SNCF Business Resiliency project & implementation

Alain Mortier, Mainframe manager alain.mortier@sncf.fr
Bertrand Delaporte, TPF System manager bertrand.delaporte@sncf.fr

AGENDA

- SNCF in brief
- Resilient architecture
- Solution components
- Logical corruption protection

SNCF IN BRIEF

History

- SNCF was born in **1938**
- First high speed commercial line (limited to 170 mph) **1981**
- Distribution system (mainly high speed passengers bookings) on TPF **1993**
- World speed train record (357 mph) **2007**
- 2 billionth high speed train passenger **2013**
- 30% of high-speed tickets sold on mobile phone **2015**

SNCF Group today

- 14 million passengers a day worldwide
- Revenues of 33,5 € billion
- 270.000 employees spread across 120 countries

RESILIENT ARCHITECTURE

Project MILESTONES

- 2010 CEO Commitment to improve SNCF DRP
- 2011 Study, architecture definition and implementation
- 2012 Dual site solution set up locally (set up in our main DC)
  Data replication to be ready for splitting the data center
  Improved recovery from DB corruption
- 2016 Availability of D/R datacenter (8km as the crow flies)
- 2017 Study and validate the solution
  Full physical & logical protection
- 2018 Dual Site Implementation

RESILIENT ARCHITECTURE

Remote datacenter location

One of our objective was to be able to use cross connectivity between the 2 datacenters
e.g. having our z13 using tapes or/and disks of the other site
↳ It does allows to fallback quickly on the remote z13 just in case…
to switch tapes on remote VTS in case of maintenance…

↳ We did request D/R site to be located 30 km away maximum from our production site…

Partners selection for the optical link

A tender was issued to select links between the 2 datacenters.
↳ *IBM-Resiliency Services* solution with ADVA FSP3000 multiplexers has been chosen
↳ SNCF own fibers has been chosen with 2 different paths (13 and 26 km)

RESILIENT ARCHITECTURE

Validation by simulation

A prototype was setup using optical fiber coils :

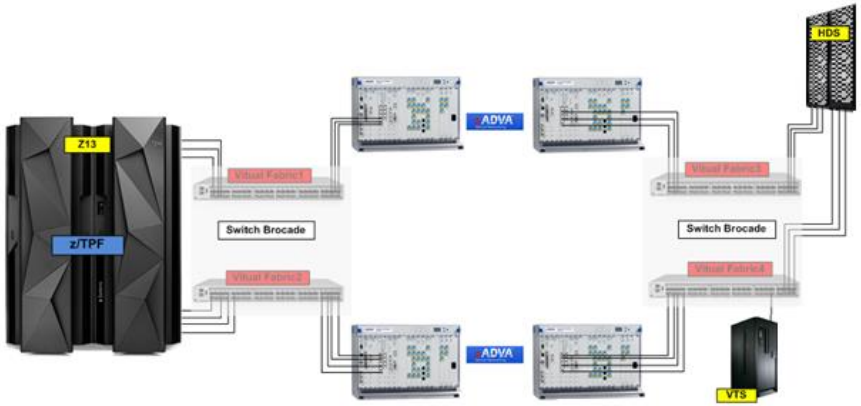

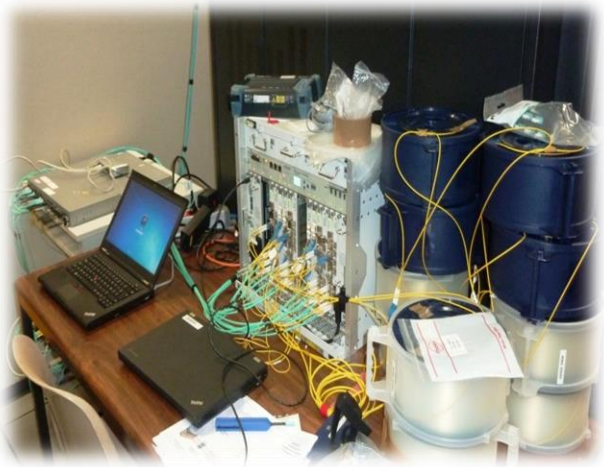

- Latency due to distance was close to the theoretical value (0,25 msec)
- z/TPF system has been resilient to a path failure/switch at ADVA level (lost paths can be recovered with a ZPATH UP for the disks…)

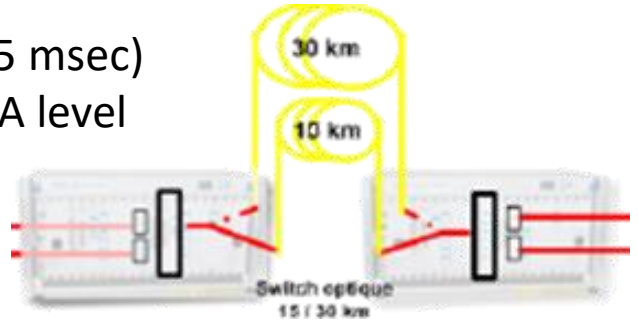

Decisions taken at that stage

Project gave us the opportunity for a technological upgrade :

- VTS TS7740 +TS3500 (robot) upgraded to 2 TS7760 in GRID mode to simplify our tape replication
- 15k HDDs controllers upgraded to SSD based disk controllers to recover 0,25 msec latency

