Back to the Future: Creating Consistent Copies at Isracard

Mike Shorkend
Isracard Group

1:30 PM on Thursday, February 7, 2013
Session Number 12504

http://www.linkedin.com/pub/mike-shorkend/0/660/3a7
mshorkend@isracard.co.il
mike@shorkend.com

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

AIX*
DB2*
HiperSockets
IBM*
IBM logo*
IMS
CICS
System z
System z9
System z10
System z114
Tivoli
WebSphere*
z/OS*
z/VM*
zSeries*

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Linux is a registered trademark of Linus Torvalds 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.
Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.
Control-M and Control-O are trademark of BMC

* All other products may be trademarks or registered trademarks of their respective companies.

Agenda

- **Introduction**
- Level 1: Synchronous Replication
- Level 2: Logical Copies
- Level 3: DRP testing
- Level 4: Third Site Copy
- Questions

Over 100,000 merchants

Over 50 million business transaction per month

Monthly turnover of 9 billion NIS

3.5 Million Cards
47 % market share

2 million card holders

About me

- *Manager, Central Infrastructures at Isracard*
- *Responsible for z/OS, z/VM, Linux(z and x), enterprise storage*
- *2 teams – Mainframe OS, Linux and Storage*
- *My background is z/OS system programming, tuning and capacity planning*
- *6 years at Isracard*

The Challenges and Triggers

- Normal threats like floods, earthquake, fire
- Geo-political specific threats like terror and cyber attacks
- In November 2008 a large Israeli financial institute had a 60 hour outage due to a logical error that was replicated to the DR site.
- Compliance
- Financial Constraints

