

What Happened to My DB2? The Top Missteps in High Availability

John B. Tobler

Senior Technical Staff Member DB2 for z/OS, IBM

Nigel Slinger

Distinguished Engineer DB2 for z/OS, IBM

Tuesday August 13, 2013 3:00PM - 4:00PM



DB2 for z/OS High Availability

- Platform for high volume OLTP workloads running mission critical applications



- DB2 is designed to help achieve 24/7/365 availability

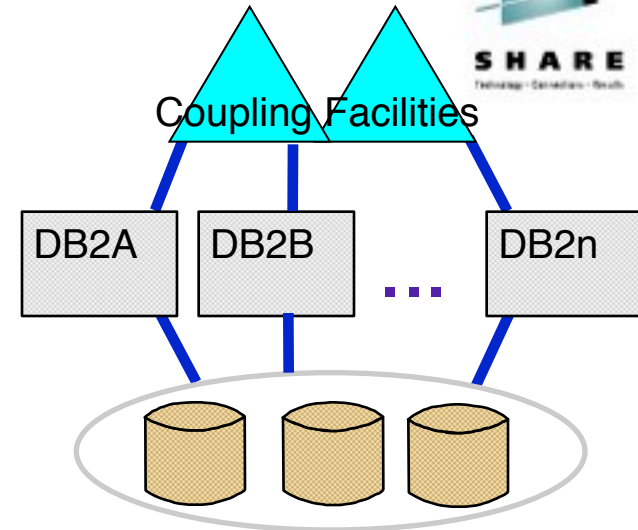
Design for High Availability

- **Data Sharing Provides Infrastructure to Eliminate Single Points of Failure**

- DB2 subsystem or z/OS LPAR
- Hiperswap for disk controller high availability
- Redundant CFs

- **Goal: Continuous availability across planned or unplanned outage of any single hardware or software element**

- Remove all causes for unplanned outages
- Build on legacy of robust, fault tolerant z/OS components
- On a failure
 - Isolate failure to lowest granularity possible
 - Automate recovery and recover fast



High Availability Is Not Achieved Without Consideration

- Experience show us there are pitfalls that can reduce your availability profile