RESILIENT ARCHITECTURE

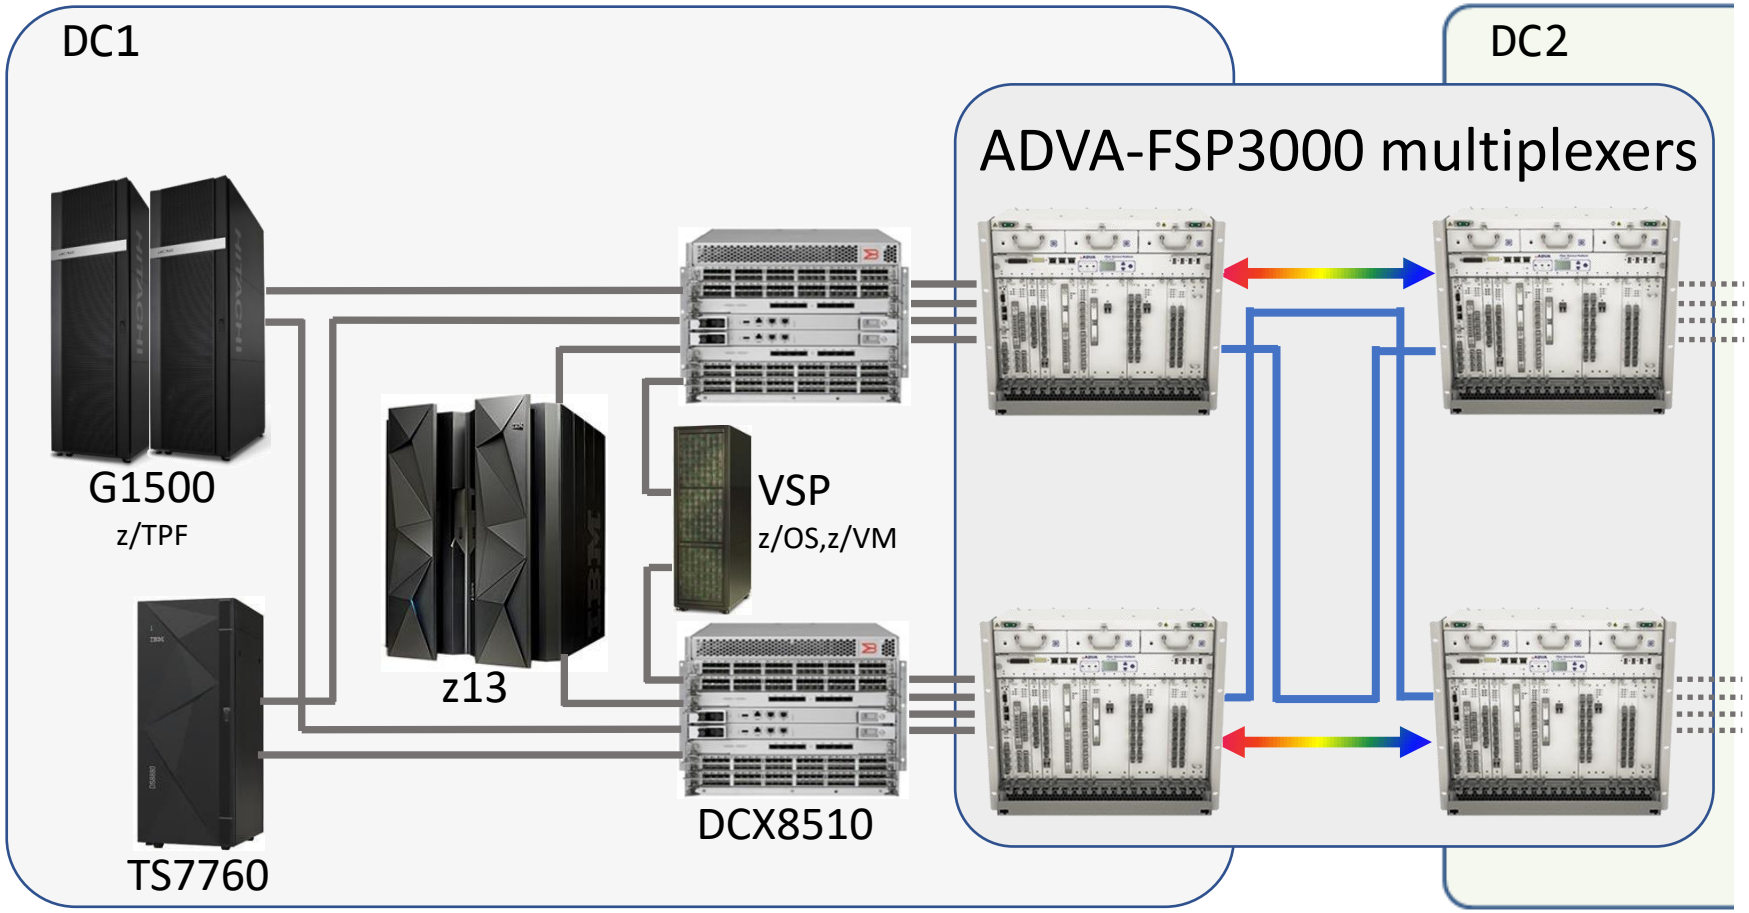

RESILIENT ARCHITECTURE

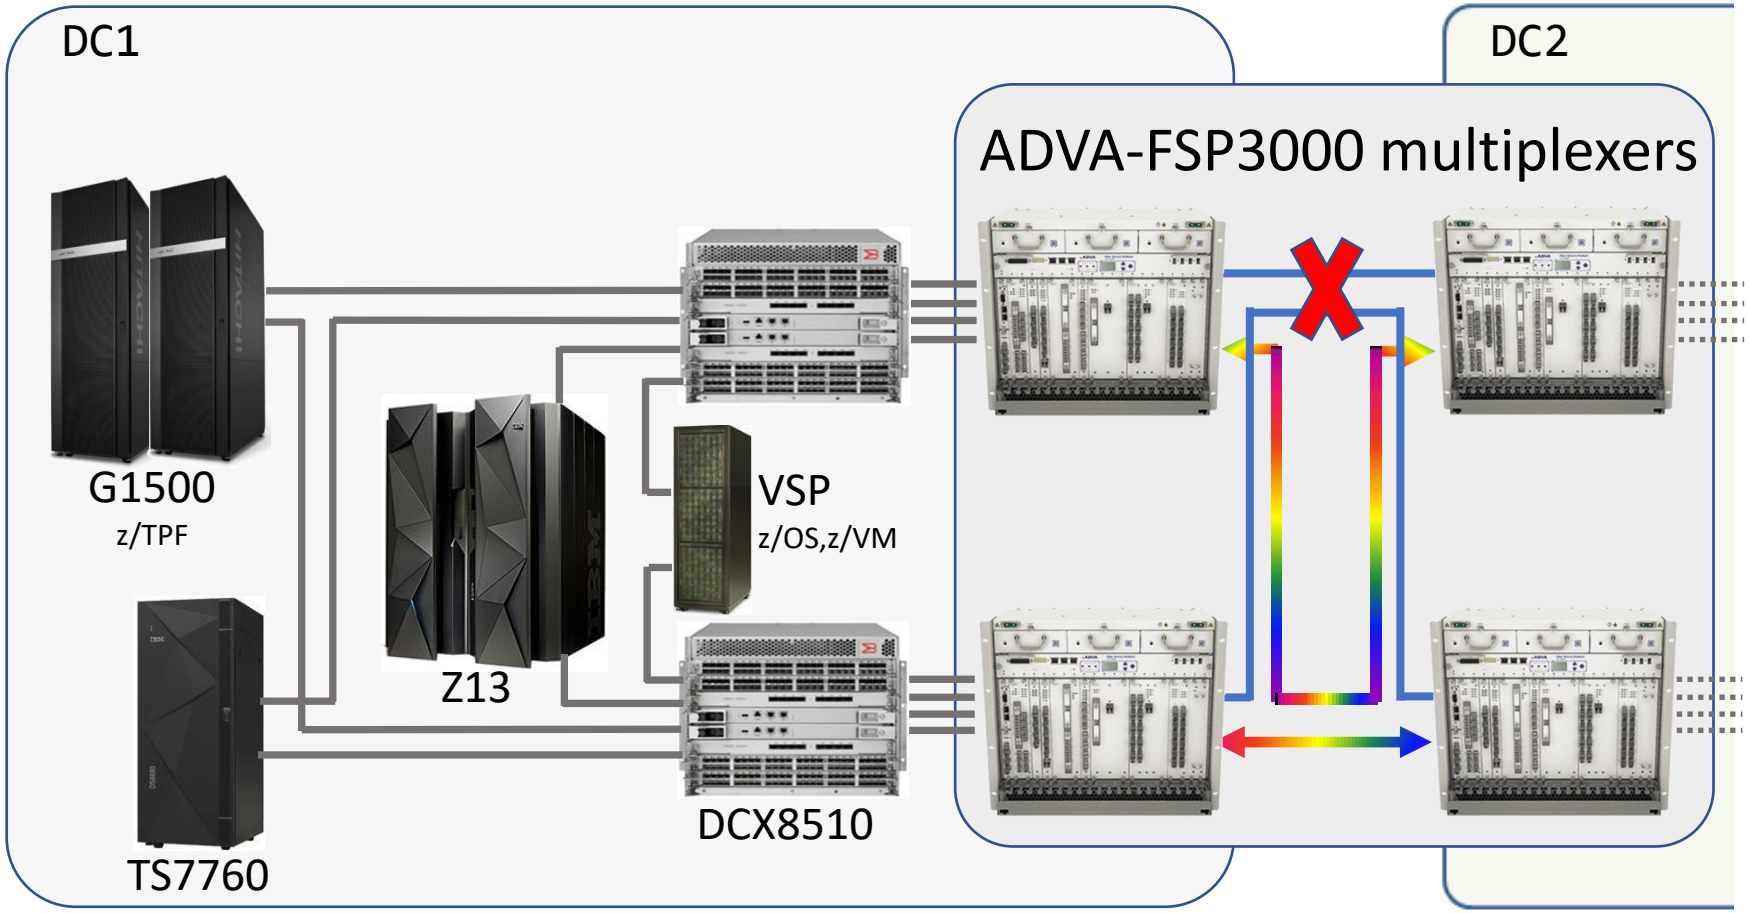

Should a path fail, switching is done at multiplexor level…

RESILIENT ARCHITECTURE

Physical installation and tests

Many different configurations, in particular with *cross* site connected equipment, were tested
(e.g. D/R testing scenarios including z13 on one site using VTS/DASD on the other site, restart on D/R site, on Image)

For each various configuration we did carry out performance measurements and hardware failure testing

⇨ Latency induced by distance was close enough to theoretical value (speed of light through glass)
⇨ We didn't encountered any show stopper

Important considerations as result of testing (none of them linked to the dual site configuration):

- Improved tape availability through GRID
- Improvement of z/TPF tape management that would be much appreciated:
  - Having an easy way to increase the size of the *tape status table* (RFE 131445)
  - Having an easy way to recover tape drives after hardware failure (path down)
  - Get access to full sense code table documentation
  - Improve interpretation and reaction by z/TPF upon an unknown sense code
- New disks performance!

