IDUG The Worldwide DB2 User Community

Session: K14

IMS and GDPS Disaster Recovery Mirroring Solutions

Glenn Galler
IBM SW IT Specialist (IMS)
Advanced Technical Support
Ann Arbor, Michigan
gallerg@us.ibm.com

IDUG 2008 North America

Experience IDUG

May 22, 2008 • 10:30 a.m. – 11:30 a.m.
Platform: System z

IDUG
The Worldwide DB2 User Community

IDUG 2008 North America

2

Trademarks.

Agenda

- IMS Business Continuity Solutions
- Concepts and Definitions
- Geographically Dispersed Parallel Sysplex (GDPS)
- Restarting IMS
- IMS CF Structures
- Testing IMS DR
- How IBM Can Help
- Further Information

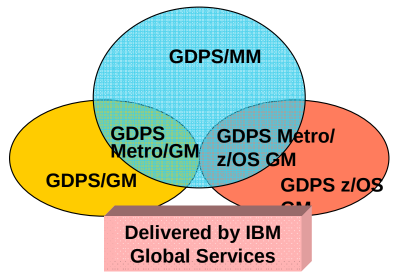

In this presentation, we will discuss the key concepts of disaster recovery as they relate to the IMS product. We will go over the important concepts and highlight the definitions that are key to understanding the IBM disaster recovery strategies.

The main focus will be how IMS works in the IBM disaster recovery solution called GDPS. We will look at the various ways GDPS can be used for synchronous and asynchronous solutions. Specifically, these solutions are dependent on the distance between the recovery sites and on how many recovery sites are in the solution.

We will also look at the key issue of providing a consistency point from which an IMS system can be emergency restarted.

If IMS is using all aspects of a Parallel Sysplex environment (ex. Datasharing), we will look at how to deal with the structures in the Coupling Facility (CF) for the various GDPS solutions. For example, for the GDPS Metro Mirror solution, it is possible to duplex some of the CF structures.

With any disaster recovery solution, the key is to test the recovery process at all of the recovery sites. This testing will ensure that all of the key datasets are present and the consistency point this is created are sufficient for emergency restarting IMS.

The IBM GDPS solutions come with several key service offerings to assist customers with configuring and implementing their disaster recovery solutions. Finally, there is a list of customer references and links to other interesting articles.

Concepts and Definitions

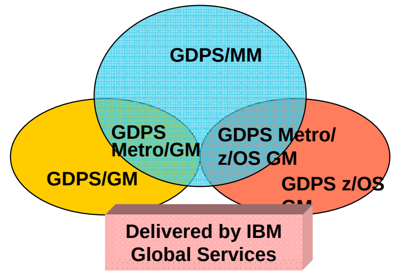

- ✓ Remote Mirroring managed by GDPS
- ✓ RTO vs. RPO
- ✓ Asynchronous vs. Synchronous
- ✓ Capacity Backup Upgrade (CBU)

In this section, we will discuss the key concepts of Disaster Recovery and some important definitions.

IMS Business Continuity Solutions

- **Geographically Dispersed Parallel Sysplex (GDPS)**
  - Manages:
    - IBM Metro Mirror (PPRC)
    - IBM Global Mirror
    - IBM z/OS Global Mirror (XRC)
  - Controls remote copy configuration and storage subsystem
  - Provides automation of sysplex operational tasks
  - Independent of applications like IMS and DB2
  - Includes IBM Services for configuration and manageability

GDPS is a multi-site management facility for IBM remote copying disaster recovery solutions. It is a combination of system code and automation and utilizes the capabilities of Parallel Sysplex technology and the storage subsystem. It manages processors, storage, and network resources.

As an integrated disaster recovery solution, it is able to handle the complexity of switching the network, the systems, and the DASD subsystems. GDPS is not part of z/OS or DFSMS. Instead, it is an automation solution using Tivoli NetView and System Automation for z/OS. GDPS supports the IBM Transaction Managers, IMS and CICS, and the Database Managers, IMS, DB2, and VSAM). In fact, GDPS is an application independent disaster recovery solution.

GDPS extends the continuous availability benefits of the Parallel Sysplex by automating the procedures required to recovery from a disaster or other failure situation. It also allows customers to manage their planned exception conditions. GDPS allows a customer to control its own continuous availability and disaster recovery goals without being dependent on third party disaster recovery providers.

GDPS is a multi-vendor solution which means it can handle vendor products that meet the right levels of the PPRC or XRC protocols. For example, both the HDS and EMC product lines support these protocols.

There are multiple service offerings available from IBM Global Services. These services are available when you license a GDPS solution. A detailed explanation of these service offerings is given later in this presentation.

Concepts and Definitions

- **Disaster recovery**
  - Process of recovering an environment after a major disaster
  - Bring to the point at which business can be conducted
- **Recovery Time Objective (RTO)**
  - How much time is available/allowed to recover the applications
  - All critical operations are up and running again
- **Recovery Point Objective (RPO)**
  - How much data can be lost in the event of disaster?
  - What is the last point-in-time when all data was current?

IDUG 2008 North America 6

Disaster Recovery is the ability to restart a production environment in a remote site following a disaster that makes the primary site unusable. Customers must determine the amount of time and the amount of data they can afford to lose when they are setting their recovery goals.

The Recovery Time Objective (RTO)is the amount of time it takes to bring up the production system in the remote site. The Recovery Point Objective (RPO) is the amount of data that can be lost in such a disaster. Typically, most institutions with critical systems on mainframe computers have a RTO of under two hours. By comparison, most UNIX environments with less critical systems have an RTO of under eight hours. The RTO and RPO objectives need to be determined on an Application-by-Application basis.

The cost of the disaster recovery solution needs to be in proportion to the business value of IT. In other words, the disaster recovery solution should not cost more than the financial loss suffered from the disaster itself.

The amount of data mirrored to a remote site and the currency of the data will greatly affect both the RTO and RPO of the remote site. If the production system can simply be emergency restarted from a recent checkpoint, there will be less data lost (RPO) and the restart will take less time (RTO). The amount of time increases dramatically for IMS when the system needs to recovery data from older image copies. In disaster recovery terms, it is always better to do a Restart than a Recovery.

