DB2 for z/OS and Storage Advanced Copy Services: Creating a DB2 System Level Backup

Using FlashCopy and/or TimeFinder on EMC Hardware

Session 16796

Brett Quinn, EMC Corporation

Tom Moulder, EMC Corporation

#SHAREorg

SHARE is an independent volunteer-run information technology association that provides **education**, professional **networking** and industry **influence**.

SHARE in Seattle 2015

Topics

- Local replication technologies
  - FlashCopy
  - TimeFinder
  - Achieving ‘space efficiency’
- DB2 System Level Backup
  - Concepts / Benefits
  - Rocket Software Utilities
  - Backup to Virtual Tape

Volume Flashcopy

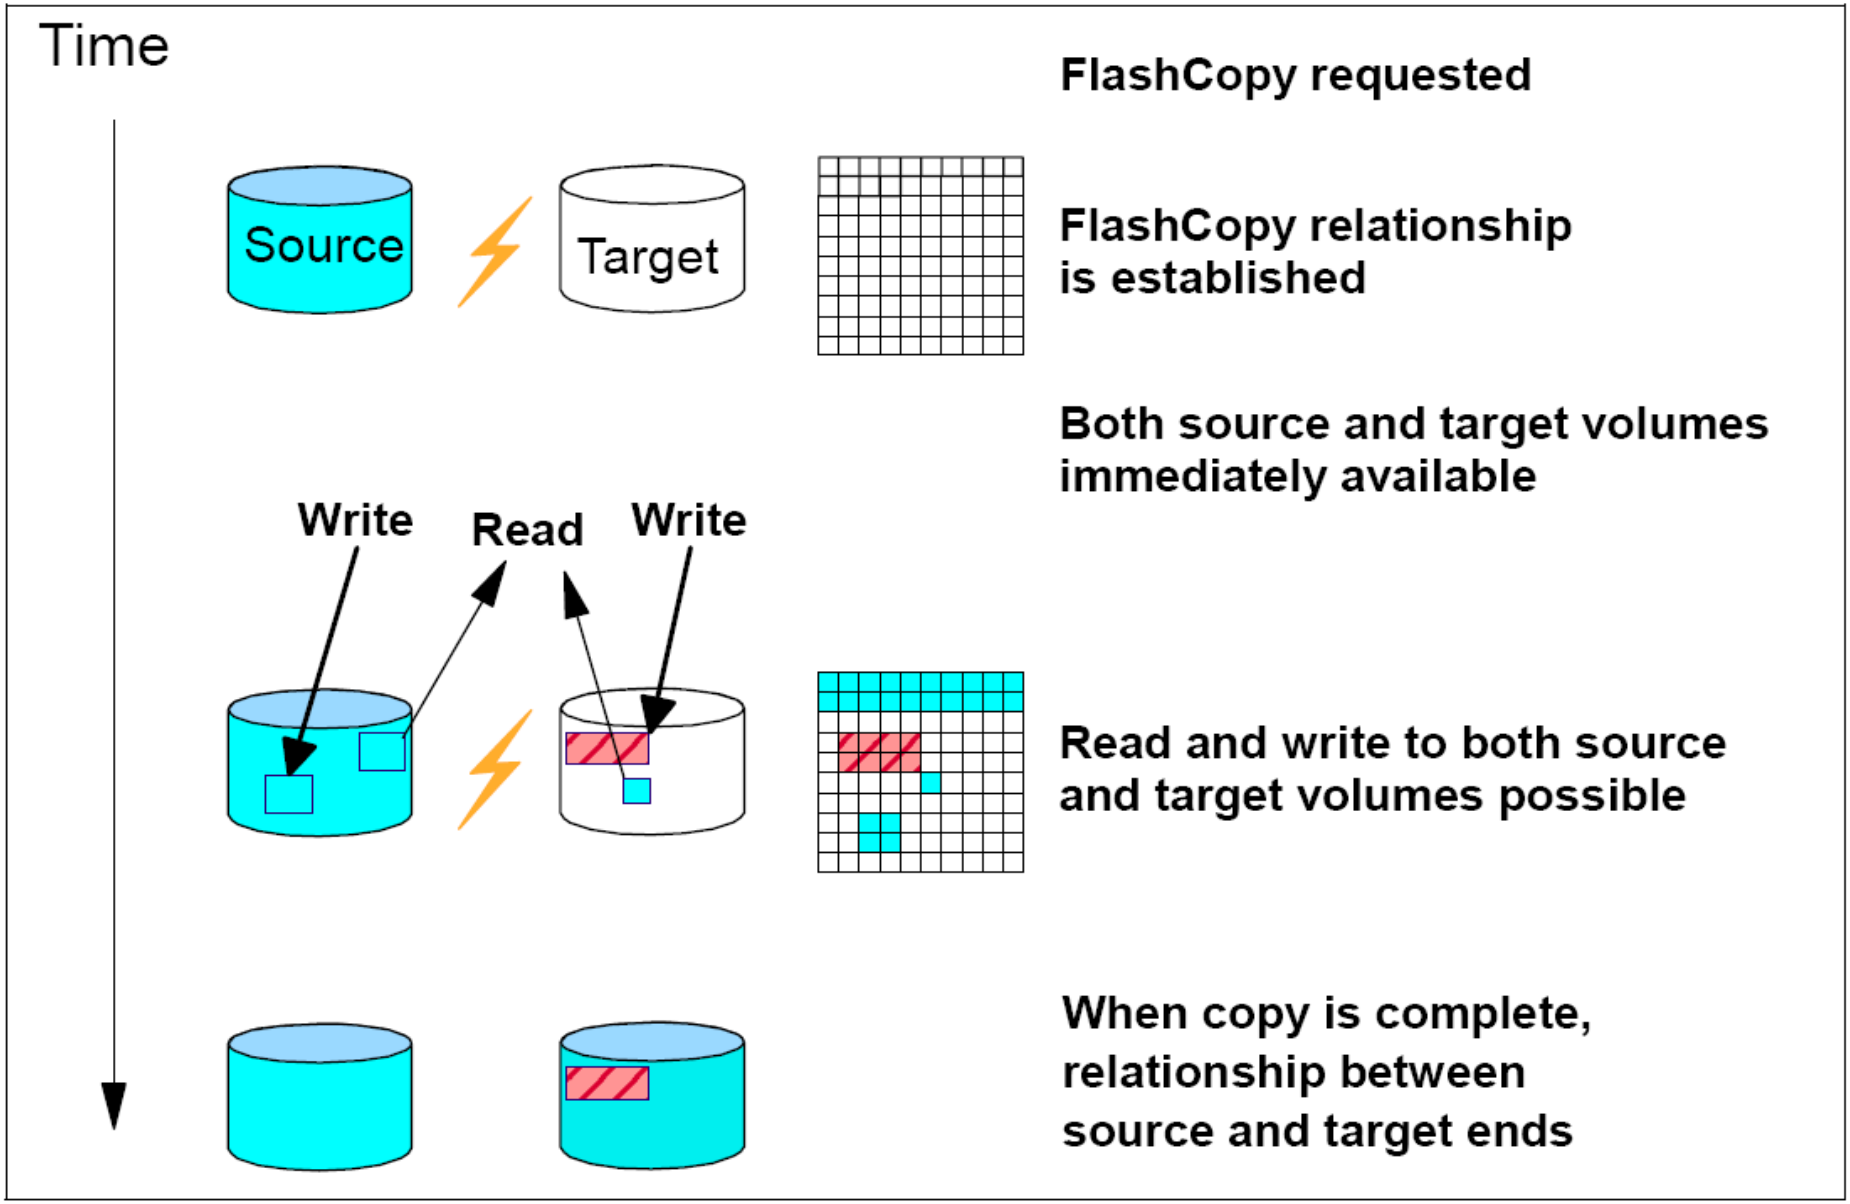

Figure 2-1 FlashCopy with background copy Source IBM Redbook: SG24-5680-04

Space Efficient FlashCopy

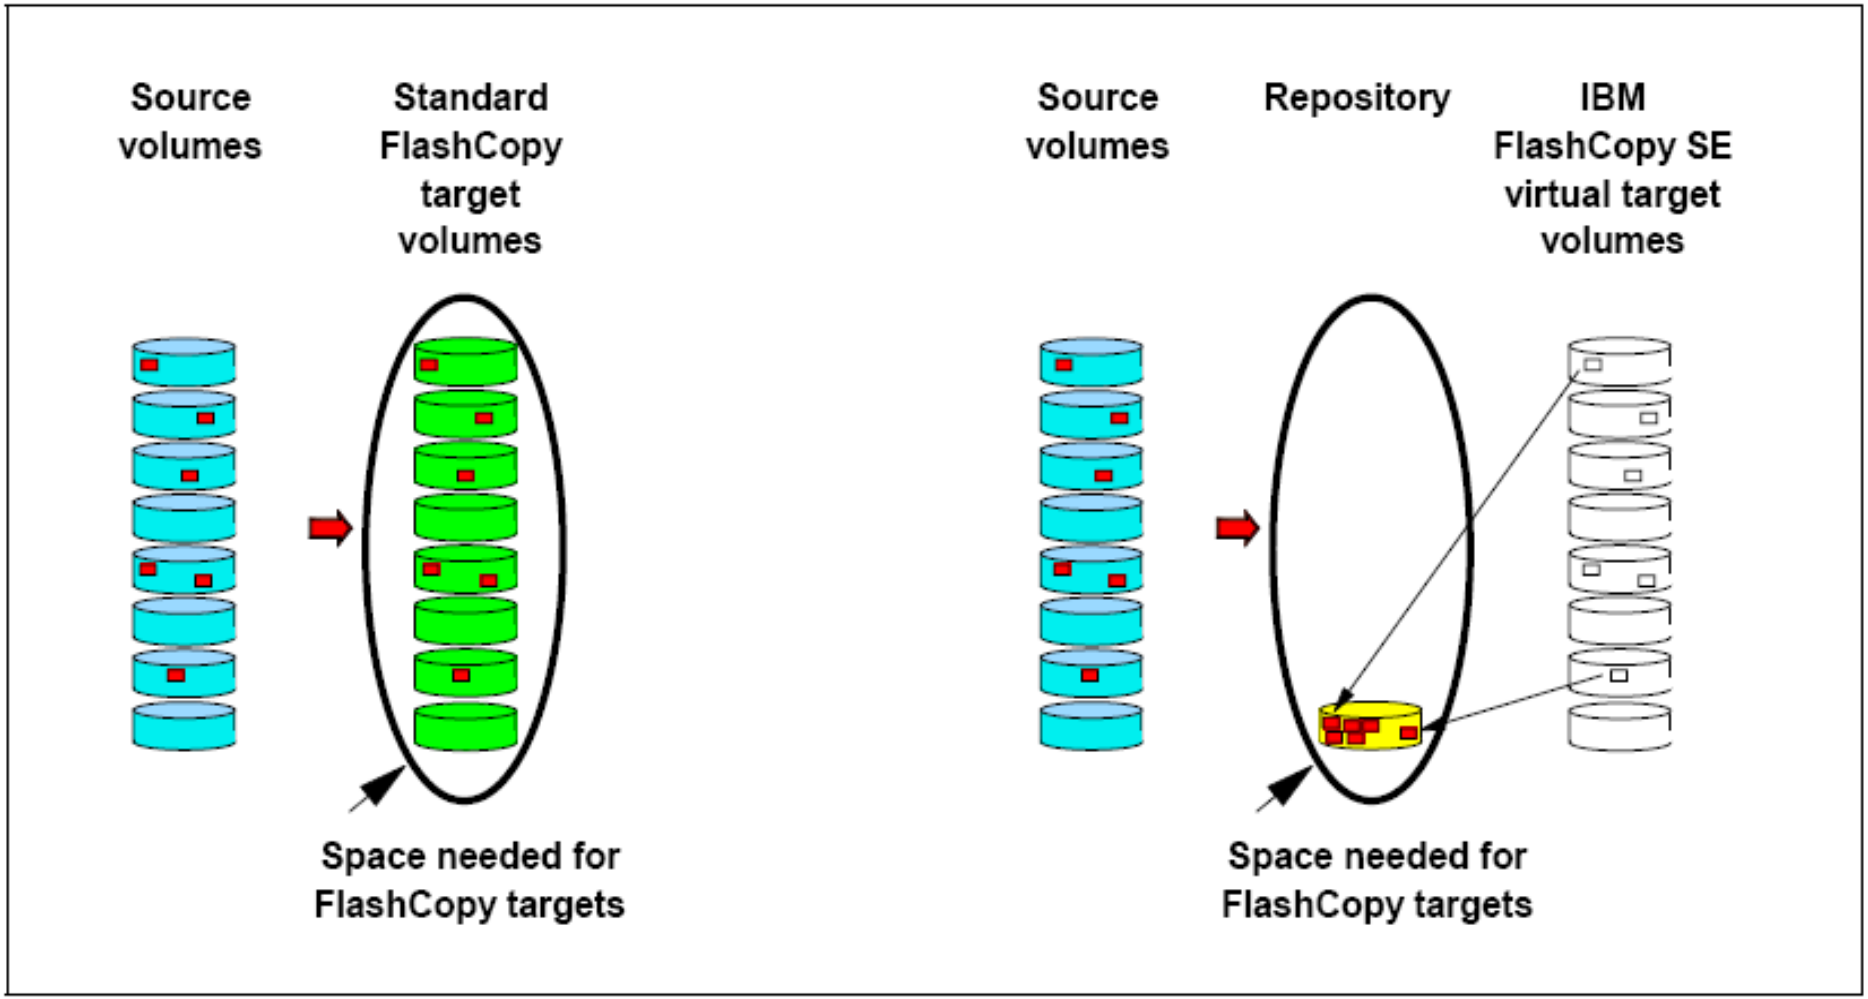

Figure 11-1 Concept of FlashCopy SE

Source: IBMREDP-4368-00

EMC TimeFinder Features

- TimeFinder/Clone
  - Up to 16 high-performance physical copies
  - Full-volume (w/precopy) and dataset level
  - Incremental resync (dataset and multi-volume)
  - Consistency Group support (**dataset and volume!**)
- TimeFinder/Snap
  - Space-saving volume snapshot images (up to 128)
  - Uses Virtual Devices (VDEVs)
  - Typically requires less than 30% additional capacity
  - Async COFW, Dynamic SAVEPOOL expansion
  - Incremental resync
  - Consistency Group support