SOLUTION COMPONENTS

ADVA FSP3000 (9U)

- 2 WDM pairs (1st is nominal on EST route, 2nd on WEST route)
- RSM/OLM module (Remote Switch Module / Optical Line Monitoring)

FIBERS

- 14 (8Gb) ISL links between DCX and ADVA
- For production data transfer we do have as many ISL links as of FICON/Fiber Channel links

DCX 8510

- 5 virtual fabrics used for flow segregation :

```
z/TPF Production -1- DASD
                 -2- Tape
                 -3- Replication (HUR)
Z/VM, z/OS, z/TPF non Production -4- DASD + Tape
                 -5- Replication (HUR+PPRC)
```

VTS 7760

- GRID replication (Copy deferred)

G1500

- FMD disks
- Inbox replication : ShadowImage
- Out-of-Box replication: HUR

SOLUTION COMPONENTS: FLASH DISKS

Our previous configuration was build around 146GB 15K rpm HDDs

The new HV G1500 are SDD based, Flash Module Drive (FMD) disks

Our production is based on 2 controllers for each site

- Data are replicated asynchronously between sites using HUR
- We do take consistent Shadow Images for local D/R regularly

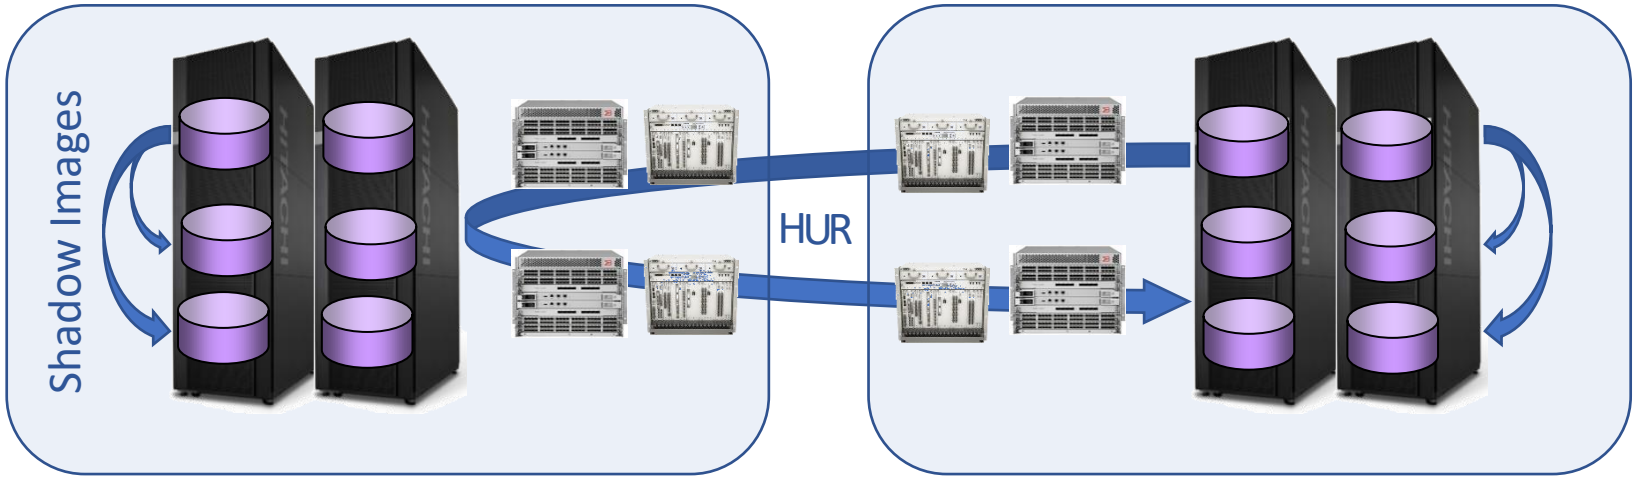

Data consistency warantee maintained between controllers without freezing I/O, thanks to:

- « At Time Split » for Shadow images
- « Extended Consistency Group » for HUR

SOLUTION COMPONENTS: FLASH DISKS

After the first tests we can say that we were… **AFRAID !**

How our legacy applications would be impacted with such low response time ?
How will VTS handle the RTX logging during batch process ?

Many tests were carried out. (we had to make some improvements to our testing tools suite!)

Major items :

- Code for Nightly File Maintenance had to be adapted
- We have reduced parallelism for some batch processes (ECB #)
- Our logging is now on 2 tapes (RTX/RTY)

Cutover in Production was done 6 months ago, we did not experience any issue since then

SOLUTION COMPONENTS: FLASH DISKS

Hardware architecture

- 4 FICON zHPF (8GB due to DCX)
- 512 GB Cache memory
- 1 array group (2D+2D) 3.2TB FMD for reservation system data + images
- 1 array group (3D+1P) for HUR journal
- 670 LDEVs defined (102 primes + 102 dupes + 408 shadow images + 58 spares)

FMD advantages vs SSD

- Hardware inline compression without impact on performance
- Extended life Time of cells : data are written compressed and binary zeros are extremely good candidates for compression and never written as is…

SOLUTION COMPONENTS: FLASH DISKS

Datacollection run concurrently with transactional traffic

	192* 15K disks R10 (2D+2D)	4* FMD disks R10 (2D+2D)
Service time	1,04 msec	0,17 msec
Device queue (mean)	14,85	4,82
LIOCB (mean use)	17,54	7,99

Batch duration

	192* 15K disks R10 (2D+2D)	4* FMD disks R10 (2D+2D)
RECOUP	1h 40min (300 ECBs)	22min (160 ECBs)
NFM	1h 10min	11min
ULCUs	28 min (200 ECBS)	6min (50 ECBs)
METRIX (stats batch)	1h	5min

With data replication active

SOLUTION COMPONENTS: FLASH DISKS

IOs per controller during RECOUP in excess of 120.000 IOPs (vs 25.000 IOPs)

