#SHAREorg

SHARE

DB2 and Advanced Copy Services: All You Need to Know

Paul Pendle

Independent Consultant

Tuesday, March 11, 2014

Session Number: 15352

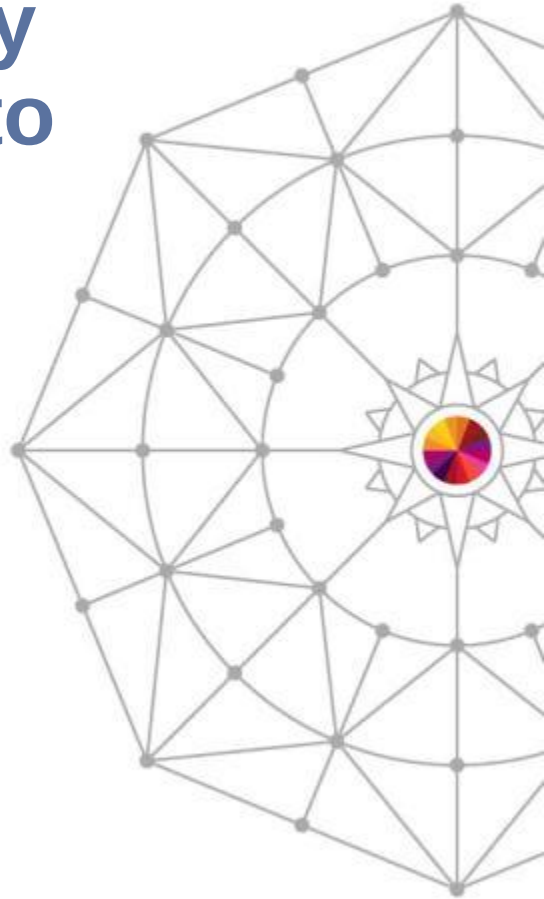

SHARE in Anaheim

Agenda

- What is/are Advanced Copy Services
  - Deep dive into FlashCopy
  - Shallow dip into PPRC/XRC
- Using FlashCopy in z/OS
  - Consistency Groups
- Using DB2 with FlashCopy
- Considerations for using DB2, FlashCopy, and PPRC/XRC
- Conclusions

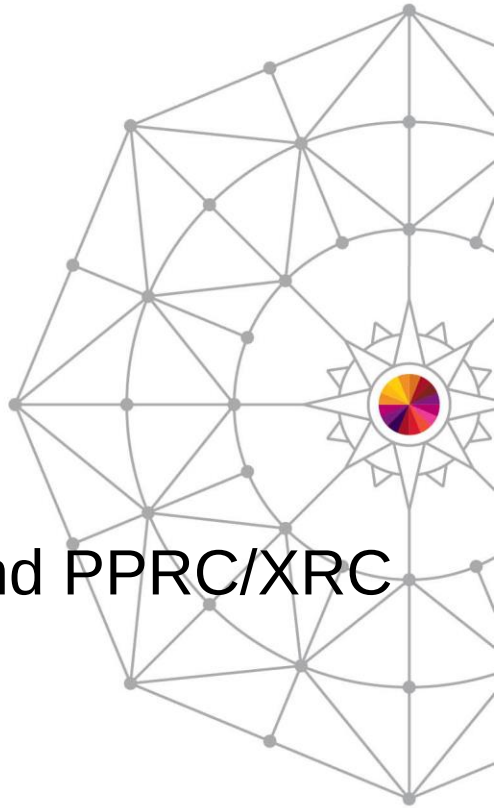

Introduction to Advanced Copy Services

- Storage array facilities that replicate data
  - Locally (FlashCopy)
    - Volume to Volume
    - Dataset to dataset
  - Long distance (PPRC, XRC)
    - Synchronously and asynchronously
    - Volume to volume

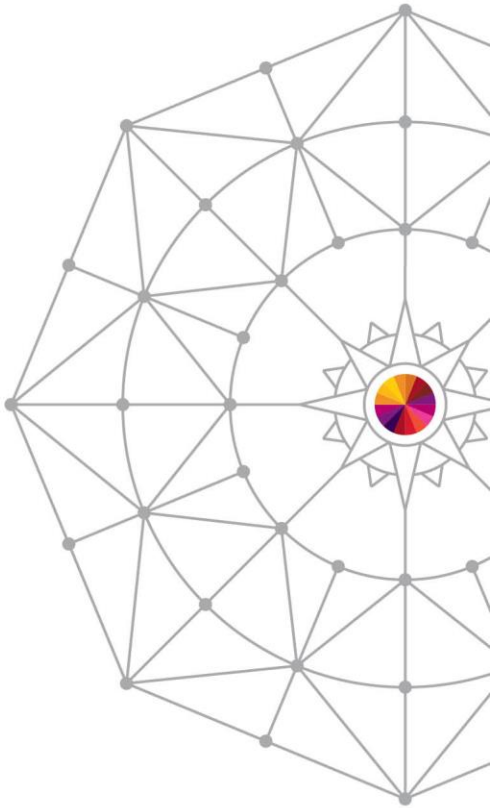

Array-based Local Replication

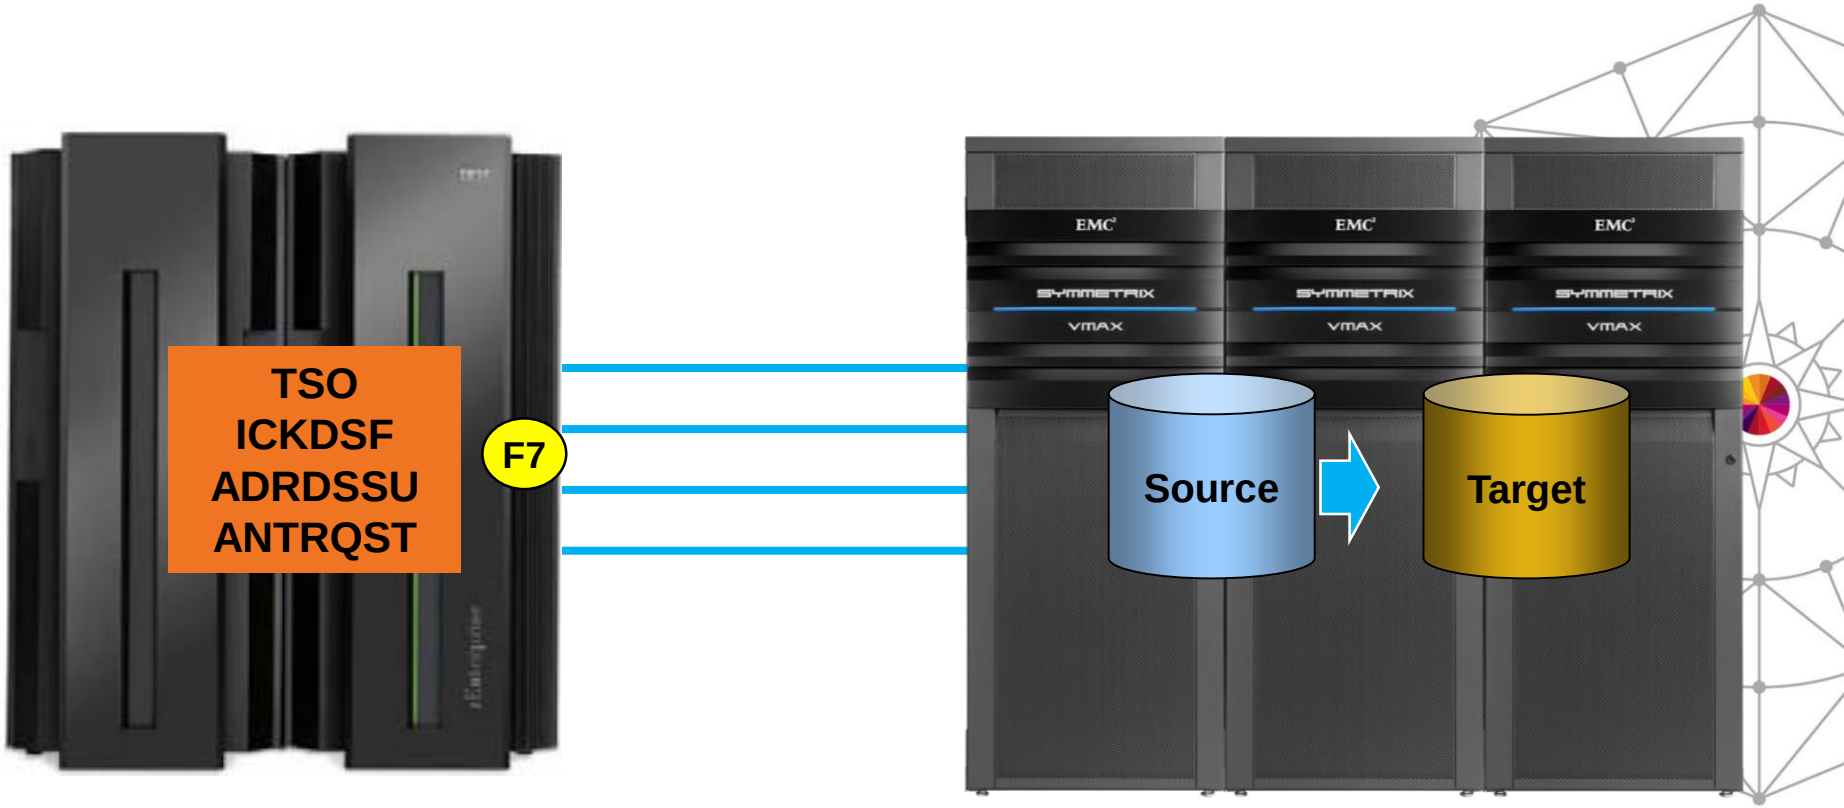

Array-based Copy Advantages

- No host CPU or I/O to perform the data movement
- Extremely fast copies
- Both dataset and volume replication
- Allows incremental copies
- Can allow space efficient copies (FlashCopy SE)
- Time-to-use very fast (before copy completion)

FlashCopy Primitives

- ESTABLISH
  - Creates a relationship between a source and target volume/dataset and in most cases copies the data from source to target
- WITHDRAW
  - Terminates a FlashCopy relationship between a source and target

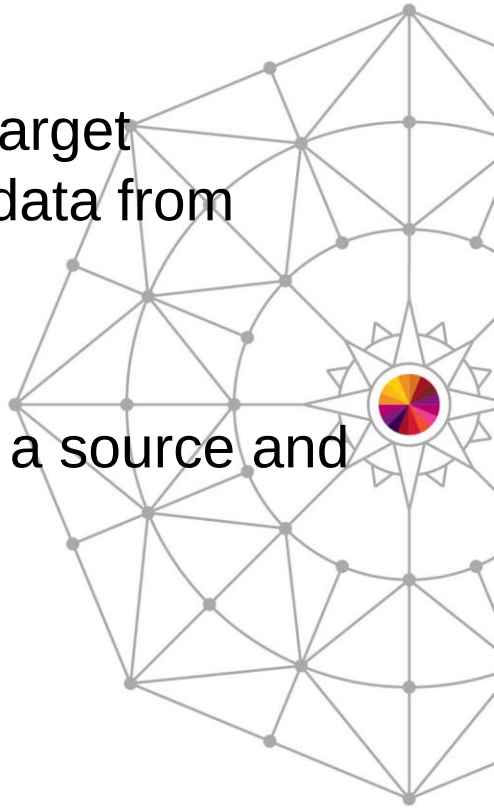

What's missing?

- RESTORE !!!!!!!
  - ESTABLISH with FAST REVERSE RESTORE ….

