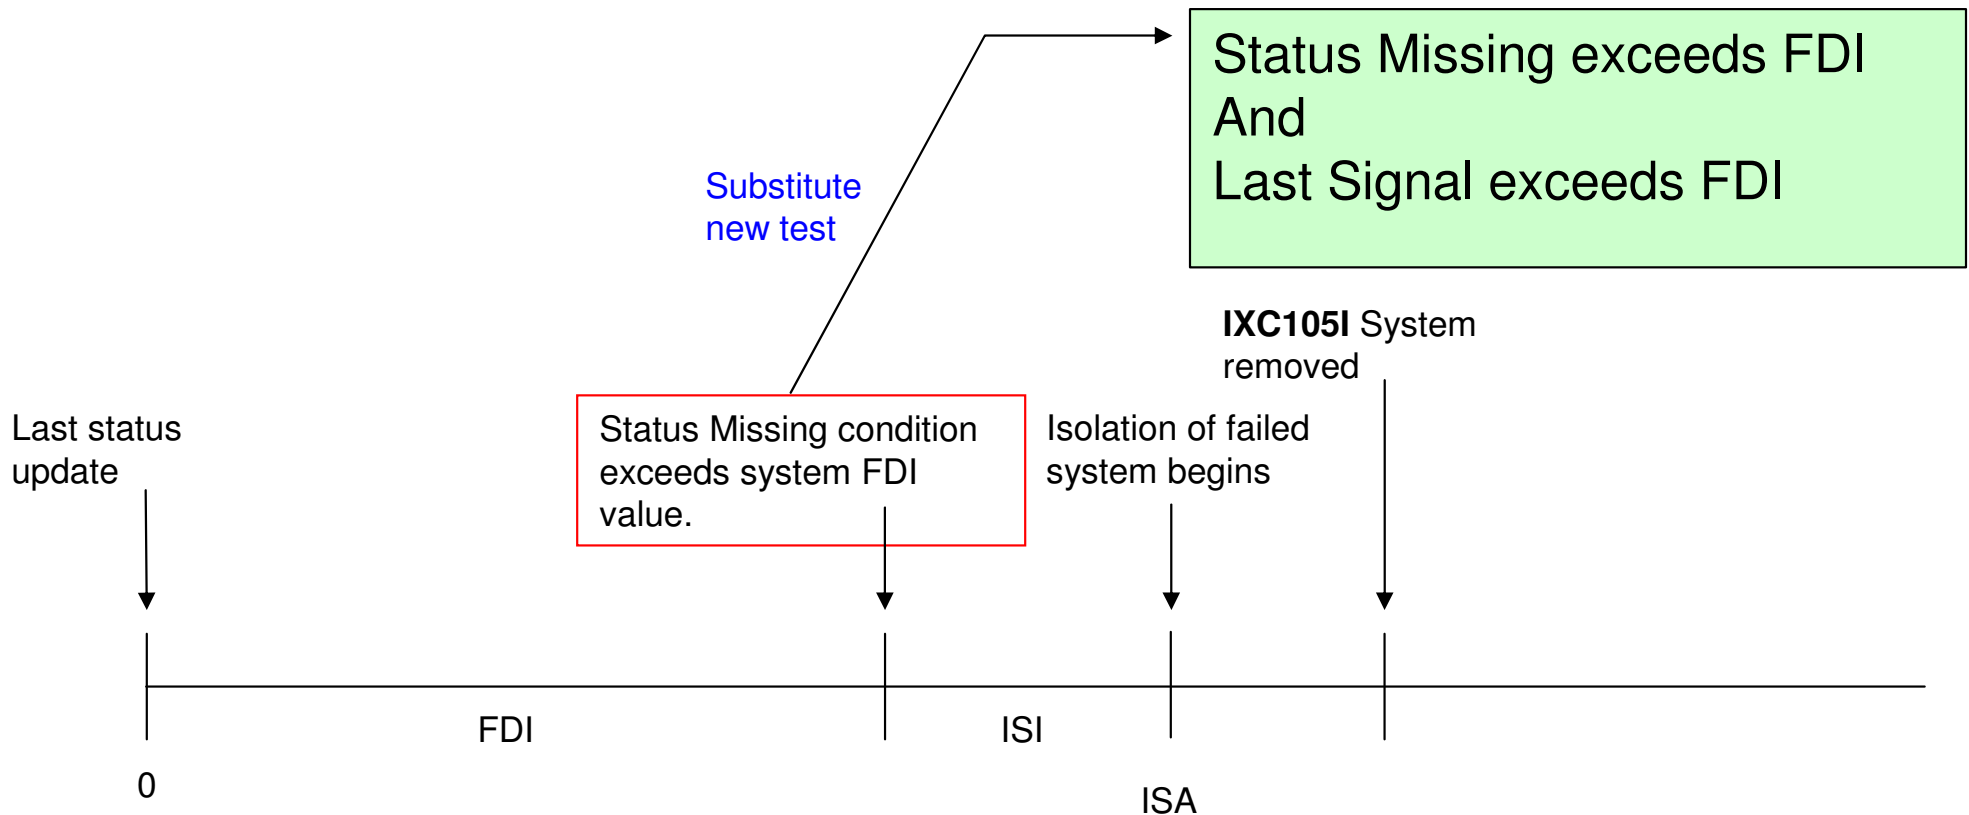


- Every 3 seconds:
- Write TOD to CDS
  - Read TOD of others\*
  - Compare to FDI
  - If appears dead, take action

\* Now both CDS and signal

- Every second as needed:
- Send TOD to peers

# SFM with More Reliable Detection



FDI = Failure Detection Interval

ISI = Indeterminate Status Interval (time value from ISOLATETIME etc)