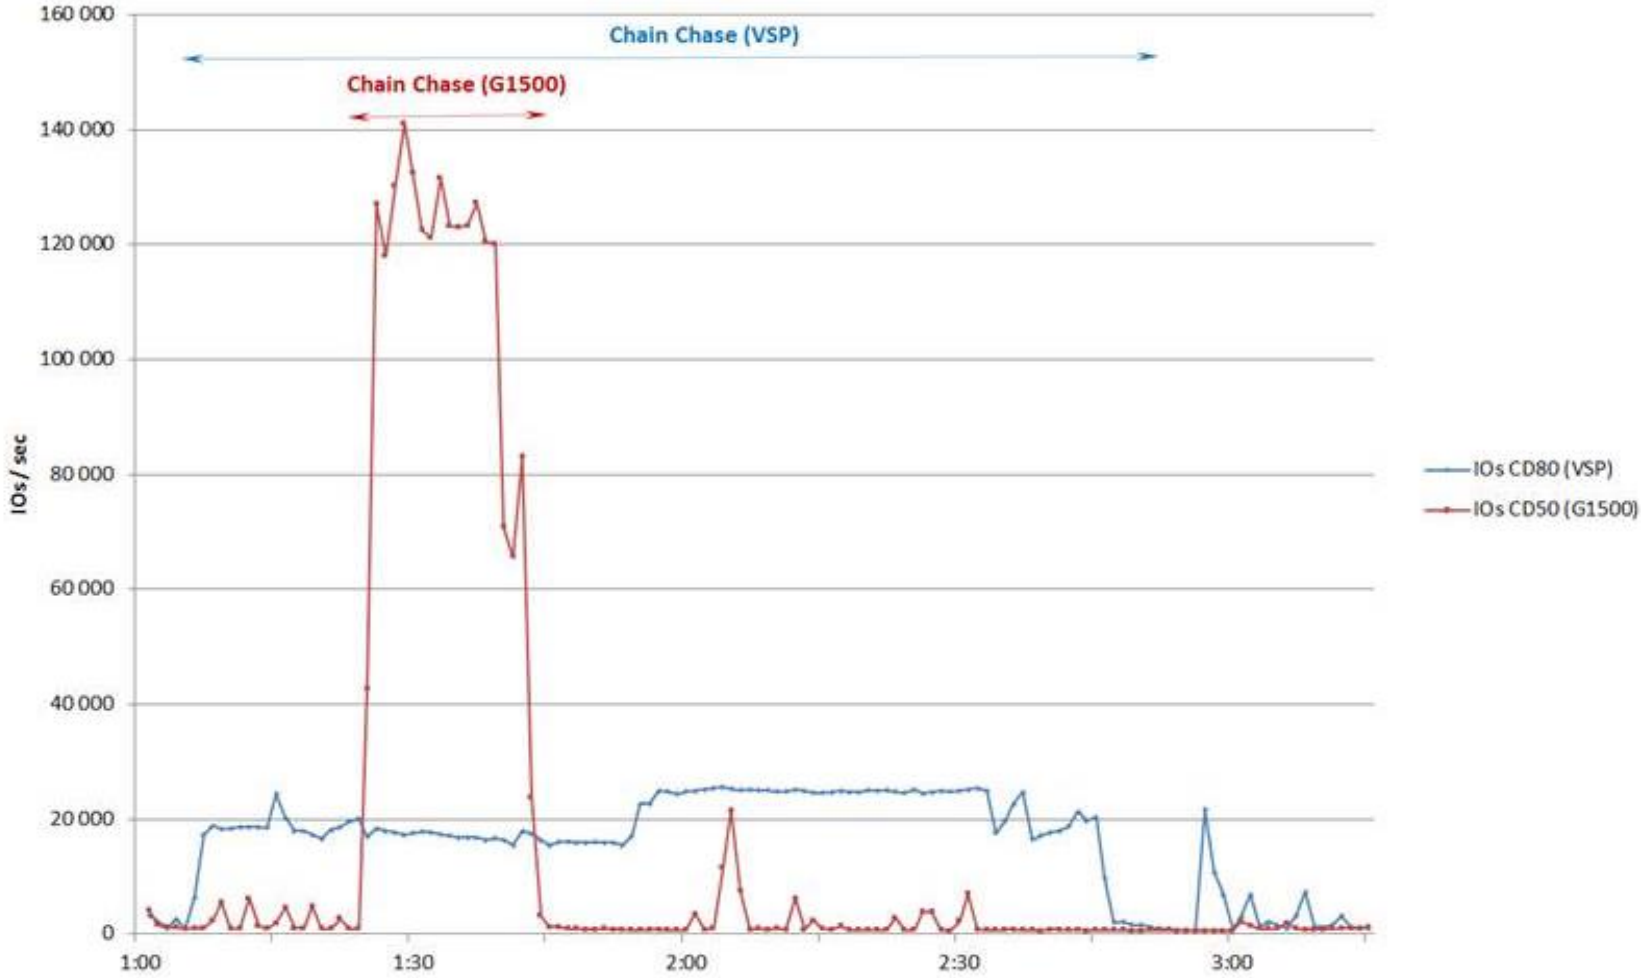

With data replication active

SOLUTION COMPONENTS: FLASH DISKS

Service Time during RECOUP remain close to 0,3msec (vs 3 to 5 msec)