FlashCopy Considerations

- Establish is three phases
  - Create the relationship (very fast)
  - Copy the tracks (optional: NOCOPY)
  - Withdraw the relationship (optional and very fast)
- Three FASTREPLICATION modes under DFDSS
  - PREFERRED (default)
  - REQUIRED
  - NONE
- Copy on first write (COFW) overhead
- Read indirection overhead

FlashCopy Relationships (Temporary)

P	P	P	P
P	P	P	P
P	P	P	P
P	P	P	P

Protection map

I	I	I	I
I	I	I	I
I	I	I	I
I	I	I	I

Indirection map

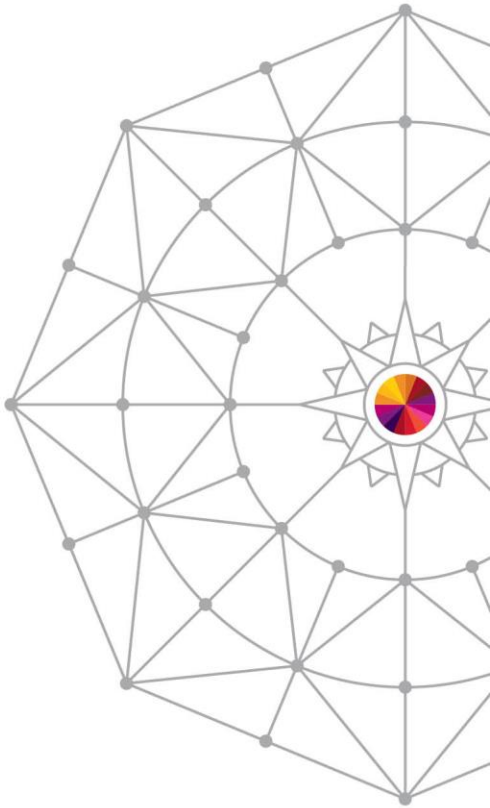

FlashCopy Relationship (Persistent)

P	P	P	P
P	P	P	P
P	P	P	P
P	P	P	P

Protection map

I	I	I	I
I	I	I	I
I	I	I	I
I	I	I	I

Indirection map

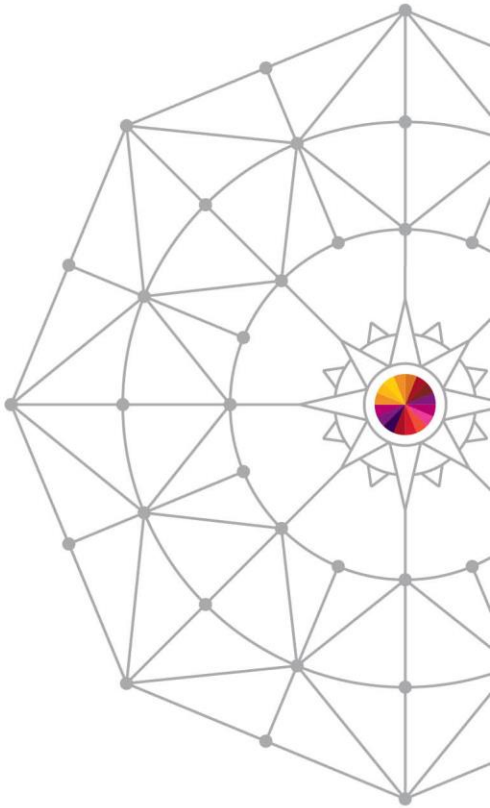

- Space Efficient
- NOCOPY
- Incremental

How Does Incremental Help?

P	P		P
	P	P	
			P
P			P

Protection map

I	I		I
	I	I	
			I
I			I

Indirection map

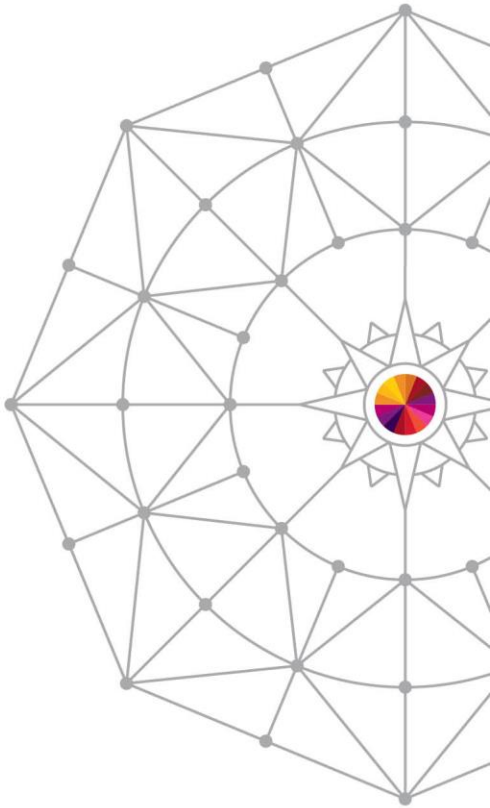

Incremental FlashCopy Considerations

- Reduced work for the storage controller
- Faster completion of background copy
- Reduced COFW impact
- Reduced indirected reads
- Only one FlashCopy session can be incremental
  - Others are withdrawn on completion

Incremental Limitation (Full Volume)