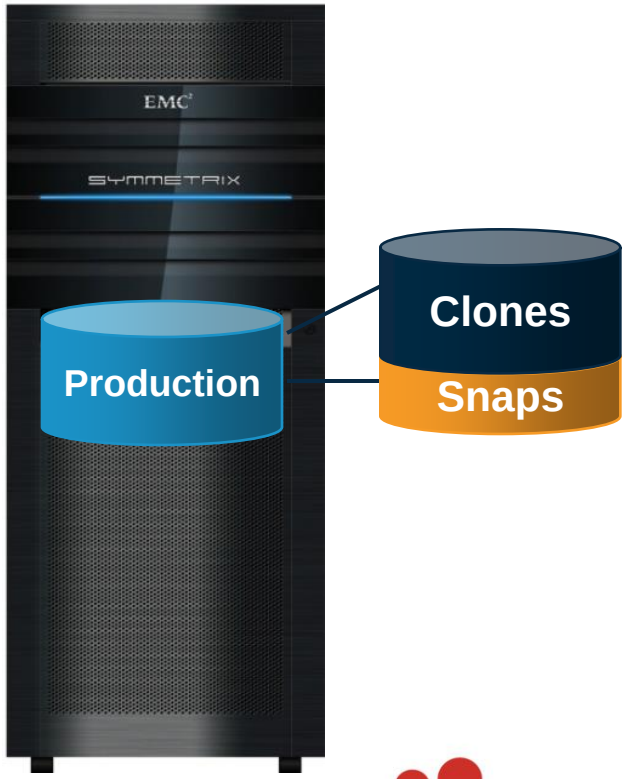

TimeFinder/Snap – Pointer-Based SNAPs

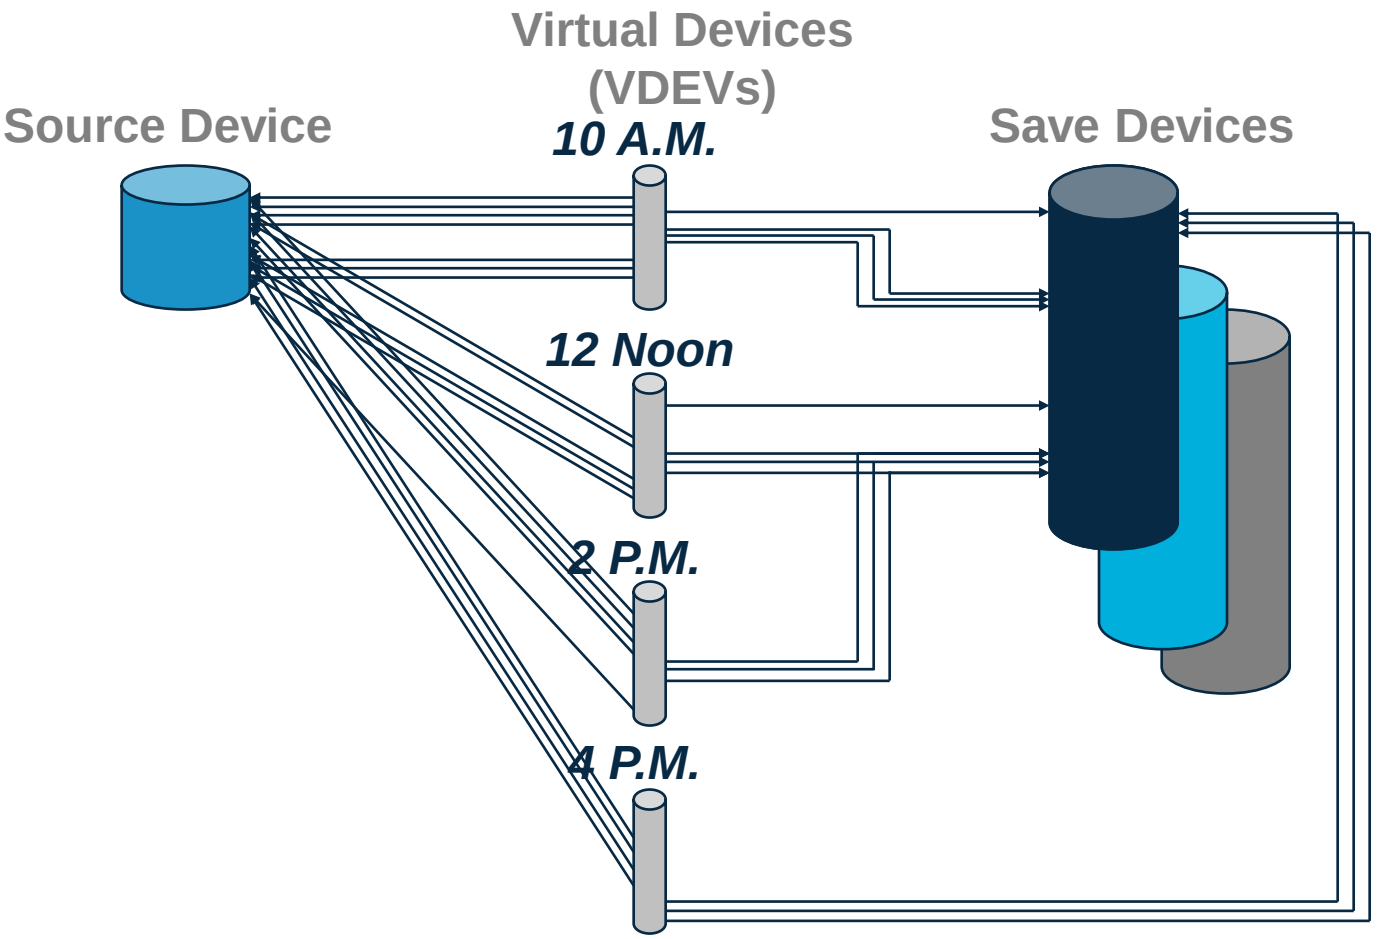

Virtual Provisioning on EMC VMAX

- Storage capacity is structured in pools defined by RAID protection, drive technology, rotational speed
- Thin devices are logical volumes that are provisioned to hosts
- Workload is spread across many disks in 12 track chunks