Isracard Infrastructure

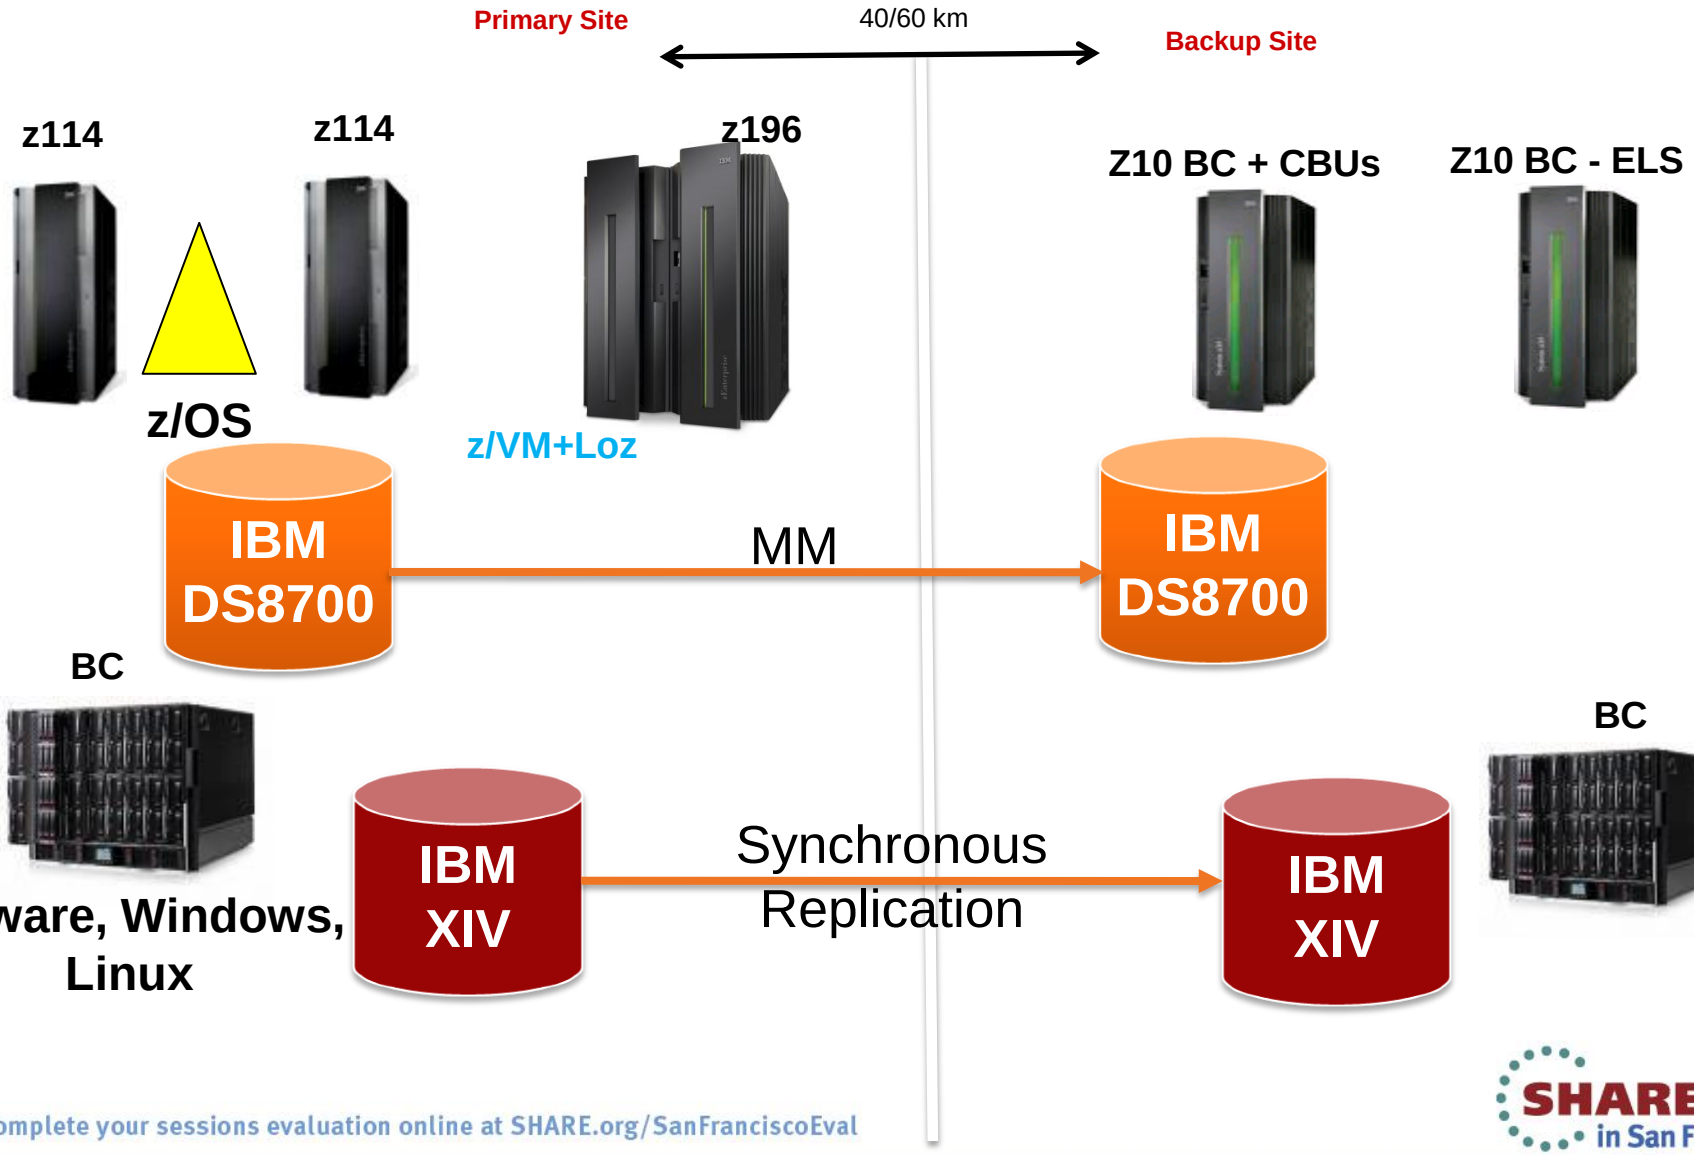

Isracard Mainframe Infrastructure

z/OS 1.13
DB2 10
CICS/TS 4.1
IMS(DBCTL) 11.1

z114

z114

z/VM 5.4
RHEL 5.6
WAS
Oracle
WMB

z196- ELS

IBM DS8700

IBM XIV

IBM VT7720

Agenda

- Introduction
- **Level 1: Synchronous Replication**
- Level 2: Logical Copies
- Level 3: DRP testing
- Level 4: Third Site Copy
- Questions

Synchronous replication

- All production DASD are replicated to the DR site using Metro Mirroring(aka PPRC).
- Managed by Tivoli Productivity Center for Replication
- Approximately 16TB (9TB allocated) on 1800 volumes
- If one pair fails, I/O is frozen and all pairs are suspended creating a write dependent consistent mirror at the DR site(deals with the ‘rolling disaster’ scenario)
- I/O is released after a suspend(the other option is a sysplex wide outage). Availability preferred over mirror update.
- Monitored by hourly jobs