Once the RTO and RPO objectives are known, the customer will be able to weigh the decisions of Synchronous vs. Asynchronous, distance between the Primary and Secondary sites, GDPS options, the number of remote sites, and many other important details of the disaster recovery solution.

Synchronous vs. Asynchronous

SYNCHRONOUS Remote Copy

Continuous Data Availability

Use when:

- ***Response time impact*** **is acceptable**
- **Within *metro distance***
- ***No data loss*** **is the objective**
- ***Fastest recovery time*** **is required**

Extended Distance Recovery

ASYNCHRONOUS Remote Copy

Use when:

- ***Smallest response time impact*** **to applications is required**
- ***Extended distance*** **disaster recovery is the objective**
- ***Minimal data loss*** **is acceptable**

The decision to use a synchronous solution or an asynchronous solution is based on the RTO and RPO objectives of the primary and secondary sites. If the secondary site is within metro distances (less than 100 KM) and the desire is to have no data loss and a very fast recovery time, and there can be a small impact to response time, then a synchronous disaster recovery solution is correct.

However, if the distance between the primary and secondary sites is not within metro distances and it is okay to have a small impact to response times and possibly a small amount of data loss, then an asynchronous disaster recovery solution would be the correct choice.

With Synchronous writes, the Acknowledgement or Device End (DE) I/O Complete is not returned to the primary system until it is received by the secondary system. This allows the primary and secondary systems to stay in synchronization with each other.

With Asynchronous writes, the Acknowledgement or Device End (DE) I/O Complete is returned immediately to the primary system without waiting for the acknowledgement to be received from the secondary site. This means the primary and secondary sites are not at the same point of synchronization.

Capacity Backup Upgrade (CBU)

- **Capacity Backup Upgrade (CBU)**
  - Temporary activation of central processors (CPs)
  - Concurrently increment the capacity of your processor
  - Alters the model and related LIC of the target machine
    - *To agreed upon model for a 90-day period of time.*
  - CPs can be added to the zSeries models
    - *With no system power-down and no re-IML/IPLs*
  - GDPS can automate the CBU activation process
    - *GDPS will "call home" using Remote Support Facility (RSF)*
      - Verify authorization and automatically unlock capacity

IDUG 2008 North America 8

The Capacity Backup Upgrade (CBU) feature allows the customer to temporarily activate additional central processors (CPs) at a secondary site for the purposes of recovering a production system during a disaster.

When the CBU is activated, IBM will alter the model of the hardware to reflect the change in processing power for a 90-day period of time. CBU allows the customer to add the additional CPs without powering down the system or having to re-IPL the machine. GDPS has the ability to automate the activation of CBU during a disaster event.

The Remote Support Facility (RSF) will be used to contact IBM to verify authorization and automatically unlock the new CPs in the system.

GDPS

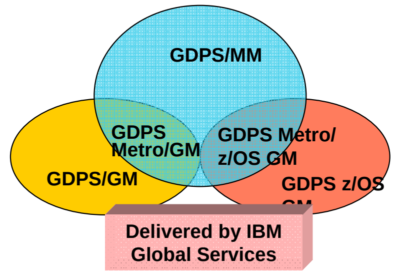

- ✓ Metro Mirror (PPRC)
- ✓ Global Mirror
- ✓ z/OS Global Mirror (XRC)
- ✓ Metro – Global Mirror
- ✓ Metro – z/OS Global Mirror

In this section, we will focus on the five most common Disaster Recovery mirroring configurations:

1. Metro Mirror (PPRC)
2. Global Mirror (Asynchronous PPRC)
3. z/OS Global Mirror (XRC)
4. Metro – Global Mirror
5. Metro – z/OS Global Mirror

GDPS/PPRC

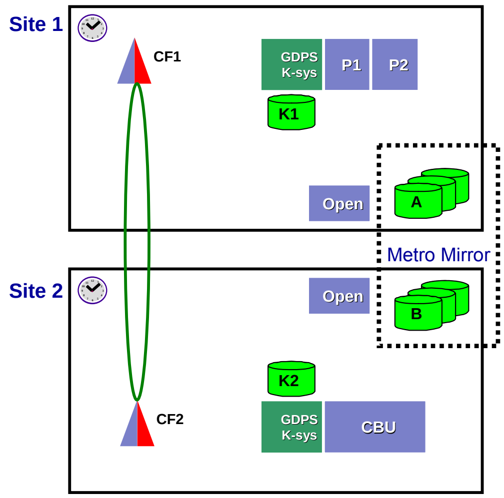

- RPO is zero
- Provides Data Consistency
- Emergency restart of IMS
- Can be used for planned outage
- Used for local CU/LSS failure with Hyperswap Manager
- Automation and Freeze policy
- Duplexing of CF Structures okay
  - GDPS handles “break-duplexing”
- Distances up to 300 km (with RPQ)

The first GDPS solution we will discuss is called GDPS/PPRC which is based on Metro Mirror technology. This was previously known as Synchronous PPRC. This solution requires the primary and secondary sites to be within a maximum of 300 KM (with an RPQ). Generally, the distance between the primary and secondary sites are much lower than 300 KM as this distance will impact the application response times.

With GPDS/PPRC, a write operation is not considered complete before it has been applied to both the primary and secondary sites. The strategy for this GDPS solution is to have both a small RTO and a small RPO. It is still necessary to use the GDPS FREEZE automation functions to provide a consistent point for the recovery.

With GDPS/PPRC and the Hyperswap Manager, it is possible to only switch to the DASD on the secondary site without switching the applications from the primary site’s production systems.

It is also possible to duplex the Coupling Facility (CF) structures on the secondary site. GDPS has functions that allow it to recognize the duplexed structures on the secondary site when there is a problem with the CF structures on the primary site.