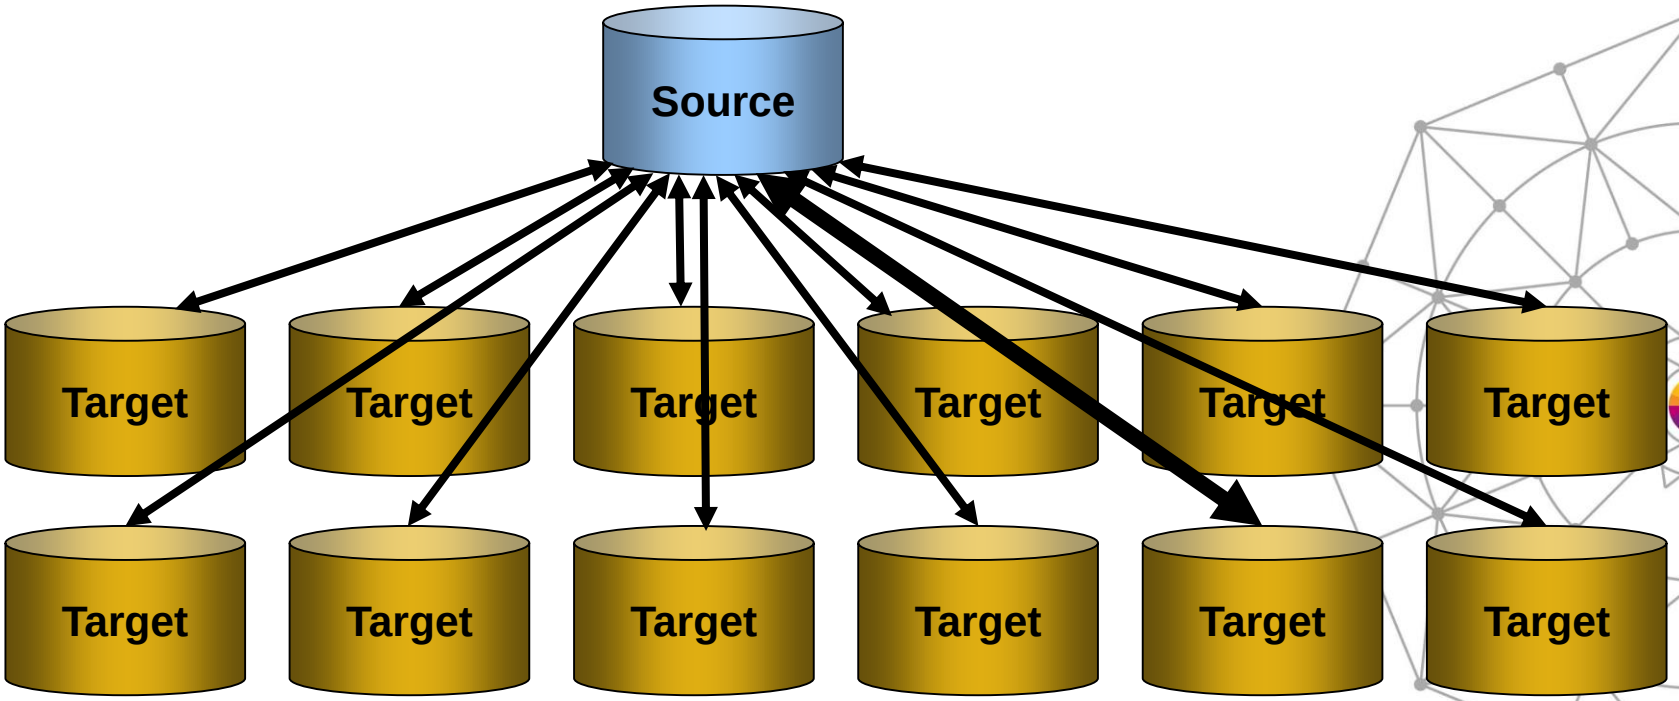

Cascading FlashCopy

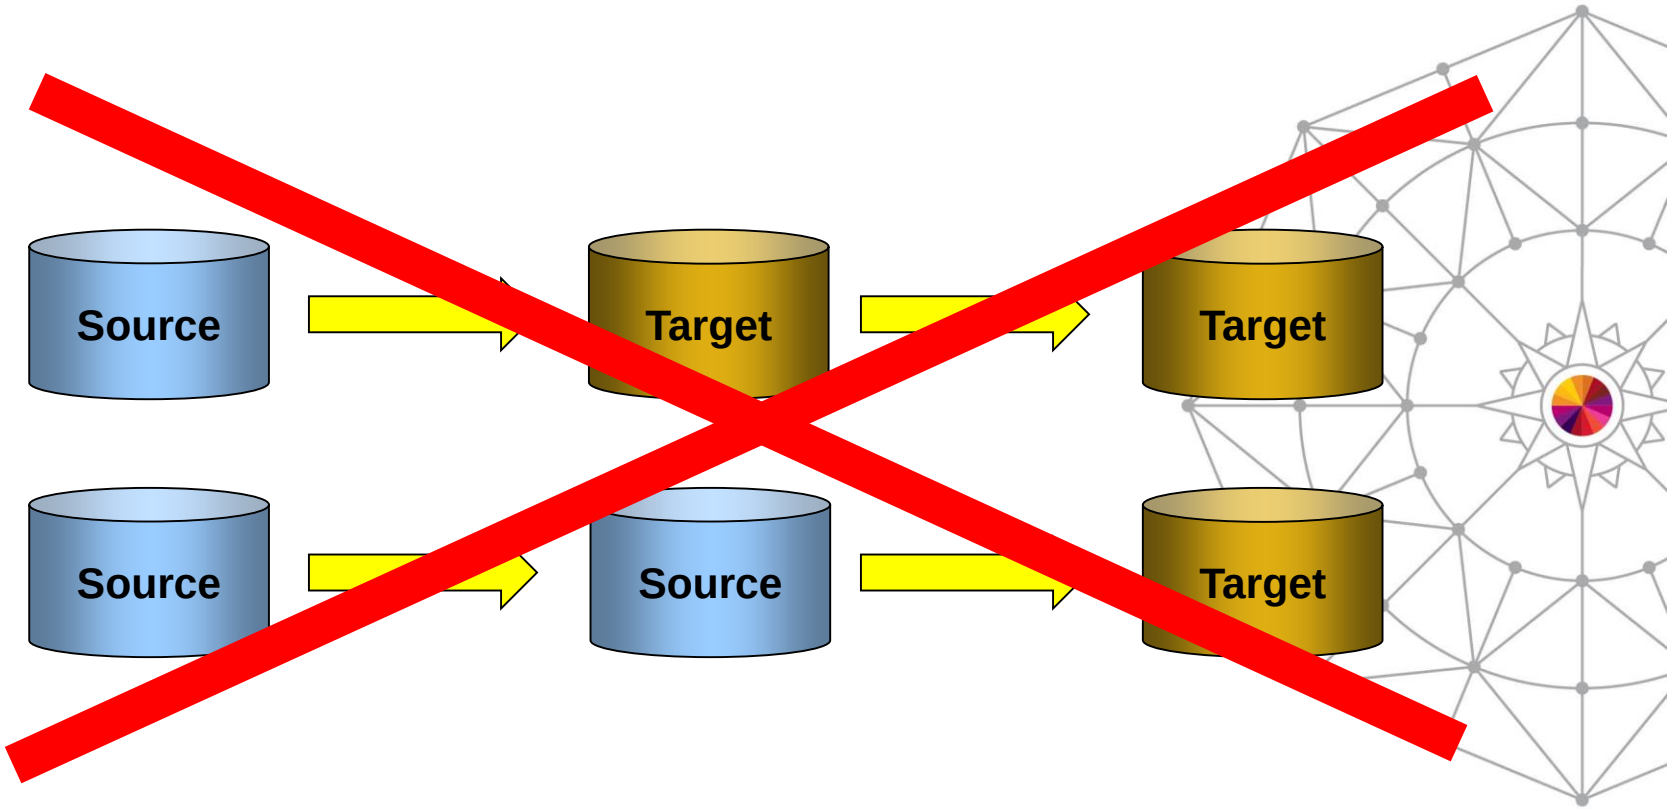

Space Efficient FlashCopy

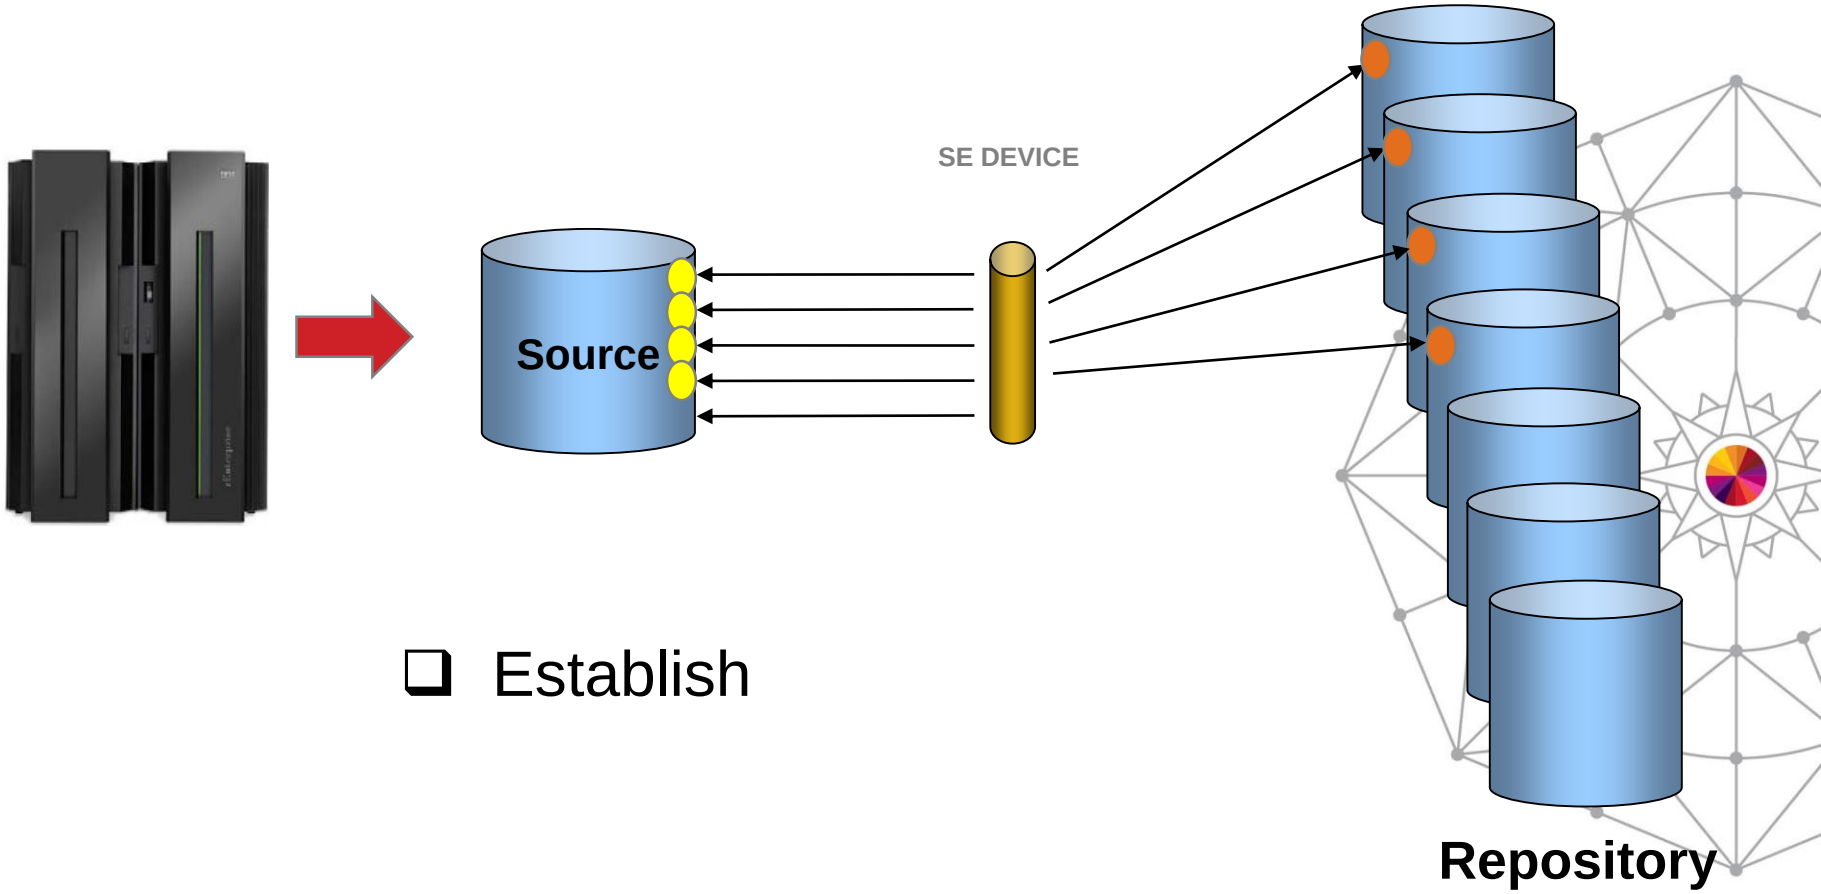

- Establish

Space Efficient FlashCopy Restore

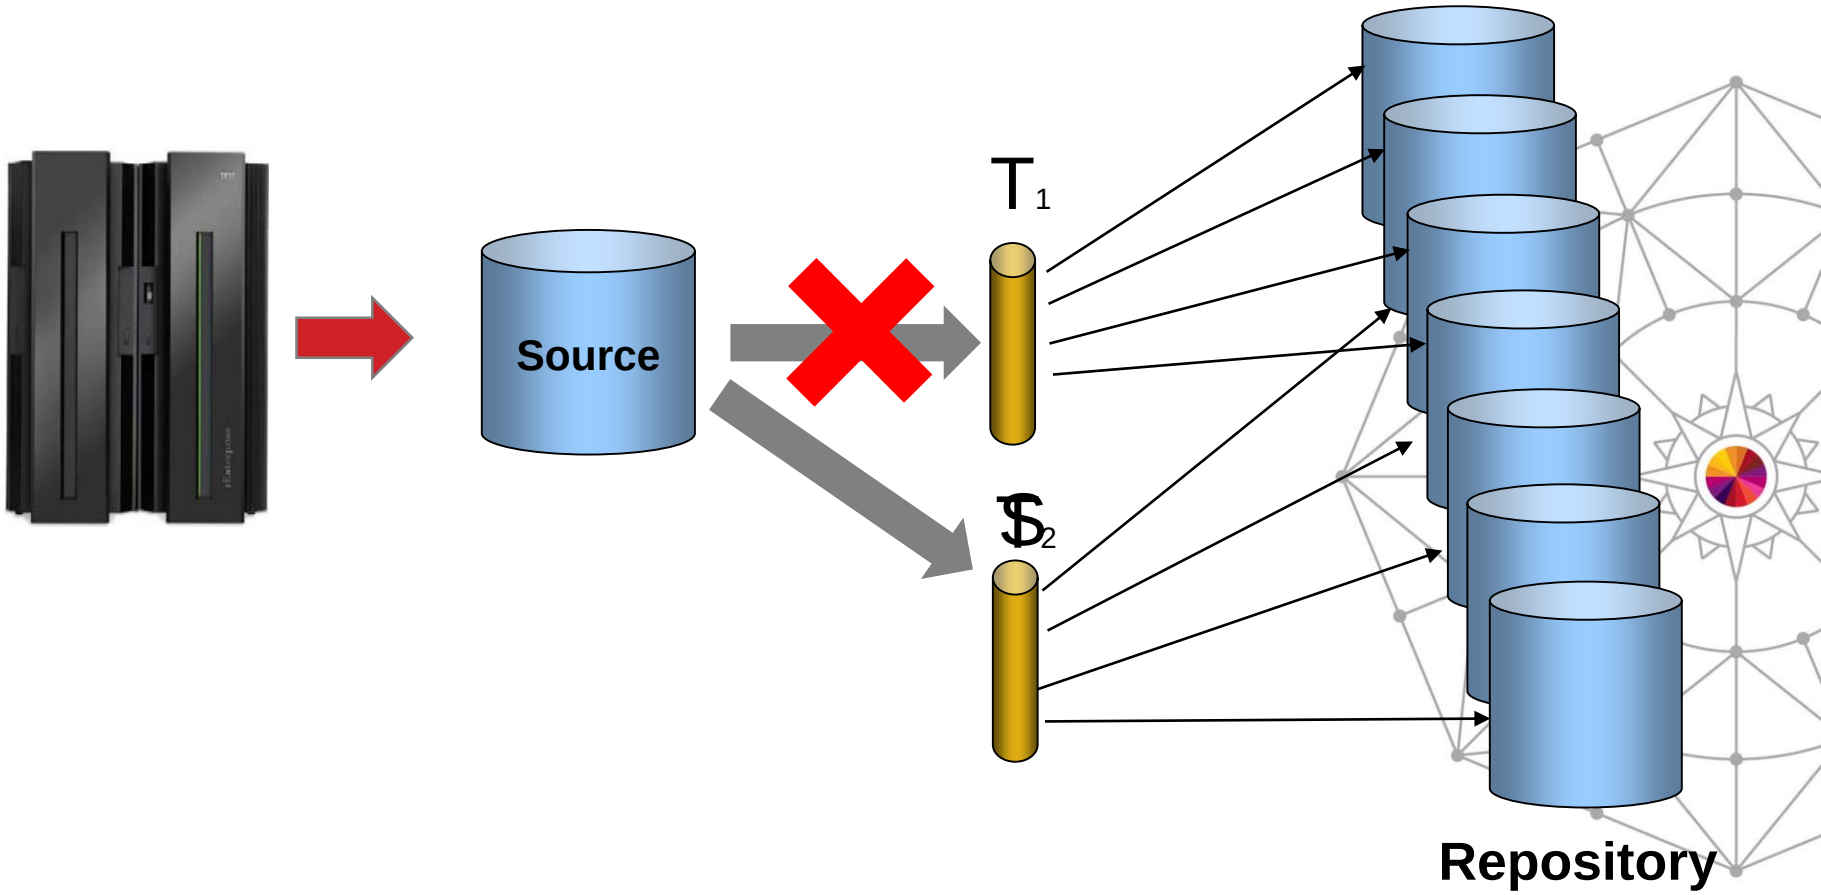

Space Efficient FlashCopy Planning

- On DS8xxx requires a separate license from basic FC
- Requires dedicated FCSE devices in the DS8xxx config
- Careful planning is required
- Has performance impact
- Change rate during life of FC session <5%
- Repository can fill up:
  - IEA499E warnings/errors
  - Repository cannot be dynamically expanded
    - Need to withdraw the relationship
    - Delete, reconfigure and redefine the repository