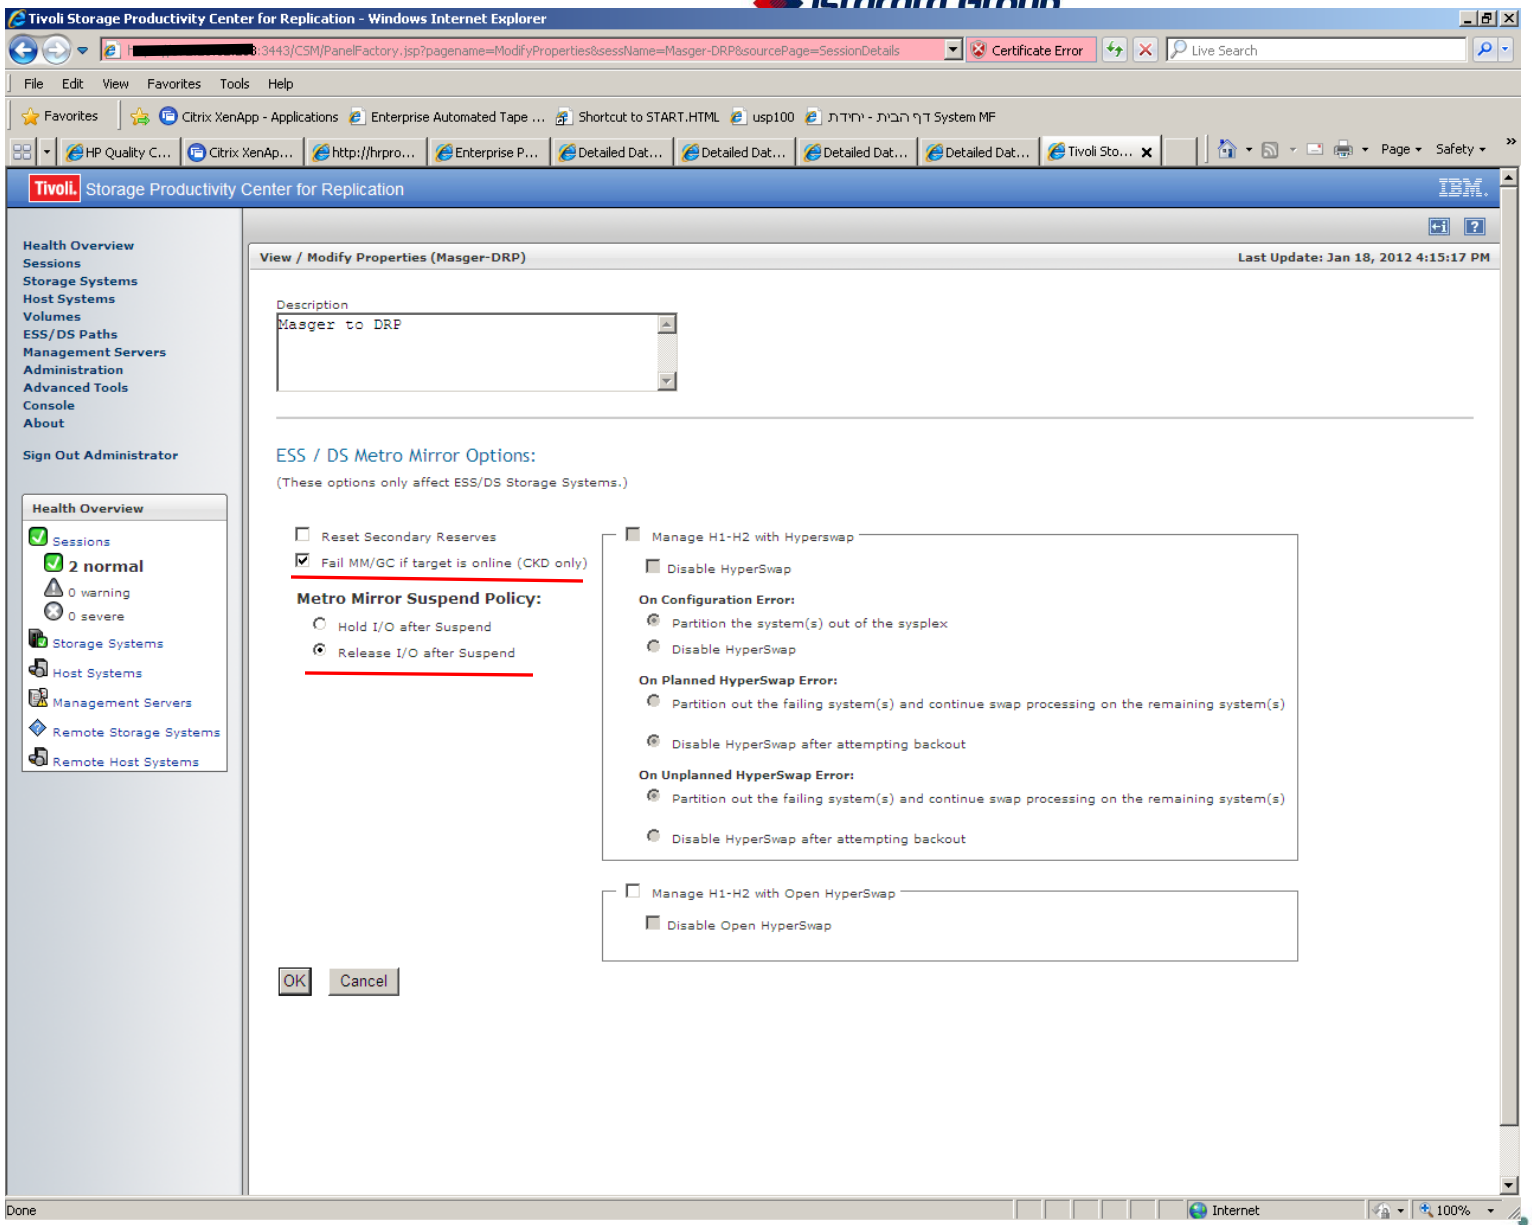

Agenda

- Introduction
- Level 1: Synchronous Replication
- **Level 2: Logical Copies**
- Level 3: DRP testing
- Level 4: Third Site Copy
- Questions

Logical Error Challenges

- If you have a software , hardware or application error that corrupts your data it gets replicated synchronously to your mirror
- Backups can help, but how do you get a consistent production copy?
- FLASH COPY is good but costly
- How do you check that your copy images are valid?

Our solution

- We take a space efficient flash copy of our production data every business day
- Two copies are kept: todays and yesterdays
- A third copy can be taken at any time(more on that later)
- After the copy is created, it is IPLed and data integrity is verified
- Only after it is verified, the previous days copy can be removed
- All automatic, using BMC/Control-M and Control-O,DSCLI and BCPii