ISA = Indeterminate Status Action

# More Reliable Detection

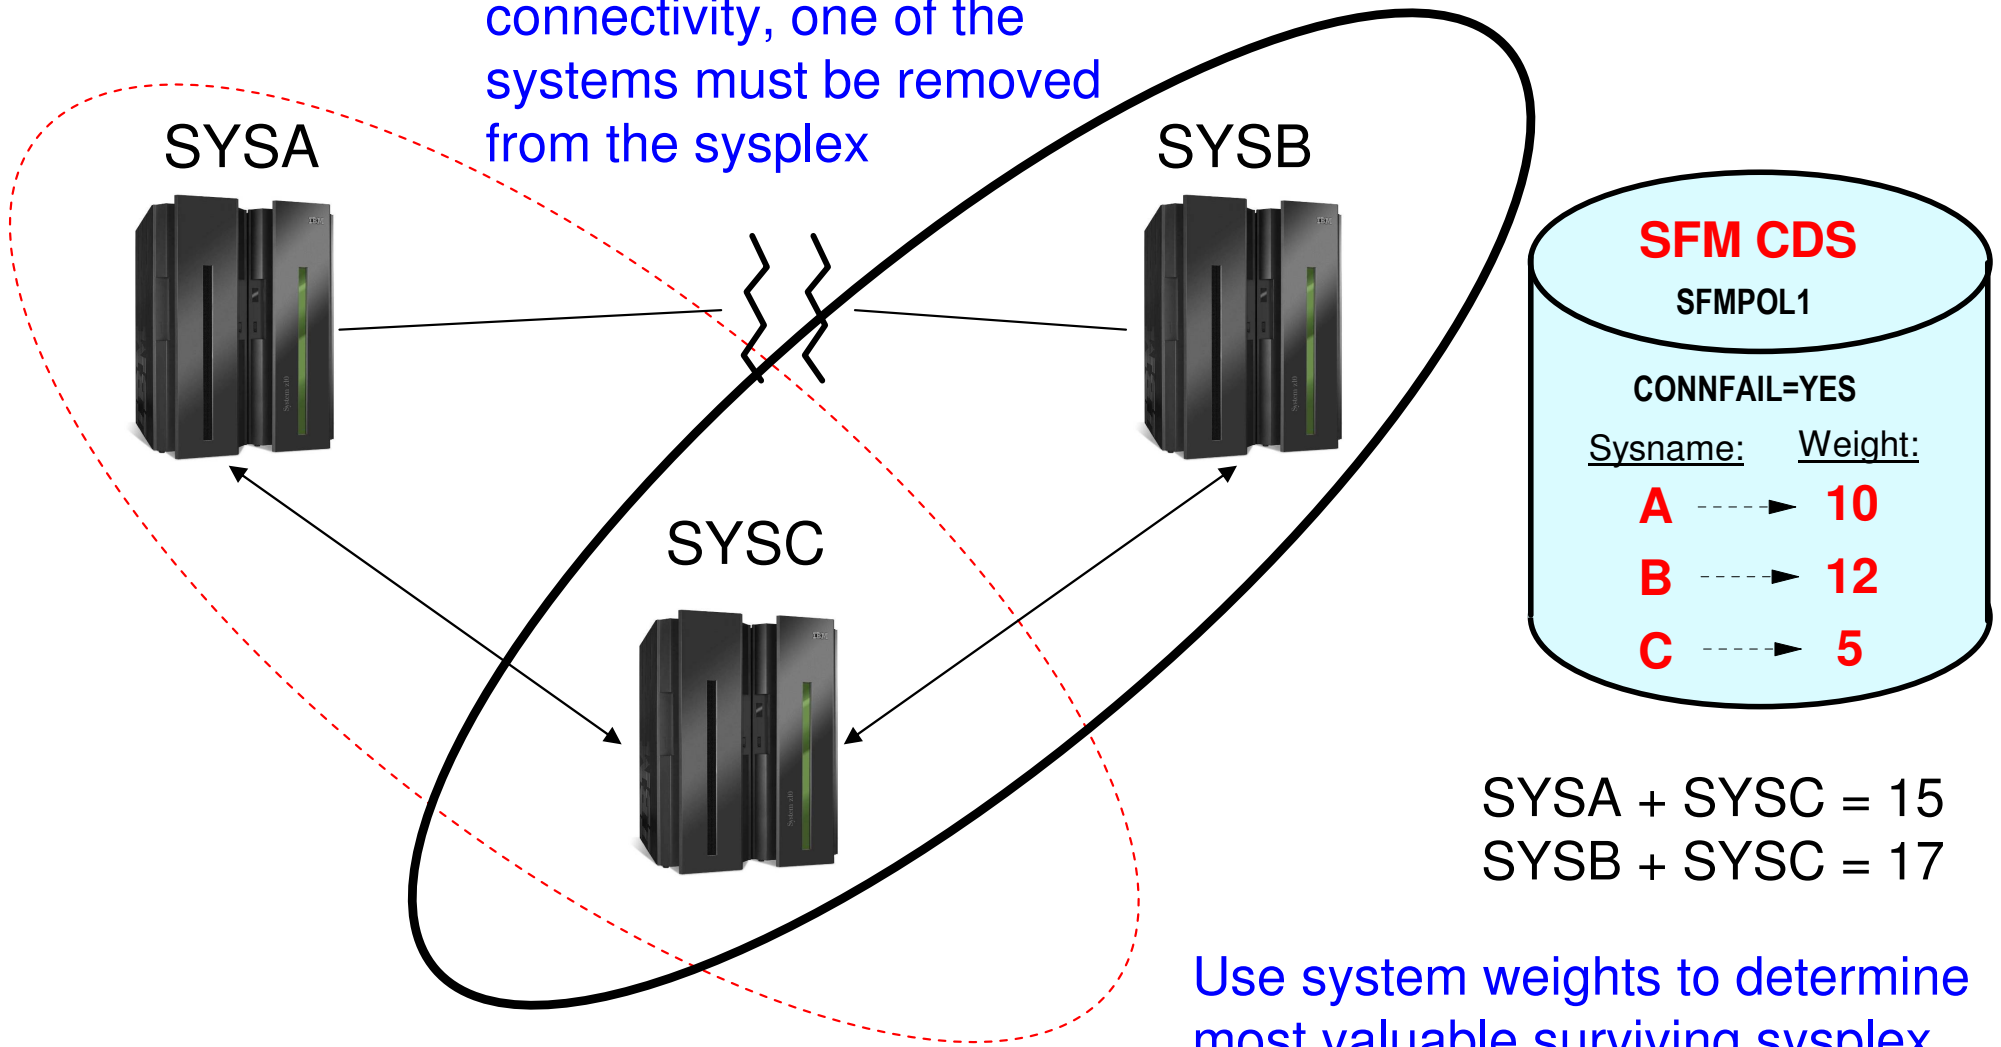
- Use of status update TODs along with the TODs from the signals proved to be a much more reliable indicator
  - When both TODs stop making progress, there is a real good chance that the system failed
  - Issues with respect to “needless” removal largely disappeared
- But what if only one of the indicators is moving?
  - Signal TODs stop but status updates continue, or
  - Status updates stop, but signal TODs continue

# Signals Stop but Status Updates Continue

- Most likely an issue with signalling paths
  - Status updates imply XCF timer DIE is running
  - Which implies signal monitor is running
  - And it would be sending fresh TODs as needed
- System status monitor ignores this case
  - Signal monitor will deal with path problems
  - Restart or stop of inoperative paths may lead to loss of signalling connectivity
- If a pair of systems does not have signal connectivity, one of them must be removed from the sysplex

# SFM and Signal Connectivity Failures

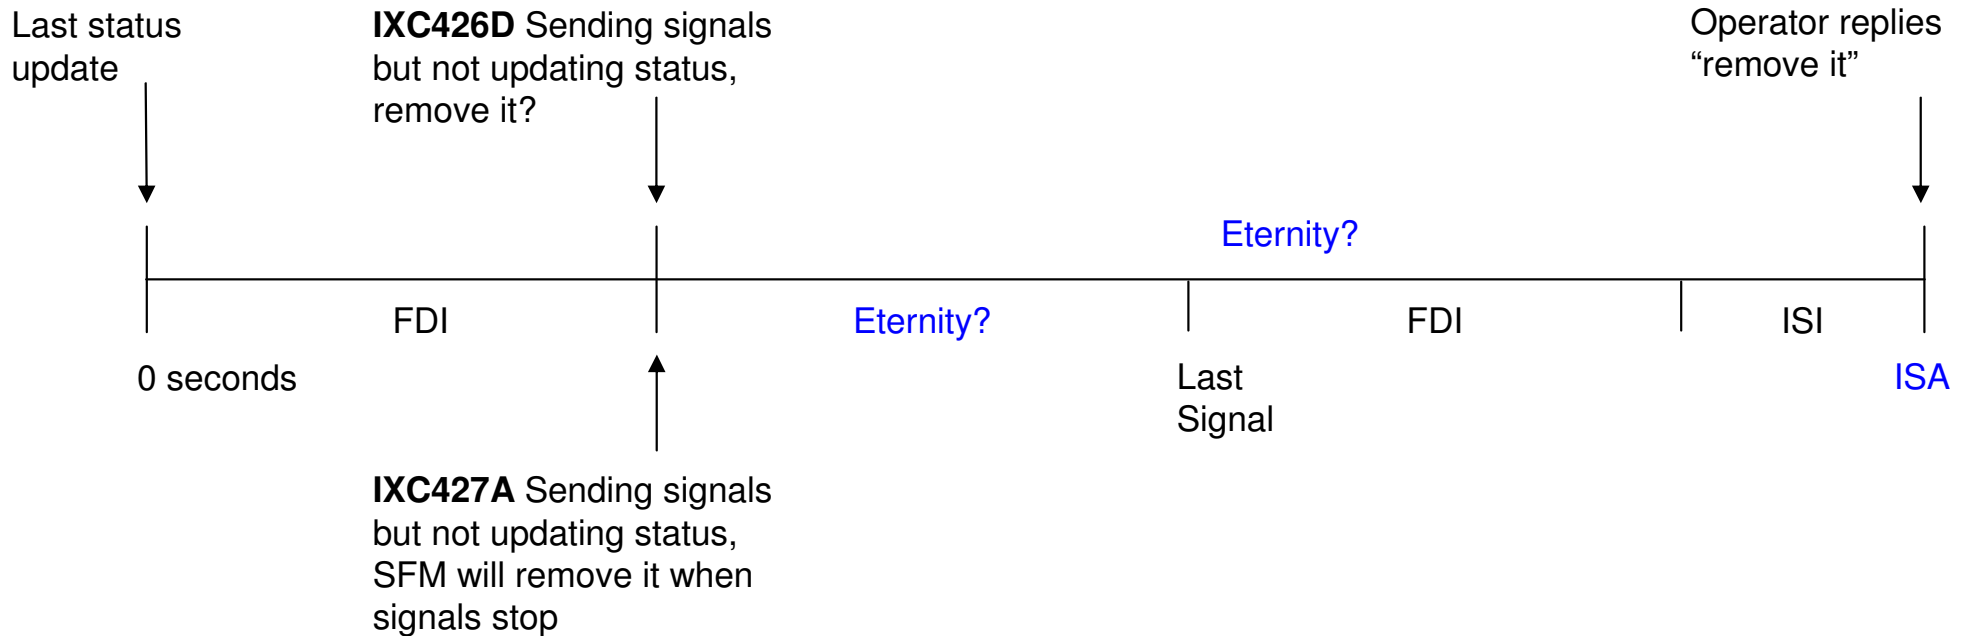
If SYSA and SYSB lose signal connectivity, one of the systems must be removed from the sysplex



# Status Updates Stop but Signals Continue

- System is not healthy
  - Often result of paging issues
  - Could be sysplex CDS issues (contention, performance, reserves, ...)
- Does not meet definition of failed
  - Not subject to automatic removal since signals imply system is still alive
  - So XCF engages the operator

# Missing Status but Sending Signals



Can we shorten eternity ?