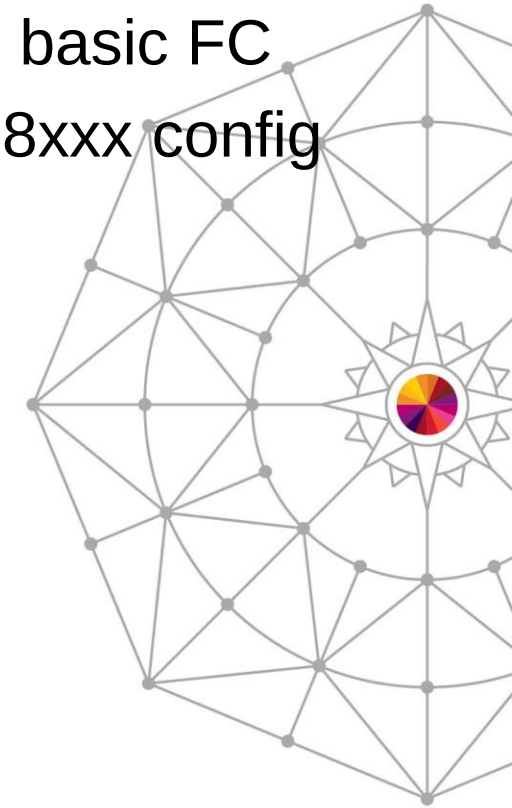

FlashCopy Consistency Groups

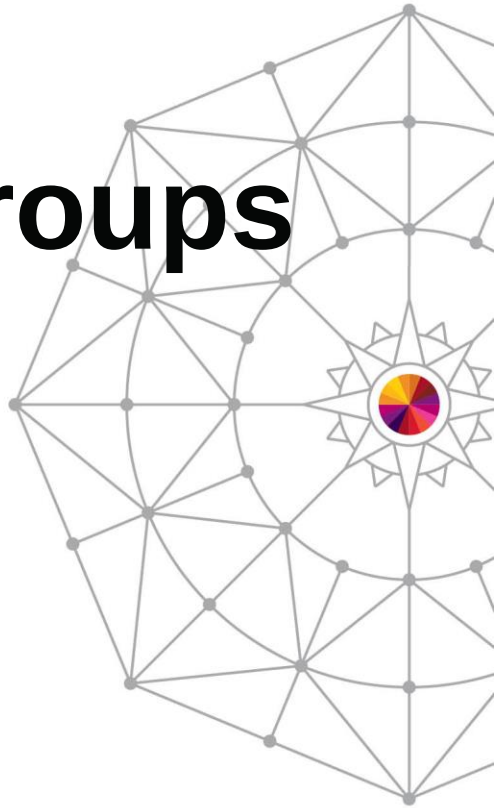

DB2 Status when replicating

For the purposes of the replicating operation there are only three states that a DB2 database can be in when you are going to copy it:

- Shut down
- In log suspend mode
- Running normally

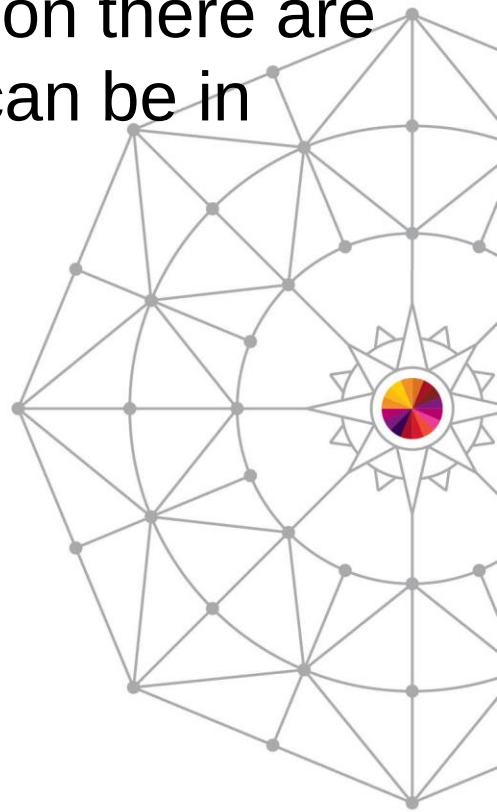

Dependent-Write Consistency

DB2 never makes a change to the database without logging the change first (ARIES principle)

Transaction Changes	Transaction state	Action on Restart
In buffer pool	Uncommitted	Discarded
On disk	Uncommitted	Undo process
In buffer pool	Committed	Redo process
On disk	Committed	No action necessary

FlashCopy Consistency

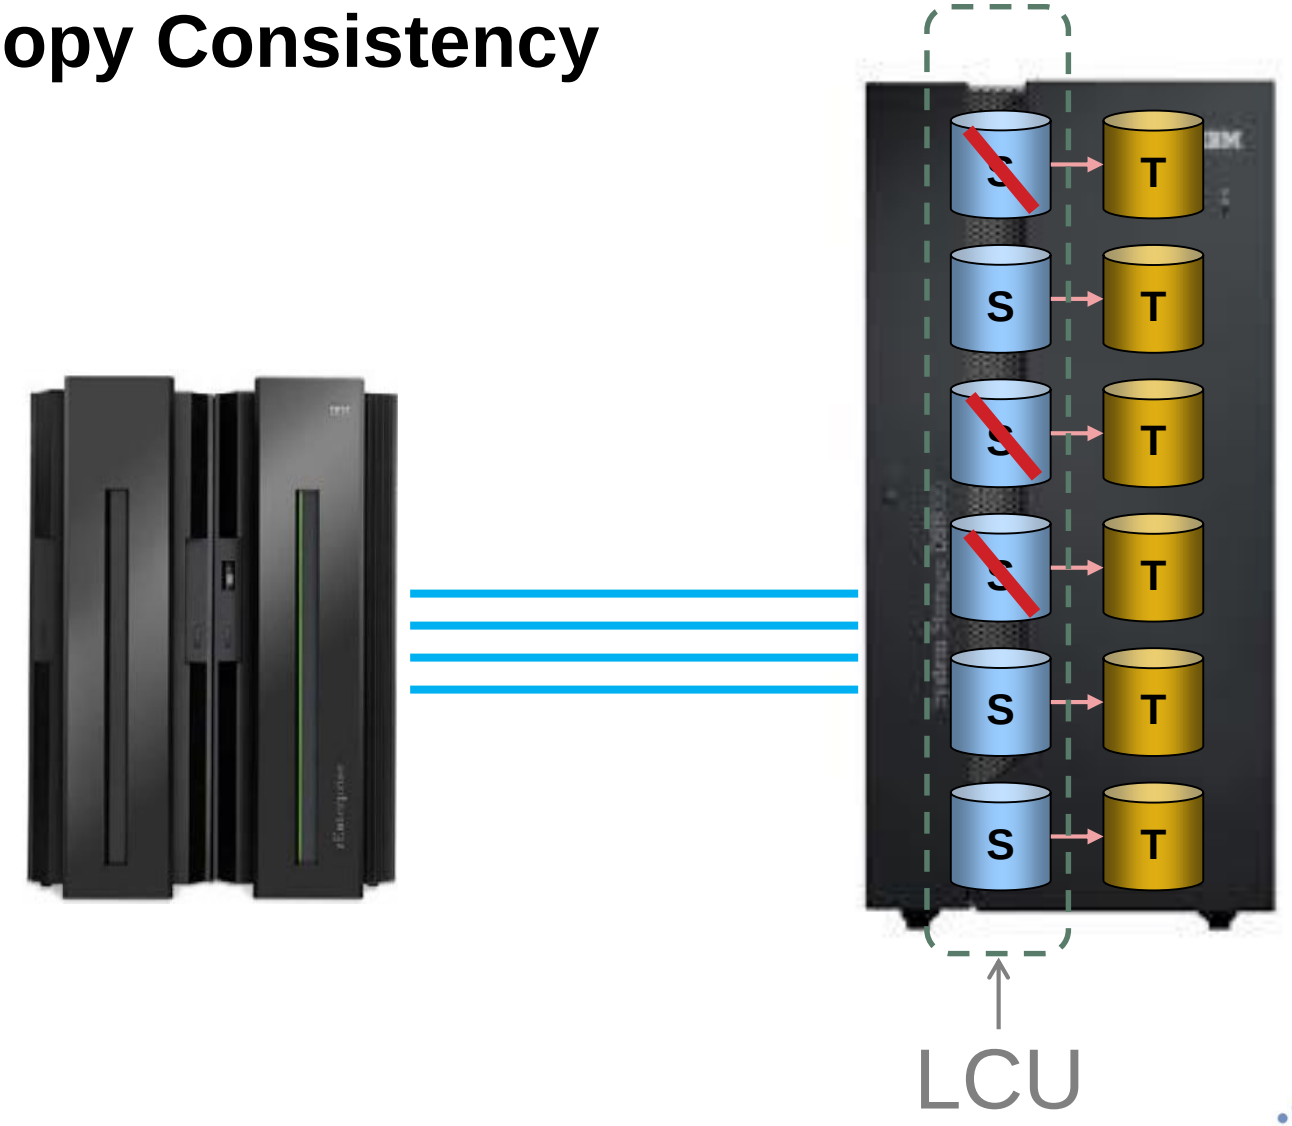

FlashCopy Outside of DB2

- Needs:
  - Consistency groups or
  - Log suspend
- Can replicate beyond the DB2 domain using Con Groups
  - IMS, Other DB2, MQ, VSAM, CICS
- Produces a restartable image

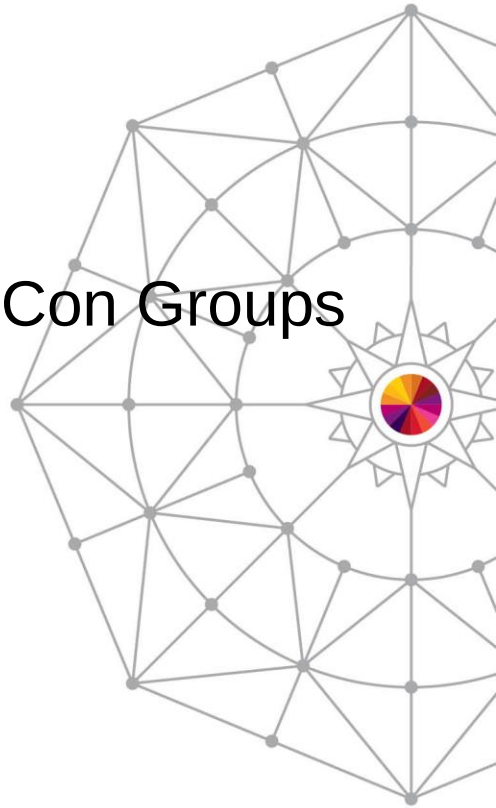

DB2 and FlashCopy

A brief history

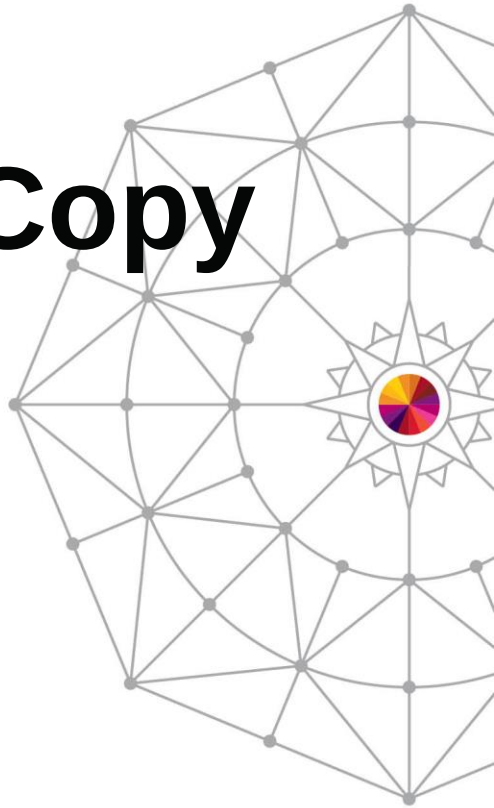

DB2 V8 and FlashCopy

- System Point-in-Time Recovery (SYSPITR)
  - Backup System utility (DSNUTILB)
  - Implemented for very large applications (e.g. SAP)
  - Only stage one (disk copy)
  - DFSMS 1.5 or later
  - Poor integration with DB2/HSM/SMS
  - Difficult to trouble-shoot
  - DB2 domain only
- FC support for CHECK INDEX SHRLEVEL CHANGE