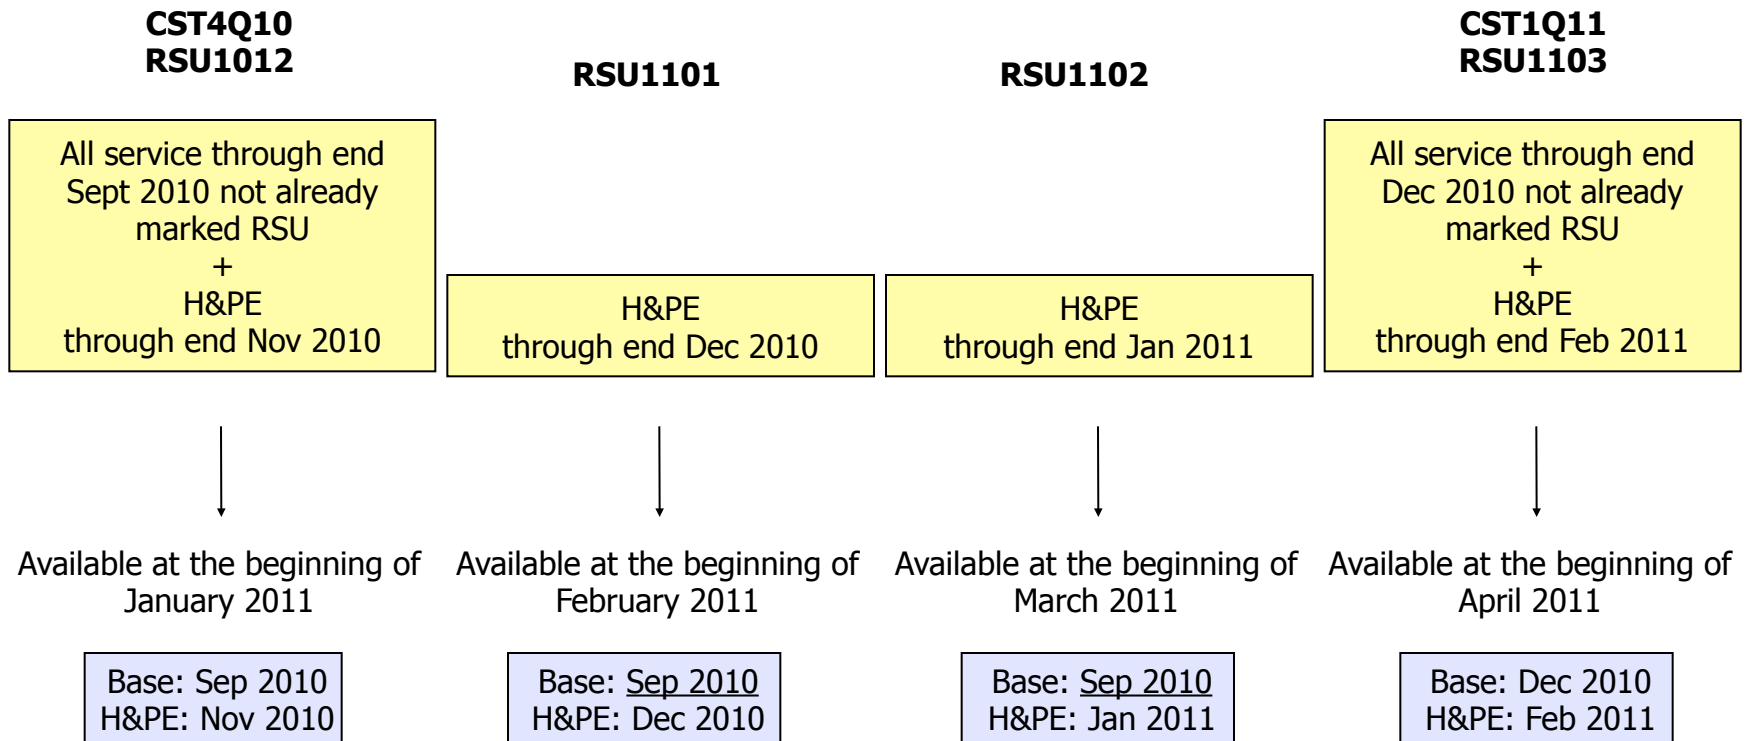


Availability Missteps

- **Maintenance Policies**
- WLM Policies
- Critical Messages
- FORCE or MEMTERM
- Configuration
- Untested DR Situations

Maintenance and RSU

- CST and RSU example



H&PE = HIPER/Security/Integrity/Pervasive PTFs + PE resolution (and associated requisites and supersedes)

Preventative Maintenance Recommendations

- As a priority implement a continuous scheduled program for applying regular DB2 preventative service using CST/RSU method
 - Paramount importance to maintain system availability and stability
- When migrating to a new release, plan on four major preventative maintenance upgrades a year and continue until world wide customer production adoption becomes trailing edge
- Pull and review Enhanced HOLDDATA on at least a weekly basis
- Pro-active checking of all HIPERs and PEs looking for critical problems e.g., data loss, INCORROUT, overlays, crashes, bad restart/recovery, etc.
- Introduce management process and procedure for expediting the apply of the most critical HIPERs after 1-2 weeks in test
- Enhance QA testing to provide better coverage and 'keep fires away' from production
- Application of preventative maintenance can become less aggressive as release becomes trailing edge
 - Perform at least two major preventative maintenance upgrades per year

Availability Missteps

- Maintenance Policies
- **WLM Policies**
- Critical Messages
- FORCE or MEMTERM
- Configuration
- Untested DR Situations

WLM Policies

- WLM controls dispatchable work on z/OS. Work is assigned to a service class. For each service class, you may define the desired performance and importance of the work.
- Poorly tuned systems can result in stalled work. Stalled work may result in a hung subsystem or even a hung sysplex.
 - Held latches
 - Held locks
 - p-lock negotiation not completing



WLM Recommendations

- The most critical work on the system should run at the highest priority level (a.k.a. SYSSTC). This includes
 - IRLM
 - TCP/IP
- DB2 MSTR and DBM1 should run with a high priority (typically importance 1)
 - DB2 monitor task in MSTR should get priority over user threads
 - p-lock negotiation, notifies, pageset opens are controlled by DBM1 priorities
- CPU critical can be considered if user work has very aggressive goals. This prevents the user work from getting a higher priority than DB2.