FDI = Failure Detection Interval  
 ISI = Indeterminate Status Interval  
 ISA = Indeterminate Status Action



# Can We Shorten Eternity?



- We have a sick but not dead situation
- Sympathy sickness will eventually occur
  - Lack of status updates suggests that sysplex CDS is not readily accessible
  - Join and Leave processing likely impacted
  - Some members record status and control info in CDS
  - Could also prevent systems from being removed from the sysplex
    - Thus preventing systems from getting back in as well

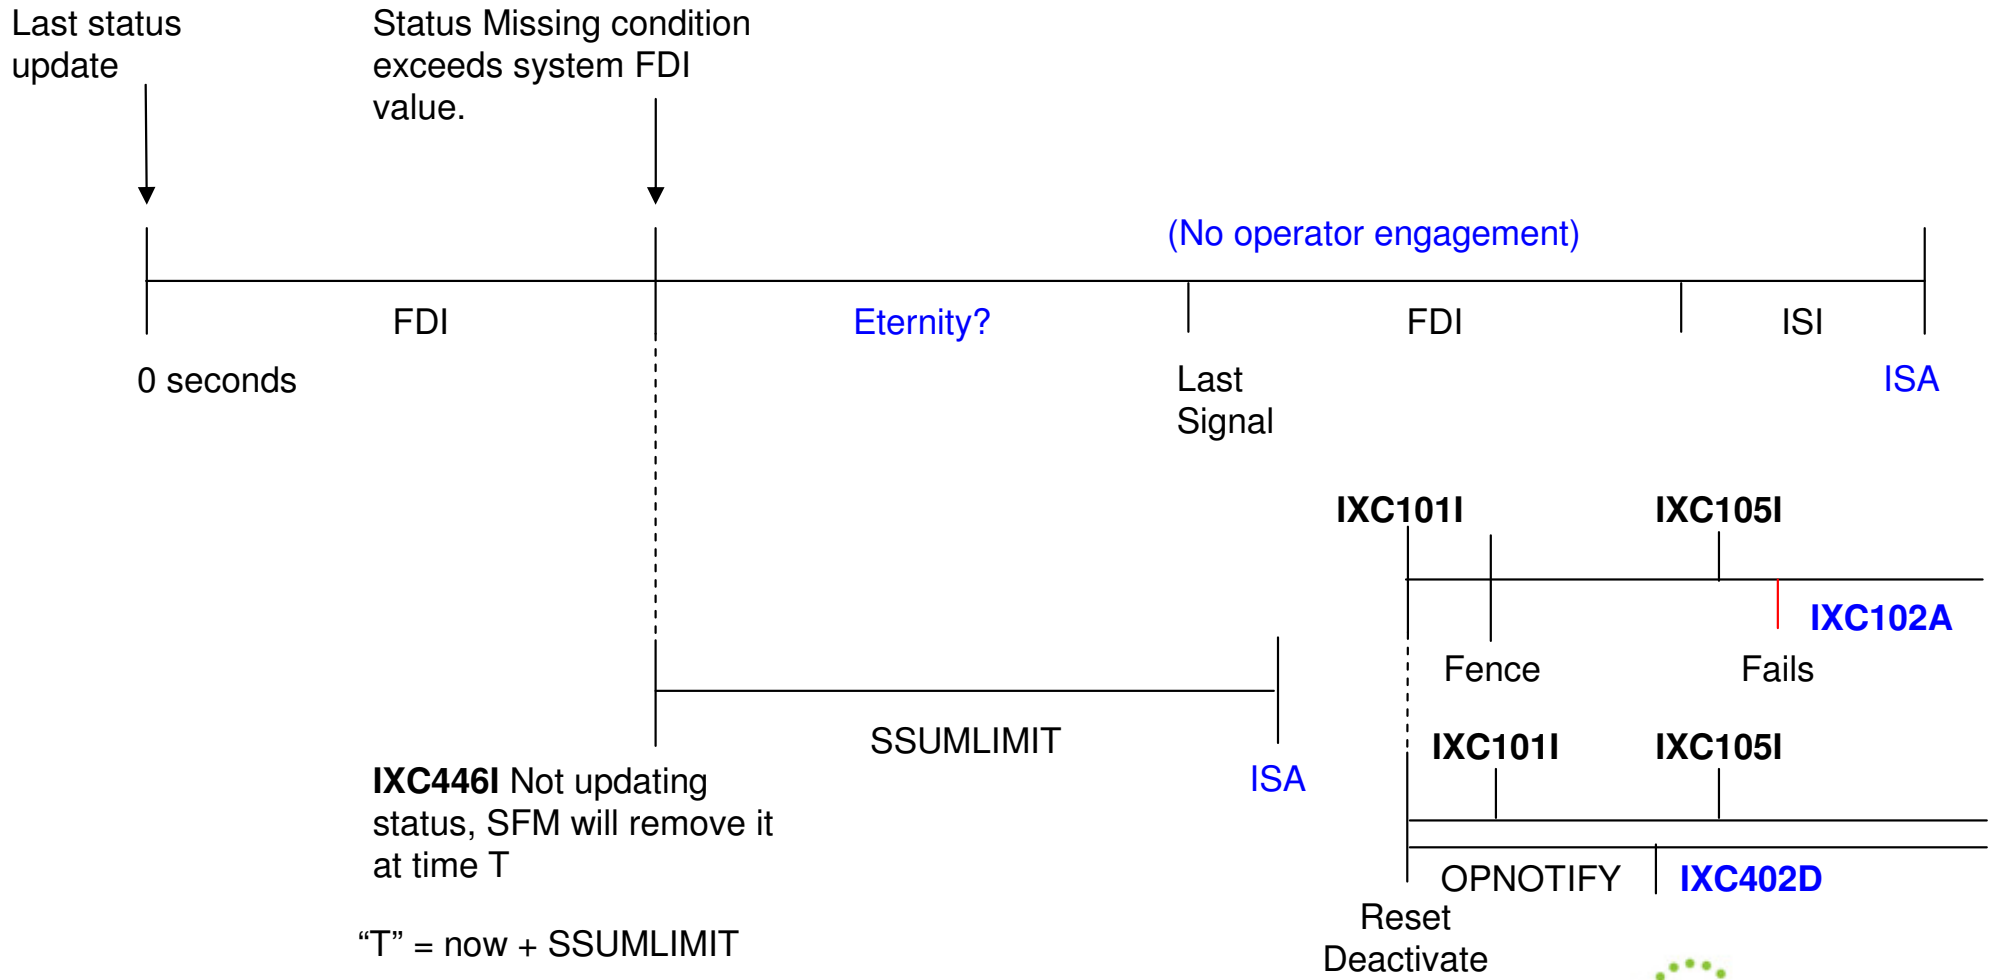
# z/OS 1.9 SSUMLIMIT



- SSUMLIMIT indicates how long a system is allowed to persist in the “not updating status but sending signals” state
  - Allows the installation to “bound” the amount of time that a sick system might impact the remainder of the sysplex
  - When the SSUMLIMIT interval expires, the system will be partitioned from the sysplex
- Not too aggressive, perhaps 15 minutes
  - Zero would be equivalent to the original status monitoring that led to “needless” removal of systems