GDPS/Global Mirror

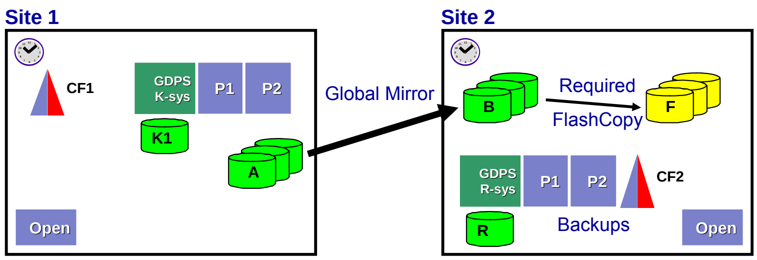

- Local site response time negligible
- Two site long distance DR and backup remote copy solution
- Consistency groups automatically established
  - Maintained every 3-5 seconds (RPO)
    - Dependent on Bandwidth
- Open system volumes can be in Consistency Group (CG)

IDUG 2008 North America 11

The GDPS/Global Mirror disaster recovery solution was previously known as Asynchronous PPRC. It is a two site solution with the primary and secondary sites are a long distance apart. In fact, the distance between the two sites is virtually unlimited and is only dependent on the capabilities of the network and the channel extension technology in place.

With GDPS/Global Mirror, it is necessary to use GDPS to form consistency groups as the data at the secondary site is not automatically in synch with the primary site. The data is applied about every 3-5 seconds so the RPO is still a relatively small amount of time. The 3-5 seconds is based on the amount of bandwidth available. Unlike IBM Metro Mirror and IBM z/OS Global Mirror, IBM Global Mirror technology requires that a FlashCopy be performed before updates for the Consistency Group are applied at the remote site.

There is support in GDPS/Global Mirror for data on Open system volumes. This data can be included in the Consistency Group allowing for a broader type of recovery at the secondary site.

It is recommended, but optional, that the secondary site use the FlashCopy technology to create a tertiary copy of the data. This is a Point-in-Time copy of the volumes and can be used for disaster recovery testing or during the application of the Consistency Group.

GDPS/XRC

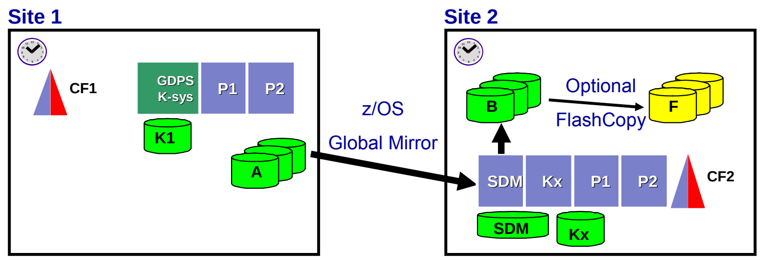

- RPO is greater than 0 (but within 1-5 seconds)
- Current RTO - Emergency restart of IMS
- Scales to any amount of DASD/distance
- Mature technology – many successes

IDUG 2008 North America 12

The GPDS/XRC which is based on the z/OS Global Mirror technology solution was previously known as Extended Remote Copy (XRC). It is a two-site solution where the two sites are a long distance apart. It uses asynchronous technology.

This solution supports only z/OS hosts. It uses a special function in the DS8000 Storage Manager called the System Data Mover (SDM). The SDM is responsible for retrieving updates from the primary disk subsystem and applying them to the secondary volumes.

The distance between the primary and secondary sites can be several thousand miles. With this distance and the fact that this is an asynchronous solution, the RPO will be greater than zero. Still, the RPO can be within 3 – 5 seconds and is dependent on the formation of the Consistency Groups (CGs) and how often the CGs are applied to the secondary volumes. The RPO for GDPS/XRC and GDPS/Global Mirror are the same and depended on available bandwidth.

With GDPS/XRC, the RTO will be fairly short as the primary and secondary sites can stay pretty current with each other. An IMS Emergency Restart from the last consistent point checkpoint is all that is required at the secondary site.

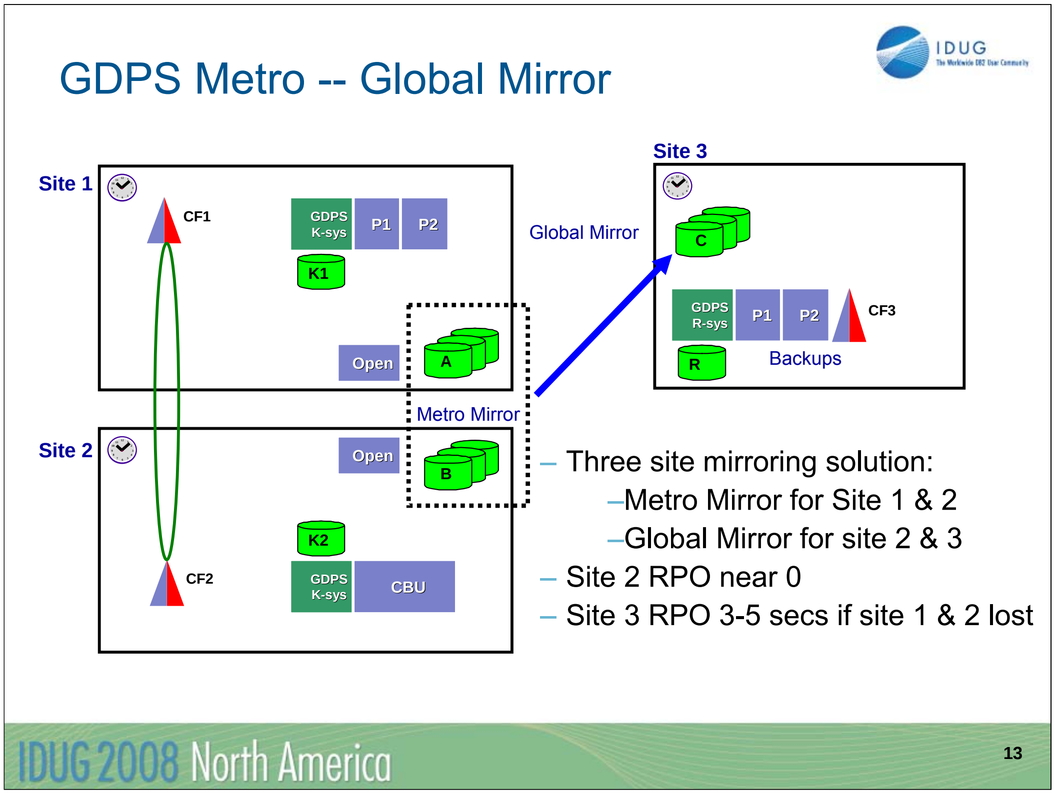

The GDPS Metro – Global Mirror solution is a three-site solution involving both the GDPS Metro Mirror and GDPS Global Mirror technologies. GDPS Metro Mirror is between the primary (site 1) and one of the secondary sites (site 2). GDPS Global Mirror is between the two secondary sites (site 2 and site 3).