WLM Policy Common Issues

- **Problem**

- Running low priority work on a highly utilized system. This is commonly batch work running at low (e.g. importance 5) priority. The work gets stalled holding critical resources -- most commonly latches.

- **Considerations**

- Be sure to enable blocked workload support (enabled by default). This is enabled via the BLWLINTHD parm in IEAOPTxx. The default is 20 seconds but may be lowered to 5 or so seconds. Note that a 5 second latch hang still may be a 'long time' and can disrupt system performance.
- Issue `-DISPLAY THREAD(*) SERVICE(WAIT)` to see long waiters. This only shows waiters suspended for > 2 minutes.
- Rely on the DB2 monitor tasks to boost work. This only works on waiters suspended for > 2 minutes.
- ★ Avoid running >>100% for long durations. Use CPU critical to avoid user work being dispatched over DB2.

WLM Policy Common Issues

- **Problem**

- Running a workload with WLM resource groups enabled, the work gets stalled holding critical resources -- most commonly latches.

- **Considerations**

- WLM provides resource group support to limit a service class to a maximum amount of processing capacity.
- Once this limit has been reached, WLM will slow resource consumption potentially making such work non-dispatchable.
- ★ Unless you have some special need to limit or protect processor capacity for a group of work, you should skip defining resource groups and let workload management manage all of the processor resources to meet performance goals.

WLM Policy Common Issues


- **Problem**

- Running a workload with LPAR soft capping enabled, work gets stalled holding critical resources -- most commonly latches.

- **Considerations**

- WLM provides LPAR level soft-capping to limit the 4 hour capacity average to a defined limit.
- If the 4 hour limit reaches that value, capacity is capped.
- ★ If LPAR soft-capping is used, make sure the defined limit can still handle the workload. If capacity is significantly limited, reaching the cap may result in performance and potentially availability problems.
- ★ Another consideration is to configure CPs offline. This reduces capacity in a more consistent fashion vs. soft capping as soft capping allows the capacity to exceed the limit as long as the 4 hour average does not.

Not So Fast...WLM and Specialty Engines



**Internal Coupling
Facility (ICF) 1997**



**Integrated Facility
for Linux® (IFL)
2000**



**System z Application
Assist Processor (zAAP)
2004**

Eligible for zAAP:

- Java execution environment
- z/OS XML System Services



**IBM System z Integrated
Information Processor
(zIIP) 2006**

Eligible for zIIP:

- DB2 remote access, BI/DW, Utilities Build Index and Sort processing, XML Parsing, RUNSTATS, BP Prefetch, Deferred Write
- z/OS XML System Services
- HiperSockets for large messages
- IPsec encryption
- z/OS Global Mirror (XRC)
- IBM GBS Scalable Architecture for Financial Reporting
- z/OS CIM Server
- Select Tivoli products

WLM and Specialty Engines

- Work eligible to run on a specialty engine may fail over to the general CP if the zIIP or zAAP is saturated.
 - Controlled by IIPHONORPRIORITY and IFAHONORPRIORITY for zIIP and zAAP respectively (default is YES with NO meaning no failover to general CP)
 - Failover to CP is additionally controlled by ZIIPAWMT and ZAAPAWMT (default in the 3 - 12 ms depending on HIPERDISPATCH)

WLM Specialty Engine Common Issues

- **Problem**

- Running > 100% on specialty engines with IIPHONORPRIORITY = NO or ZIIPAWMT >> default value (same for zAAP equivalents) and work becomes stalled holding a critical resource waiting for zIIP to be available.

- **Considerations**

- Use ZIIPAWMT default or set IIPHONORPRIORITY = YES

NOTE: DB2 v11 offloads all system SRBs to zIIP except p-lock engines. This is disabled completely if IIPHONORPRIORITY = NO!

- ★ Avoid running with honor priority set to NO or with large failover time with specialty engine CPU > 100% for long durations

Availability Missteps

- Maintenance Policies
- WLM Policies
- **Critical Messages**
- FORCE or MEMTERM
- Configuration
- Untested DR Situations

Critical Messages

- The goal of DB2 is to be a self-healing, automated database. There are some cases where action or special attention may be needed.
- There are many messages deemed important (see Appendix A) and below are a few highlighted messages.
- **Problem #1:**
 - DB2 is out or almost out of ACTIVE log space and may hang the system as logging comes to a halt.
 - **DSNJ110E** and **DSNJ111E** (J111E is non-scrollable indicating no active logs are available!)
 - **Considerations:** There can be different causes:
 - Insufficient active log space
 - Excessive use of -ARCHIVE LOG command
 - Offload is not progressing or failed (e.g. **DSNJ115I**)

Critical Messages Continued...