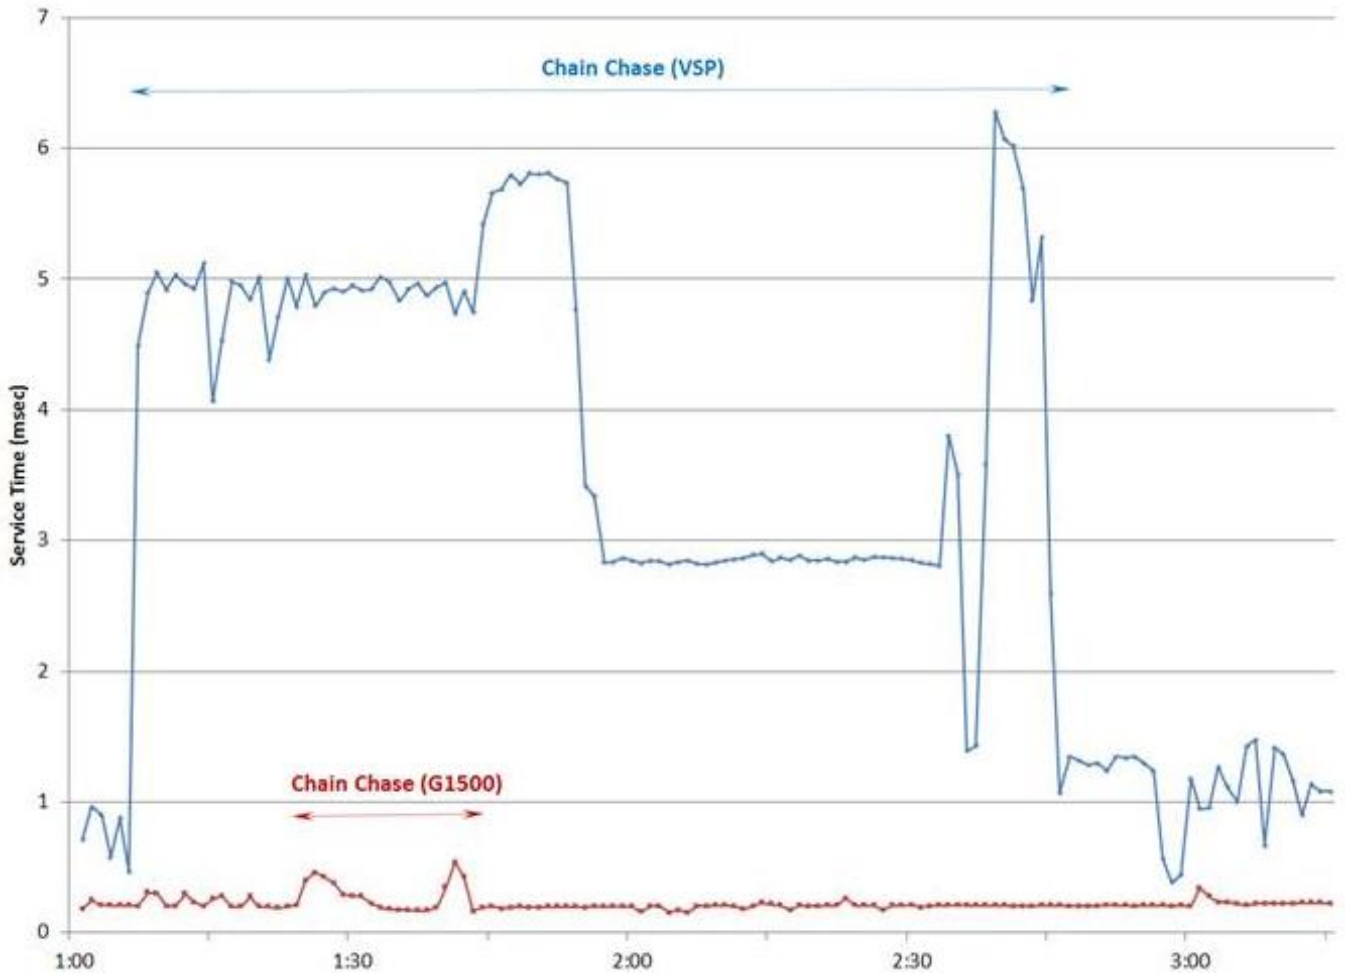

With data replication active

SOLUTION COMPONENTS: FLASH DISKS

Impact of distance on Service Time

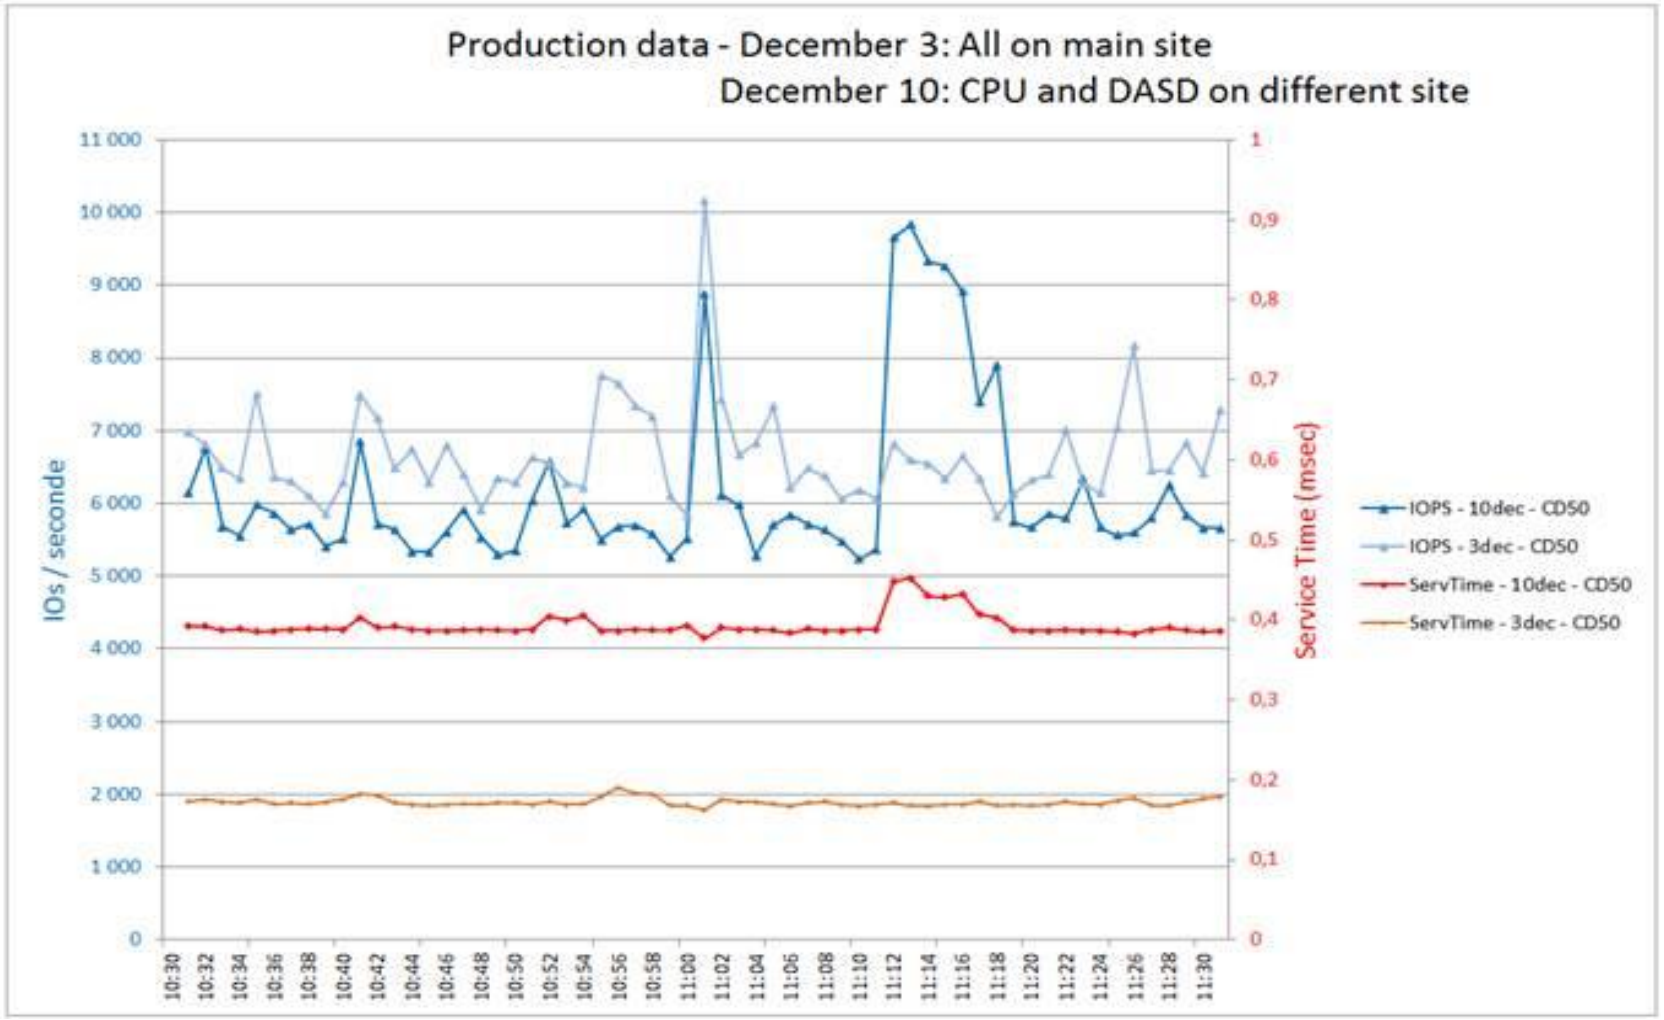

With data replication active

SNCF REPLICATION MONITORING & AUTOMATION TOOL

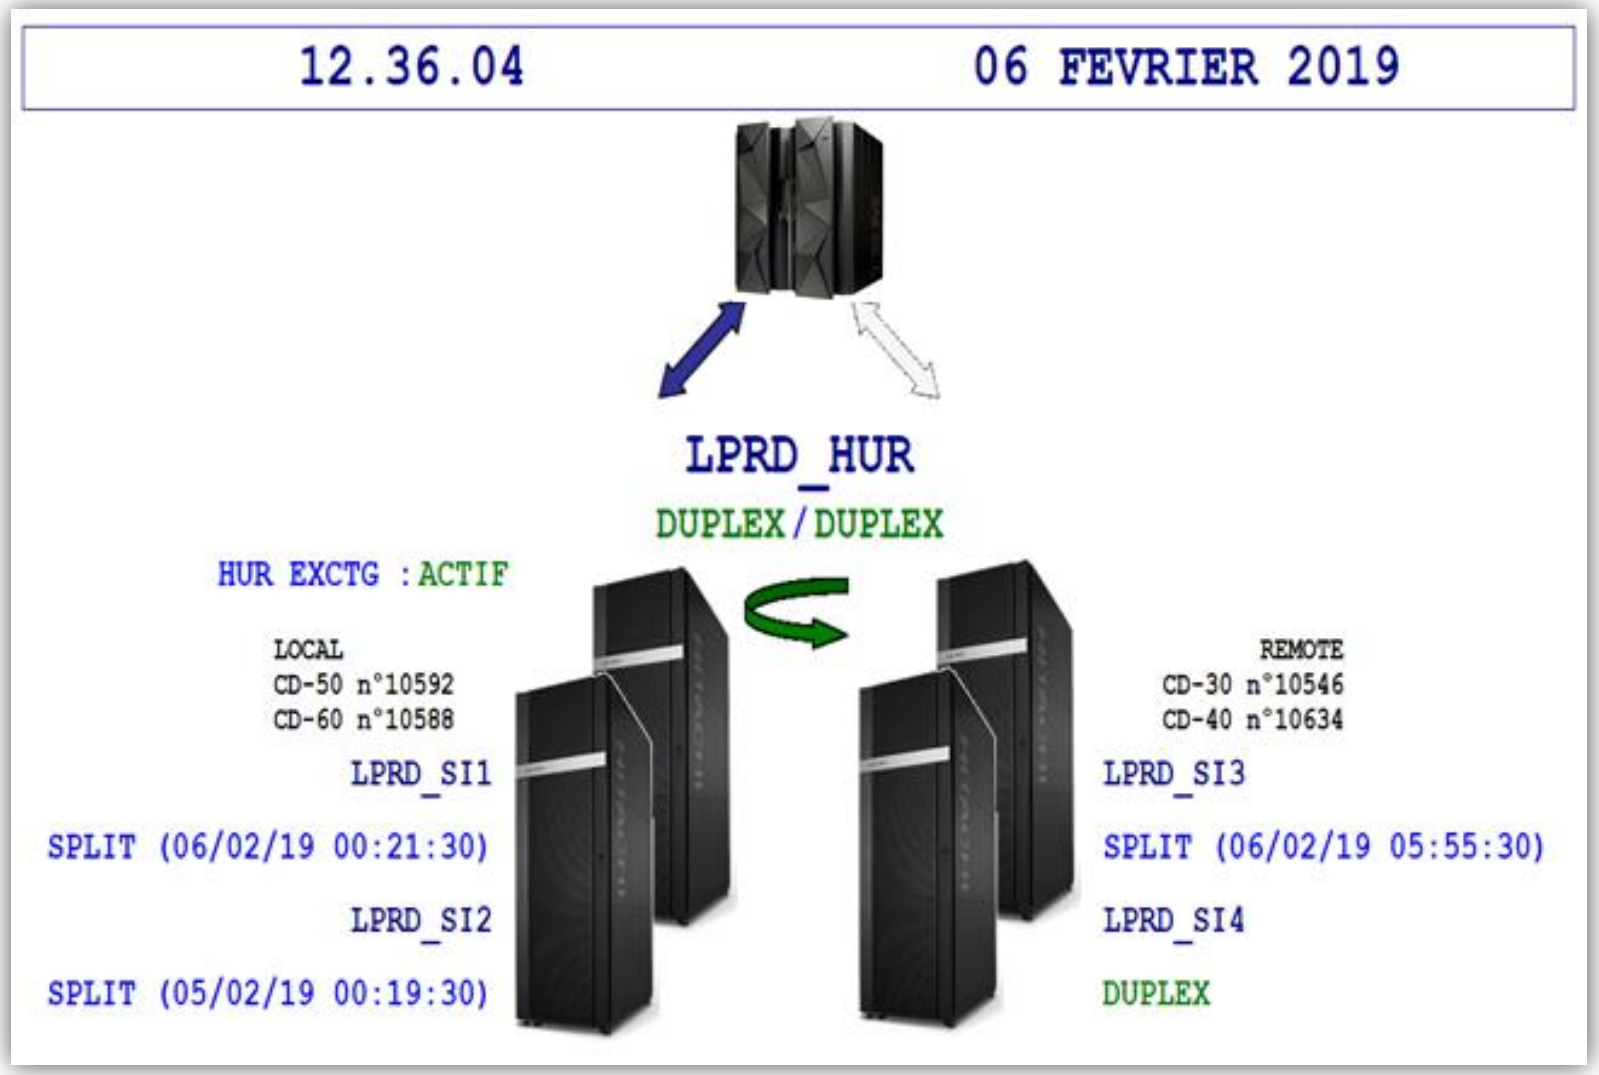

LOGICAL CORRUPTION PROTECTION

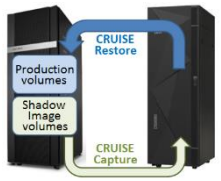

CRUISE improvement

After having implemented Shadow Image (2012) we did test a CRUISE modification using FDRSC macro.
-> Presented at the 2013 TUG: “SNCF Recovery Plan with HDS Shadow image and HUR”

FDRSC does allow to read a TPF record out of an offline ShadowImage target volume.