There is synchronous mirroring for GDPS Metro Mirror (site 1 and site 2) and asynchronous mirroring between the two secondary sites (site 2 and site 3).

The RPO for the secondary (site 2) is near zero if the primary (site 1) goes away since they are using the GDPS/PPRC synchronous mirroring technology.

The RPO for the secondary (site 3) is within 3 – 5 seconds if both the primary (site 1) and secondary (site 2) go away since the data at site 2 may be behind the primary when it goes away.

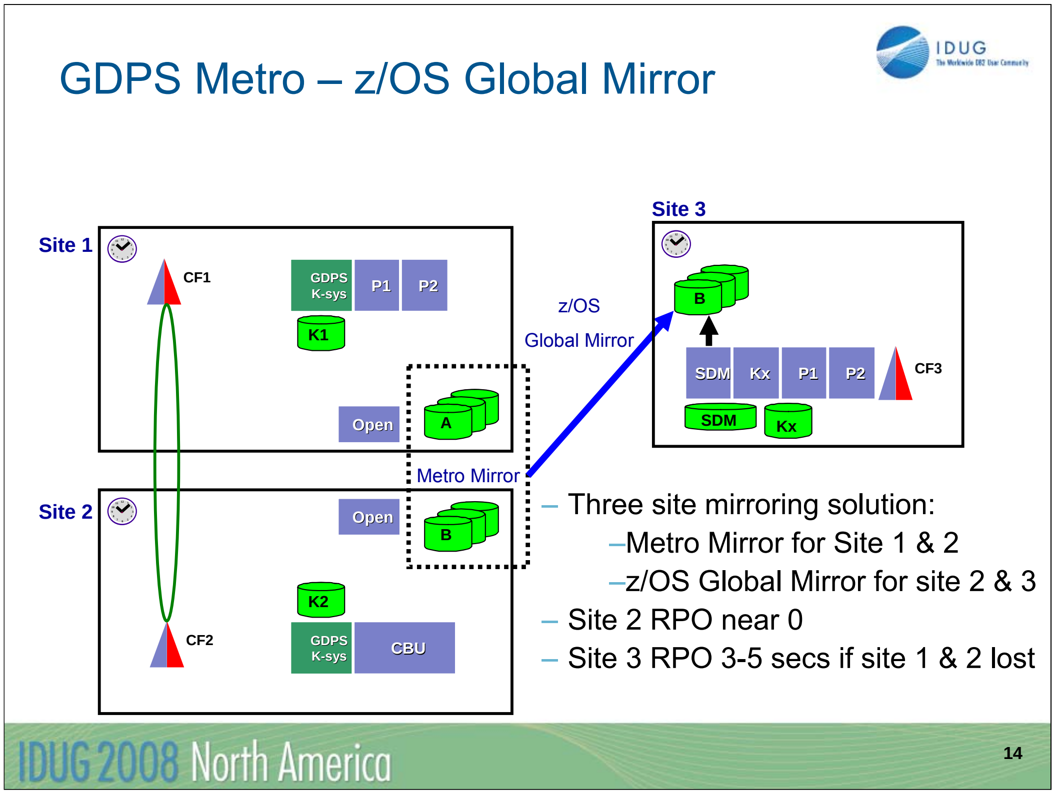

The GDPS Metro – z/OS Global Mirror solution is a three-site solution involving both the GDPS Metro Mirror and GDPS z/OS Global Mirror technologies. GDPS Metro Mirror is between the primary (site 1) and one of the secondary sites (site 2). GDPS z/OS Global Mirror is between the two secondary sites (site 2 and site 3).

There is synchronous mirroring for GDPS Metro Mirror (site 1 and site 2) and asynchronous mirroring between the two secondary sites (site 2 and site 3).

The RPO for the secondary (site 2) is zero if the primary (site 1) goes away since they are using the GDPS Metro Mirror synchronous mirroring technology.

The RPO for the secondary (site 3) is near zero if the primary (site 1) goes away because it is fairly close to the synchronous copy at the other secondary (site 2).

The RPO for the secondary (site 3) is within 3 – 5 seconds if the other secondary (site 2) goes away since the data at site 2 may be behind the primary when it goes away.

Restarting IMS

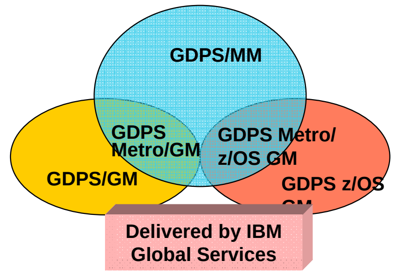

- ✓ Dependent Writes
- ✓ Consistency Groups

In this section, we will cover the key concept of Dependent Writes and Consistency Groups. This section is very important because without a consistent point in time at the secondary site, IMS does not have the ability to restart and recover data.

Dependent Writes and Consistency Groups

- **Importance of the Consistency Point in IMS**
  - Emergency Restart needs a consistent point for all data
  - Failures may not occur at the same point for all systems
  - IMS has events that are chained together
    - DASD Writes that are dependent on each other
    - Time order of the write is crucial
    - Examples:
      - Database Update & Log Record Update
      - Index & Data components

The importance of the Consistency Group (CG) in disaster recovery of database applications can not be emphasized enough. IMS needs a consistent point in time and a checkpoint in order to do an emergency restart. The CG provides the secondary site with the point in time when all updates are available and can be applied to the secondary site.

In IMS, there is the concept of Dependent Writes that must be observed in any remote mirroring solution. A Dependent Write is a set of writes where one write is dependent on the fact that the previous write has completed. Generally, it is an operation that is either completed together or not completed at all.