SYSPITR Timeline

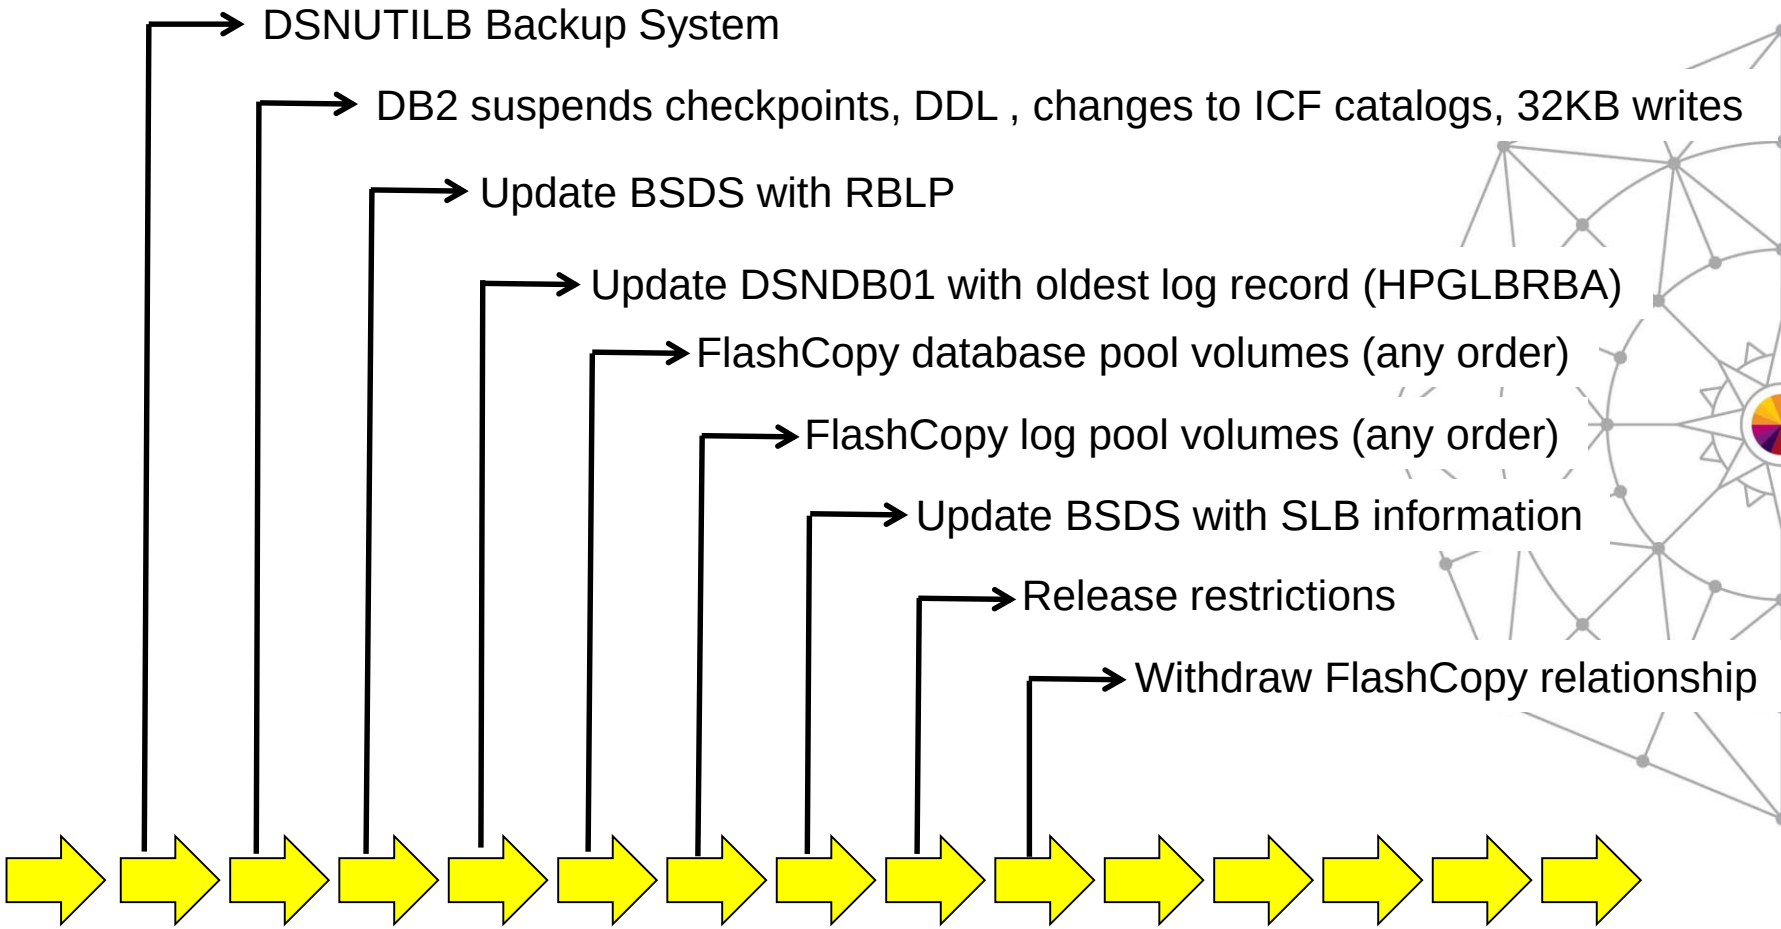

DB2 V9 and FlashCopy

- SYSPITR
  - Dataset restore from SLB
  - Still using SMS/HSM
  - Capability to dump target volumes
    - DUMP/DUMPONLY
  - Some ICF catalog issues for datasets being moved
- FC support for CHECK INDEX SHRLEVEL CHANGE
- FC support for CHECK LOB SHRLEVEL CHANGE
- DB2 domain only

DB2 V10 and FlashCopy

- SYSPITR
  - ICF catalog issues resolved (mostly)
  - Incremental Flash enabled
    - BACKUP SYSTEM ESTABLISH FCINCREMENTAL
  - Still using HSM/SMS
- DB2 domain only

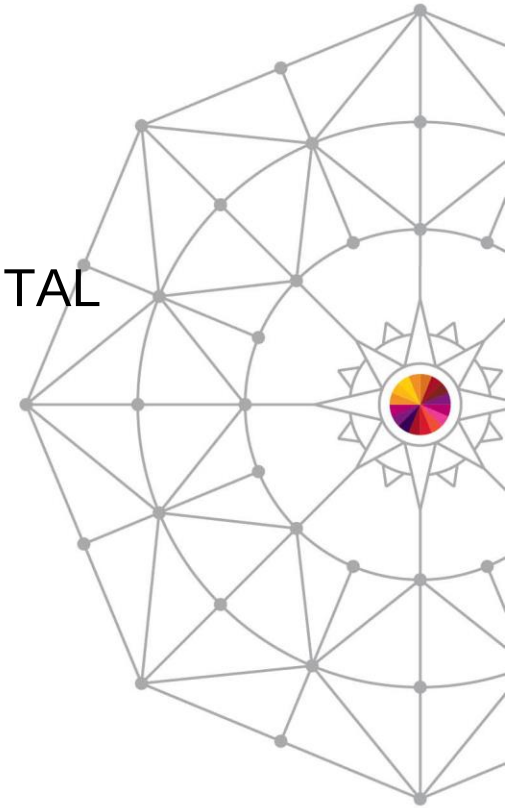

DB2 Full Volume Cloning

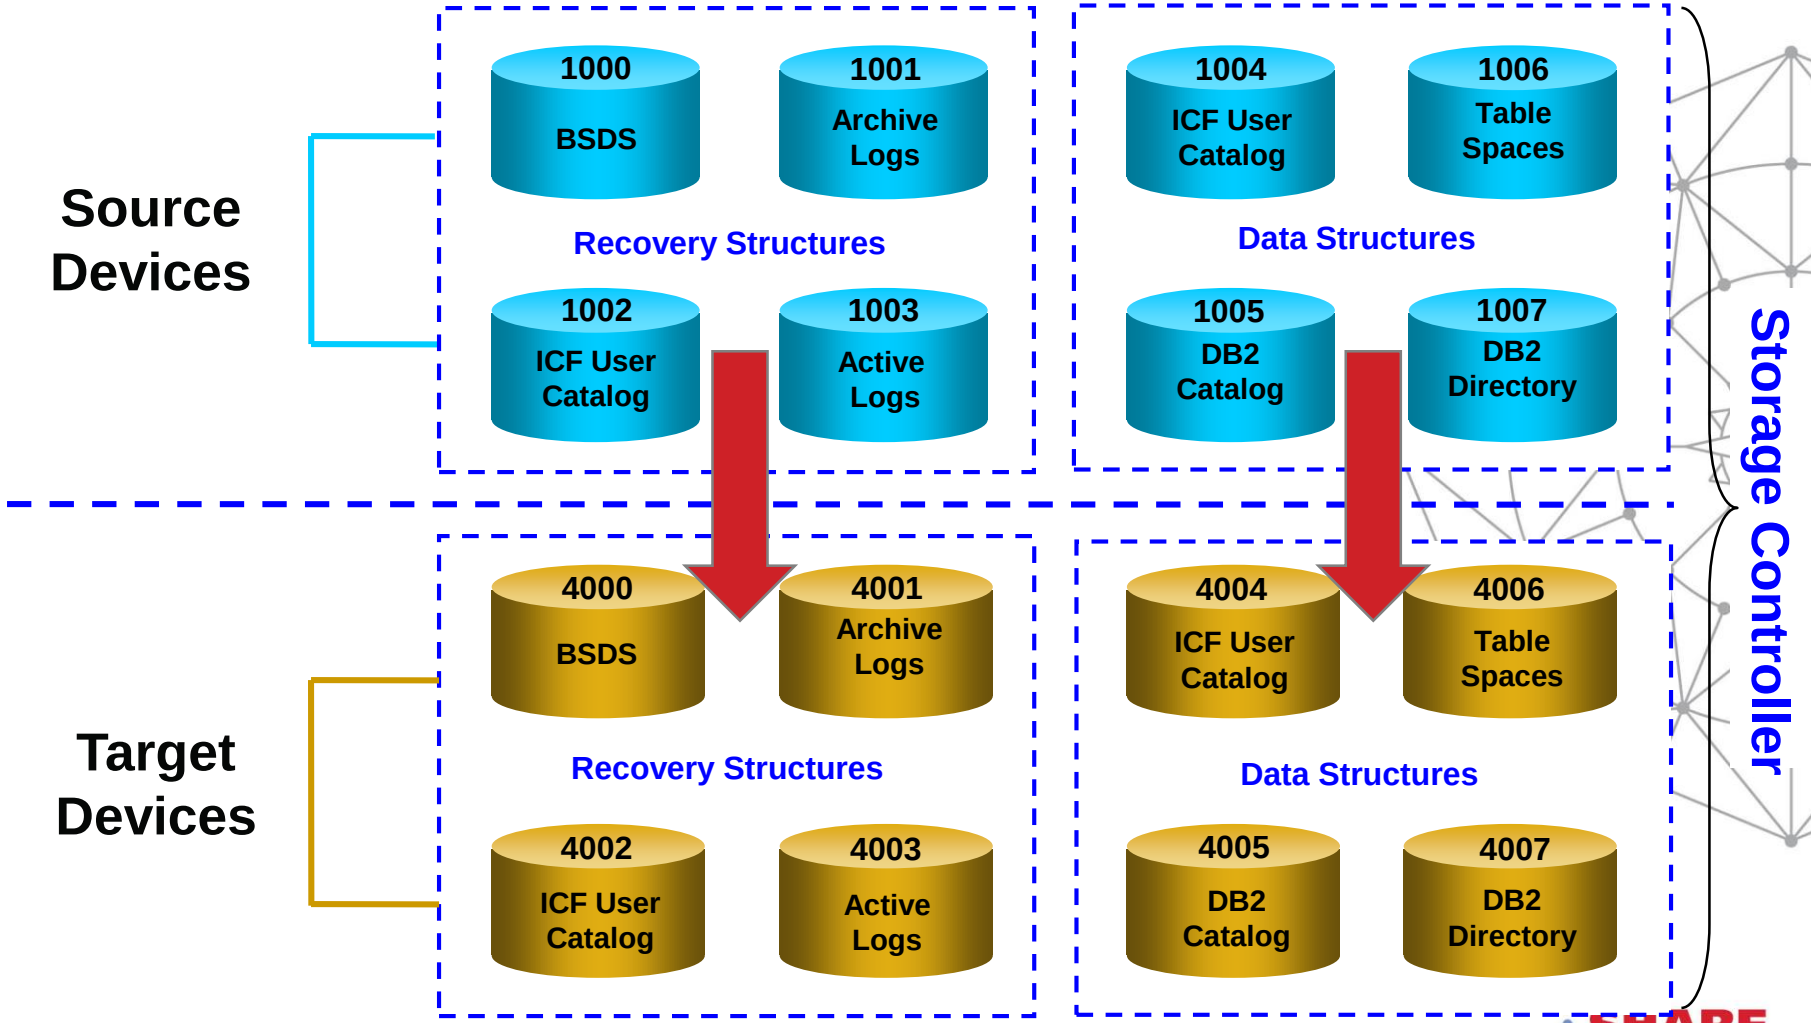

DB2 and Other FlashCopy Functions

- ZPARM
  - CHECK_FASTREPLICATION= (PREFERRED/REQUIRED)
- Some utilities can use FlashCopy
  - CHECK DATA
  - CHECK INDEX
  - CHECK LOB
  - REORG TS
  - REBUILD INDEX
  - REORG INDEX
  - COPY
  - LOAD
- FCIC in DB2 V10
  - With consistency!