Following a successful testing, we did ask IBM to provide a CRUISE user exit in order to implement FDRSC.
After submitting a RFE, we had the chance to join an “IBM Sponsor user program”…

We also did ask Hitachi Vantara for some FDRSC efficiency improvements…

Put15 did bring us APAR PH04143

3 additional user exits for the CRUISE utility:

- CASE 13 of dfuex.mac — where FDRSC can be coded
- CRUUSR_BEFORE_ALL_CHAIN_CHASES — to set environment before to start
- CRUUSR_AFTER_ALL_CHAIN_CHASES — to reset environment

LOGICAL CORRUPTION PROTECTION

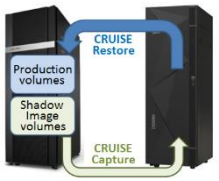

CRUISE and alternate find

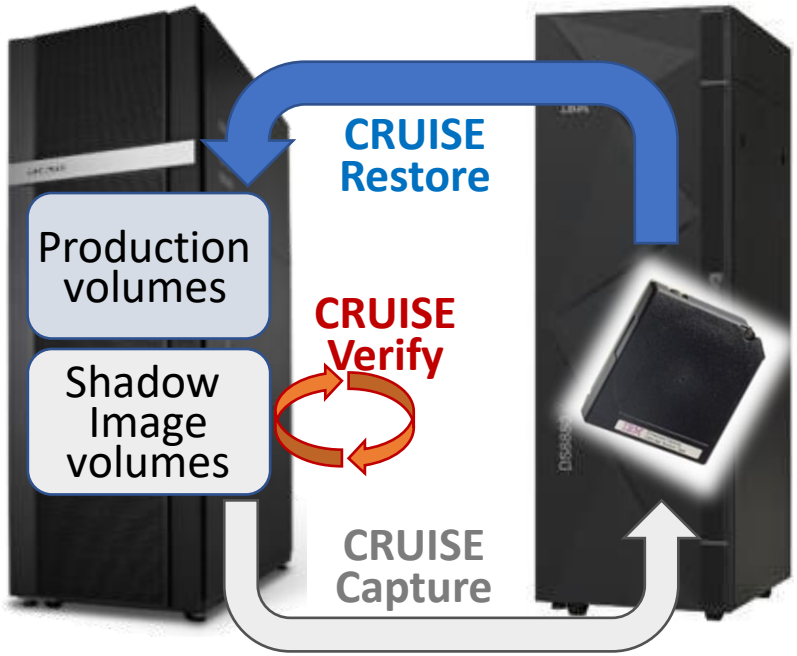

Business benefits

- Image validation before using it to restart
- Quick restore of a single corrupted database
- No need to write a utility per database

Works with classical database having a DBDEF

↳ Our classical RECOUP descriptors have been rewritten in DBDEF

LOGICAL CORRUPTION PROTECTION

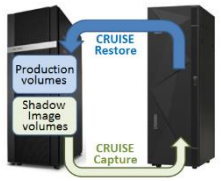

CRUISE SETUP

ZFCRU SETUP FUN-CAP REF-FR23RR **ALTFIND**-YES **UEX**-ShadowImage2

ALTFIND if YES then CASE 13 is called to do the FDRSC
UEX allows to specify the volumes to be read (Shadow Image name)

```
> ZFCRU LOCK-VERC820
> ZFCRU SETUP FUN-VER REF-C820 ALTFIND-YES UEXIT-LCOV_SI1 STA-YES
FCRU0000I 11.35.19 PARAMETER TABLE DISPLAY FOR VERC820
  FUNCTION                  VER
  . . .
  AUTOMATIC PAUSE           NO
  ALTERNATE FIND USER EXIT  YES
  REFERENCE ID              C820
  USER EXIT TEXT             L C O V _ S I 1
FCRU0000I 11.35.19 END OF DISPLAY
> ZFCRU UNLOCK-VERC820
```

- **IBM** for the *Products Sponsor User and Early Release* program
  - Daniel Jacobs and Chris Filachek for their support

- **Hitachi Vantara** who did assists us throughout the project
  - Gary Spencer who has always been available for us
  - Alain Beauregard for his great technical expertise

TPF
Users Group

FLASH DISKS CONTROLLERS

SHADOW IMAGE “AT TIME SPLIT” PROCESS

- All group pairs are split at same time
- Consistency groups are spread across controllers
- No need to freeze IOs to capture image
- Split happens upon first IO after the set time
- Pairs that didn’t receive any IOs will be split at time-out

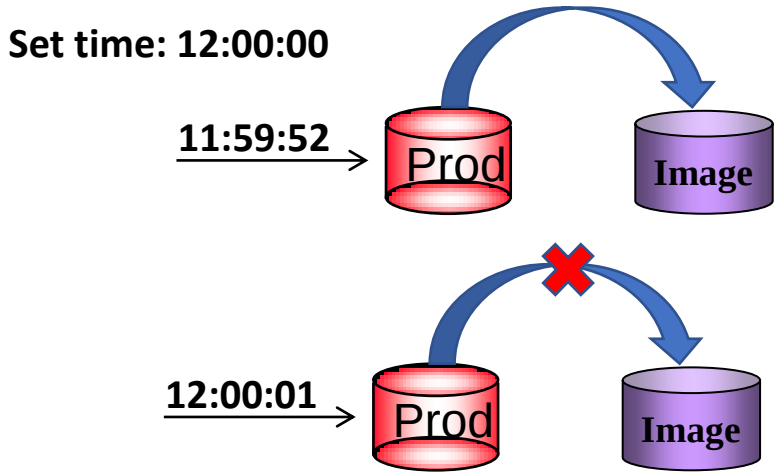

FLASH DISKS CONTROLLERS

EXTENDED CONSISTENCY GROUP PROCESS

- The write timestamp is part of the host I/O and used by HUR to maintain consistency between HUR journals
- HUR regularly checks the latest write time stamp for each journal group of the EXCTG:
  - In Journal group 1, the latest time stamp is 15:00
  - In Journal group 2, the latest time stamp is 15:02
  - In Journal group 3, the latest time stamp is 15:03
  - In Journal group 4, the latest time stamp is 15:04
- Takes the oldest time stamp: 15:00
- Then it does allow on all controllers restore of all journal data that have a time stamp up to 15:00