There are many examples in Computer Science of Dependent Writes. In a programming sense, there is code to move an element from one chain to another chain, freeing up space on the first chain. There are programming instructions that allow this to happen without interruption. Without these types of instructions, an interrupt could occur allowing two tasks to acquire the same element.

In the database area, there are operations where the both the data and index components need to be updated in a certain order. If only one of the updates is made, the database can be out of synchronization.

A key example of Dependent Writes is given on the next page.

Dependent Writes

- **Committed Database Update**

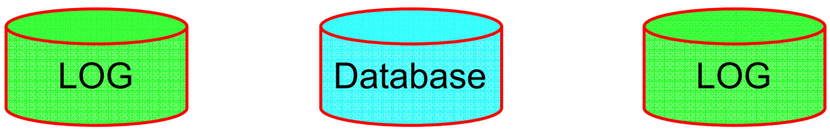

(1) Log "Before Image" (2) Update Database (3) Log "After Image"

Good Sequence of Writes
- (1)
- (1) and (2)
- (1), (2) and (3)

Bad Sequence of Writes
- (1) and (3) only

This example is the classic reason why it is so important to ensure Dependent Writes are applied correctly on the secondary site and why it is important to create Consistency Groups at the secondary site.

In this example, there are databases and there are logs. Typically, IMS will write a log record (1) to indicate that IMS is about to update a database. Then, the database will be updated (2). Finally, the log will be updated (3) to indicate that the database update was applied. The two writes to the log and the one write to the database are dependent on each other.

When the Dependent Writes are completed in the proper order, then a failure that occurs after the first log update (1) is applied is not a problem because IMS knows that the database update was not applied.

If the failure occurs after the first write to the log (1) and the update to the database (2) occurred, but before the second write to the log (3), then IMS knows to backout the database update. This is also not a problem.

If the failure occurs after the first log update (1), the update is applied to the database (2) and the second write to the log (3), then the update was successfully applied to the database and does not need to be backed out.

However, if the order of the Dependent Writes is not preserved and the first and second log writes (1 and 3) were done without the update to the database (3), then IMS will assume the update was applied correctly even though the database was never updated. This situation is avoided with Dependent Writes.

GDPS Metro Mirror: Consistency

- **(1) Consistency Group (CG)**
  - Set of volumes hold IMS datasets and logs
  - When failure occurs for a volume in CG
    - Extended Long Busy (ELB) for failing volume
    - Writes for other volumes are held
    - After ELB completes, held writes continue
- **(2) GDPS Freeze Automation**
  - FREEZE and GO:
    - Writes continue at Primary even if failing at Secondary
  - FREEZE and STOP:
    - Writes are frozen at Primary and Secondary

IDUG 2008 North America 18

The Consistency Group (CG) and the GDPS FREEZE automation function are integral parts of GDPS Metro Mirror for providing a consistent point at the secondary site. The CG is defined at the secondary site to include all of the volumes that are used by IMS. For example, this would include the database and the logs and other key datasets. When all of the writes are received at the secondary site for the CG, they are applied to the secondary volumes.

When a write fails for a volume in the CG at the secondary site, the write operation receives an Extended Long Busy (ELB) for the failing volume. All writes for other volumes in the CG are held while the ELB is in existence. When the ELB completes, which may occur after a default period of time, the other writes are allowed to continue.

The GDPS FREEZE function works hand-in-hand with the failing write of the volume at the secondary site. There are two FREEZE options: 1) FREEZE and GO, and 2) FREEZE and STOP.

With FREEZE and GO, a failing write at the secondary system will not prevent writes at the primary site from continuing even though they can not be applied to the secondary volumes. This option allows production work to continue even if there is a failure at the disaster recovery site. The GO option is used by the majority of customers.

With FREEZE and STOP, the writes on the primary system have to stop until the secondary site's problems have been corrected. The STOP option ensures the secondary site is always as current as the primary site. It also keeps production work from continuing when there is a problem at the secondary site. If the RPO is zero, then the STOP option must be used.

GDPS Global Mirror: Consistency

- **(1) Primary creates Out-of-Sync Bit Maps:**
  - Shows tracks with new data
- **(2) Consistency Groups**
  - FlashCopy is required for IBM Global Mirror
    - FlashCopy is taken before changes are applied
  - Creates Change Recording bitmap from out-of-sync bitmap
    - After writes to out-of-sync bitmap have stopped.
  - Coordinated across all systems in Global Mirror session.
  - Tracks in Change Recording bitmap updated at Secondary

IDUG 2008 North America

19

With GDPS Global Mirror, the data is “pushed” to the secondary site. The primary site creates an Out-of-Sync (OOS) bit map showing all of the tracks that have been updated with new data. The secondary site waits until all writes to the OOS bit map have been stopped and then the secondary site creates its own Change Recording (CR) bit map. This process is coordinated across all disk subsystems in the Global Mirror session.

Unlike the IBM Metro Mirror and IBM z/OS Global Mirror technologies, IBM Global Mirror requires the FlashCopy support on the secondary site. A FlashCopy is taken before any changes in the CR are applied. When the FlashCopy is compolete, the tracks that have been identified in the CR bit map as having changed are updated.

A Consistency Group (CG) is formed to apply these updates when the data is available from all of the disk subsystems in the session. If some of the disk subsystems are unavailable, the Global Mirror will skip a cycle and apply the CG updates when all of the disk subsystems are available.

GDPS/XRC: Consistency

- **(1) Primary writes are timestamped for Secondary**
  - Timestamps are stored in Side File cache
- **(2) System Data Mover (SDM)**
  - Global Mirror uses SDM to “pull” the data to Secondary
  - SDM uses timestamps from Side File to order writes
  - SDM selects min of max timestamps with all before it.
  - Consistency Group is journalled
    - Applied to Secondary volumes.

With GDPS/XRC, the data is “pulled” to the secondary site by the System Data Mover (SDM). The primary site will timestamp all writes using the sysplex timer. When the SDM receives the timestamped writes, it applies them to a Side File Cache.