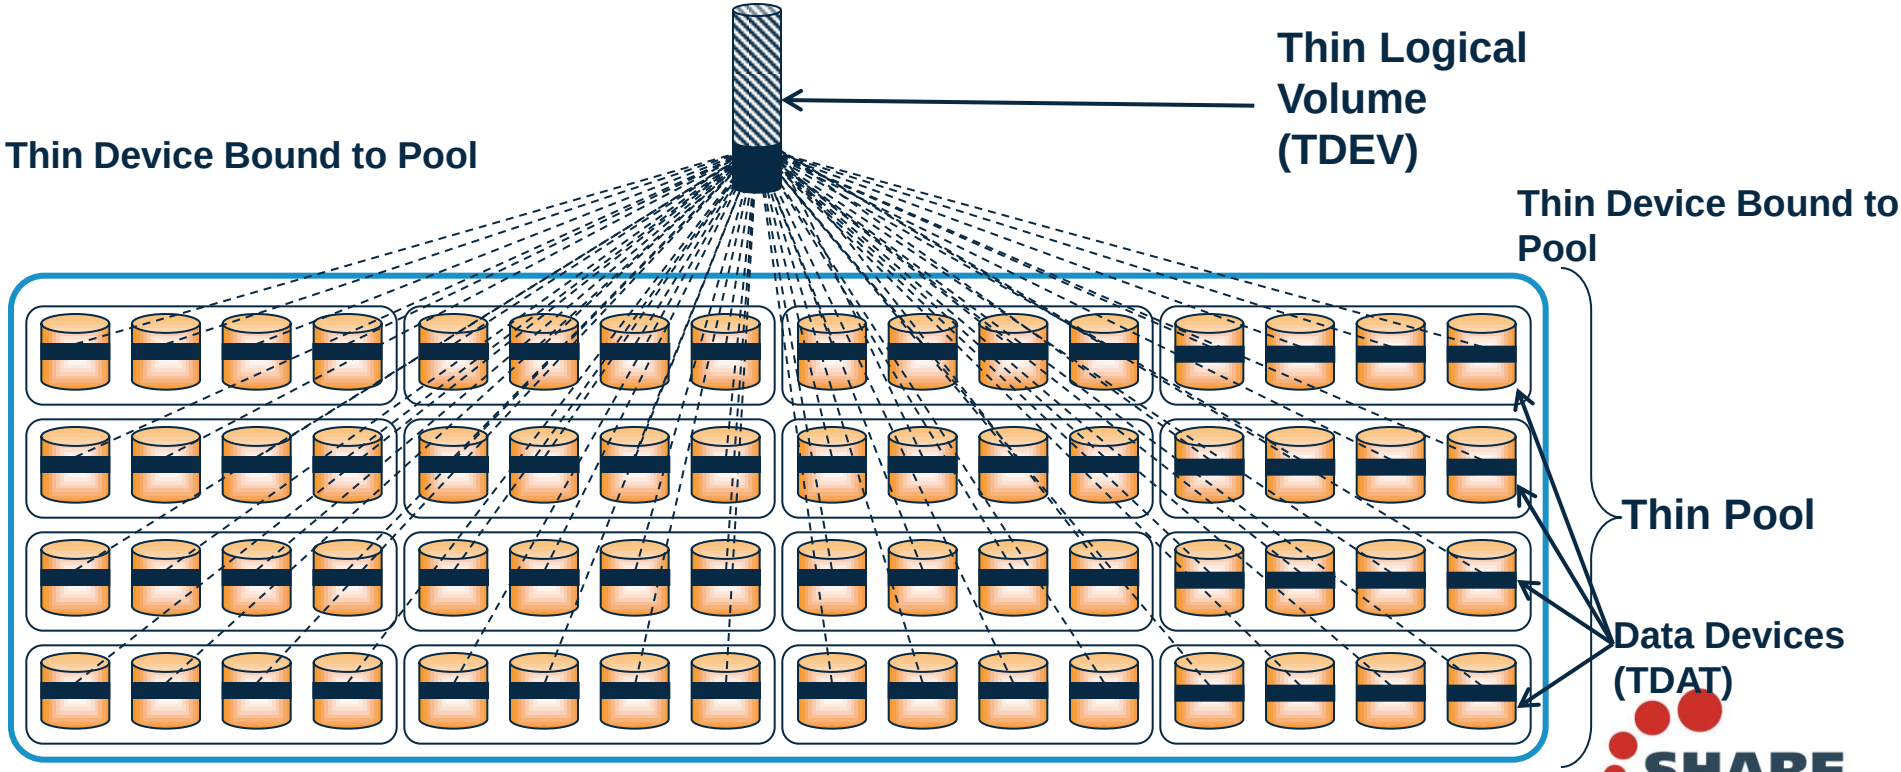

TimeFinder/Clone to Thin Devices

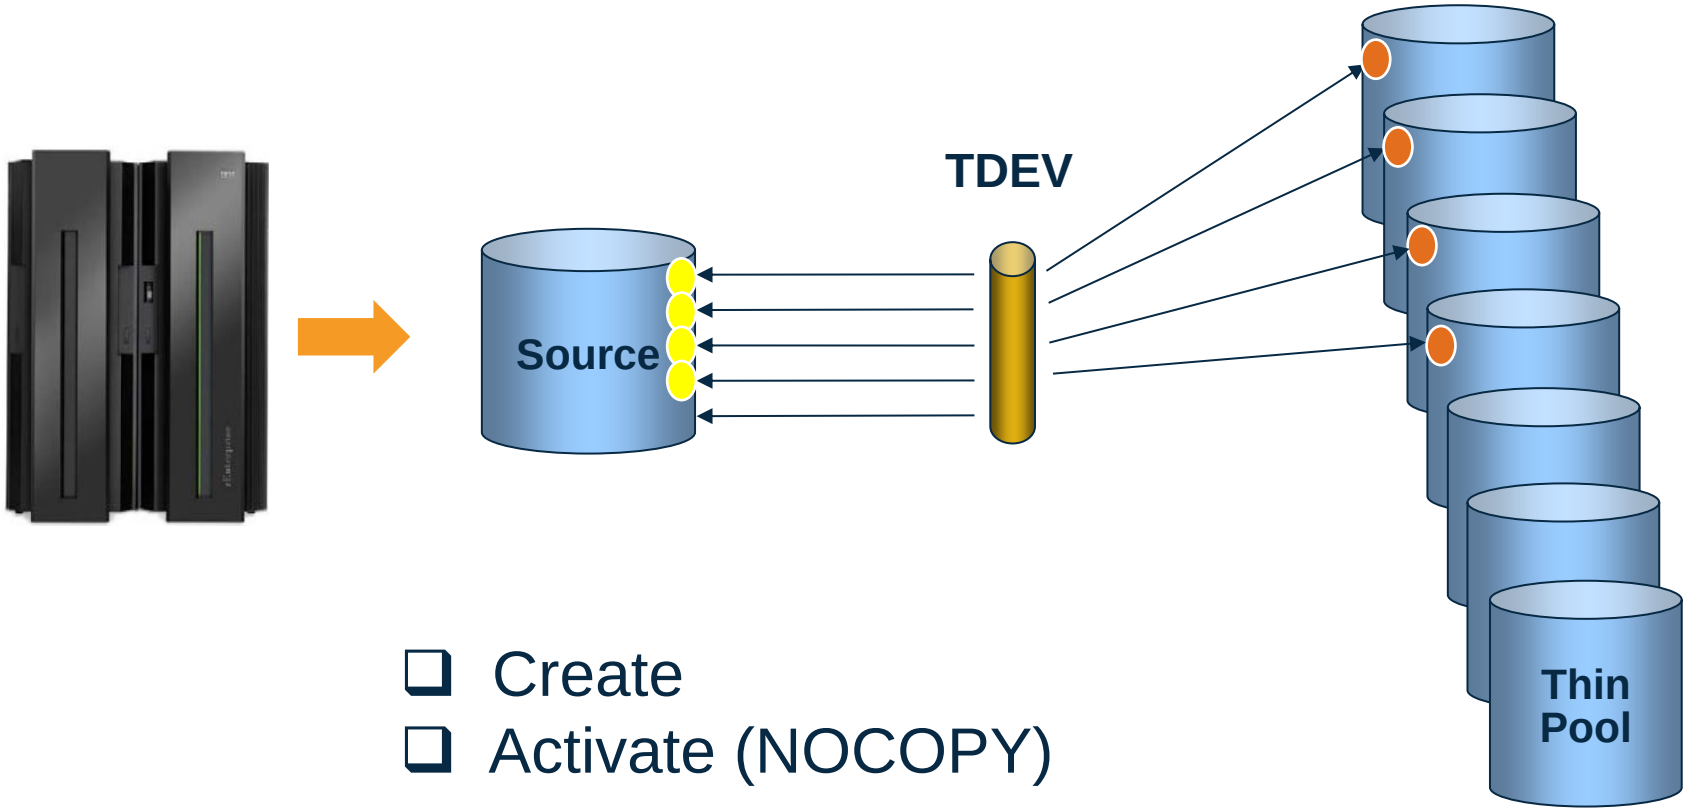

- Create
- Activate (NOCOPY)

TimeFinder/Clone and Virtual Provisioning: COPY vs NOCOPY

- TF/CLONE NOCOPY to THIN target results in space efficient copy
- Target can be in shared or dedicated VP pool (repository-like) that can be over-provisioned

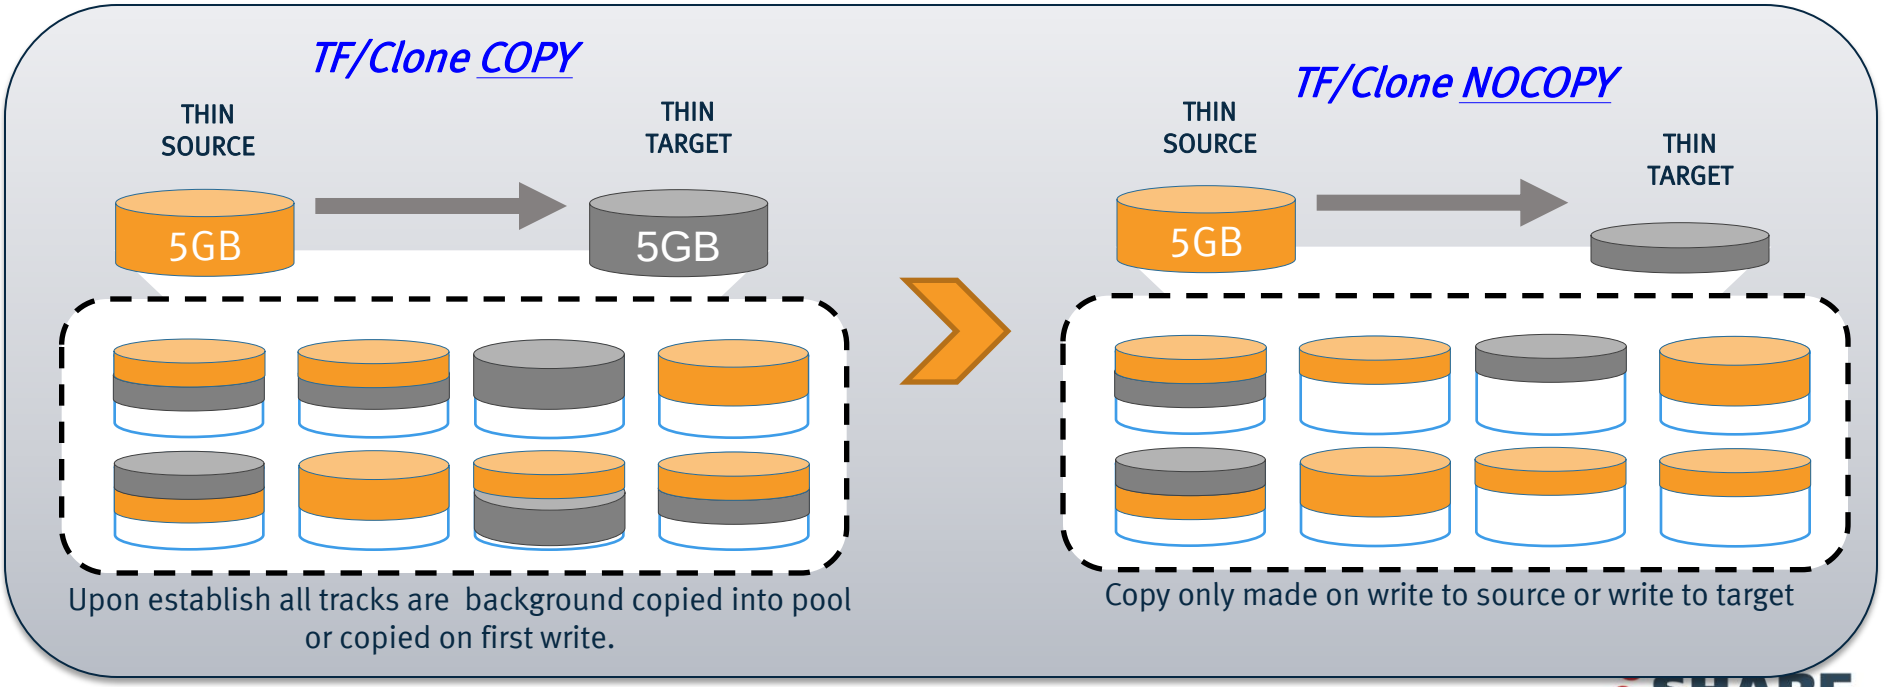

Achieving Space Efficient FlashCopy with EMC VMAX

- Use of Symmetrix Virtual Provisioning provides equivalent functionality to Space Efficient FlashCopy
- FlashCopy parameter 'SETGTOK' allows, but does not require, the target volume to be an SE volume
  - Symmetrix devices do not present as Space Efficient devices
- However, using a Symmetrix thin device as a FlashCopy target **and specifying NOCOPY**, results in a space efficient copy

FlashCopy on VMAX with Virtual Provisioning: Copy vs NoCopy

- FlashCopy NOCOPY to THIN target results in space efficient copy
- Target can be in shared or dedicated VP pool (repository-like) that can be over-provisioned