- **Problem #2:**
 - p-lock negotiation is not progressing which can lead to a sysplex hang
 - IXL045E and perhaps DXR167I message is issued
 - **Considerations:** If pervasive this should be investigated. The problem may be dispatching priority of DBM1, a logging issue, etc. Dumps may be required for additional analysis.
- **Problem #3:**
 - DBM1 is nearly out of 31-bit virtual storage (largely a v9 and prior issue) which may result in performance issue or loss of a DB2.
 - DSNV508I message is issued
 - **Considerations:** Member may be constrained virtually but this also may be normal for predictable workloads. Other considerations are RELEASE(COMMIT/DEALLOCATE) and application behavior.

Critical Messages Continued...

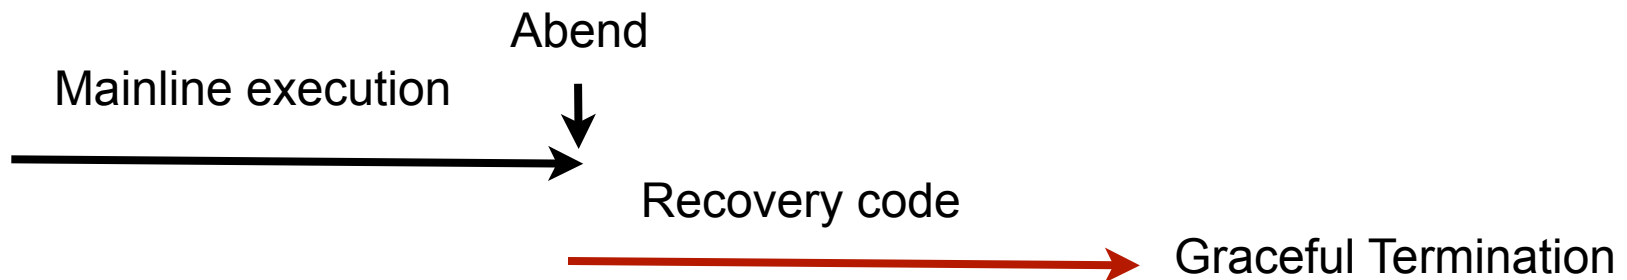
- **Problem #4:**
 - The system is consuming a large amount of real storage and DB2 may crash or the LPAR may be lost.
 - **DSNS003I** message (non-scrollable message indicating >80% of REALSTORAGE_MAX is consumed)
 - Additionally may see **IRA200E** or **IRA201E** messages indicating > 70% and 85% of aux storage is consumed.
 - **Considerations:** DSNS003I may indicate a storage problem within DB2. Use monitors to help analyze. The LPAR may also be underprovisioned and additional memory may be needed.

Availability Missteps

- Maintenance Policies
- WLM Policies
- Critical Messages
- **FORCE or MEMTERM**
- Configuration
- Untested DR Situations


FORCE or MEMTERM

- DB2 exploits z/OS recovery services in the form of FRRs, ESTAEs, etc. for robust fault tolerance
 - User initiated ABENDs should be tolerated. Examples are -CANCEL THREAD (ABEND04E RC00E50013), canceling an address space or detaching a TCB (ABEND13E, ABEND33E, ABEND222, etc.)
 - DB2 recovery gets control and releases critical resources (i.e., latches, locks, storage)



FORCE or MEMTERM

- When an address space MEMTERMs, functional recovery does not get a chance to execute.



- Because recovery does not execute, DB2 may terminate if a critical resource is held (RC00F30801)
- MEMTERMs can be explicit or implicit
- Common causes are applications that leak storage (ABEND878), aggressive termination of address space via FORCE command