At a certain point in time, the SDM will collect all of the timestamped writes in the Side File Cache and it will select the minimum timestamp from all of the maximum timestamps in the Consistency Group. This minimum timestamp acts as the cut-off point for the group of updates. This collection of writes in the Consistency Group is logged in the journal and then applied to the secondary site.

One SDM can handle approximately 2000 volumes. SDMs can also be clustered with other SDMs to handle very large z/OS environments.

IMS CF Structures

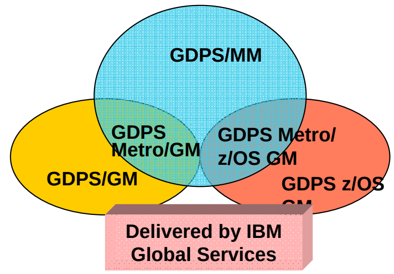

- ✓ Duplexed vs. Non-Duplexed

In this section, we will discuss how the CF Structures can be duplexed in a Metro Mirror (PPRC) configuration. We will also discuss which structures would benefit from duplexing and which structures would be better to rebuild at the remote site.

GDPS and Duplexing CF Structures

- **GDPS V3.3, later (GDPS Enhanced Recovery Support)**
  - Structures can be duplexed across two sites
    - Structures in same site as secondary disk
      - Retained in break duplexing condition
  - Applicable to FREEZE and STOP policy
    - Reduces recovery times associated with recovering CF data
- **Secondary CF Structures agrees with secondary disks**
  - Some failure scenarios require reallocating/rebuilding structures.
    - Dependent on type of failure at primary site

With GDPS V3.3 (GA in January, 2006), IBM introduced the function called "GDPS Enhanced Recovery" in conjunction with new function in z/OS. This capability allows structures that are duplexed across two sites to be retained at the secondary site where the secondary disk resides in a break duplexing condition. This function, which is only applicable with FREEZE and STOP policies, reduces recovery times associated with recovering CF resident data.

The GDPS enhanced recovery support requires z/OS APAR OA11719 which is available back to z/OS V1.5.

Although duplexing is occurring across two sites, there are cases where the surviving CF structure at the secondary site may not be usable. To be usable, it must be consistent with the secondary copy of the mirrored disks for recovery purposes. If it is unusable, it will need to be allocated and rebuilt at the secondary site during emergency restart.

GDPS MM, GM, z/OS GM: CF Structures

- **GDPS/PPRC**
  - Primary and Secondary distance allows for Duplexing
    - Determine the structures where duplexing is useful
  - Logical configuration of Primary and Secondary must be same
    - If Primary is data sharing, then Secondary is too
- **GDPS/Global Mirror, GDPS/XRC**
  - Distance between sites usually prevents Duplexing
    - Note: Distances do not need to be large
  - All CF structures are allocated during Emergency Restart

The GDPS/PPRC solution allows duplexing of CF structures where the structures are duplexed across the primary and secondary sites. For GDSP/Global Mirror and GDPS/XRC, the distances between the primary and secondary sites usually prevent the duplexing of CF structures between the two sites. For these disaster recovery solutions with long distances between the primary and secondary sites, the CF structures must be allocated and rebuilt which can be a time-consuming event.

The configuration of the primary site must be maintained at the secondary site. That is, if the primary site is operating a parallel sysplex with data sharing, then the secondary site must also do data sharing in a parallel sysplex environment.

GDPS Metro Mirror: IMS CF Structures

- **No Duplexing:**
  - OSAM and VSAM Buffer Pools
    - Stored on DASD when data is committed
    - Secondary Site:
      - All buffers are invalid and structure is rebuilt
  - IRLM Locks
    - Secondary site:
      - Restart backs out inflight transactions to release locks
      - IRLM rebuilds lock structure as empty

In IMS, there are CF structures which are good candidates for duplexing and there are CF structures which are not and should be rebuilt at the secondary site. The CF Structures which are not good candidates are: 1) OSAM and VSAM Buffer Pools, and 2) IRLM Locks.

When the OSAM and VSAM Buffer Pools structures are committed and updated in the CF, the changes are also applied to disk storage. For this reason, the OSAM and VSAM buffers should not be duplexed. Instead, they should all be invalidated and rebuilt as needed at the secondary site.

The IRLM Lock CF structure is also a poor candidate for duplexing. IMS Emergency Restart will automatically back out all inflight transactions and release any IRLM locks allowing IRLM to rebuild the lock structure as an empty CF structure.