# Missing Status but Sending Signals with SSUMLIMIT

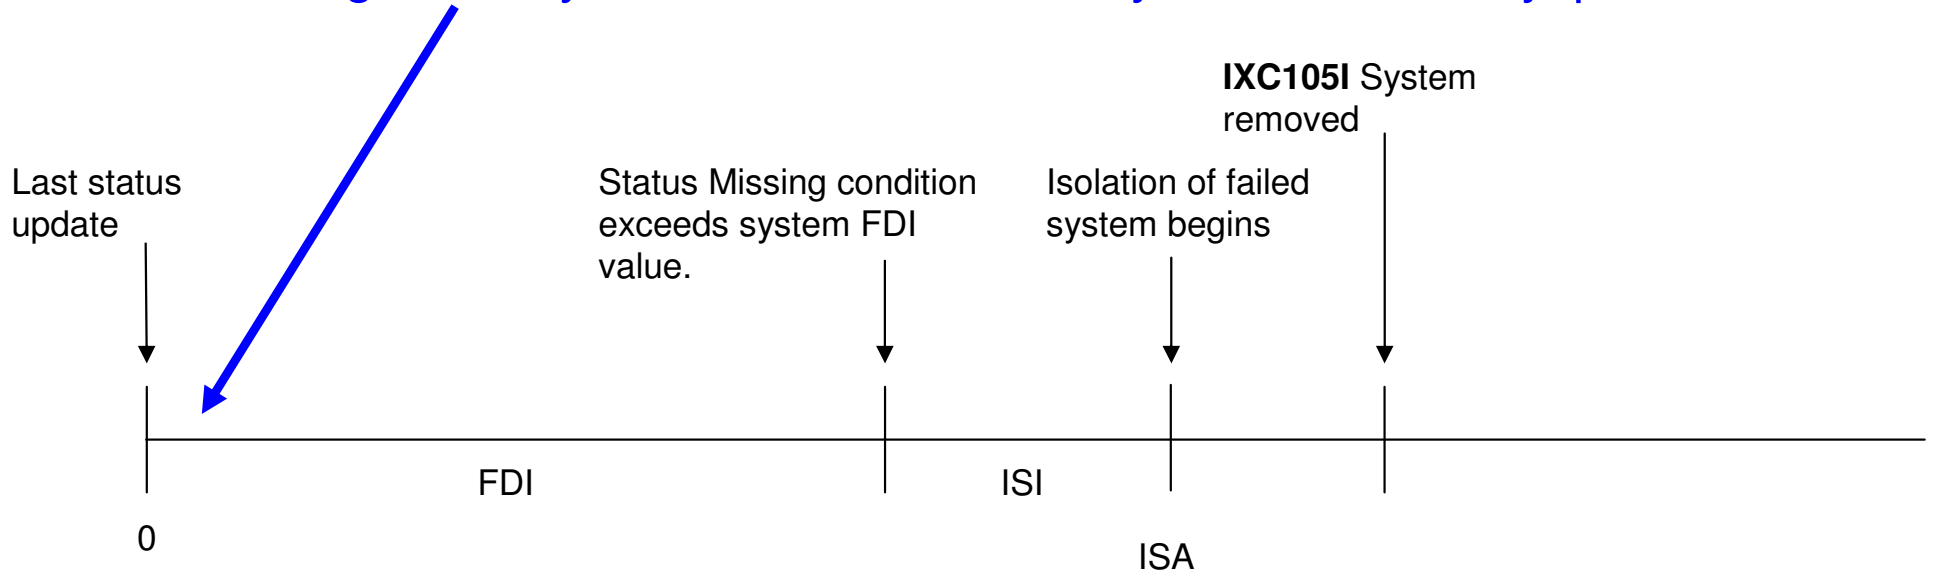
- Race Conditions
- Downstream from ISA, there could be operator engagement



# “Not only merely dead, but really most sincerely dead”

A system that is not updating status and not sending signals may well be dead


If so, SFM could immediately proceed to remove the system from the sysplex



FDI = Failure Detection Interval

# 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

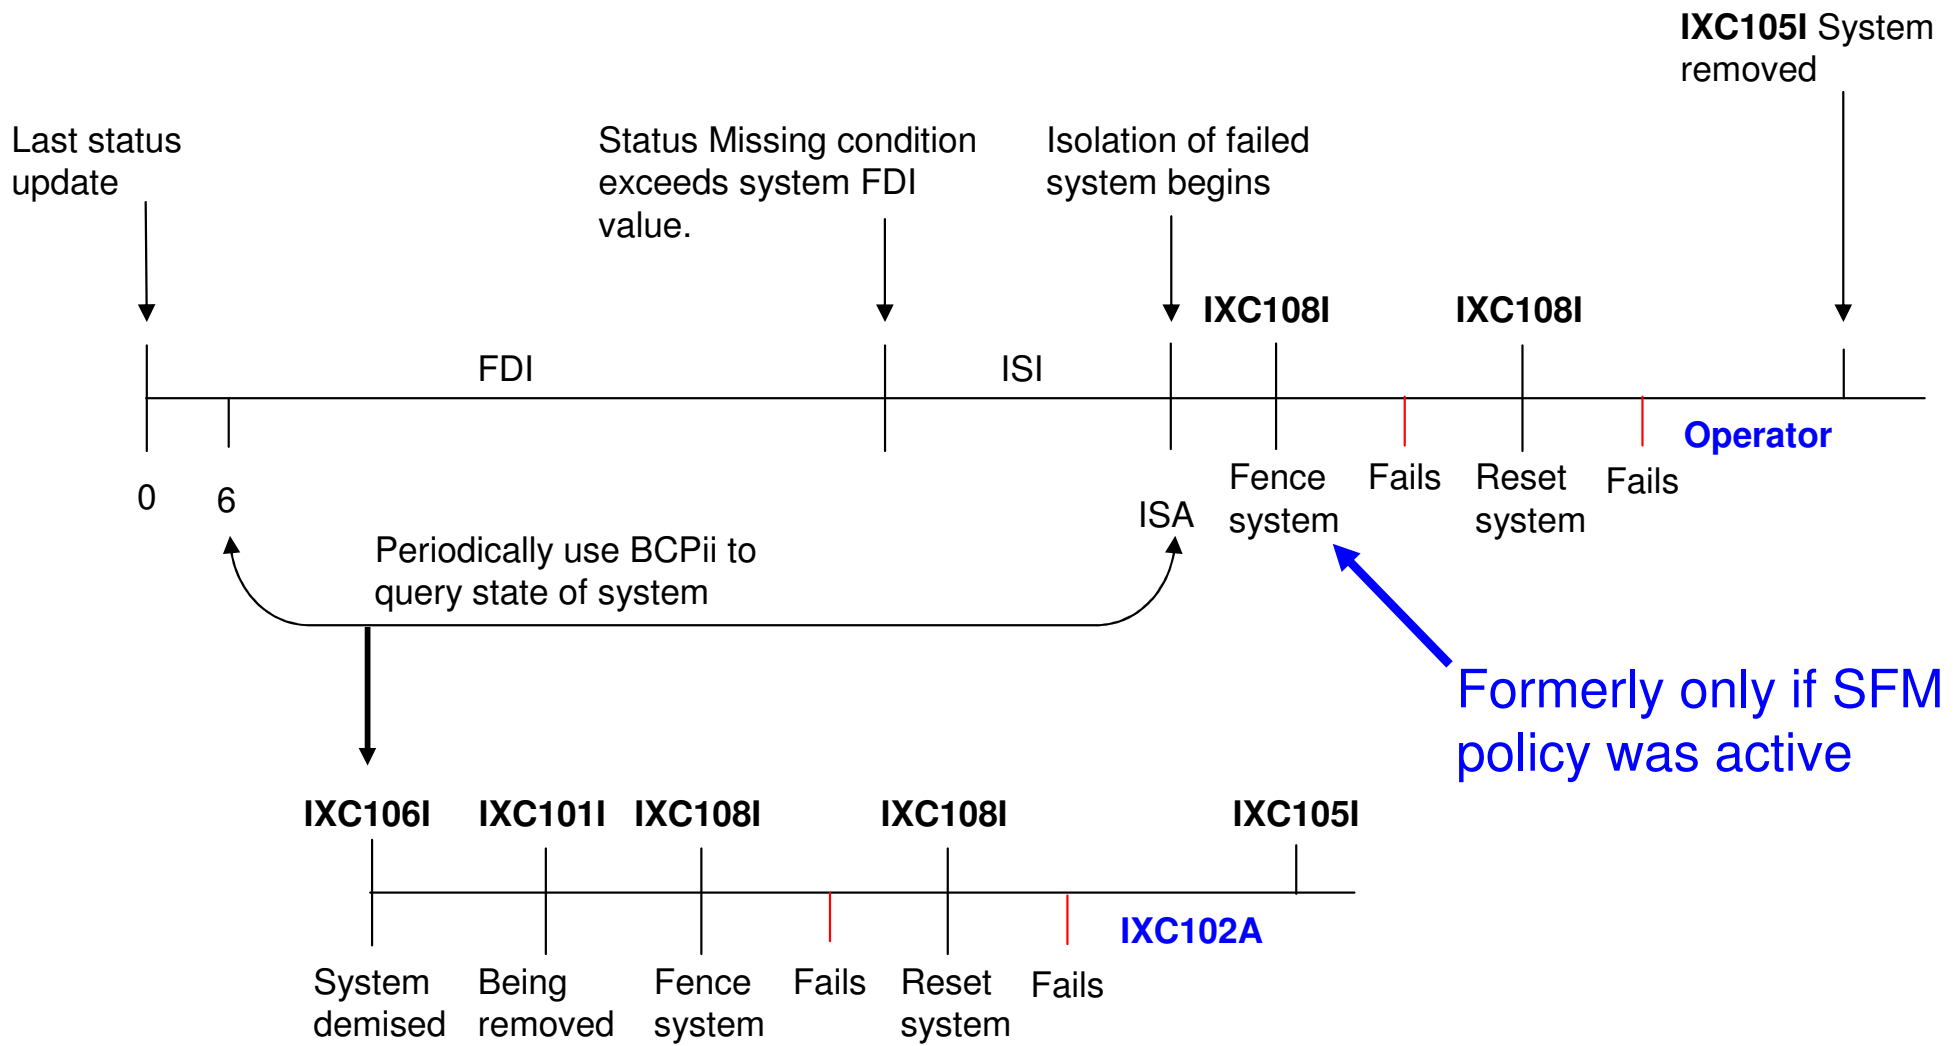
Requires operational  
SE and HMC network