FORCE or MEMTERM

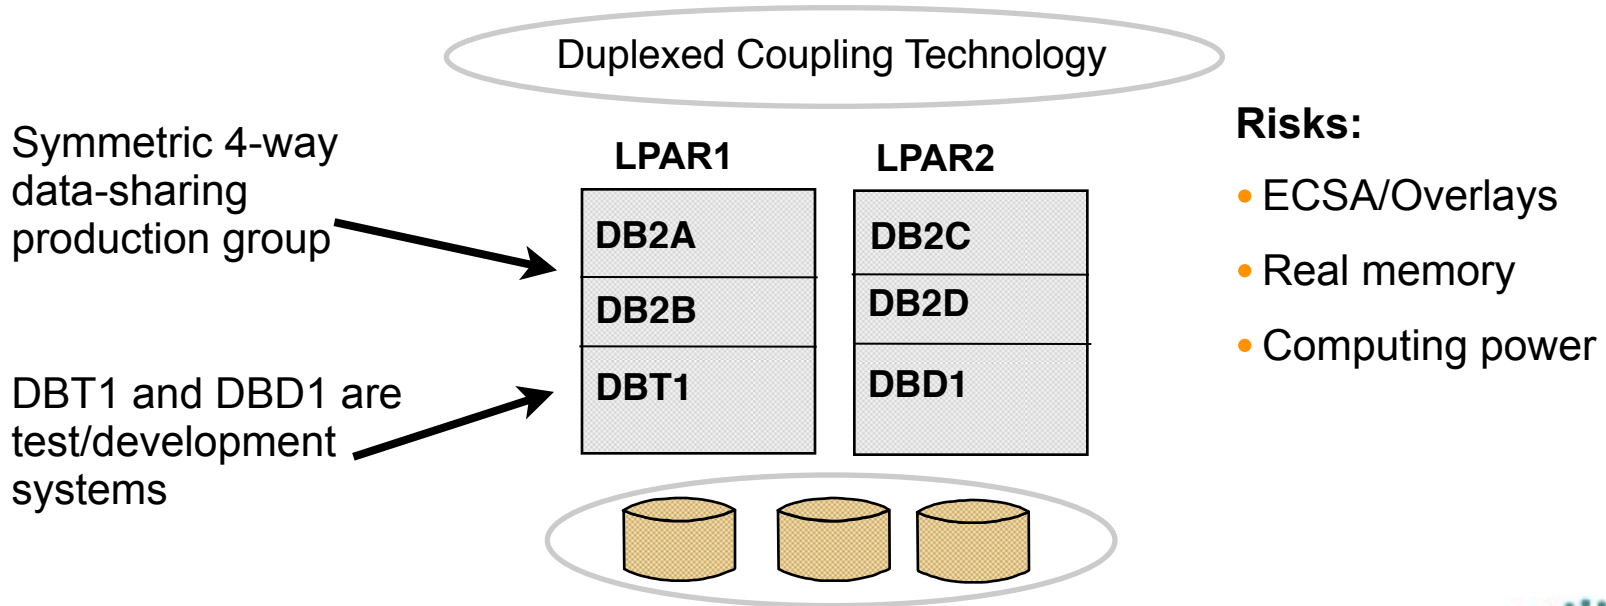
- ★ Use the appropriate hammer to terminate work
 - -CANCEL THREAD() -- this is a best effort cancel that may take a few seconds to be accepted
 - Cancel the address space
 - Note: If a cancel is accepted, abort processing must complete before termination. Abort processing can be long duration and a force will not expedite that processing and may jeopardize the subsystem. -DISPLAY THREAD TYPE(SYSTEM) will show threads that are in abort/deallocation processing after a cancel termination (**DSNV504I**)
 - If the work appears to be stuck, consider a force of the address space with the awareness that DB2 may terminate.

Availability Missteps

- Maintenance Policies
- WLM Policies
- Critical Messages
- FORCE or MEMTERM
- **Configuration**
- Untested DR Situations