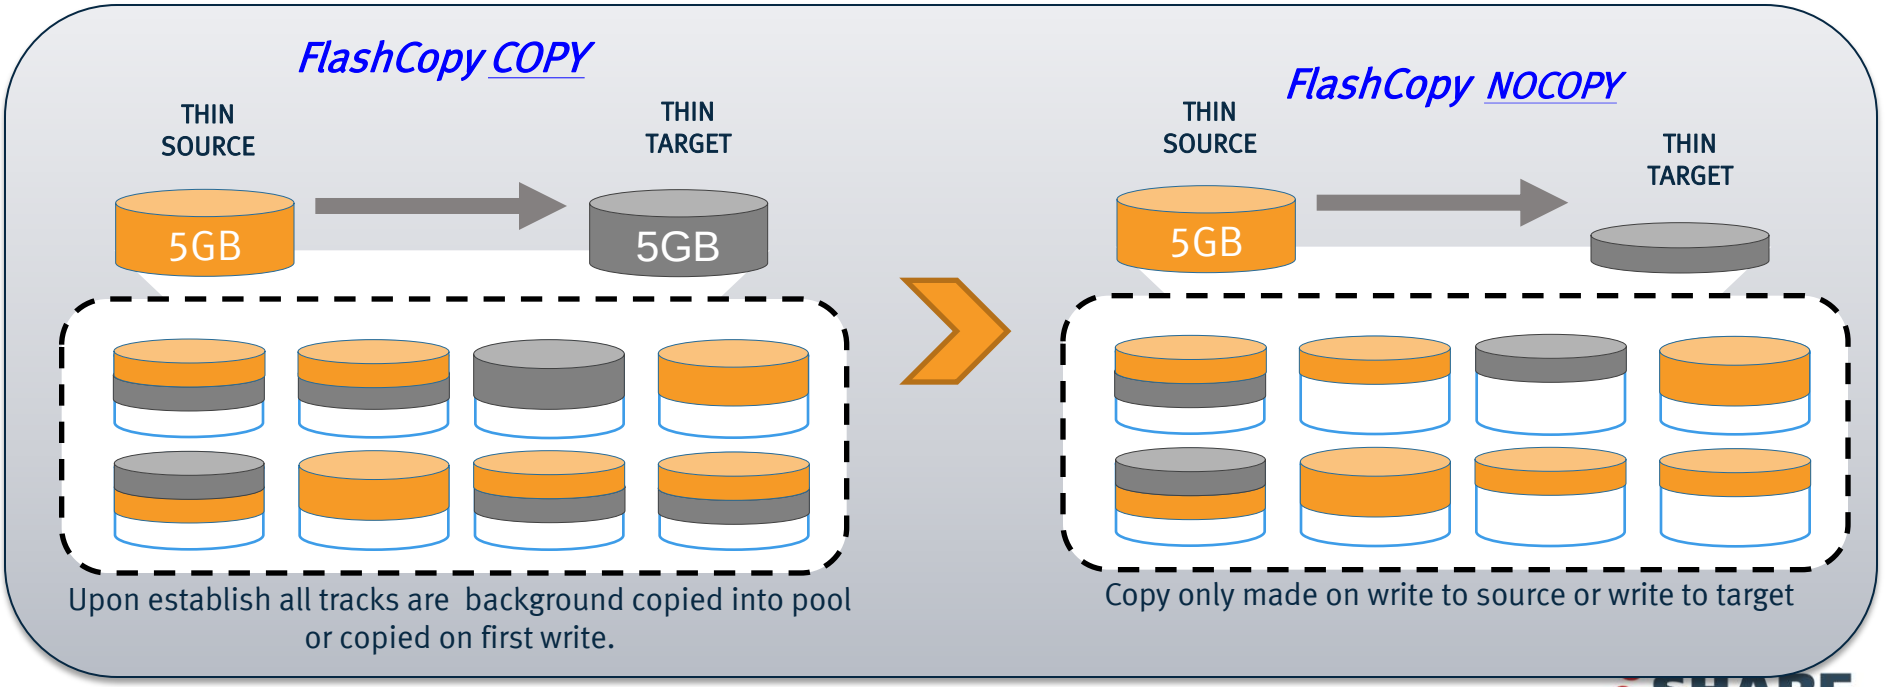

Space Efficient FlashCopy Comparison

	Space Efficient FC devices	EMC thin devices
Convert to full copy	No	Yes
Incremental possible	No	Yes
Special devices required	Yes	Thin
Repository resizable	No	Yes (thin pool)
Standard FC possible	No	Yes

SE Flashcopy 'equivalence' using VMAX thin devices: Consideration

- Devices do not present as SE FlashCopy devices
  - No reporting of SE devices in
    - TSO FC QUERY
    - ICKDSF FC QUERY RELATIONS

DB2 System Level Backup

Preparing the Configuration

- SMS managed DB2 volumes are required
- Two types of Storage Groups required
  - System Level Data Sets
    - BSDS,Logs (Active and Archive)
    - Dedicated ICF catalog for these data sets
  - Application Level Data Sets
    - DB2 Catalog, Tablespaces/Indexspaces
    - Dedicated ICF catalog for these data sets
- **ONLY** DB2 data sets on these volumes

SMS Storage Groups

- Target devices may be in HSM Copy Pools
- Target devices may be in ordinary Storage Groups
- When using the “BACKUP SYSTEM” and “RESTORE SYSTEM” DB2 commands Copy Pools are required.
- Recovery Expert allows either