# 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 prevent human error that can cause data corruption

# Unresponsive System with BCPii



FDI = Failure Detection Interval



# z/OS 1.11 SFM with BCPii



- With BCPii, XCF can know system is dead, and:
  - Bypass the Failure Detection Interval (FDI)
  - Bypass the Indeterminate Status Interval (ISI)
  - Bypass the cleanup interval
  - Reset the system even if fencing fails
  - Avoid IXC102A, IXC402D and IXC426D manual intervention
  - Validate “down” to help avoid data corruption

**Helps improve availability**

# 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
  - **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

May need  
MCL  
Fixes !

# “Sick But Not Dead” Refinements



- For cases where it can be known that system is dead, the FDI, ISI, and SSUMLIMIT intervals are irrelevant
- Remain relevant for “sick but not dead” cases
  - Including cases where BCPii cannot ascertain the state of the system
- We will now explore additional refinements that:
  - Reduce “needless” removal
  - Improve “needed” removal

# Refinement to Avoid “Needless” Removal

- FDI needs to be short enough to recognize unresponsive systems before sympathy sickness gets too severe
- Yet long enough to allow the system to overcome “normal” stalls and hangs
- Historically,  $FDI = 2 * spintime + 5$ 
  - Want to allow time for system to recover from an excessive spin condition
  - “2” worked pretty well since first action of ABEND was usually sufficient to break out of the spin
  - But not always ...

# z/OS 1.11 XCF FDI Consistency



- Enforces consistency between the system Failure Detection Interval (FDI) and the excessive spin parameters
  - $FDI = (N+1) * spintime + 5$
- Allows system to perform full range of spin recovery actions before it gets removed from the sysplex
- Avoids false removal of system for a recoverable situation

# z/OS 1.11 XCF FDI Consistency ...



D XCF,C

```
IXC357I 15.12.46 DISPLAY XCF Effective Values E SYS=D13ID71
SYSTEM D13ID71 DATA
```

INTERVAL	OPNOTIFY	MAXMSG	CLEANUP	RETRY	CLASSLEN
165	170	3000	60	10	956

SSUM ACTION	SSUM INTERVAL	SSUM LIMIT	WEIGHT	MEMSTALLTIME
PROMPT	165	N/A	N/A	N/A

```
PARMLIB USER INTERVAL: 60
DERIVED SPIN INTERVAL: 165
SETXCF USER OPNOTIFY: + 5
```

**User FDI**  
**Spin FDI**  
**User OpNotify**  
 - Absolute  
 - Relative

< - - - snip - - - >

OPTIONAL FUNCTION STATUS:

FUNCTION NAME	STATUS	DEFAULT
DUPLEXCF16	ENABLED	DISABLED
SYSSTATDETECT	ENABLED	ENABLED
<b>USERINTERVAL</b>	<b>DISABLED</b>	<b>DISABLED</b>

**Switch**

# Is Spin FDI too long?



- If system is truly dead
  - If detected via BCPii, FDI is irrelevant
  - If BCPii cannot ascertain, detection is elongated
- If system is sick but not dead
  - No status updates, sending signals
    - SSUMLIMIT is key, FDI is “irrelevant”
  - No status updates, not sending signals
    - For spin loops, spin FDI is the desired value
    - If not spin loop, detection is elongated

Even without BCPii, likely OK. But watch and adjust as needed or if concerned. Once BCPii set up, should be rare.