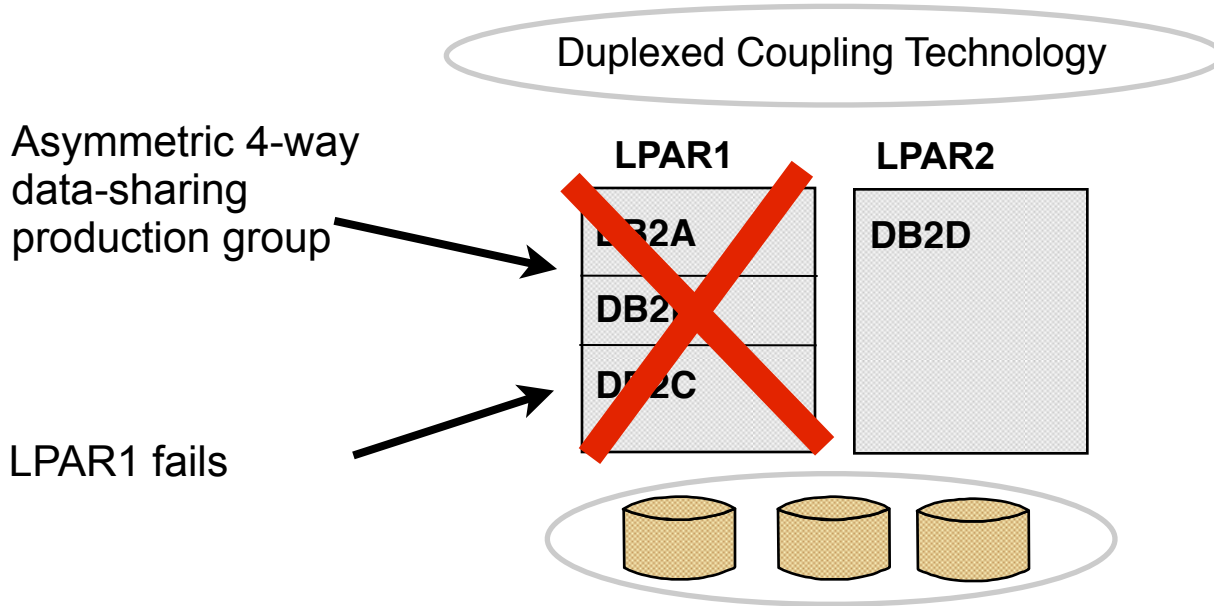
Design for High Availability?

- Data Sharing provides infrastructure to eliminate single points of failure
- The following configurations are some that may lessen the availability profile of a production system
- **Configuration #1:**



Design for High Availability?

- Configuration #2:

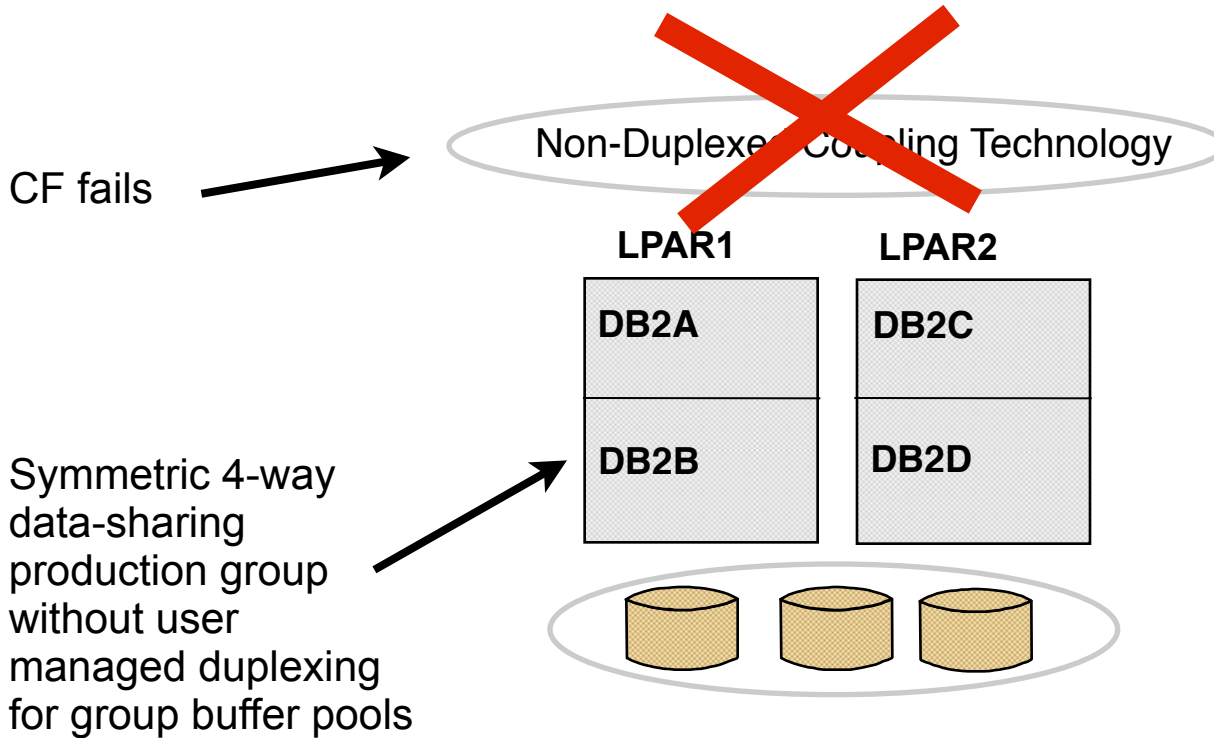


Risks:

- Assymmetric design may result in group wide outage as the remaining members cannot handle the workload

Design for High Availability?