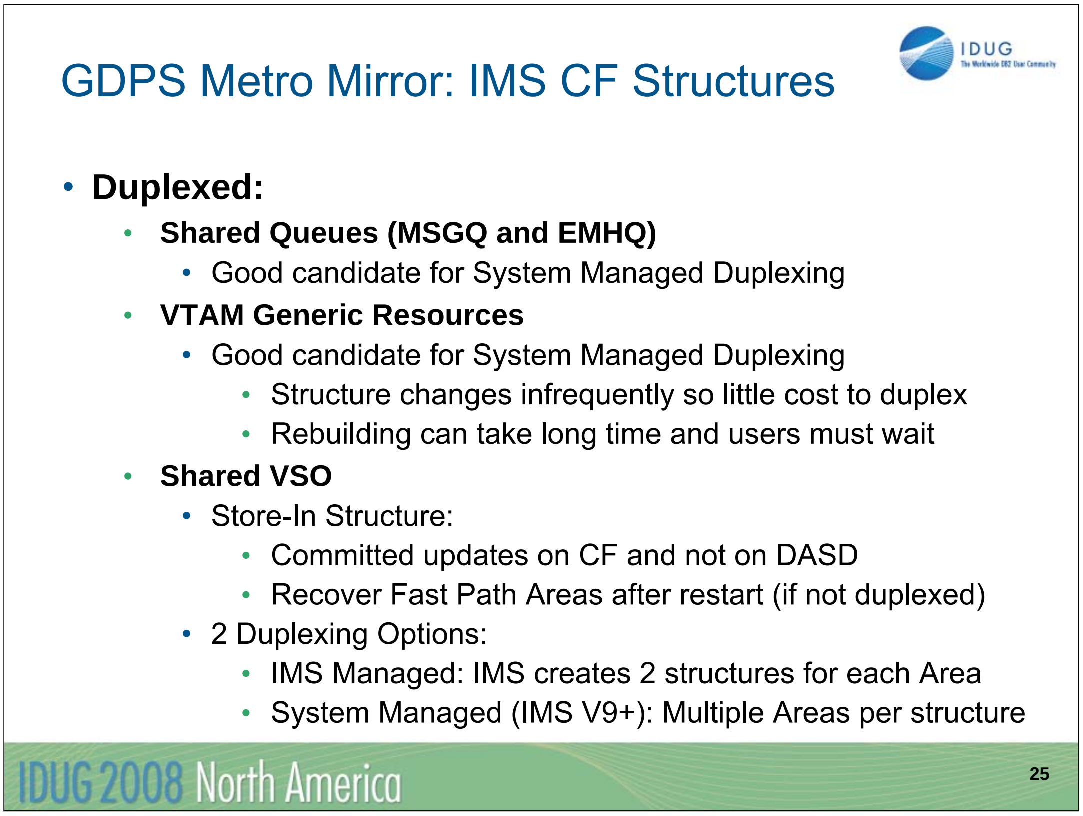

The IMS CF Structures which are good candidates for duplexing are: 1) Shared Queues (MSGQ and EMHQ), 2) VTAM Generic Resources, and 3) Shared VSO.

The Shared Message Queues (MSGQ and EMHQ) CF structures are good candidates for System Managed Duplexing because it avoids rebuilding them at the secondary site.

The VTAM Generic Resources CF structure is another good candidate for System Managed Duplexing because it changes infrequently so there is very little cost to duplexing. Also, rebuilding the structure can take a long time to complete.

The Shared VSO CF structures use the Store-In technology. This means that committed updates may be on the CF and not on disk storage. In a disaster, the Fast Path Areas would need to be recovered after restart if they are not duplexed.

There are two duplexing options for the Shared VSO Fast Path CF structures. Prior to IMS V9, each Shared VSO area needed to be in its own CF structure. IMS provided the option to manage the duplexing of these VSO Shared areas. With IMS V9 and above, there could be more than one VSO Shared areas in a CF structure. This allowed these later VSO Shared areas to be duplexed and managed by System Managed Duplexing.

Testing IMS

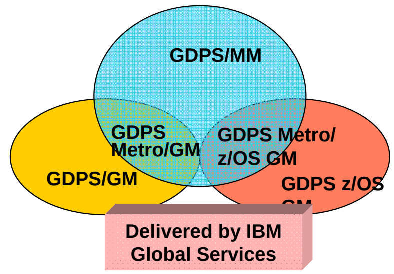

✓FlashCopy

✓IMS Datasets

In this section, we will focus on the use of FlashCopy for Global Mirror and for testing at the remote site.

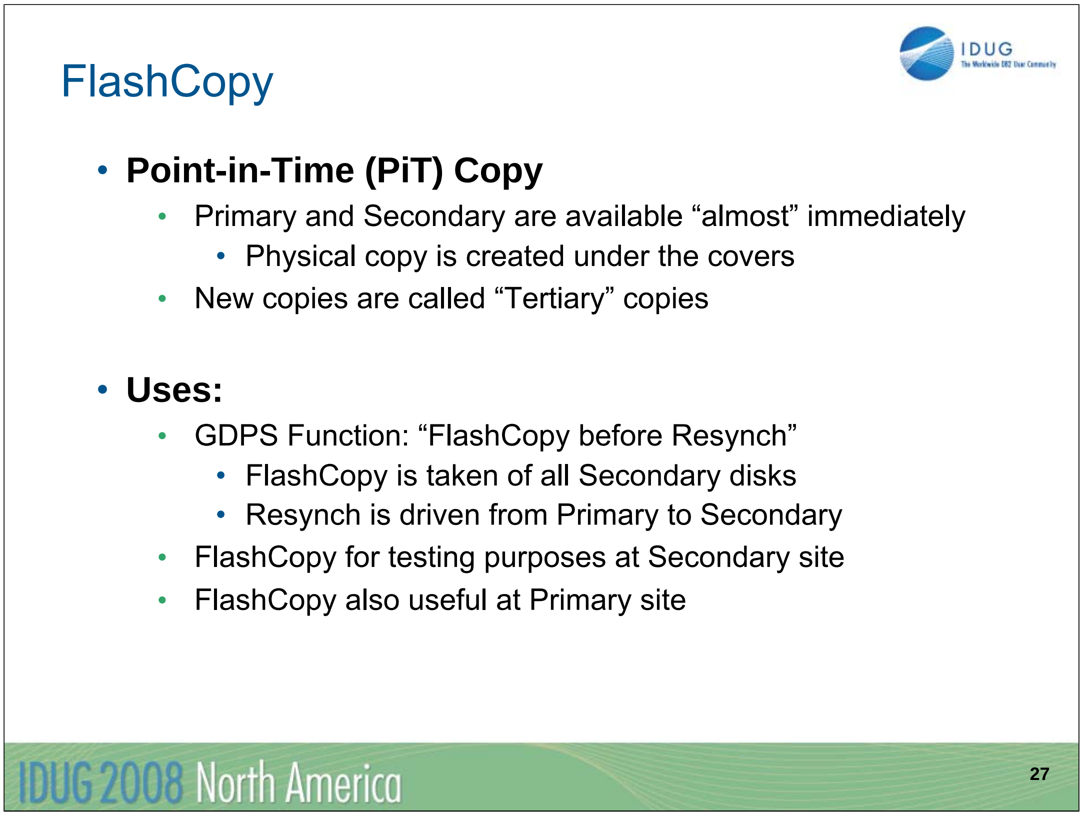

FlashCopy provides a Point-in-Time (PiT) copy of the volume data. After FlashCopy, the volumes are almost immediately available. The physical copy occurs under the covers. The new copies are called "Tertiary" copies.