SSUMLIMIT is tens of minutes. Dominates time to resolution

Seldom goes beyond ABEND

Probably Rare

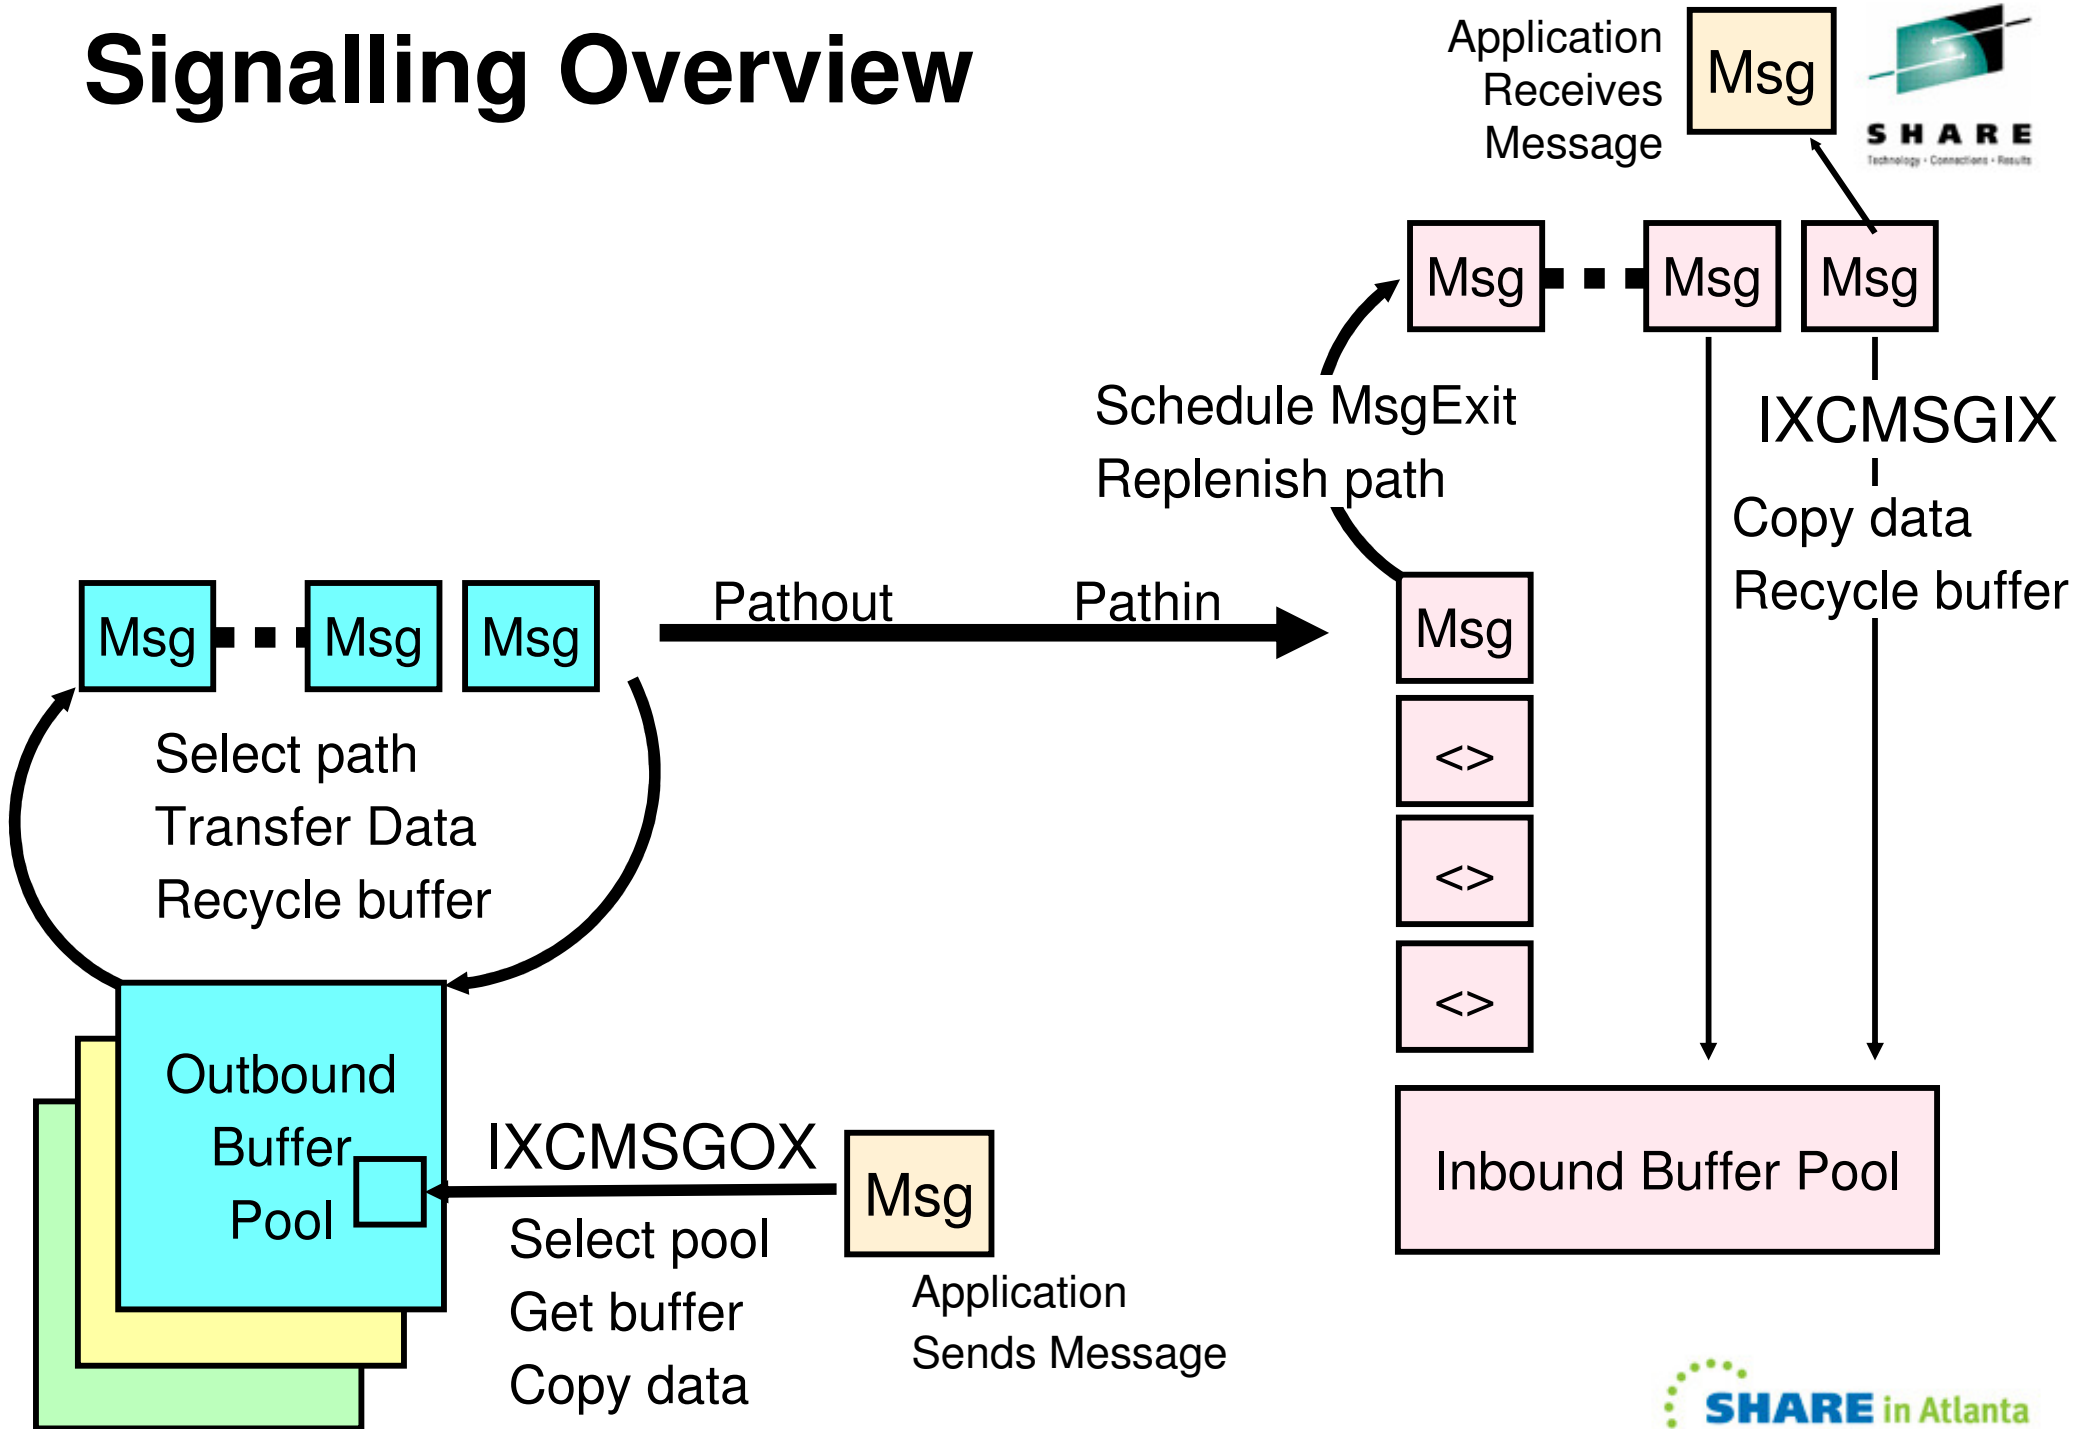


# Refinements for “Sick But Not Dead”



- Signalling Sympathy Sickness
- Unresponsive Critical Members
- Unresponsive CF structure connectors

# Signalling Overview



# z/OS 1.8 Signalling Sympathy Sickness



- XCF detects and surfaces inter-system signalling sympathy sickness caused by stalled group member(s)
- SFM policy MEMSTALLTIME specification determines how long XCF should wait before taking action to resolve the problem
- After expiration, the stalled member is terminated
  - For GRS, XCF, or Consoles, implies system termination
- Provides a backstop that can take automatic action in case your automation or manual procedures fail to resolve the issue

# Signalling Sympathy Sickness Indicators

## Impacted System

- D XCF,G... shows stalls
- IXC467I Restart stalled I/O

### Stalled Members

## Culprit System

- D XCF,G... shows stalls
- IXC431I member stalled
- ABEND 00C 020F0006
- IXC430E stalled members