- **Configuration #3:**

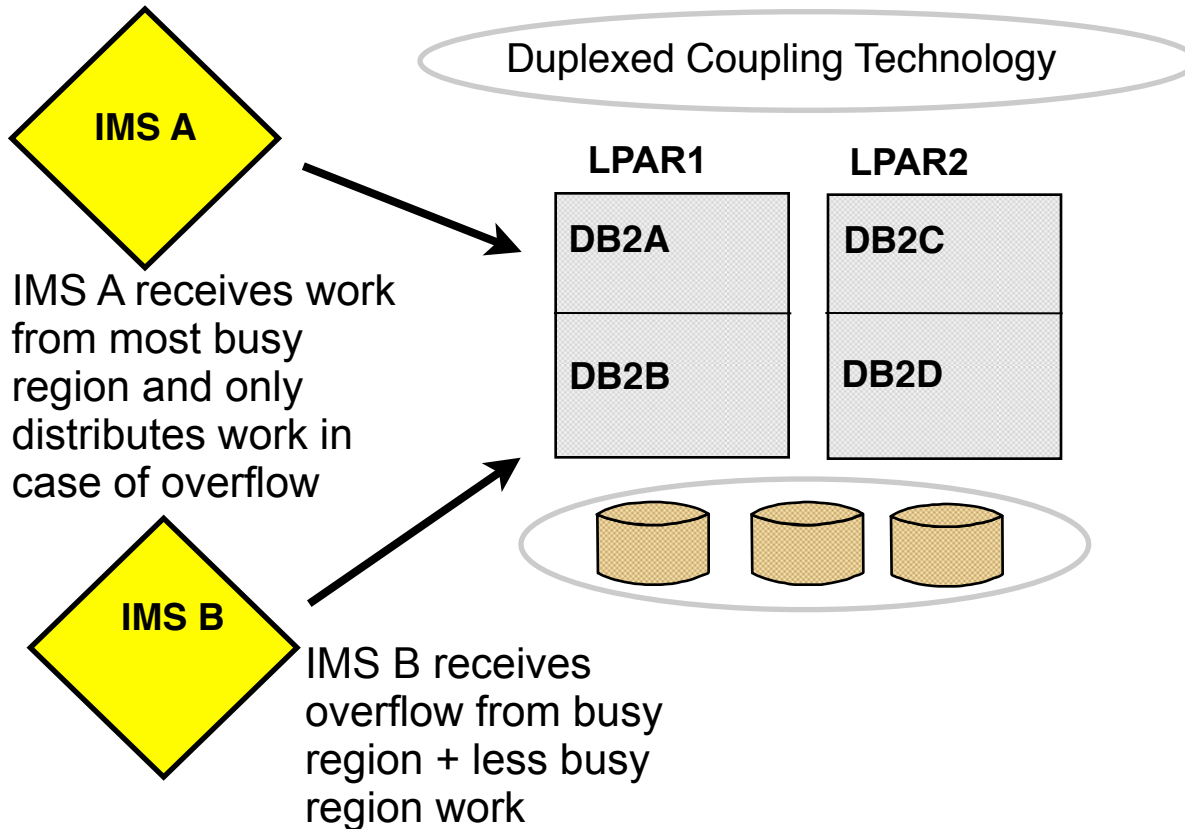


Risks:

- Group outage because group buffer pools cannot be rebuilt from existing members
- Note: SCA and the lock structure can be rebuilt

Design for High Availability?

- **Configuration #4:**



Risks:

- Sudden contention in the CF due to 2 members now accessing the same objects
- Affinities should at the very least be understood (i.e., IMS A -> DB2A). Considerations may be member cluster, better workload balancing, data partitioning etc.)