Building Blocks(1/3):Space Efficient Flash Copy

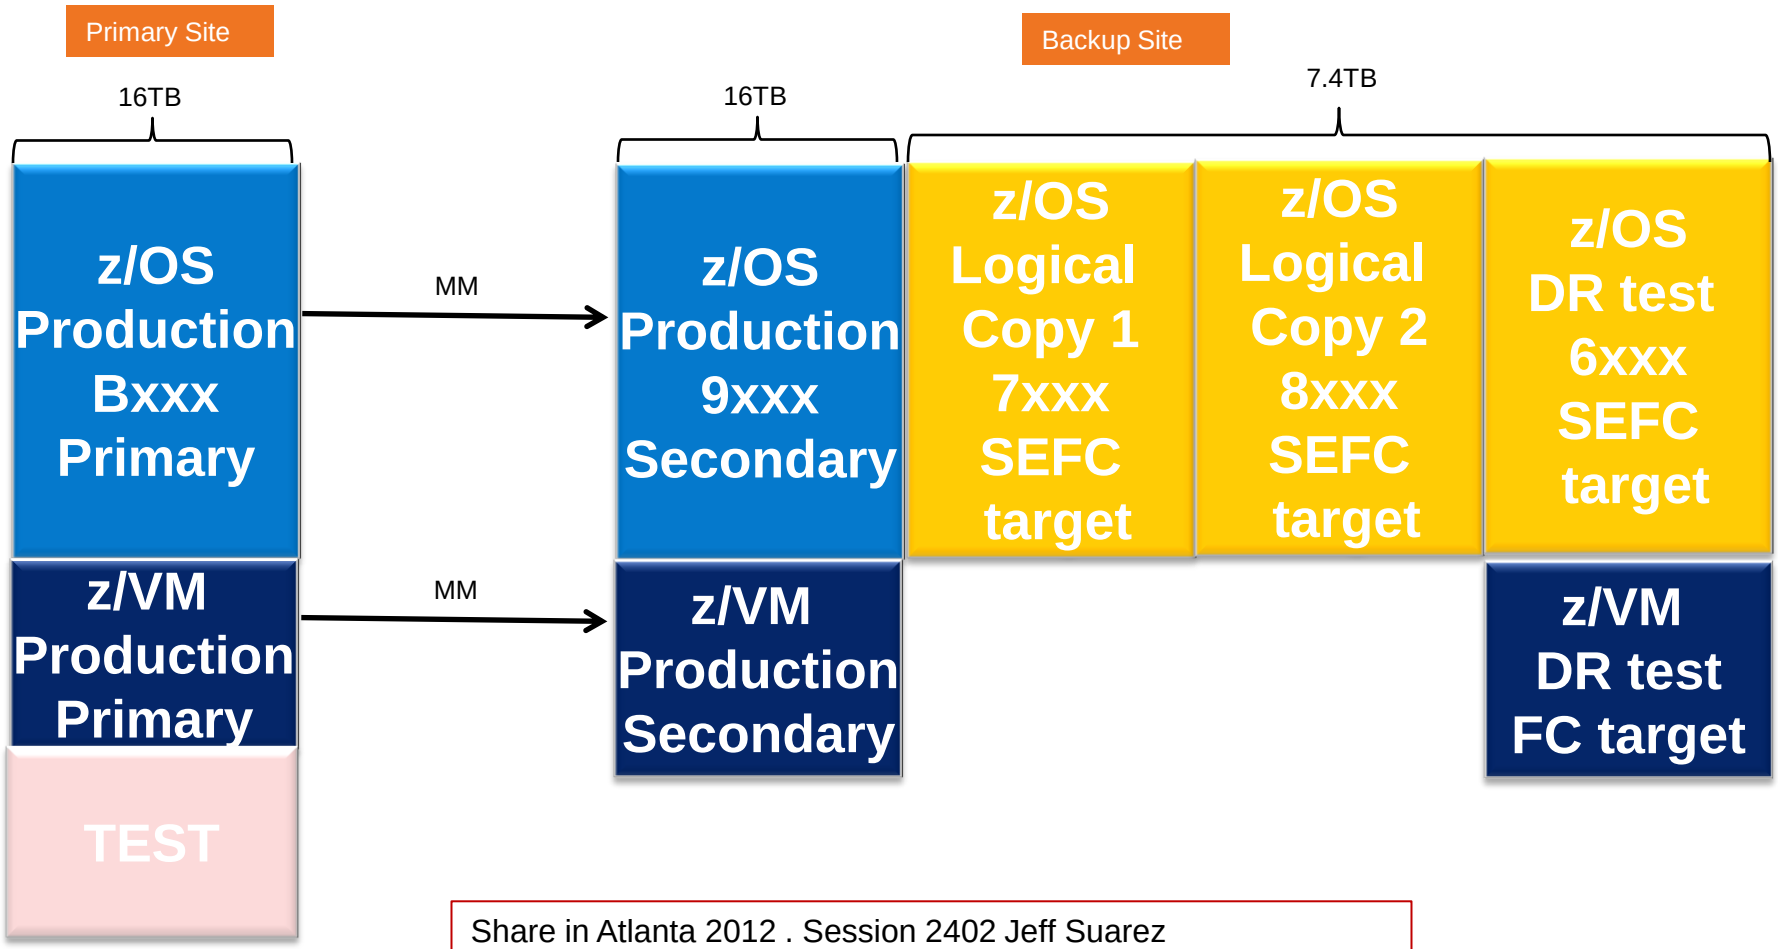

Share in Atlanta 2012 . Session 2402 Jeff Suarez
Share in Austin 2009. Session 3080. Linda Gundy

Building Blocks(2/3):BCPii

What is BCPii?

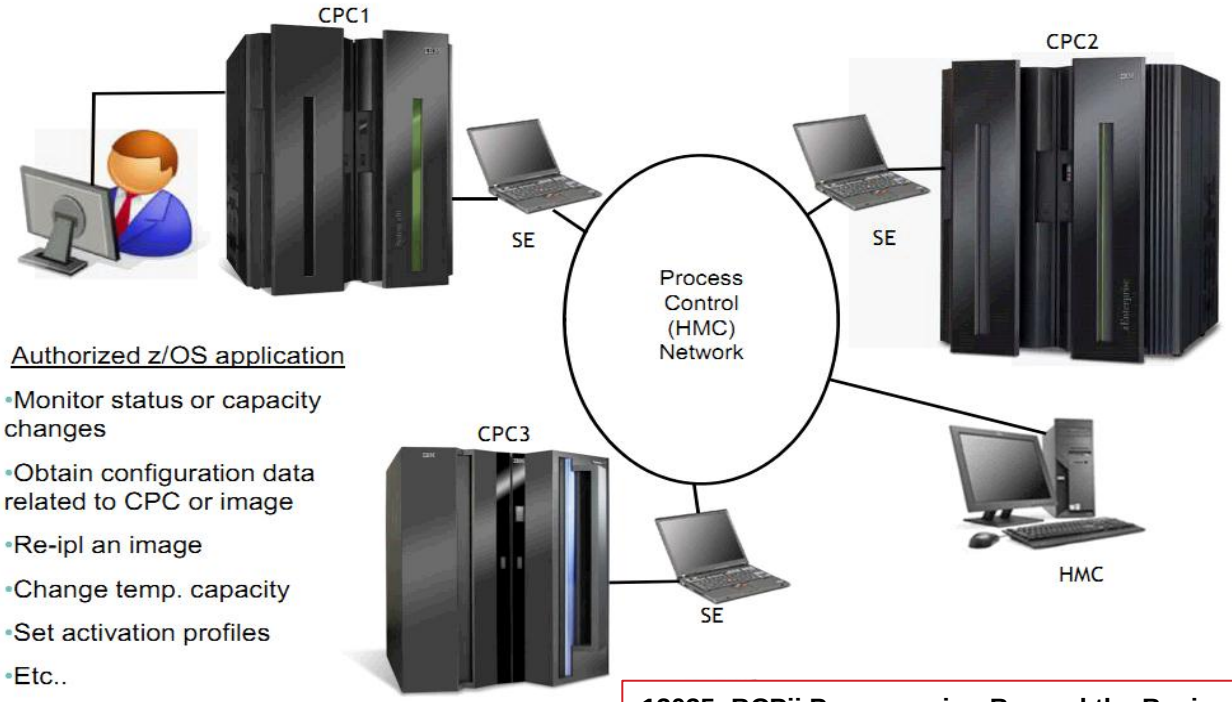

Authorized z/OS application

- Monitor status or capacity changes
- Obtain configuration data related to CPC or image
- Re-ipl an image
- Change temp. capacity
- Set activation profiles
- Etc..

13035: BCPii Programming Beyond the Basics for the z/OS System Programmer
Steve Warren

Building Blocks(3/3)

- Control-M – z/OS and distributed scheduling
- Control-O – z/OS Automation
- DSCLI – command line interface to the SSPC(System Storage Productivity Center)

The Big Picture

- When the job(on z/OS) indicating end of day processing has finished, a condition is raised by Control-M.
- This condition causes a script to be run on the SSPC that creates the flash
- When the script ends, Control-M raises a condition that causes a job on z/OS to run that activates the coupling facility and the z/OS image at the DR site.
- Another job monitors the IPL message log

Creating the New Flash

```
cd \
cd "C:\Program Files\IBM\dscli"
echo ****************************************
date /t
time /t
dscli -cfg "C:\Program Files\IBM\dscli\profile\DS8700_DRP.profile" -hmc1 xxx.xx.52.211 -script c:\ControlM\rmflashmankal1.script >c:\ControlM\rmflashmankal1.log
dscli -cfg "C:\Program Files\IBM\dscli\profile\DS8700_DRP.profile" -hmc1 xxx.xx.52.211 -script c:\ControlM\flashmankal1.script >c:\ControlM\flashmankal1.log
dscli -cfg "C:\Program Files\IBM\dscli\profile\DS8700_DRP.profile" -hmc1 xxx.xx.52.211 -script c:\ControlM\unfreezeflashmankal1.script >c:\ControlM\unfreezeflashmankal1.log
exit
```

flashmankal1.script

```
mkflash -freeze -tgtse -nocp -seqnum 01 9000-9018:7000-7018 9100-911F:7100-711F 9400-942F:7400-742F 9500-
952F:7500-752F 9600-963B:7600-763B 9700-973B:7700-773B 9800-98C7:7800-78C7 9900-99C7:7900-79C7 9A00-9AC7:7A00
-7AC7 9B00-9BC7:7B00-7BC7 9C00-9CC7:7C00-7CC7 9D00-9DC7:7D00-7DC7 9E00-9EB4:7E00-7EB4 9F00-9FEF:7F00-7FEF
unfreezeflash 90 91 94 95 96 97 98 99 9A 9B 9C 9D 9E 9F
lsflash -l 9000-901F
lsflash -l 9100-911F
lsflash -l 9400-942F
lsflash -l 9500-952F
lsflash -l 9600-963B
lsflash -l 9700-973B
lsflash -l 9800-98C7
lsflash -l 9900-99C7
lsflash -l 9A00-9AC7
lsflash -l 9B00-9BC7
lsflash -l 9C00-9CC7
lsflash -l 9D00-9DC7
lsflash -l 9E00-9EB4
lsflash -l 9F00-9FEF
```