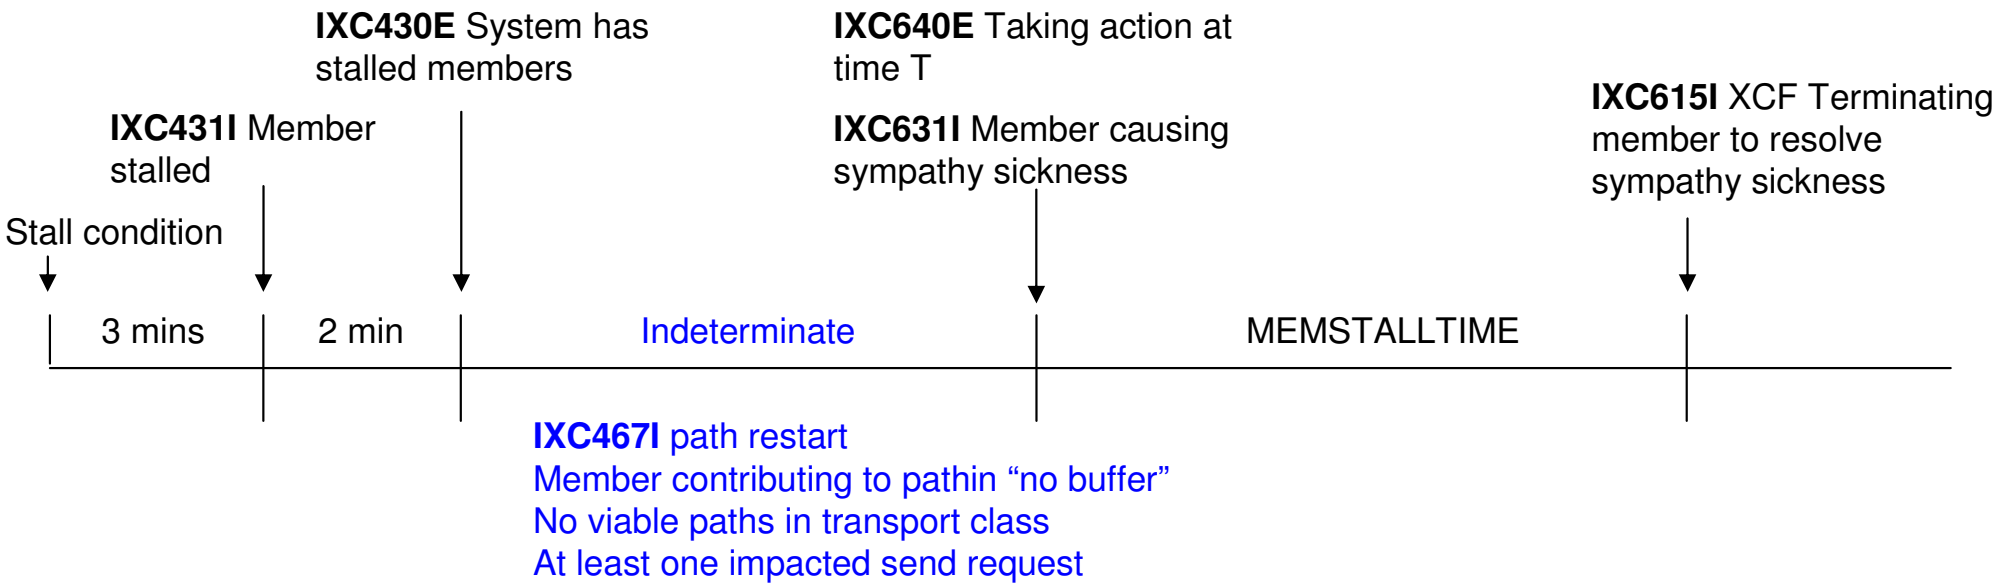
- IXC440E impacted
- ### Sympathy Sickness

- IXC631I mem causing SS
- IXC640E if/when to act
- ABEND 00C 020F000C

If SFM allowed  
to take action

- ABEND 00C 020F000D
- IXC615I terminating
  - ABEND 00C 00000160
  - Wait State 0A2 rsn 160

# SFM for Signalling Sympathy Sickness



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

# z/OS 1.12 Critical Members



- A system may appear to be healthy with respect to XCF system status monitoring
  - Updating status in sysplex CDS and sending signals
- But is the system actually performing useful work?
- There may be critical functions that are non-operational
- Which in effect makes the system unusable, and perhaps induces sympathy sickness elsewhere in the sysplex
- Action should be taken to restore the system to normal operation

# z/OS 1.12 Critical Members ...



- Member Impairment
  - A member is **confirmed** to be impaired when its status exit indicates “status missing”
  - A member is **deemed** to be impaired if it is stalled with no signs of activity
- XCF now surfaces impairment for all members

# z/OS 1.12 Critical Members ...



- A Critical Member is a member of an XCF group that identifies itself as “critical” when joining its group
- If critical member is impaired long enough, XCF will eventually terminate the member
  - Per member’s specification: task, space, or system
  - MEMSTALLTIME determines “long enough”
- GRS is a “system critical member”