There are multiple purposes for FlashCopy in a GDPS disaster recovery environment. For example, a FlashCopy can be done just prior to a resynchronization between the primary and secondary sites. The GDPS function "FlashCopy Before Resynch" provides this exact functionality and it causes the secondary site to take a FlashCopy before applying the new updates. This is the same as backing up the data before the new updates are applied. The primary site drives the resynchronization, so when the resynch is completed, the primary can issue the FCWITHDRAW command to remove the relationship between the primary and secondary sites.

FlashCopy can also be used for testing the data at the secondary site. The Tertiary copy of the data at the secondary site can be used to test IMS Emergency Restart.

It is also useful to have FlashCopy at the primary site since at some point in time, it may be necessary to switch the production workload back to the primary site. FlashCopy can be used to create a Tertiary copy of the data at the primary site before resynchronization occurs between the secondary and primary sites.

Testing at the DR Site

- **Use FlashCopy to create Tertiary copy of data**
- **Emergency Restart IMS from last Consistent Point**
  - Verify all datasets are available for restart
    - Assume Primary site is no longer available
  - Consider:
    - Proclibs, Maclibs
    - User exits, Randomizing routines
    - FORMAT
    - DBDLIB, PSBLIB, PGMLIB, ACBLIB
    - OLCSTAT
    - Etc.
  - Test multiple types of JCL
    - Utilities, IMS, Batch, etc.

IDUG 2008 North America

28

It is very important to test the data at the secondary site. To do this, the customer needs to do an IMS Emergency Restart from the last consistent checkpoint in the Recon record. FlashCopy can be used to create a Tertiary copy of the data for testing purposes. With this alternate copy of the data, the customer can be sure they will not be disturbing the real secondary copy of the data during testing.

During testing, the customer should ensure they have all of the important datasets required to restart their production system. There are several important datasets and here are a few examples: Proclibs, Maclibs, User Exits, Randomizing Routines, FORMAT, DBDLIB, PSBLIB, PGMLIB, ACBLIB, OLCSTAT, and several others.

When testing on the secondary site, it is best to try all of the different types of JCL (ex. Utilities, Batch, Control Region, etc.) that is used on the production system. One strategy may be to log all of the events that occur on the primary site and try to recreate these on the secondary site.

Further Information

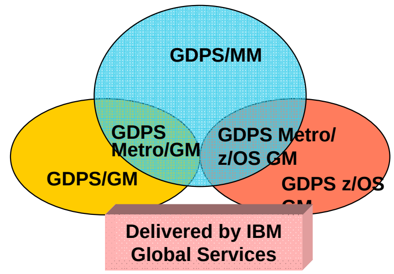

- ✓ Publications
- ✓ Websites

In this section, we will look at some documents, presentations and websites that are available for reading about GDPS and Mirroring Disaster Recovery solutions.

Additional Information

- **Detailed GDPS Presentation and Information e-mail:**
  - gdps@us.ibm.com
  - www.ibm.com/systems/z/gdps
- **White Papers:**
  - Business Continuity Considerations and IBM eServer zSeries
  - GDPS - Ultimate e-business Availability Solution (GF22-5114)
- **Publications:**
  - GDPS Family of Offerings Introduction to Concepts and Capabilities (SG24-6374)
  - TS Disaster Recovery Solutions Redbook (SG24-6547)
  - z/OS Advanced Copy Services (SC35-0428)
  - ESS Copy Services on zSeries Redpiece (SG24-5680)
  - ESS Copy Services on Open Redpiece (SG24-5757)

Here are some articles form the Web that may also help.

The Alinean ROI **Report** - January 2004:

www.alinean.com/newsletters/2004-1-Jan.asp

How To Quantify **Downtime:**

www.itmanagementnews.com/2004/0311.html

Research shows application **downtime** is costly:
www.itweb.co.za/office/compuwaresa/0406070830.htm

Calculating the **cost** of **downtime** – Computerworld:

www.computerworld.com/securitytopics/ security/recovery/story/0,10801,91961,00.html

Most firms cannot count **cost** of IT **downtime** | The Register

www.theregister.co.uk/2004/04/19/app_downtime_survey/

How to calculate and convey the true **cost** of **downtime:**

techrepublic.com.com/5100-6313-1038783.html

Redbooks

- **SG24-6547**
  - Business Continuity: Part 1 Planning Guide
- **SG24-6548**
  - Business Continuity: Part 2 Solutions Overview
- **SG24-6374**
  - GDPS Family of Offerings Introduction to Concepts and Capabilities

These are good Redbooks on the topic of Business Continuity.

Summary

- IMS Business Continuity Solutions
- Concepts and Definitions
- Geographically Dispersed Parallel Sysplex (GDPS)
- Restarting IMS
- IMS CF Structures
- Testing IMS DR
- How IBM Can Help
- Further Information

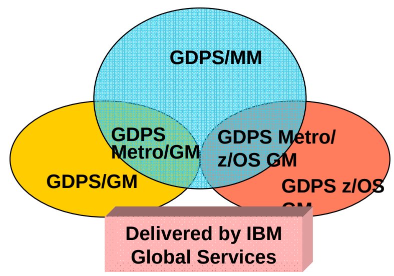

In this presentation, we discussed the key concepts of disaster recovery as they relate to the IMS product. We will go over the important concepts and highlight the definitions that are key to understanding the IBM disaster recovery strategies.

The main focus will be how IMS works in the IBM disaster recovery solution called GDPS. We will look at the various ways GDPS can be used for synchronous and asynchronous solutions. Specifically, these solutions are dependent on the distance between the recovery sites and on how many recovery sites are in the solution.

We will also look at the key issue of providing a consistency point from which an IMS system can be emergency restarted.

If IMS is using all aspects of a Parallel Sysplex environment (ex. Datasharing), we will look at how to deal with the structures in the Coupling Facility (CF) for the various GDPS solutions. For example, for the GDPS Metro Mirror solution, it is possible to duplex some of the CF structures.

With any disaster recovery solution, the key is to test the recovery process at all of the recovery sites. This testing will ensure that all of the key datasets are present and the consistency point this is created are sufficient for emergency restarting IMS.

The IBM GDPS solutions come with several key service offerings to assist customers with configuring and implementing their disaster recovery solutions. Finally, there is a list of customer references and links to other interesting articles.