DB2 Full Volume Cloning

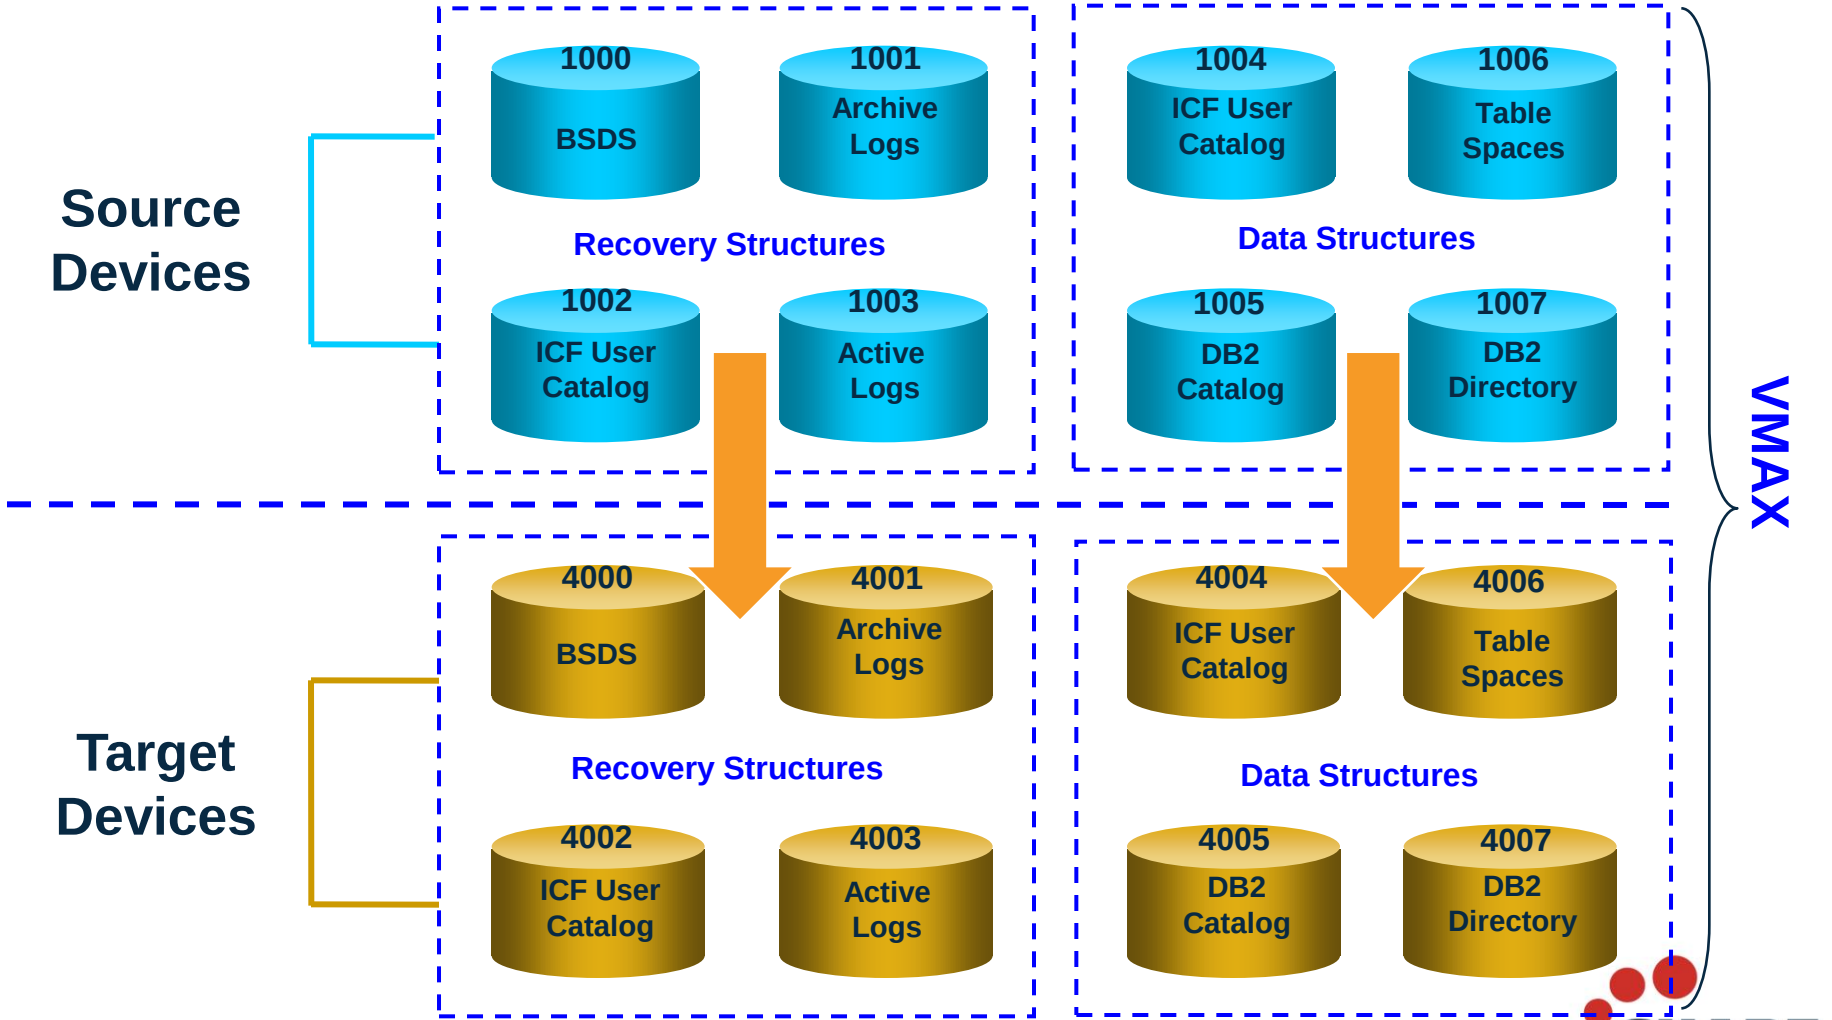

Source Devices

- May be Thick or Thin
- May be R1 or R2 (Primary or Secondary)
- Consistency across all DB2 volumes is required
  - VMAX can use Enginuity Consistency Assist feature with volumes in a Consistency Group
  - DB2 commands allow for consistency
    - Set Log Suspend before Clone/FlashCopy
    - Set Log Resume after Clone/FlashCopy

Target Devices

- May be Thick, Thin (TDEV) or Virtual Device (VDEV)
- May be R1 or R2 (Primary or Secondary)
- Must be in the same array as the source
- Must be equal number of volumes as the source
- Must be the same geometry as the source