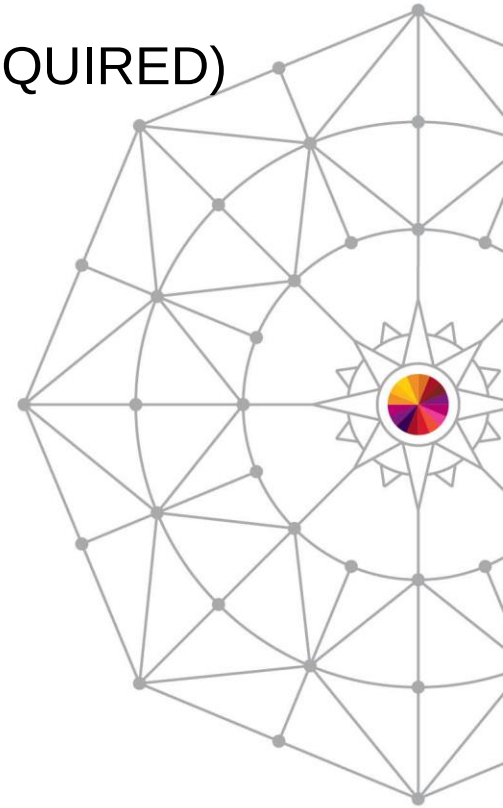

FlashCopy, DB2, and Remote Replication

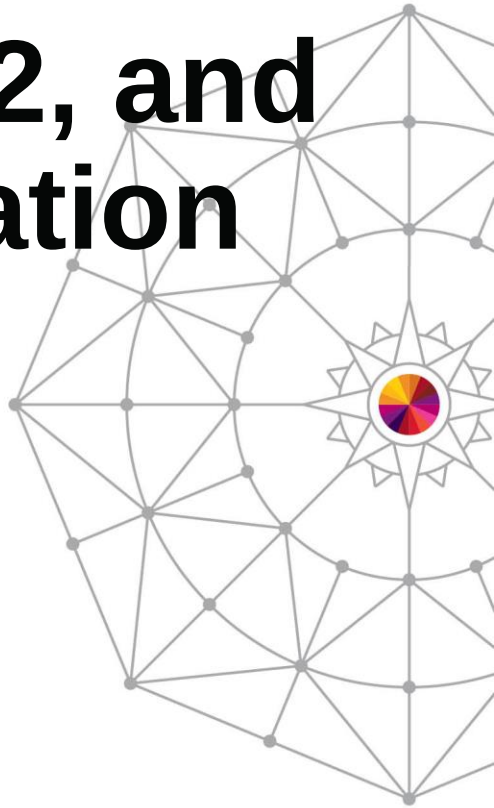

Array-based Remote Replication

Synchronous (Metro Mirror aka PPRC)

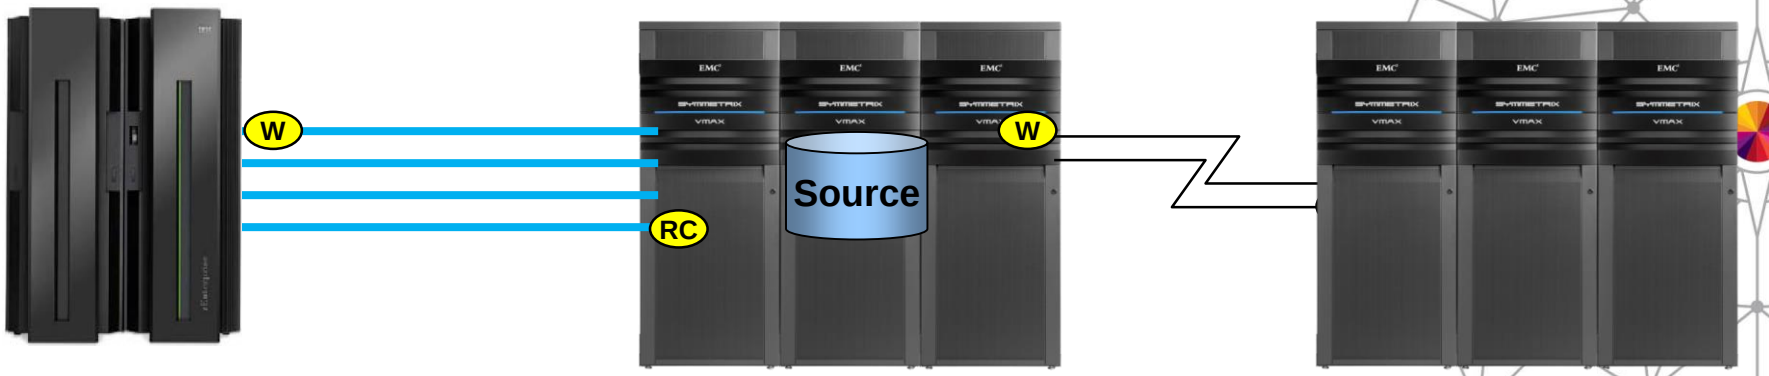

Considerations for Synchronous Replication

- Write performance
  - Log writes
    - VSAM striping can help mitigate some latency
  - Synchronous buffer pool writes
  - Page steal writes
  - Increased DISC time
- Database consistency
- Distance impact on writes
  - Every 200km miles adds 2ms additional latency (round trip)

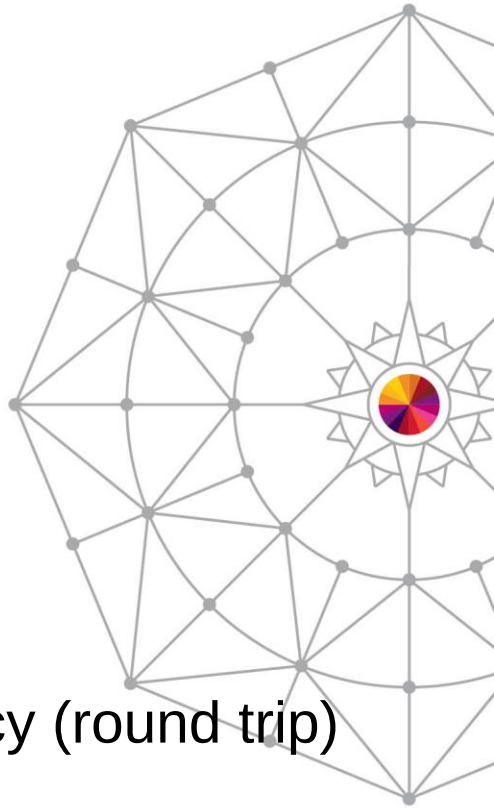

Array-based Remote Replication

Asynchronous (z/OS Global Mirror aka XRC)

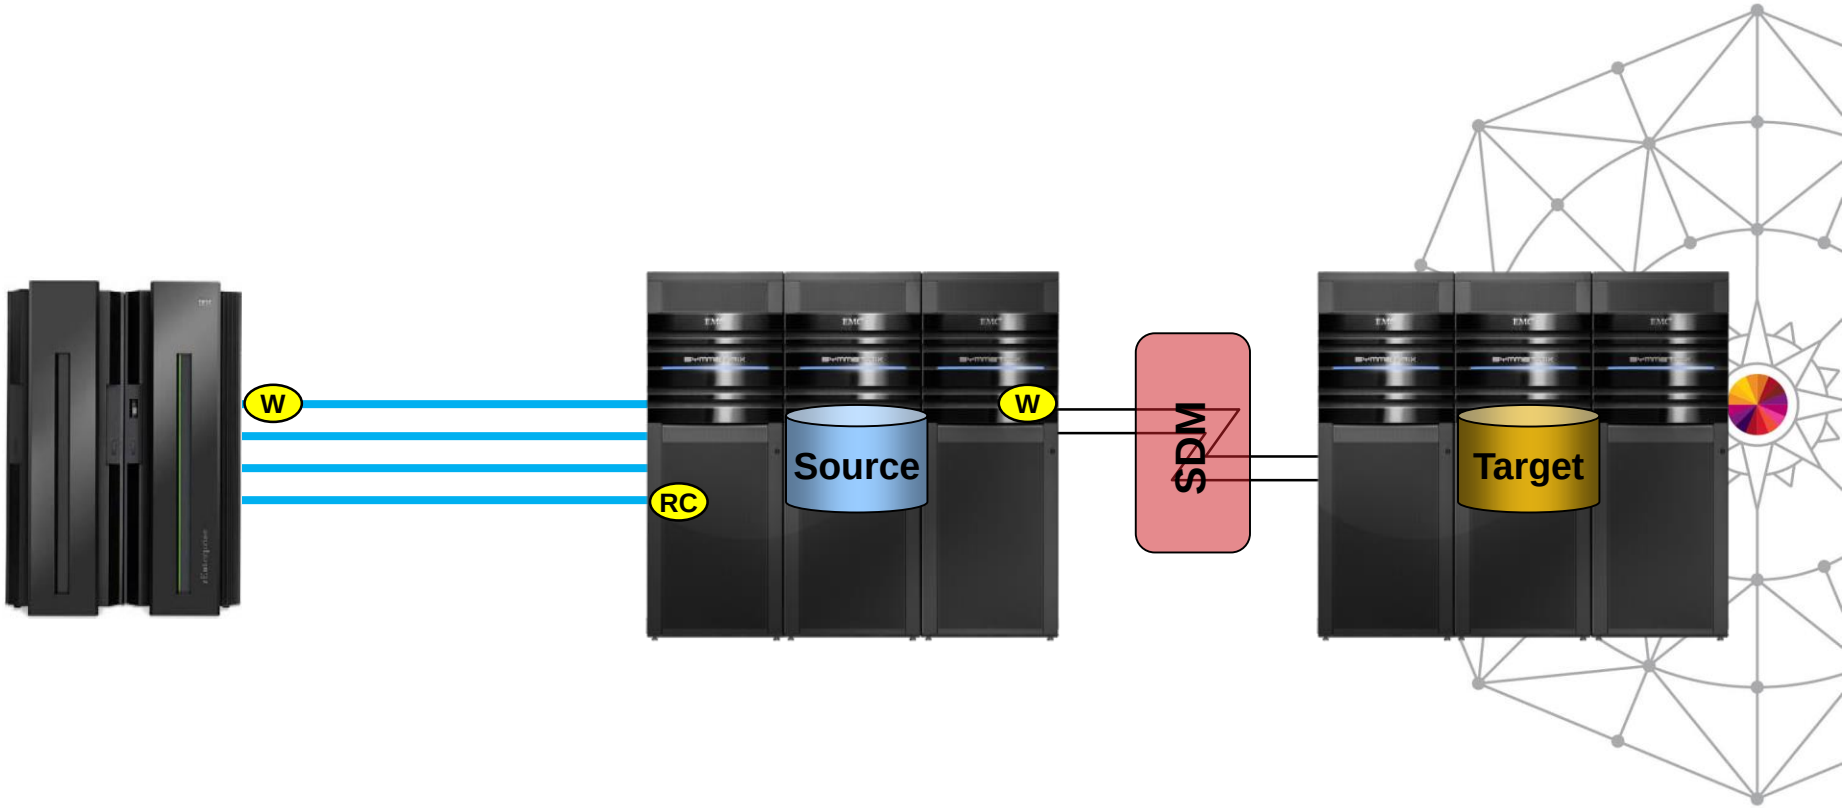

Considerations for Asynchronous Replication XRC

- Host write performance largely unaffected
- z/OS timestamps the writes
  - Writes are accumulated in a side file
  - System data mover on the target side reads the side file
- Some data loss is expected