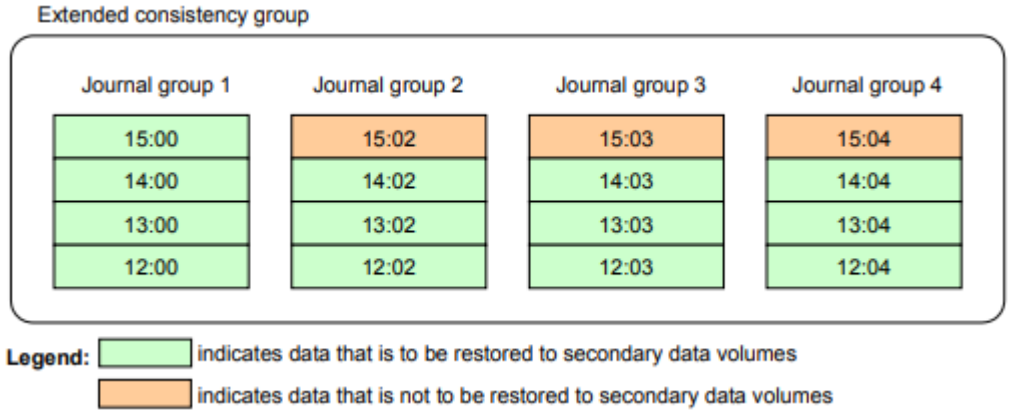

HOW TO RESTART ON A SHADOW IMAGE

One CTK0 defined per database area

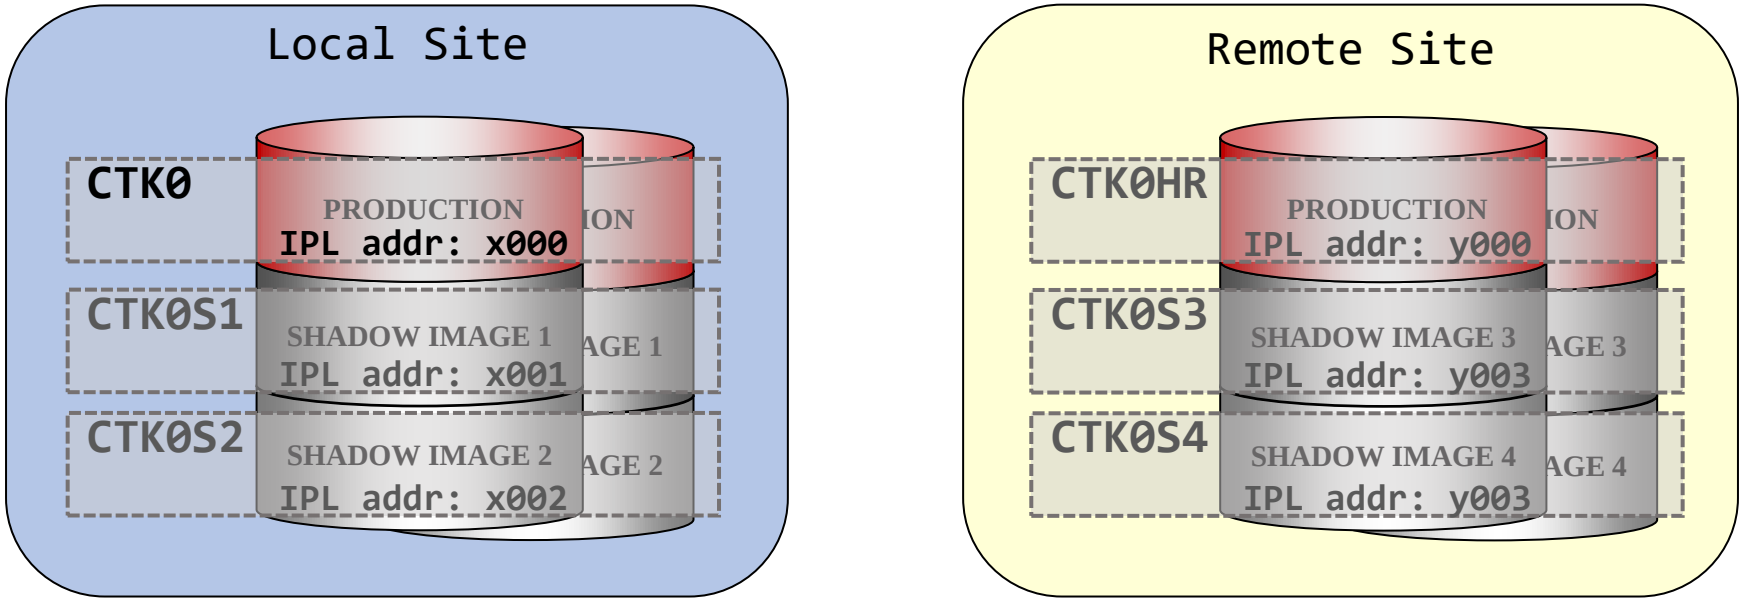

Restart process:

- 1st IPL on a GF pack built with the *database area* CTK0 wanted
- 2nd IPL on the *database area* IPL address
- RTX/RTY tapes can be injected to reduce RPO (Recovery Point Objective)

SHADOW IMAGES STRATEGY

A compromise : recent image (⇘RTO) / older image (crawling corruption)

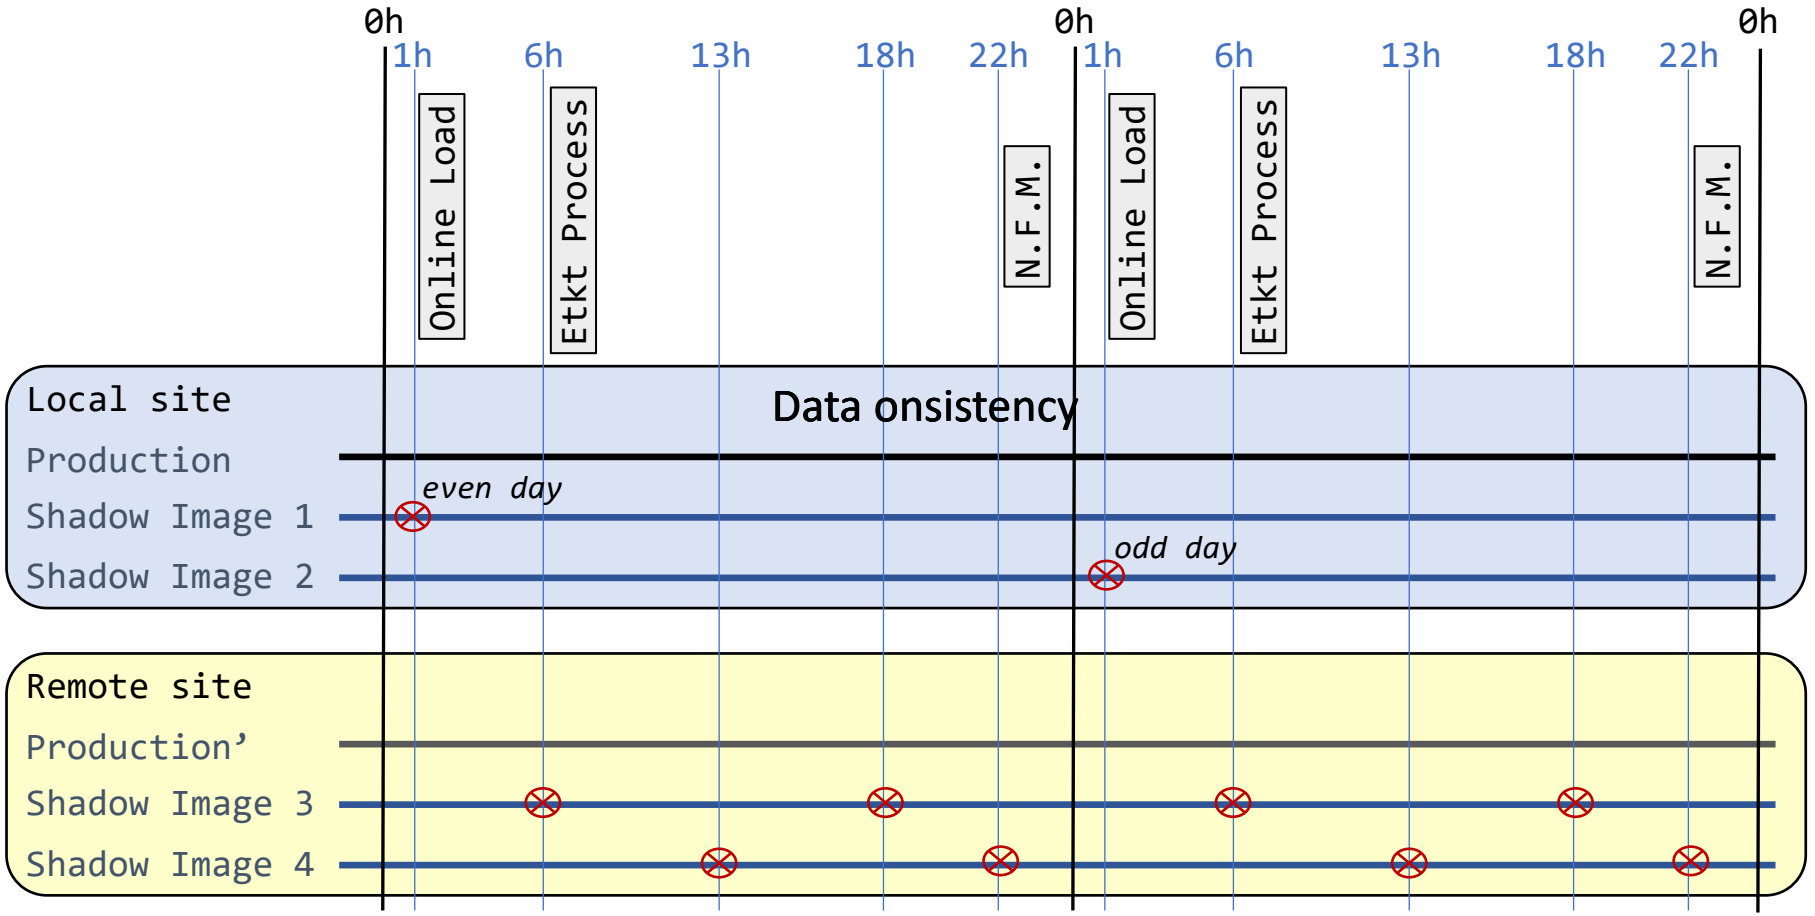

Introduction : ShadowImage

- ShadowImage At Time Split :
  - Suspend of each SI pair based on timestamp
  - Allow multi-controller consistency