FlashCopy Backup

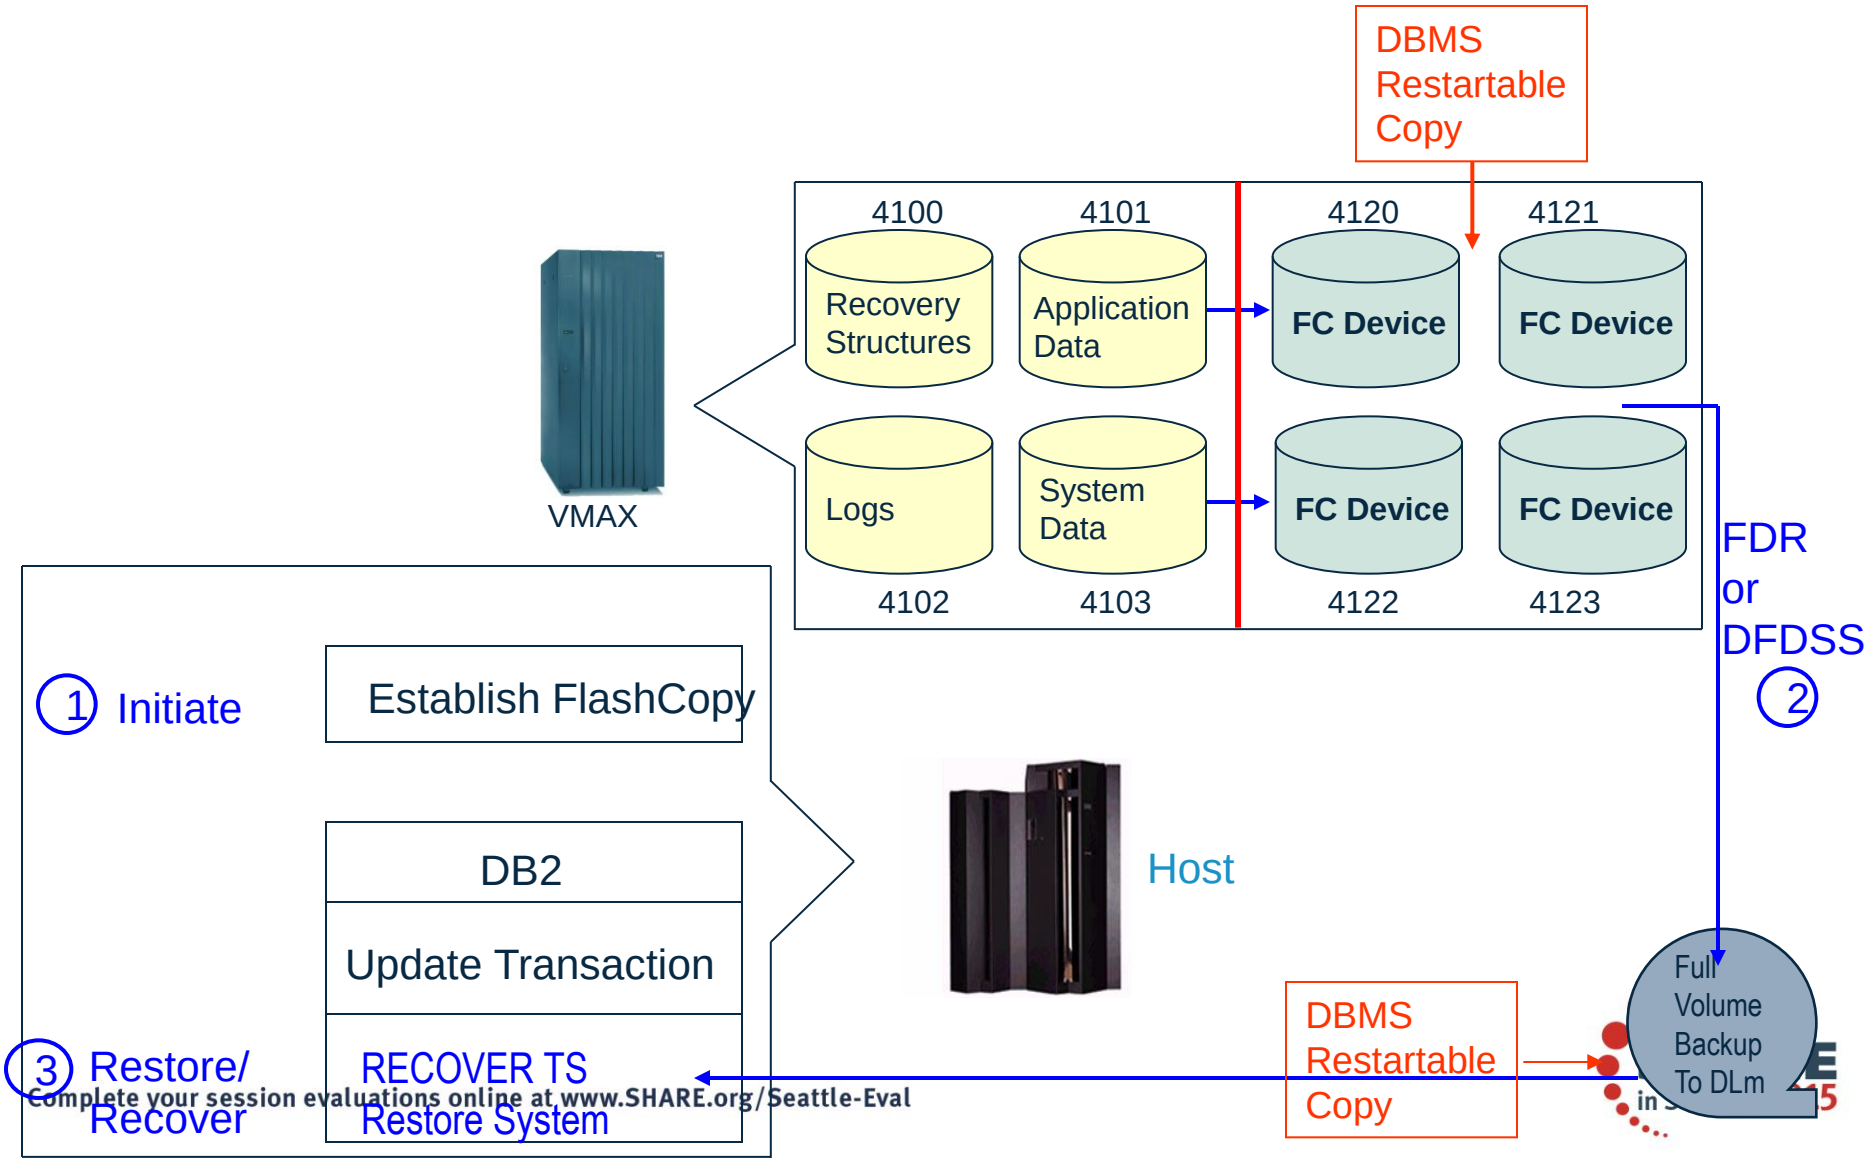

Clone Backup

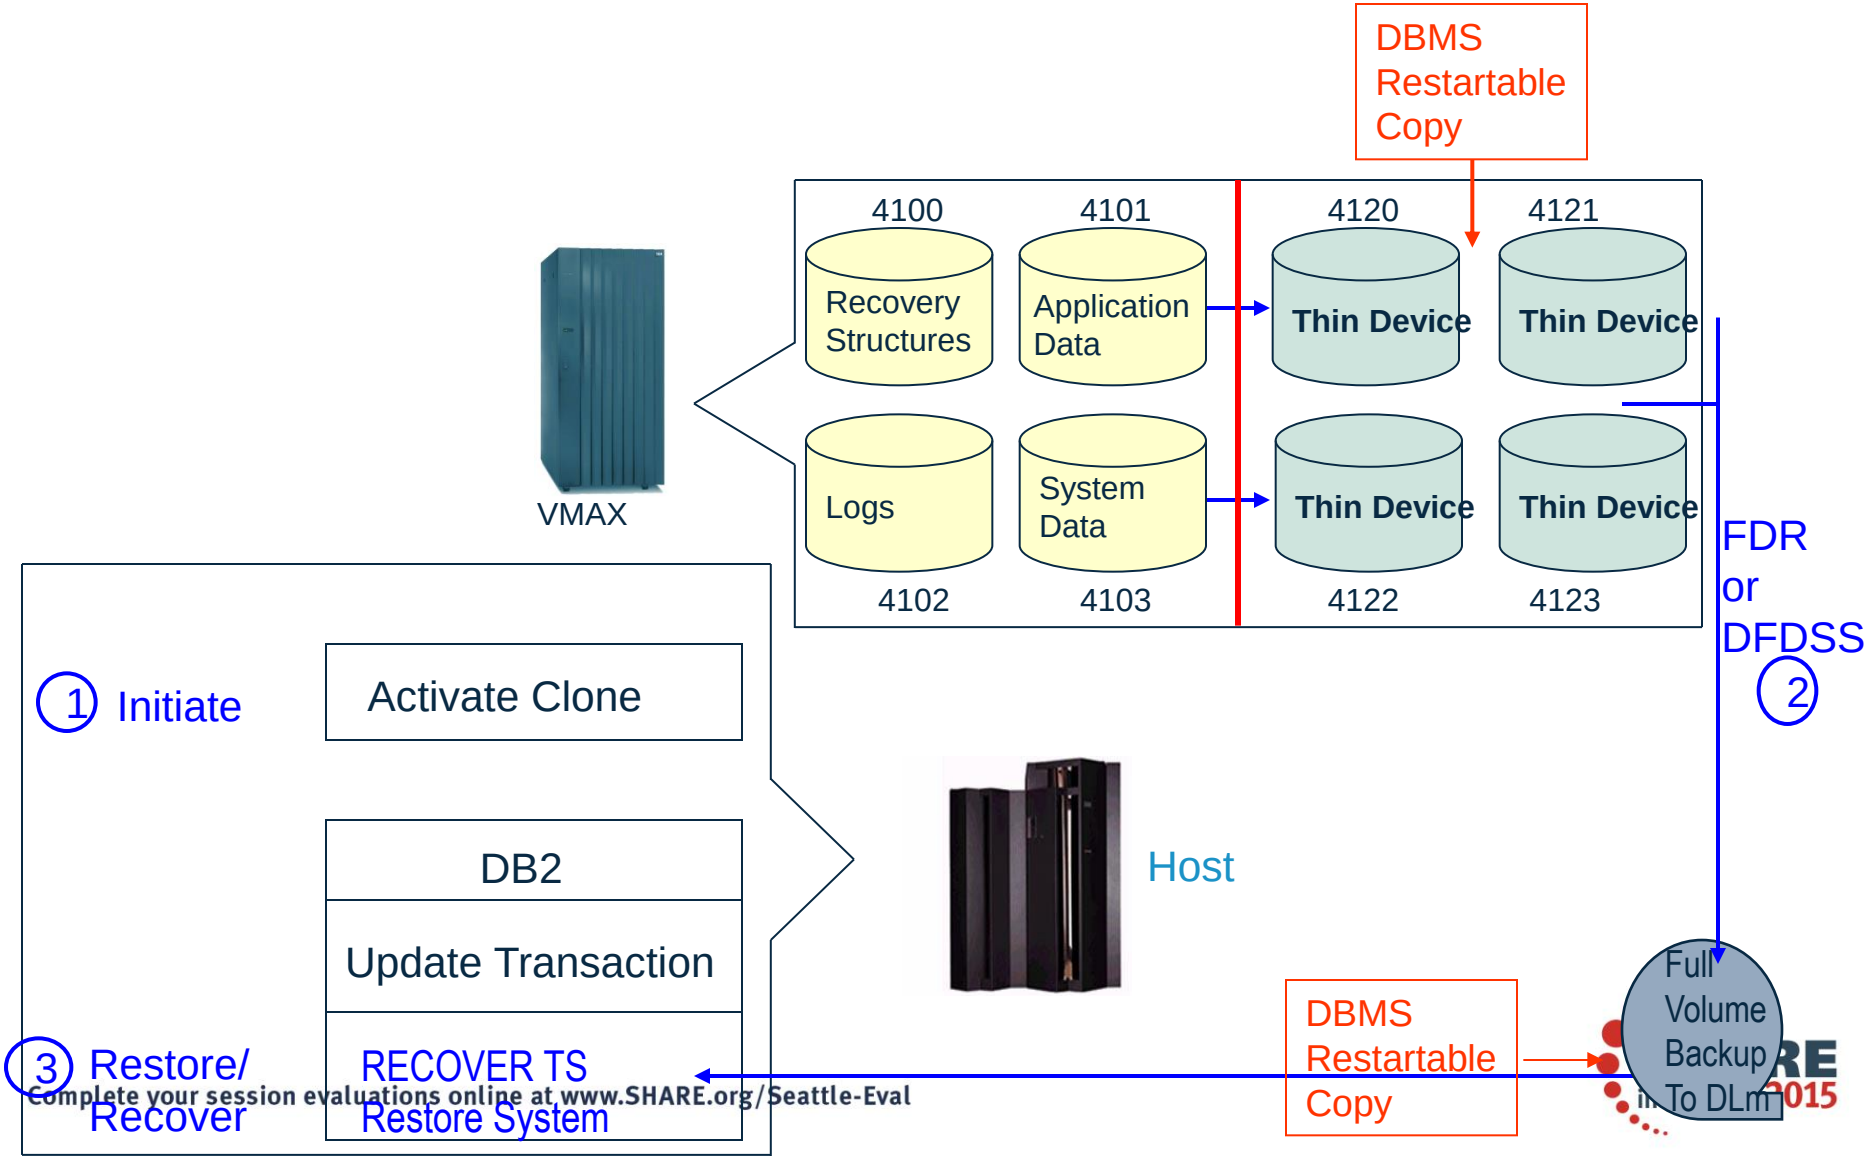

DB2 System Backup With Thin Devices

- Use thin devices in SMS backup copy pools
- Specify VERSIONS=0 on copy pool backup definition
  - Generates a NOCOPY parameter on the FlashCopy
- Changed tracks are stored in the thin pool
- Withdraw operation occurs after HSM dumps the volume

Rocket Software Products

- Rocket Backup and Recovery for DB2
- Fully integrates with both FlashCopy and TimeFinder/Clone and TimeFinder/Snap backups with DB2
- Supports thin and thick devices as backup targets
- Resold by IBM as a part of DB2 Recovery Expert

Profile Specifications for Rocket

```
MAINSTAR  V2R2 -------- Source Stogroup Selection ------- 2015/02/11  15:49:13
Option  ===>                                                Scroll ===> PAGE

Line Commands: I - Enter  D - Delete

--------------------------------------------------------------------------------
 Creator: TMOULD1    Name: DB0P                                SSID: DBS1
 Share Option: U  (Upd,View,No)    Description: PRODUCTION PROFILE

--------------- Enter Storage Groups ---------------          Row 1 of 3
Cmd  Stogroup

     DBSDSG
     DBSLSG
****************************** Bottom of Data ********************************
```

Storage Groups that are not Copy Pools

Profile Choices for Snapshot

```
+------------- Enter New Backup Profile Options --------------+
|                                                              |
|    Creator           TMOULD1                                 |
|                                                              |
|    Profile Name                                              |
|                                                              |
|    Description                                               |
|                                                              |
|    DB2 SSID          DBS1   (? for system list)              |
|                                                              |
|    Backup Method     S   (Bcv/Snap/Flash/Db2/dfsmsdss(L))    |
|                                                              |
|    Source/           M   (Auto discover/pool mapping,        |
|    Target Mapping            Stogroup discover/pool mapping, |
|                              Manual)                         |
|                                                              |
|    Update Option     U   (Update, View only, No access)      |
|                                                              |
|                                                              |
+--------------------------------------------------------------+
```