For Consistency

Target is Space efficient

Copy on Write

Automated IPL

- Submit job that activates the coupling facility
- Submit job that listens on console traffic (of the IPLing image)
- Submit job that activates the z/OS image(load on activation set)
- Respond to WTORs using the listener job using CONTROL/O
- ControlO/Cosmos takes over the IPL process when it can
- When the system is up , run a CICS transaction (using the MODIFY command) to verify data integrity

Automated IPL

```
_SDSF OUTPUT DISPLAY S16LISTN JOB06908  DSID     2 LINE  INVALID COMMAND
 COMMAND INPUT ===>                                          SCROLL ===> CSR
******************************** TOP OF DATA ***********************************
                    J E S 2  J O B  L O G  --  S Y S T E M  S Y S E  --  N O D E

04.14.11 JOB06908 ---- SUNDAY,     12 FEB 2012 ----
04.14.11 JOB06908  IRR010I  USERID DCONOP   IS ASSIGNED TO THIS JOB.
04.14.11 JOB06908  ICH70001I DCONOP   LAST ACCESS AT 04:14:11 ON SUNDAY, FEBRUAR
04.14.11 JOB06908  $HASP373 S16LISTN STARTED - INIT IT   - CLASS Z - SYS SYSE
04.14.11 JOB06908  IEF403I S16LISTN - STARTED - TIME=04.14.11
04.14.11 JOB06908  BCPIILSN Starting
04.14.11 JOB06908  BCPIILSN Opening Log file
04.14.11 JOB06908  BCPIILSN Preparing
04.14.11 JOB06908  BCPIIXEQ Driven
04.14.12 JOB06908  BCPIIXEQ Main
04.14.12 JOB06908  BCPIIXEQ Prepare
04.14.12 JOB06908  BCPIILSN Conecting to target CPC
04.14.12 JOB06908  BCPIIXEQ Driven
04.14.12 JOB06908  BCPIIXEQ Main
04.14.12 JOB06908  BCPIIXEQ Connect
04.14.12 JOB06908  BCPIILSN Conecting to target LPAR
04.14.12 JOB06908  BCPIIXEQ Driven
04.14.12 JOB06908  BCPIIXEQ Main
04.14.12 JOB06908  BCPIIXEQ Connect
04.14.12 JOB06908  BCPIIXEQ Driven
04.14.12 JOB06908  BCPIIXEQ Main
04.14.12 JOB06908  BCPIIXEQ Event
04.14.12 JOB06908  BCPIILSN to ask when to stop ...
04.14.12 JOB06908 *81 BCPIILSN Deactive BCPii Listener? Y to deactivate
04.19.33 JOB06908 *MANKAL1 - IXC420D - I
04.29.43 JOB06908 *MANKAL1 - COSMOS - 5
04.30.11 JOB06908  MANKAL - *MASTER* IS UP
04.30.11 JOB06908  MANKAL - PCAUTH__ IS UP
04.30.11 JOB06908  MANKAL - RASP____ IS UP
04.30.12 JOB06908  MANKAL - TRACE___ IS UP
04.30.12 JOB06908  MANKAL - DUMPSRV_ IS UP
04.30.12 JOB06908  MANKAL - XCFAS___ IS UP
04.30.13 JOB06908  MANKAL - GRS_____ IS UP
04.30.13 JOB06908  MANKAL - SMSPDSE_ IS UP
04.30.13 JOB06908  MANKAL - CONSOLE_ IS UP
```

Automated IPL

```
SDSF OUTPUT DISPLAY SI6LISTN JOB06908  DSID   102 LINE 0        COLUMNS 02- 81
COMMAND INPUT ===>                                            SCROLL ===> CSR
******************************** TOP OF DATA ********************************
2012043 02.18.34.24     IEA371I SYS0.IPLPARM ON DEVICE 7BB0 SELECTED FOR IPL PARA
2012043 02.18.34.27     IEA246I LOAD   ID M1 SELECTED
2012043 02.18.34.30     IEA246I NUCLST ID 00 SELECTED
2012043 02.18.34.33     IEA519I IODF DSN = IODF.IODFD3
2012043 02.18.34.35     IEA520I CONFIGURATION ID = SYSIM1   . IODF DEVICE NUMBER =
2012043 02.18.34.38     IEA091I NUCLEUS 1 SELECTED
2012043 02.18.44.35   IEA370I MASTER CATALOG SELECTED IS CATALOG.MASTER.SYSI
2012043 02.18.44.55   IEA009I SYMBOLIC DEFINITIONS WILL BE READ FROM:
2001177 02.18.44.58           IEASYM00
2001177 02.18.44.61           IEASYM1I
2001177 02.18.44.63           IEASYMSI
2012043 02.18.44.82 *IEA247I USING IEASYSGB FOR z/OS 01.11.00 HBB7760
2012043 02.18.44.92   IEA007I STATIC SYSTEM SYMBOL VALUES
2001181 02.18.44.94           &SYSALVL.   = "2"
2001181 02.18.44.97           &SYSCLONE.  = "1I"
2001182 02.18.45.00           &SYSNAME.   = "SYSI"
2001182 02.18.45.02           &SYSPLEX.   = "PLX1"
2001182 02.18.45.05           &SYSR1.     = "NSR101"
2001183 02.18.45.08           &BCMVER.    = "V6R3M0"
2001183 02.18.45.11           &BHPFTPIP.  = "SIV"
2001183 02.18.45.14           &BHPOSAAD.  = "032"
2001184 02.18.45.17           &BMCCVER.   = "V2R2M25"
2001184 02.18.45.19           &BMCDVER.   = "V10R1MK"
2001184 02.18.45.22           &BMCEVER.   = "V2R2M26"
                .
                .
                .
2012043 04.19.30.89   IXC404I SYSTEM(S) ACTIVE OR IPLING: SYSI      SYSE      SYSD
2012043 04.19.30.92   IXC419I SYSTEM(S) NOT SYNCHRONIZED: SYSI      SYSE      SYSD
2012043 04.19.30.95 *    IXC420D REPLY I TO INITIALIZE SYSPLEX PLX1, OR R TO REIN
2012043 04.19.34.39   IEE600I REPLY TO 00 IS;I
```