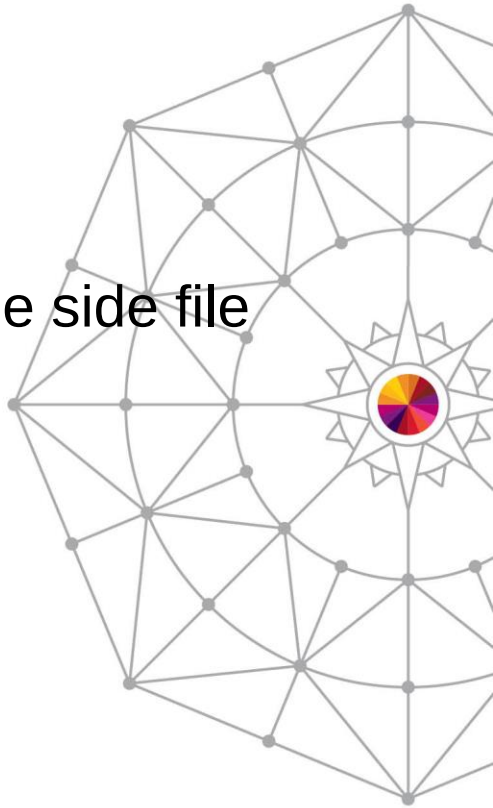

Array-based Remote Replication

Synchronous (PPRC)

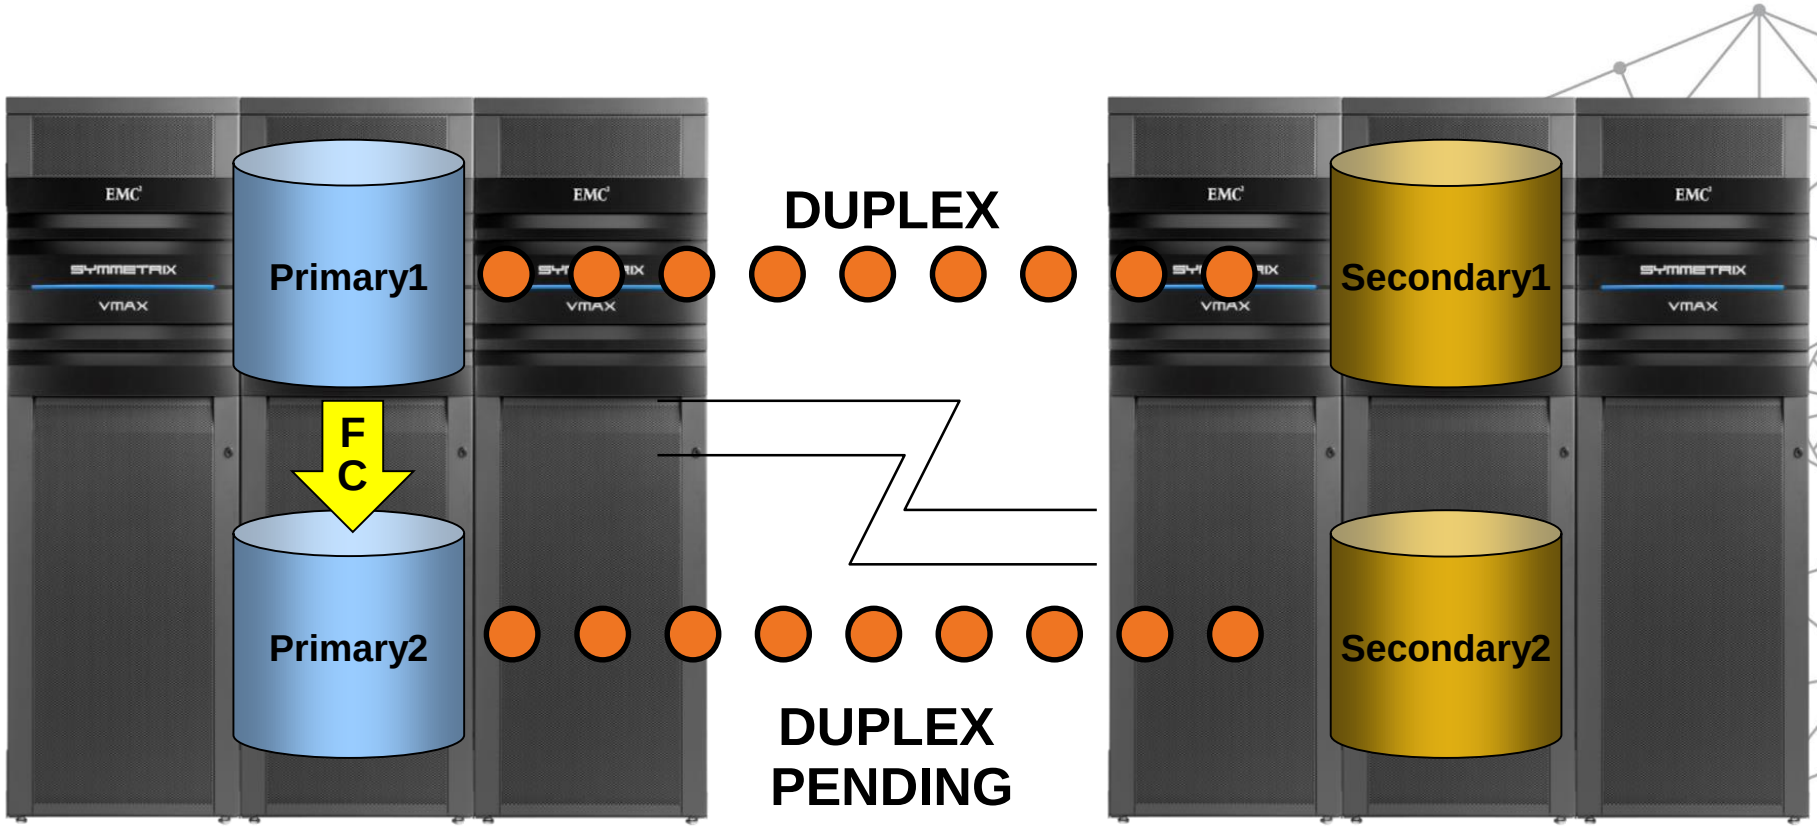

Duplex Pending Status

- Generally not acceptable in PPRC installations
- Secondary site may not be recoverable
- HyperSwap event will fail

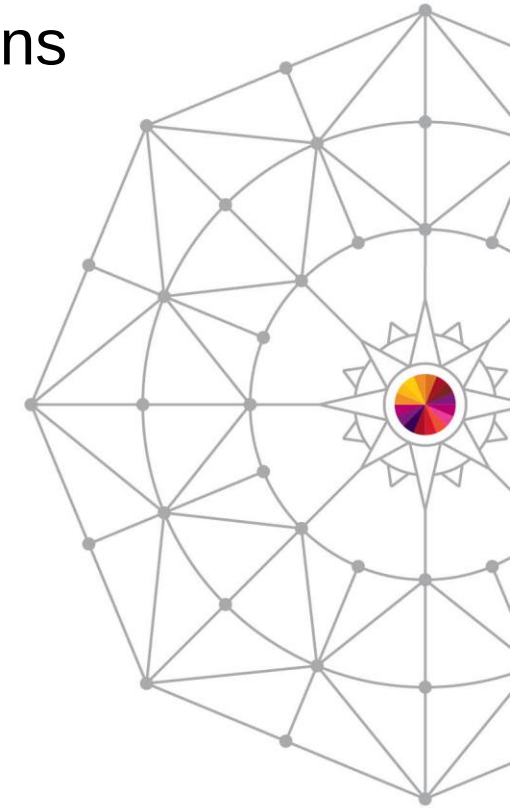

Solutions to Duplex Pending Issue

- SYSPITR
  - Set up COPY POOLs outside of PPRC
- DB2 utilities
  - Point temporary files at non-PPRC SMS storage group
    - ZPARM UTIL_TEMP_STORCLAS