# z/OS 1.12 Critical Members ...

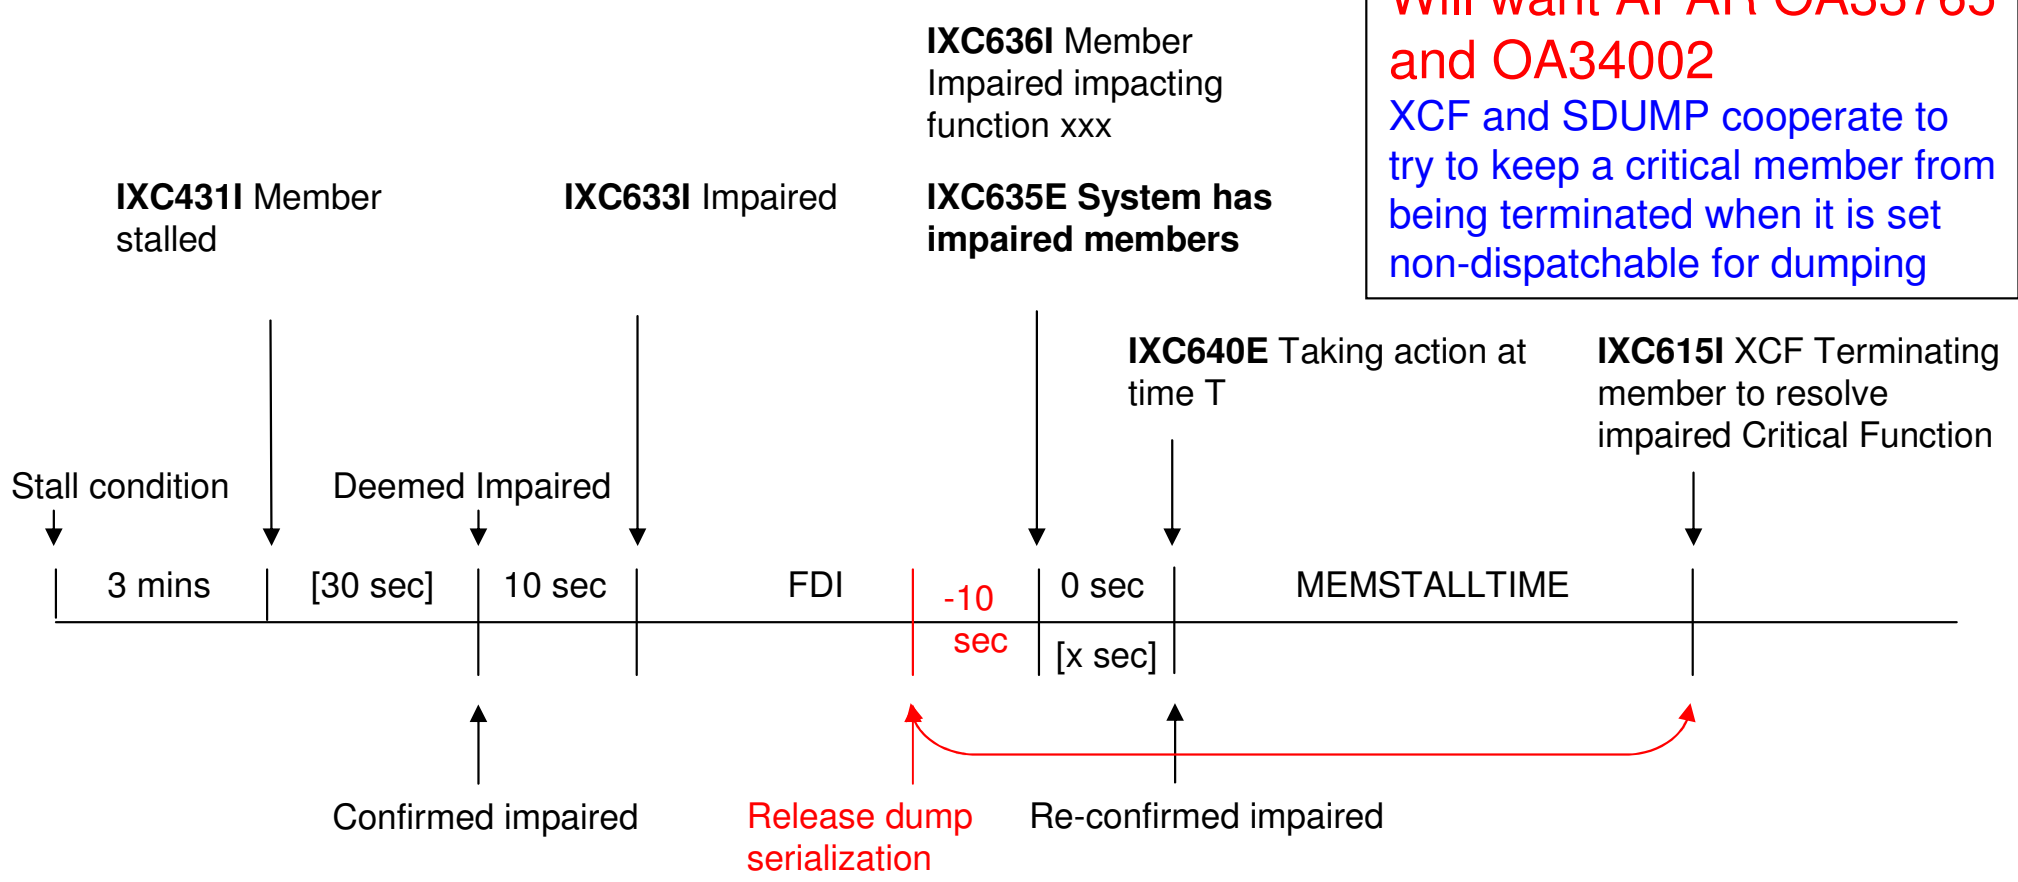


- Key Messages
  - IXC633I “member is impaired”
  - IXC634I “member no longer impaired”
  - **IXC635E “system has impaired members”**
  - IXC636I “impaired member impacting function”

# SFM for Impaired Critical Members



**Will want APAR OA33765 and OA34002**  
 XCF and SDUMP cooperate to try to keep a critical member from being terminated when it is set non-dispatchable for dumping



Stall condition = At least 1 exit stalled for 30 seconds or work item on head of queue for 30 seconds

Deemed impaired = IXC431I issued and for last 30 seconds, either all scheduled user exits stalled or no user exits scheduled

# Unresponsive Structure Connectors



- Connectors to CF structures need to participate in various processes and respond to relevant events
- XES monitors the connectors to ensure that they are responding in a timely fashion
- If not, XES issues messages (IXL040E, IXL041E) to report the unresponsive connector
- Users of the structure may hang until the offending connector responds or is terminated

# z/OS 1.12 CFSTRHANGTIME



- CFSTRHANGTIME
  - An SFM Policy specification to indicate how long the system should allow a structure hang condition to persist before taking corrective action(s) to remedy the situation
  - CFSTRHANGTIME(NO) is the default
- Corrective actions may include:
  - Stopping rebuild
  - Forcing the user to disconnect (signal structures only)
  - Terminating the connector task, address space, or system

# z/OS 1.12 CFSTRHANGTIME ...

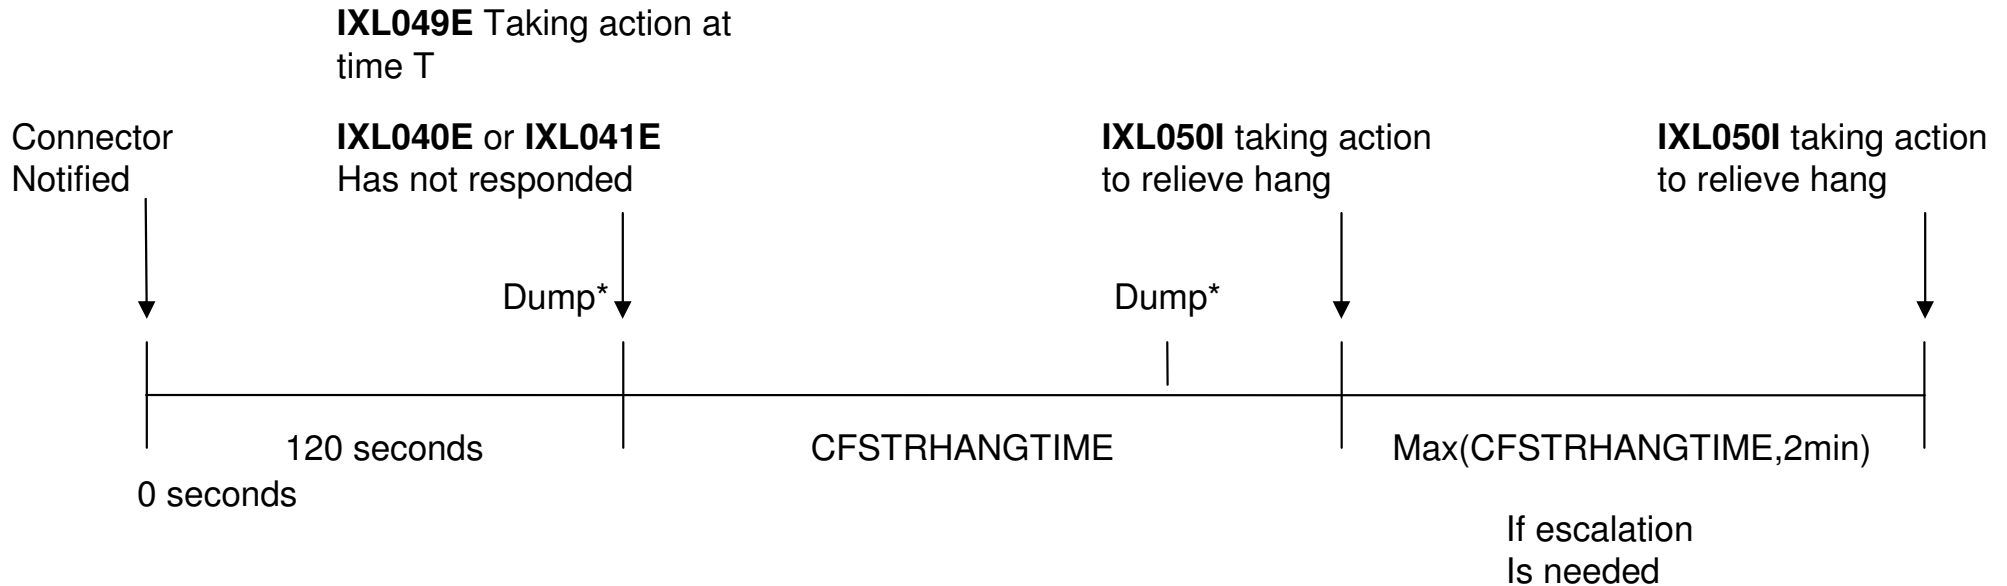


## Messages

```
IXL049E HANG RESOLUTION ACTION FOR CONNECTOR NAME: conname  
TO STRUCTURE strname, JOBNAME: jobname, ASID: asid:  
actiontext
```

```
IXL050I CONNECTOR NAME: conname TO STRUCTURE strname,  
JOBNAME: jobname, ASID: asid  
HAS NOT PROVIDED A REQUIRED RESPONSE AFTER noresponsetime  
SECONDS.  
TERMINATING termtarget TO RELIEVE THE HANG.
```

# SFM for Unresponsive Connectors



Dump\* = Originally, dump taken either when hang announced xor just prior to termination.  
With latest service/release, dump taken always and only when hang is announced

# SFM Suggestions



If using GDPS, use their recommendations

- **Enable SFM with BCPii**
- **SFM Policy Specifications**
  - ISOLATETIME(0) -All releases
  - SSUMLIMIT(900) -z/OS 1.9
  - MEMSTALLTIME(300) -z/OS 1.8 and 1.12
  - CFSTRHANGTIME(900) -z/OS 1.12 (worth watching)
  - CONNFAIL(YES) -All releases, YES is default
- **COUPLExx Specifications**
  - INTERVAL (omit for default) -spin FDI z/OS 1.11
  - OPNOTIFY -All releases, your call

# Summary

- Failing to deal with an unresponsive system in a timely manner can cause sympathy sickness
- Appropriate configuration of the Sysplex Failure Management (SFM) policy allows the systems in the sysplex to automatically take corrective action to resolve sympathy sickness problems when your manual procedures and automation fail to resolve them in a timely manner
- **Enable SFM with BCPii**



# Other Sources of Information



- *MVS Setting Up a Sysplex (SA22-7625)*
- *MVS System Commands (SA22-7627)*
- *MVS System Messages IXC-IZP (SA22-7640)*
- *MVS Initialization and Tuning Reference (SA22-7592)*
- *MVS Programming: Callable Services for High Level Languages (SA22-7613)*
- *Redbook: System z Parallel Sysplex Best Practices (SG24-7817)*