Automated IPL

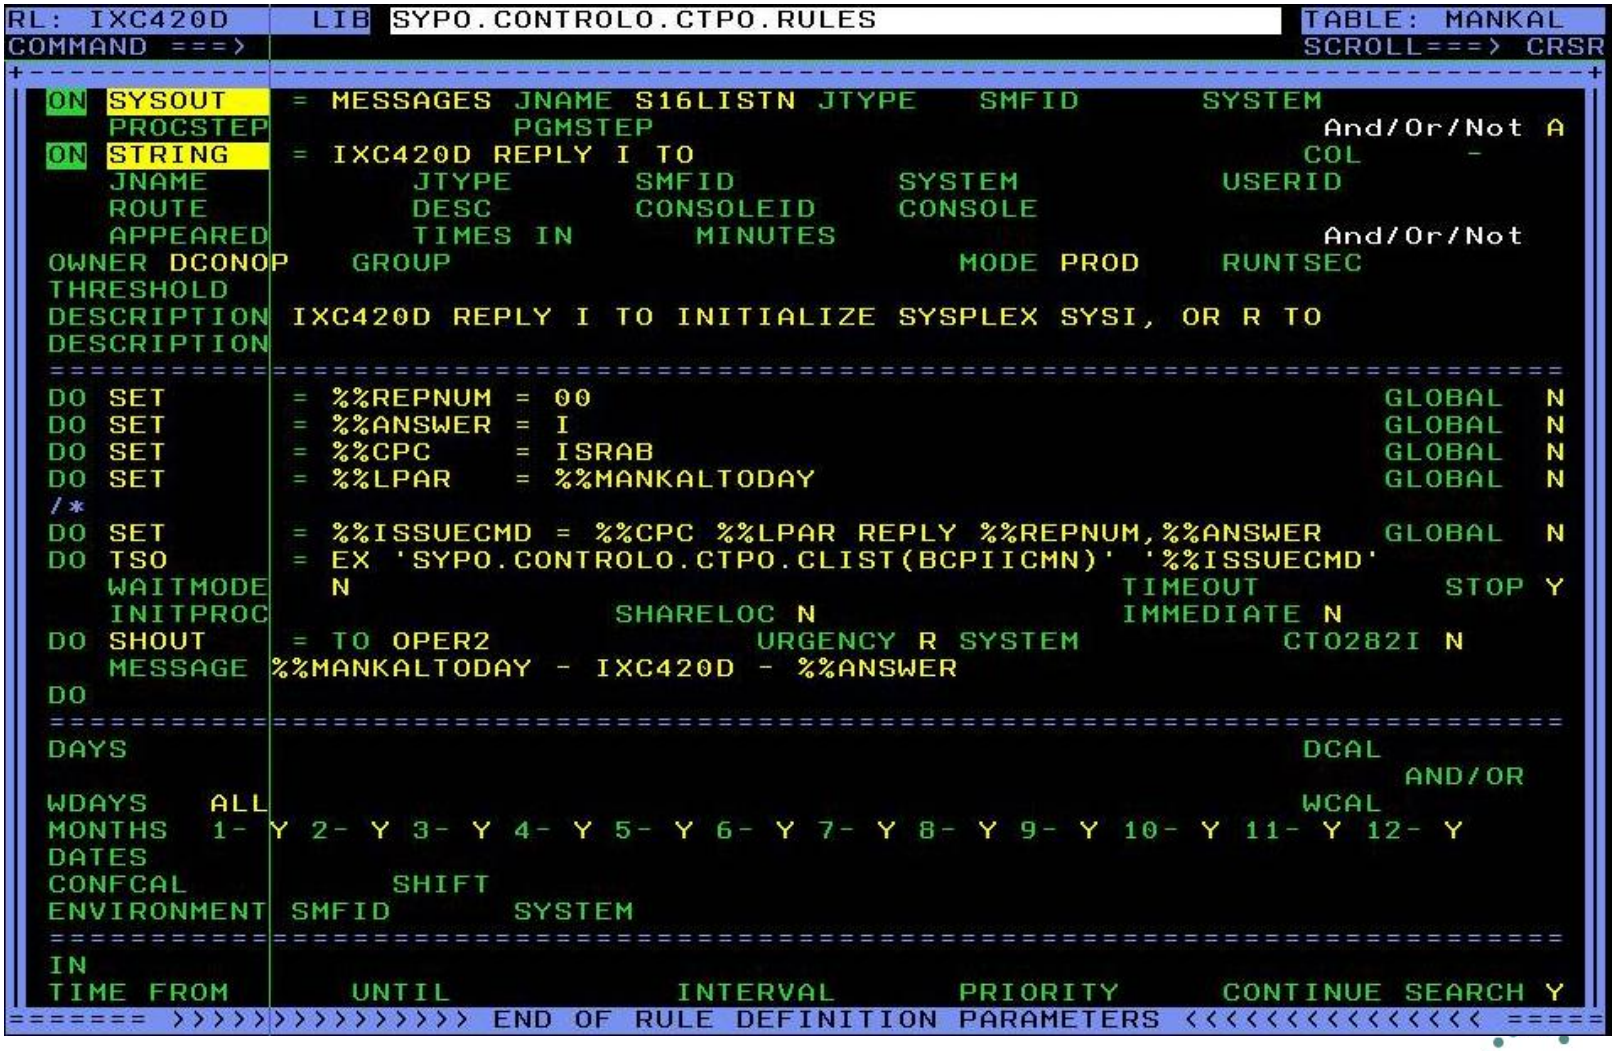

Automated IPL

```
----------------------- CONTROL-O COSMOS OBJECT STATUS ------< D >--------(OC)
COMMAND ===>                                                 SCROLL===> CRSR
O OBJECT   CURRENT  DESIRED  CLASS     MODE     STATUS
  CICSPIN1 DOWN     DOWN     CICS      FREE     STEADY DOWN
  CICSPTR1 DOWN     UP       CICS      FREE     CD=DU PR=IMSPDBC CICSPIS1
  CICSPIS1 DOWN     DOWN     CICS      FREE     STEADY DOWN
  CICSPIS2 DOWN     DOWN     CICS      FREE     STEADY DOWN
  CICSPHR1 DOWN     DOWN     CICS      FREE     STEADY DOWN
  CICSPBT1 DOWN     DOWN     CICS      FREE     STEADY DOWN
  CICSPMQ1 DOWN     DOWN     CICS      FREE     STEADY DOWN
  CICSPPXI DOWN     DOWN     STANDARD  FREE     STEADY DOWN
====== >>>>>>>>>>>>>>>     NO MORE ENTRIES IN THE LIST     <<<<<<<<<<<<<<< ======
```

```
06.10.09 JOB06908 *MANKAL - ********************************************
06.10.09 JOB06908 *MANKAL - * THE CICS CHECK OF MANKAL ENDED OK   *
06.10.09 JOB06908 *MANKAL - ********************************************
```