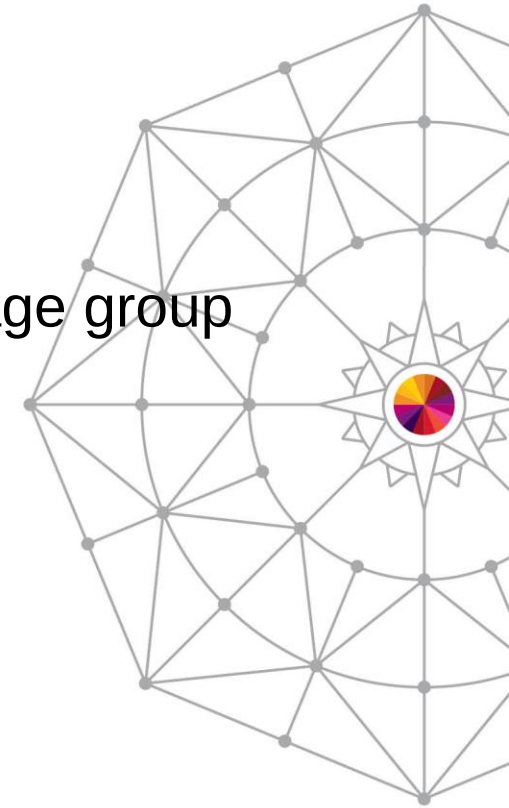

Preserve Mirror and PPRC

AKA Remote Pair FlashCopy

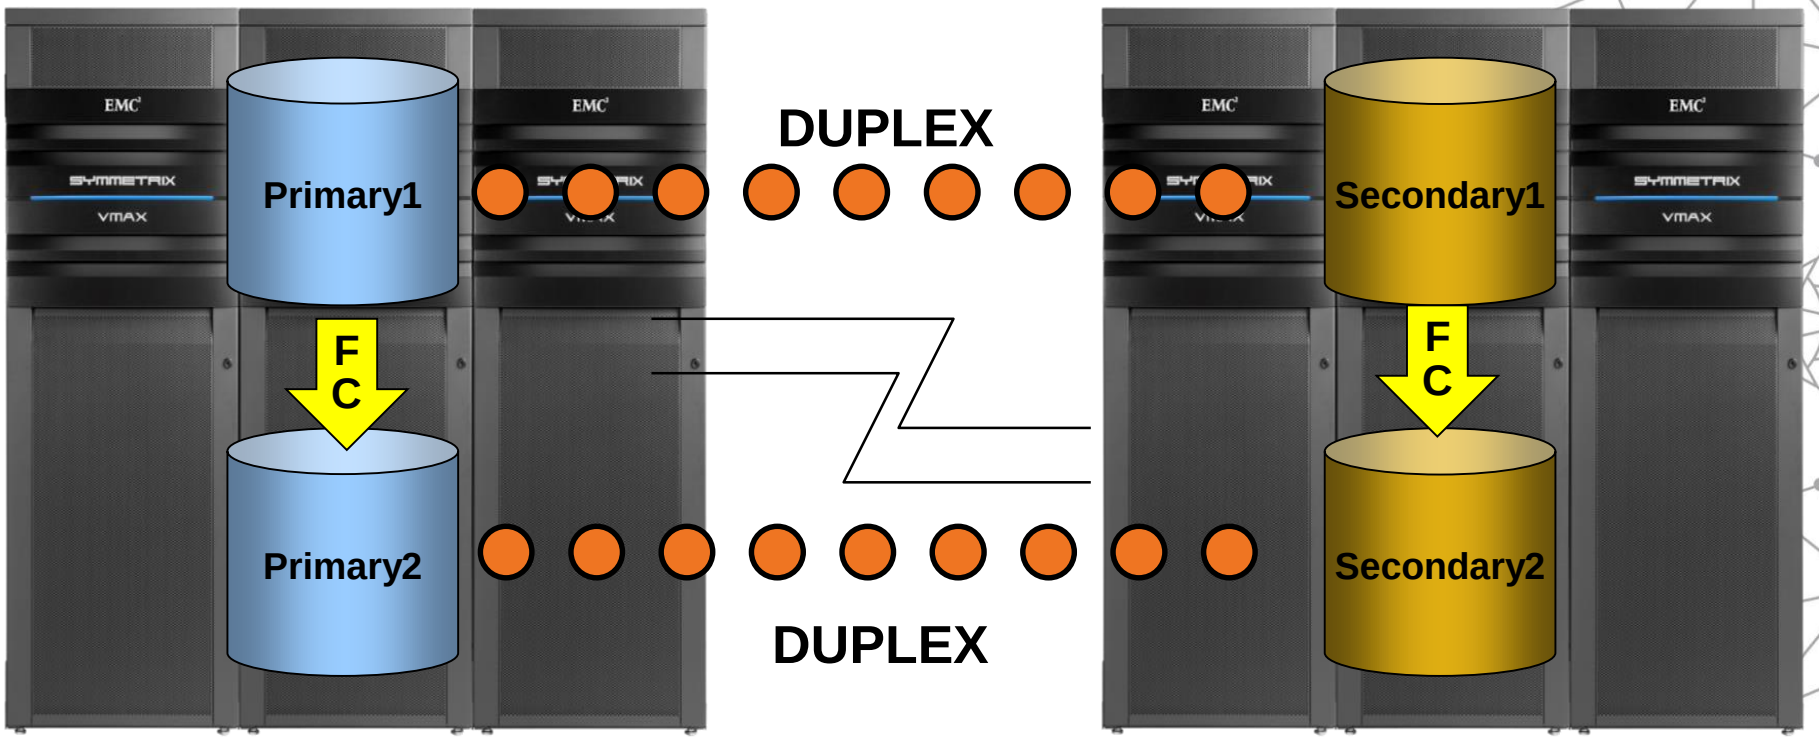

Preserve Mirror Considerations

- Requires specific microcode levels and APARs
- FlashCopy Source and Targets must be PPRC Primary volumes
- HyperSwap will work when Preserve Mirror is used
- Copy work performed by FlashCopy does not have to be transmitted across the network

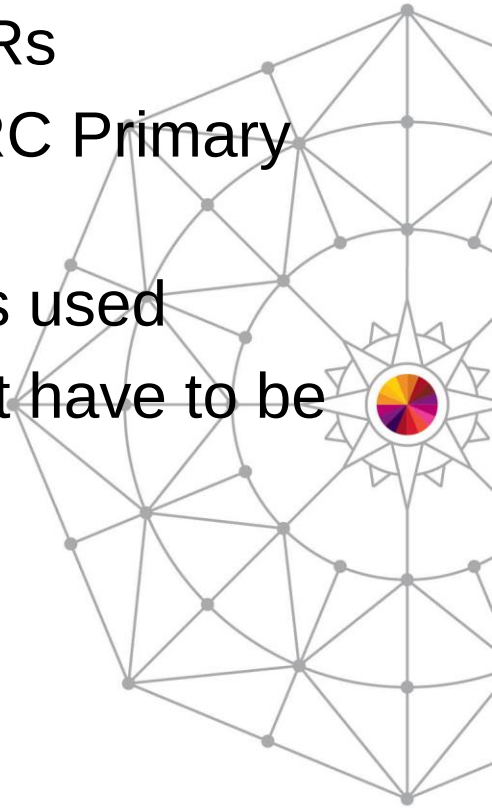

DSNZPARM FLASHCOPY_PPRC

- Values
  - NONE results in DUPLEX PENDING
  - PREFERRED means DUPLEX PENDING is allowed
  - REQUIRED means DUPLEX PENDING is not allowed (GDPS)
  - Blank no preserve mirror command sent to DFSMSdss
- Applies to DB2 Utilities
  - CHECK DATA/INDEX/LOB
  - REORG TABLESPACE/INDEX
  - REBUILD INDEX
  - LOAD
  - COPY/RECOVER

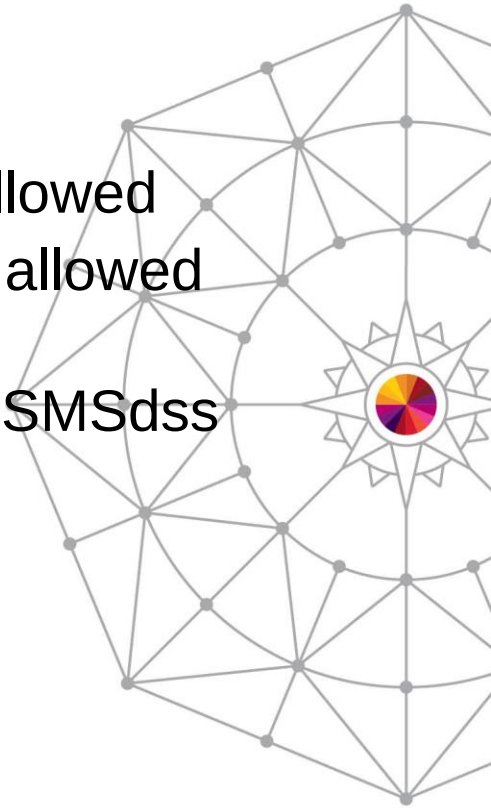

DB2 Utilities using DFSMShsm and Preserve Mirror

- No ZPARM settings oddly
- Preserve Mirror attribute set at the Copy Pool level in SMS
  - FRBACKUP (NO|PN|PP|PR)
  - FRRECOV (NO|PN|PP|PR)

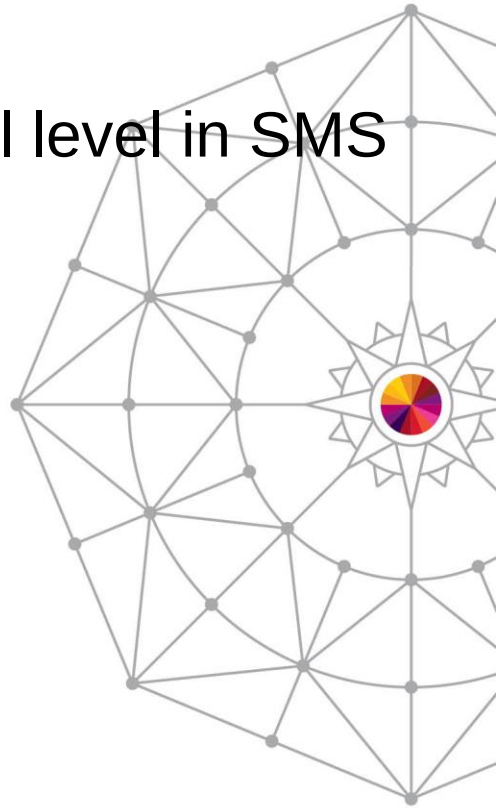

FlashCopy, DB2, and XRC

- FlashCopy to a XRC source is not allowed
  - Consider how z/OS timestamps the writes
- BACKUP SYSTEM
  - Keep copy pools separated from XRC
  - Restore System cannot use FlashCopy
    - Or disable mirroring
  - Object recovery slow replication
- DB2 utilities
  - Point temporary files at non-XRC SMS storage group
    - ZPARM UTIL_TEMP_STORCLAS

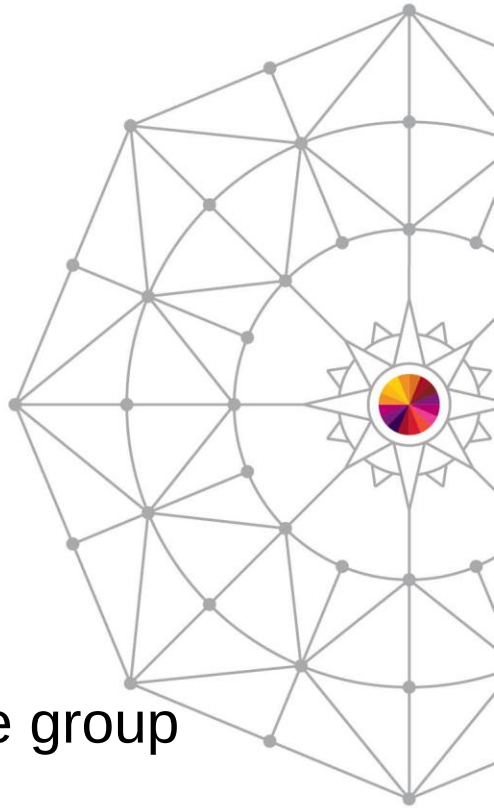

Drivers for Selecting the ACS Solution

- Restartable or recoverable
- Federated with other applications/platforms
  - z/OS
  - DB2 LUW
  - MQ
- Space efficient copies
- DB2 availability
- Volume or dataset
- Object restore requirements
- Integration with DB2 recovery
- Considerations for remote replication

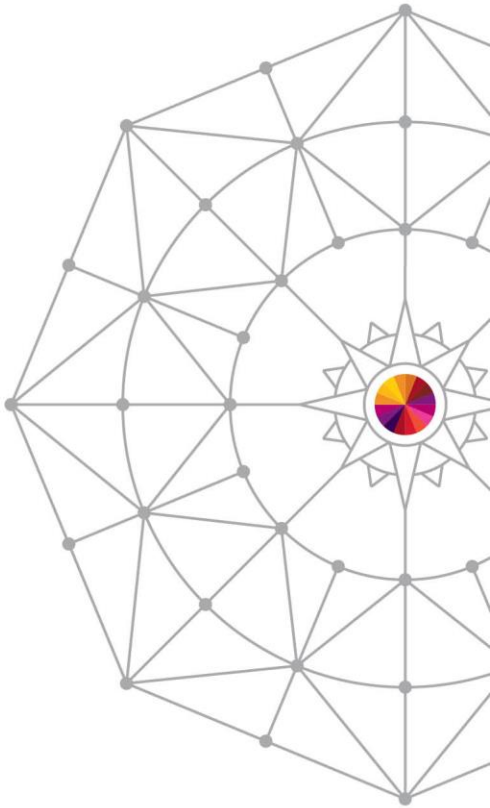

Conclusions

- FlashCopy replication can enhance:
  - System backup and restore
  - Object backup and restore
- Care must be taken to deploy the correct options for a given situation
- More care is necessary when using remote replication
- Collaborate closely with system programmers and storage administrators for success

#SHAREorg

SHARE

DB2 and Advanced Copy Services: All You Need to Know

Paul Pendle

Independent Consultant

Tuesday, March 11, 2014

Session Number: 15352

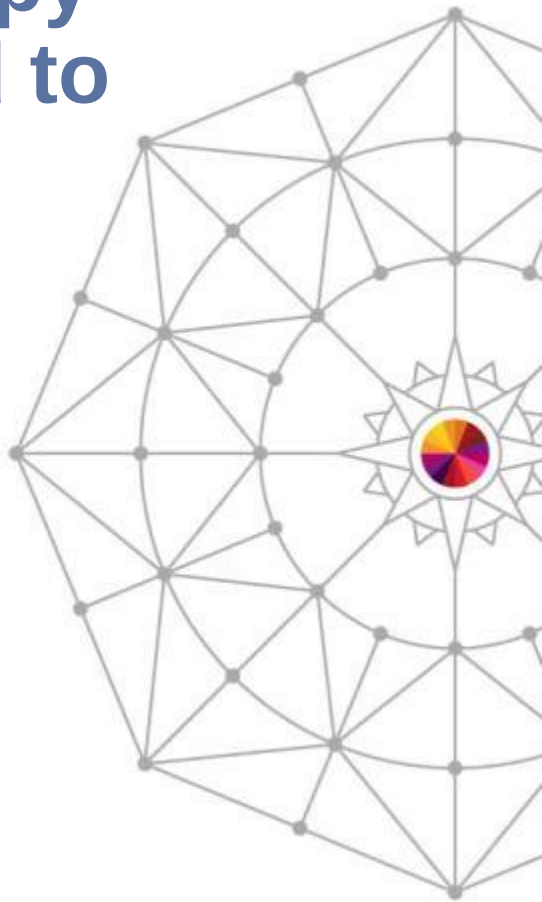

SHARE in Anaheim