Chosen at Profile creation and can't be changed afterwards

Automatic Mapping

<---DB2 Volume--->			<-Target-->		<--------Data Types-------->					
Volser	Ucb#	Devtyp	Volser	Ucb#	Obj	OCat	ALog	ACat	RLog	RCat
DBS001	4500	3390-9	BKP025	4528	No	No	Yes	Yes	Yes	Yes
DBS002	4501	3390-9	BKP026	4529	No	No	Yes	No	Yes	No
DBS003	4502	3390-9	BKP027	452A	No	No	No	No	Yes	No
DBS004	4503	3390-9	BKP028	452B	No	No	No	No	Yes	No
DBS005	4504	3390-9	BKP013	451C	Yes	Yes	No	No	No	No
DBS006	4505	3390-9	BKP014	451D	Yes	No	No	No	No	No
DBS007	4506	3390-9	BKP015	451E	Yes	No	No	No	No	No
DBS008	4507	3390-9	BKP016	451F	Yes	No	No	No	No	No
DBS009	4508	3390-9	BKP017	4520	Yes	No	No	No	No	No
DBS010	4509	3390-9	BKP018	4521	Yes	No	No	No	No	No
DBS011	450A	3390-9	BKP019	4522	Yes	No	No	No	No	No
DBS012	450B	3390-9	BKP020	4523	Yes	No	No	No	No	No
DBS013	450C	3390-9	BKP021	4524	Yes	No	No	No	No	No
DBS014	450D	3390-9	BKP022	4525	Yes	No	No	No	No	No
DBS015	450E	3390-9	BKP023	4526	Yes	No	No	No	No	No
DBS016	450F	3390-9	BKP024	4527	Yes	No	No	No	No	No

Time Required for Flash

```
Utility Executed:..... Backup
Profile Name:..... TMOULD1.DB0P
DB2 Subsystem:..... DBS1
DB2 Version:..... 1010
Backup Type:..... Flash Copy
Backup Contains:..... Object Data and Log Data
Partial Backup:..... No
Nbr of Volumes:..... 0016
Backup RBA:..... 000001619270
Last Checkpoint RBA:..... 000001614346
HPGRBLP RBA:..... 000001616DC0
Backup Date:..... 11/17/2014
Backup Time:..... 2014-11-17-14.25.13.667532
Consistency Method:..... Flash Consistency Group
Supports Object Restore:.. Yes
I/O Suspend Time:..... 2014-11-17-14.25.13.666161
I/O Resume Time:..... 2014-11-17-14.25.13.774279
```

Backup Elapsed:..... 00.10 Seconds

To Log Suspend or Not That is the Question

- Log Suspend in DB2?
  - Disk Hardware (FlashCopy Consistency Group or TF/Clone Consistency function) provides a “Dependent Write Consistent” (aka power failure consistent) restartable image of the DB2 subsystem
  - What could go wrong?
  - Got Referential Integrity?
    - DB2 performs separate writes for Parent/Child Updates
    - I have seen the DASD Hardware hold writes from the host in between the Parent and Child updates
    - After a subsystem was restored with no logs applied the Parent and Child were inconsistent
  - If you do not perform a Log Suspend, then plan on applying all logs.

To Image Copy or Not

- With DB2 V10 an SLB supports recovery of individual objects
- Savings if Image Copies were suspended
  - If you use Virtual Tapes, the capacity requirements are minimized
  - If you use Physical Tape, you need fewer tapes
  - If you use Disk, the capacity requirements are minimized
    - If you replicate the disk pool to another site, network capacity requirements are minimized

Offload to a EMC Disk Library for mainframe (DLm)

- Define the DLm tape devices as 3590-1
- In Recovery Expert choose number of tasks based on tape devices available for use
- Choose Local and Remote Options (LP/RP)
- Choose Data Set Name options and Unit Names to direct output to the desired Volsers
- Choose Generations to determine rentention
- Choose to Delete aged generations

Offload Options

```
Enter the Offload options to associate with this Backup profile:

 Local Primary             ==> N   (Yes/No/Update)
 Local Backup              ==> N   (Yes/No/Update)
 Recovery Site Primary     ==> N   (Yes/No/Update)
 Recovery Site Backup      ==> N   (Yes/No/Update)

 Offload Generations       ==> 01  (1 - 99)
 Delete Aged Backup files  ==> Y   (Yes/No)
 Compress Data             ==> N   (Yes/No)
 Data Mover                ==> D   (Dfsmsdss, Fdr, or fdrInstant)
 Encrypt Data              ==> N   (Yes/No/Update)
 Number of Tasks           ==> 02  (1 - 99)
```

Offload Example

- The following example used these options
  - Number of Tasks = 2
    - One for Local Backup Copy
    - One for Remote Backup Copy
  - Controls the use of Volsers
    - Higher the number uses more Volsers takes less time
    - Lower the number uses less Volsers takes more time

Offload to DLm Example

```
1                      J E S 2  J O B  L O G  --  S Y S T E M  T S T 1  --  N
0
 12.27.13 JOB02654 ---- WEDNESDAY, 09 OCT 2013 ----
 12.27.13 JOB02654  IRR010I  USERID *******  IS ASSIGNED TO THIS JOB.
 12.27.14 JOB02654  ICH70001I *******  LAST ACCESS AT 12:17:57 ON WEDNESDAY, OCT
 12.27.14 JOB02654  $HASP373 DBRBACK  STARTED - INIT A    - CLASS A - SYS TST2
.....
 12.28.32 JOB02654 *IEC501A M A400,PRIVAT,SL,COMP,DBRBACK,DBRJOFFL,DBRTDS.LP.TES
 12.28.32 JOB02654  IEC705I TAPE ON A400,T00059,SL,COMP,DBRBACK,DBRJOFFL,DBRTDS.
 12.28.32 JOB02654 *IEC501A M A440,PRIVAT,SL,COMP,DBRBACK,DBRJOFFL,DBRTDS.RP.TES
 12.28.32 JOB02654  IEC705I TAPE ON A440,B00032,SL,COMP,DBRBACK,DBRJOFFL,DBRTDS.
.....
 13.00.01 JOB02654  -DBRREBU1             00    638    153       .01       .00
 13.00.01 JOB02654  IEF404I DBRBACK - ENDED - TIME=13.00.01
 13.00.01 JOB02654  -DBRBACK  ENDED.  NAME-******************    TOTAL TCB CPU TI
 13.00.01 JOB02654  $HASP395 DBRBACK  ENDED
0------ JES2 JOB STATISTICS ------
-  09 OCT 2013 JOB EXECUTION DATE
-          134 CARDS READ
-        1,379 SYSOUT PRINT RECORDS
-            0 SYSOUT PUNCH RECORDS
-           89 SYSOUT SPOOL KBYTES
-   32.79 MINUTES EXECUTION TIME
.....
```

Offload to DLm Example

```
1                    Database Backup and Recovery for DB2
                         Volume Offload Detail Report
.....
    <---DB2--->                                                    <----Tape---->
    Volser Ucb# Type Offloaded to Filename                         FileSeq Volser
    ------ ---- ---- ---------------------------------------------- ------- ------
    DBTV01 5B16 LP   DBRTDS.LP.TEST.D2013282.T122832.DBTV01         001     T00059
                RP   DBRTDS.RP.TEST.D2013282.T122834.DBTV01         001     B00032
.....
1PAGE 0001     5695-DF175  DFSMSDSS V1R13.0 DATA SET SERVICES     2013.282 12:28
-  DUMP FULL -
        INDDNAME( DSDBTV01 ) -
        OUTDDNAME( SYS00023 , SYS00024 ) -
        OPTIMIZE( 4 ) -
        ALLEXCP ALLDATA(*)
 ADR101I (R/I)-RI01 (01), TASKID 001 HAS BEEN ASSIGNED TO COMMAND 'DUMP '
 ADR109I (R/I)-RI01 (01), 2013.282 12:28:32 INITIAL SCAN OF USER CONTROL STATEME
 ADR050I (001)-PRIME(01), DFSMSDSS INVOKED VIA APPLICATION INTERFACE
 ADR016I (001)-PRIME(01), RACF LOGGING OPTION IN EFFECT FOR THIS TASK
0ADR006I (001)-STEND(01), 2013.282 12:28:32 EXECUTION BEGINS
0ADR006I (001)-STEND(02), 2013.282 12:28:54 EXECUTION ENDS
0ADR013I (001)-CLTSK(01), 2013.282 12:28:54 TASK COMPLETED WITH RETURN CODE 0000
```

Native DB2 Utilities

- DB2 searches the EDT for any UNIT specification
- DLm UNIT names should be in the EDT or these utilities will fail

Summary

- DB2 System Level Backup is completely supported on EMC Hardware using either FlashCopy or TimeFinder
- VMAX has options for reducing capacity requirements for target devices using both FlashCopy and TimeFinder
- Rocket Software utilities exploit both Flashcopy and TimeFinder to produce DB2 System Level Backups
- EMC DLm provides for all requirements when offloading the target devices to tape volsers