Logical Copy - Recap

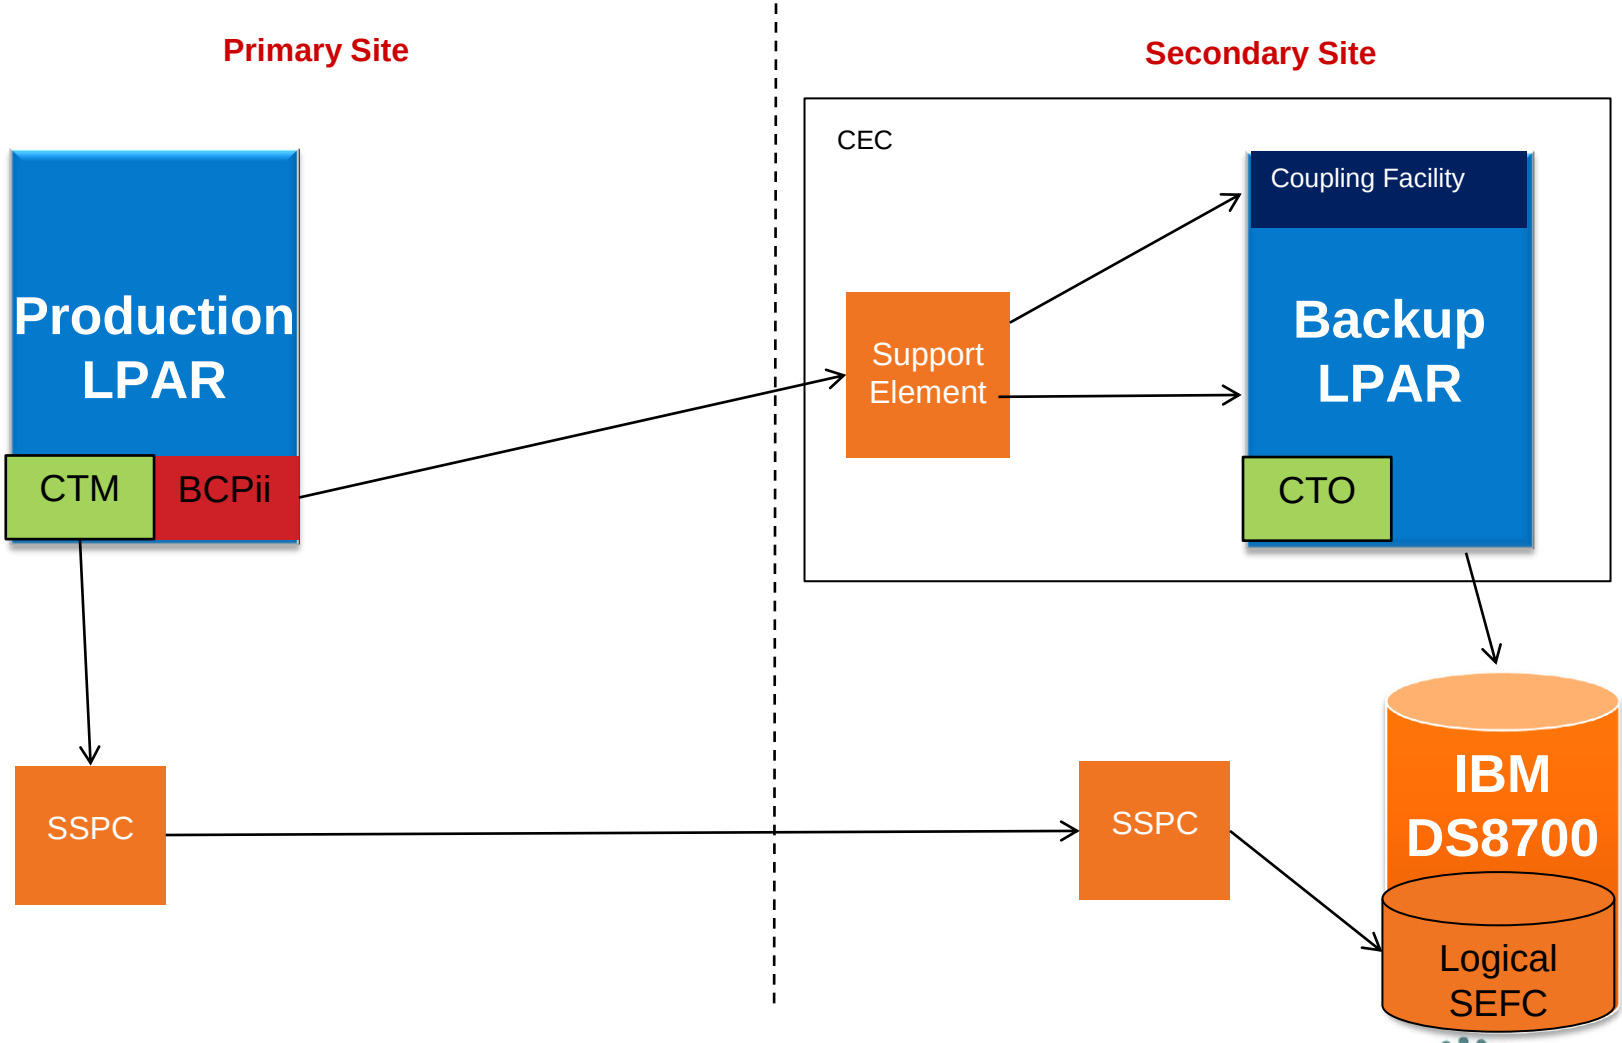

Logical Copy – side benefits

- A DR test every day!
- A true production environment which can be used to test new versions of software
- Improves MTTR – picks up errors at IPL time

SEFC– the downside

- SEFC impacts PPRC latency
- SEFC performance is impacted (affects DR tests)
- If we ever need to use it, we will not IPL directly from the copy. We will have to restore some or all of our data to the primary volumes

BCPii gotcha

- We had a problem responding to WTORs early in an IPL
- You need to set the HWI_CMD_OSCMD_PRIORITYTYPE field to HWI_CMD_PRIORITY

Agenda

- Introduction
- Level 1: Synchronous Replication
- Level 2: Logical Copies
- **Level 3: DRP Testing**
- Level 4: Third Site Copy
- Questions

DRP testing – the limitations

- We do not use the secondary PPRC volumes for DR testing
- We never stop the mirroring
- The User DR site and the IT DR center are 30km apart

DRP testing - How do we do it?