Availability Missteps

- Maintenance Policies
- WLM Policies
- Critical Messages
- FORCE or MEMTERM
- Configuration
- **Untested DR Situations**

Untested DR Situations

- DR is needed in the case of catastrophic failures
- Broad ranging subject that requires a deep dive but some basic steps should be followed
 - Make sure the DR process is tested. This is most often dictated by government standards. A recent real-world example had different JCL for the DB2 address spaces at the DR site such that a deadlock resulted on the BSDS and DB2 would not start.
 - Make sure all required jobs are available to participants in the DR. Copies in my personal directory can become a problem if I am on vacation.
- Disaster Recovery Redbook
<http://www.redbooks.ibm.com/redbooks/pdfs/sg246370.pdf>

Thanks!

jtobler@us.ibm.com

nigel@us.ibm.com

Appendix A

Message	
DSNI012I	PAGE LOGICALLY BROKEN
DSNI014I	DATA IN USE DURING ABEND
DSNI031I	LOCK ESCALATION
DSNJ004I	LOG IN SINGLE MODE
DSNJ100I	ERROR OPENING BSDS
DSNJ103I	LOG ALLOCATION ERROR
DSNJ107I	READ ERROR ON BSDS
DSNJ108I	WRITE ERROR ON BSDS
DSNJ110E	LAST COPY _n ACTIVE LOG DATA SET IS _{nnn} PERCENT FULL
DSNJ111E	OUT OF SPACE IN ACTIVE LOG DATA SETS
DSNJ114I	ERROR ON ARCHIVE DATA SET
DSNJ115I	OFFLOAD FAILED, COULD NOT ALLOCATE AN ARCHIVE DATA SET
DSNJ125I	ERROR DUMPING BSDS
DSNJ128I	LOG OFFLOAD TASK FAILED
DSNL008I	DDF ABNORMAL TERMINATION COMPLETE
DSNL030I	DDF PROCESSING FAILURE FOR LUWID
DSNL501I	CNOS PROCESSING FAILED FOR LU
DSNP002I	DEFINE FAILED FOR <i>data-set-name</i>
DSNP007I	EXTEND FAILED FOR <i>data-set-name</i>
DSNP011I	MEDIA MANAGER SERVICES ERROR FOR <i>data-set-name</i>
DSNP031I	CURRENT SPACE HAS <i>cc</i> EXTENTS FOR <i>data-set-name</i>
DSNR035I	WARNING - UNCOMMITTED UR AFTER _{nn} CHECKPOINTS
DSNX906I	PROCEDURE OR FUNCTION <i>name</i> TERMINATED ABNORMALLY

Appendix A Continued...

Message	
DXR142E	THE LOCK STRUCTURE wwwwww IS zzz% IN USE
DXR167E	IRLM HAS DETECTED A DELAY IN COMPLETION OF vvvvvv PROCESS
DXR170I	THE LOCK STRUCTURE wwwwww IS zz% IN USE
DSNB250E	A PAGE RANGE WAS ADDED TO THE LOGICAL PAGE LIST
DSNB260I	WARNING - A READER HAS BEEN RUNNING FOR xx MINUTES
DSNB319A	SHORTAGE OF SPACE IN GROUP BUFFER POOL
DSNB325A	CRITICAL SHORTAGE OF SPACE IN GROUP BUFFER POOL
DSNB331I	REBUILD STARTED FOR GROUP BUFFER POOL
DSNB743I	DUPLEXING IS BEING STOPPED FOR GROUP BUFFER POOL - FALLING BACK TO
DSNB744I	DUPLEXING IS BEING STOPPED FOR GROUP BUFFER POOL - SWITCHING TO
DSNI006I	ASYNCHRONOUS GRECP OR LPL RECOVERY IN PROGRESS
DSNJ008E	nn OF mm ACTIVE LOGS ARE FULL. ssname NEEDS ARCHIVE SCRATCH. REPLY Y
DSNJ017E	WARNING - OFFLOAD TASK HAS BEEN ACTIVE SINCE xxx AND MAY HAVE
DSNJ031I	WARNING - UNCOMMITTED UR HAS WRITTEN xxx LOG RECORDS
DSNT318I	TIMED OUT DUE TO A P-LOCK HELD BY MEMBER xxxx
DSNV508I	DSNVMON - DB2 DBM1 BELOW-THE-BAR STORAGE NOTIFICATION