- We take snapshots of our production secondary copies and use them
  - For z/OS it is another SEFC set
  - For zVM it is a FC set
  - For the distributed environment we use XIV snapshots
  - The VTL does not support snapshots, but we can read the production tapes. Scratches are taken from a special pool
- All communication between the primary site and the DR site is disconnected
- Synchronous replication for the DS8700 and XIV continues
- A test runs for about 36 hours

Agenda

- Introduction
- Level 1: Synchronous Replication
- Level 2: Logical Copies
- Level 3: DRP Testing
- **Level 4: Third copy**
- Questions

Third Copy

- The financial regulation laws require that we have a third copy(that is, at neither of our sites) of our data at a secured location
- The assumption is that this copy will be used if both sites are permanently unavailable.
- Every Friday morning we bring up an LPAR at our primary site that reads that mornings logical copy and dumps it to a TS3500.
- Cartridges and reports are exported and sent off site
- Another LPAR is needed because you can’t bring the logical copies online (same VOLSERS as the production)

Third Copy

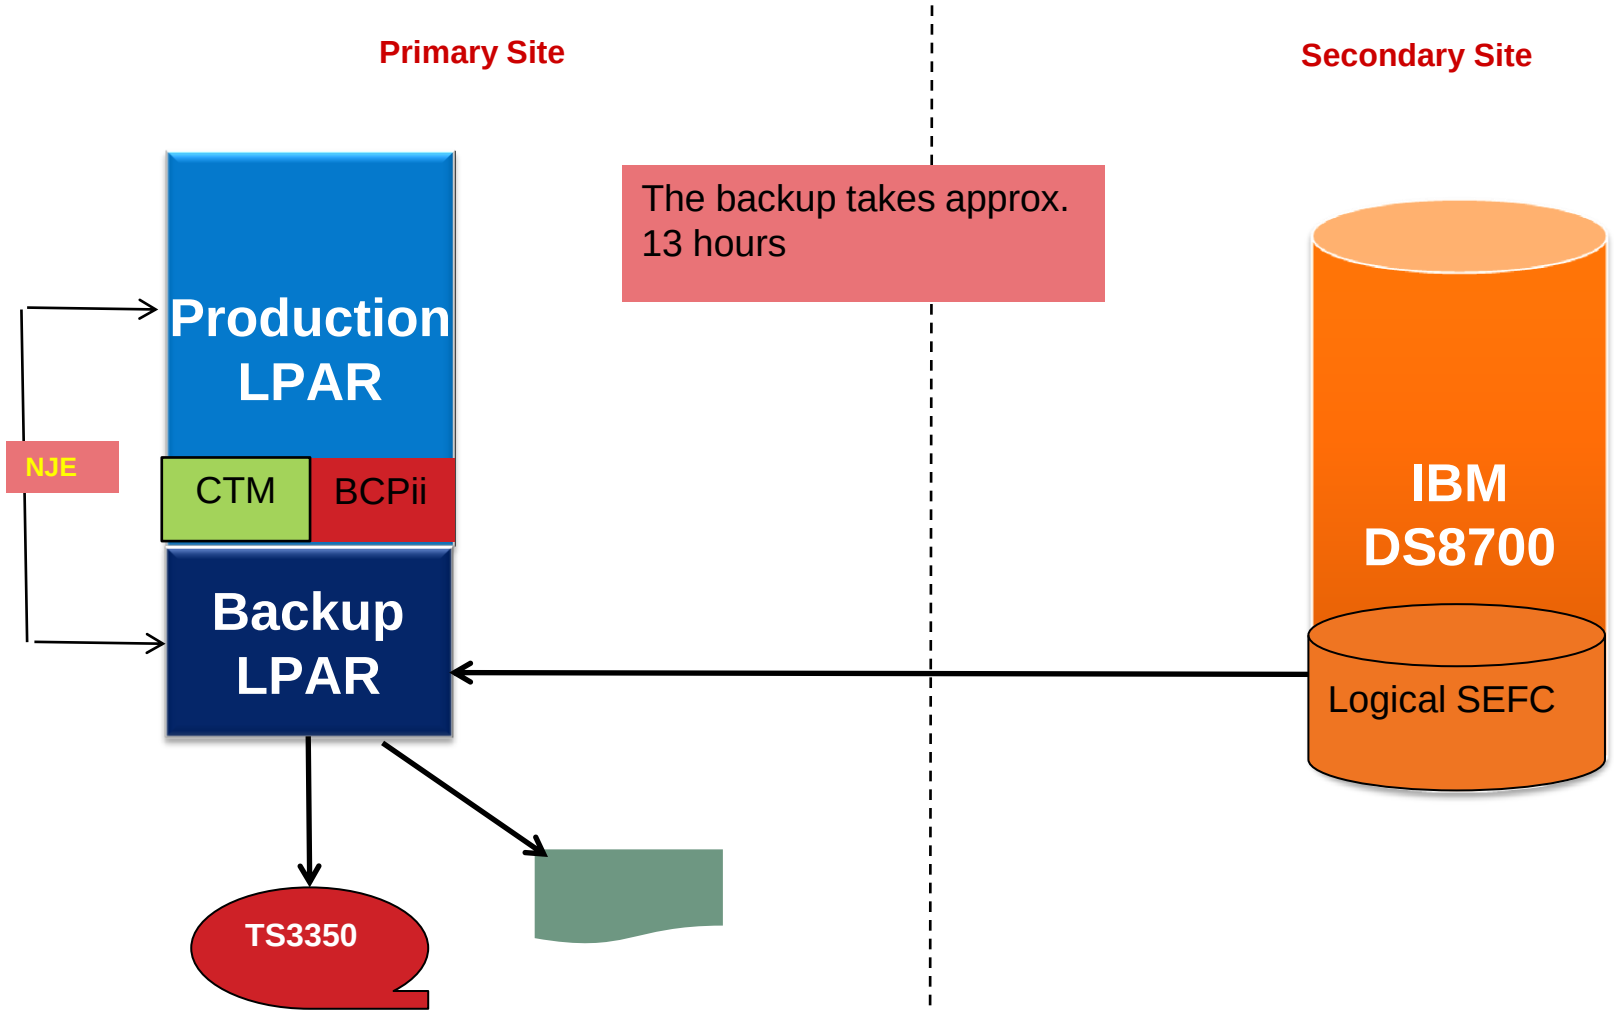

Third Copy - Output

- Cartridges that contain:
  - Our production data
  - Rexx and edit macros to customize the restore jobs at the new (unknown) site
- Hardcopy documentation
  - Requirements – Hardware, software
  - Inventory reports(created dynamically for each copy)
    - VOLSER to dataset mapping
    - Catalog structure

Next Steps

- Implement Hyperswap with TPC on z/OS
- For third copy – add stand alone IPL cartridge and test at a third site
- Distributed environment – implement logical copy

Summary

Scenario	Protection
Primary site DS8700 failure	Metro Mirror Copy
Primary site complete failure	MM copy + Backup CEC
Logical error that gets mirrored	Logical Copy
Both sites fail	Third copy

Questions ?