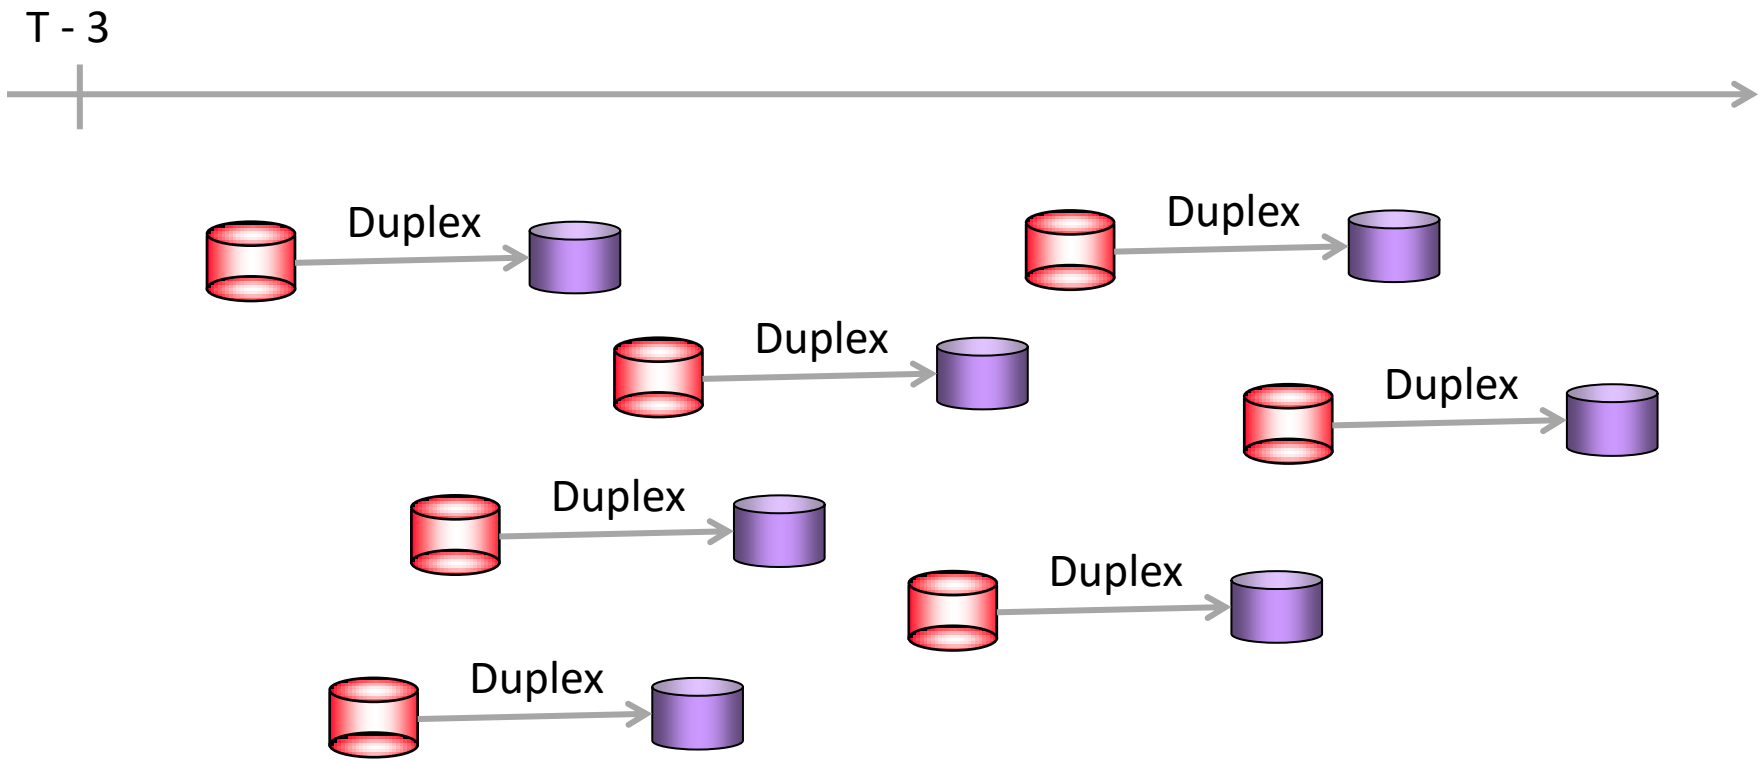

Introduction : ShadowImage

- ShadowImage At Time Split :
  - Suspend of each SI pair based on timestamp
  - Allow multi-controller consistency

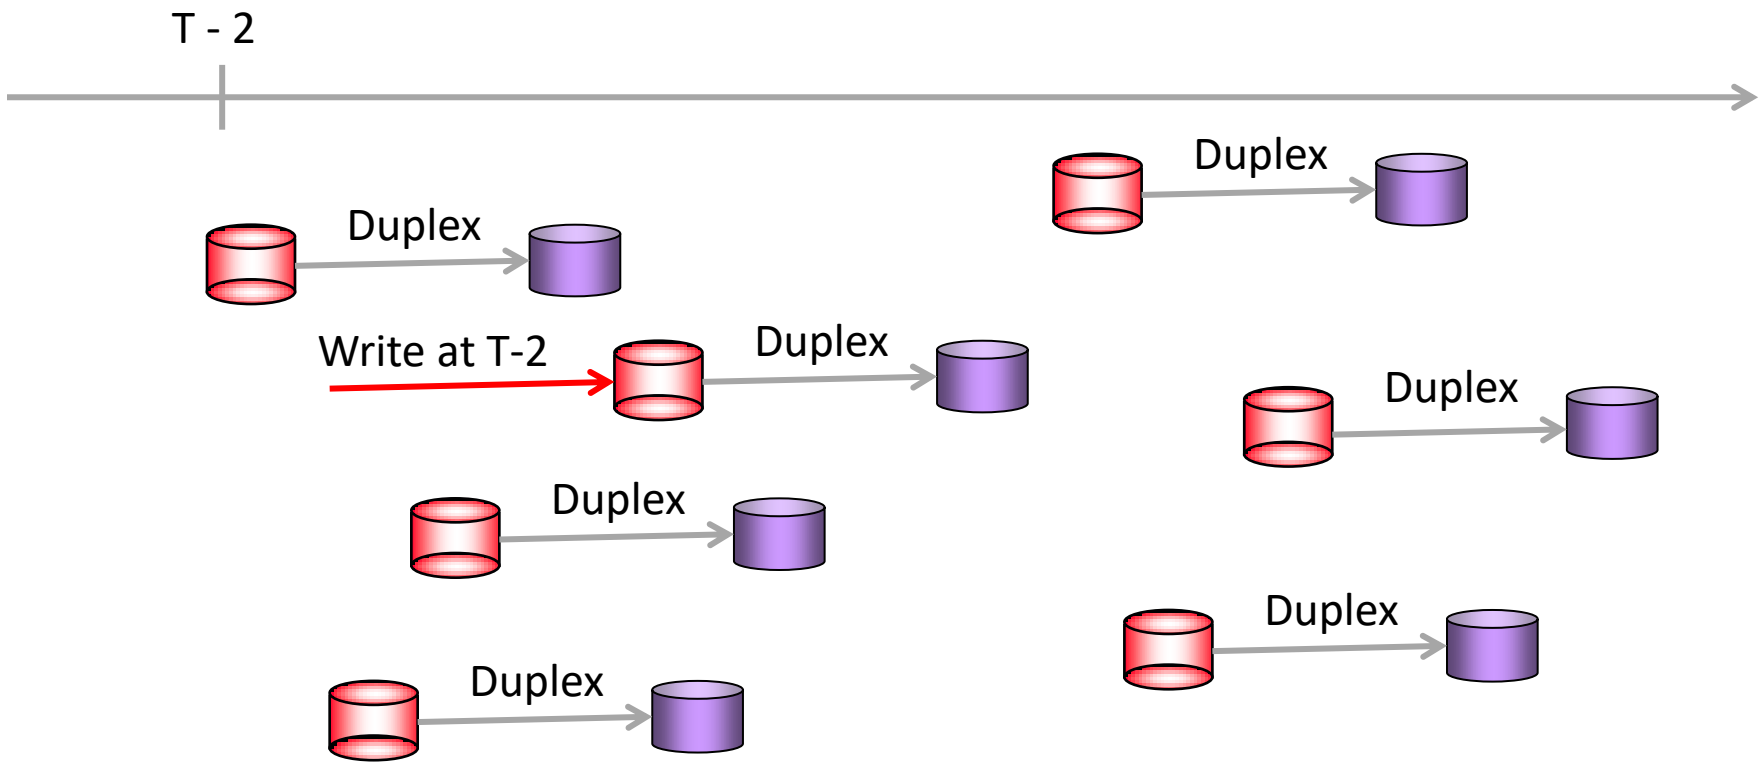

Introduction : ShadowImage

- ShadowImage At Time Split :
  - Suspend of each SI pair based on timestamp
  - Allow multi-controller consistency

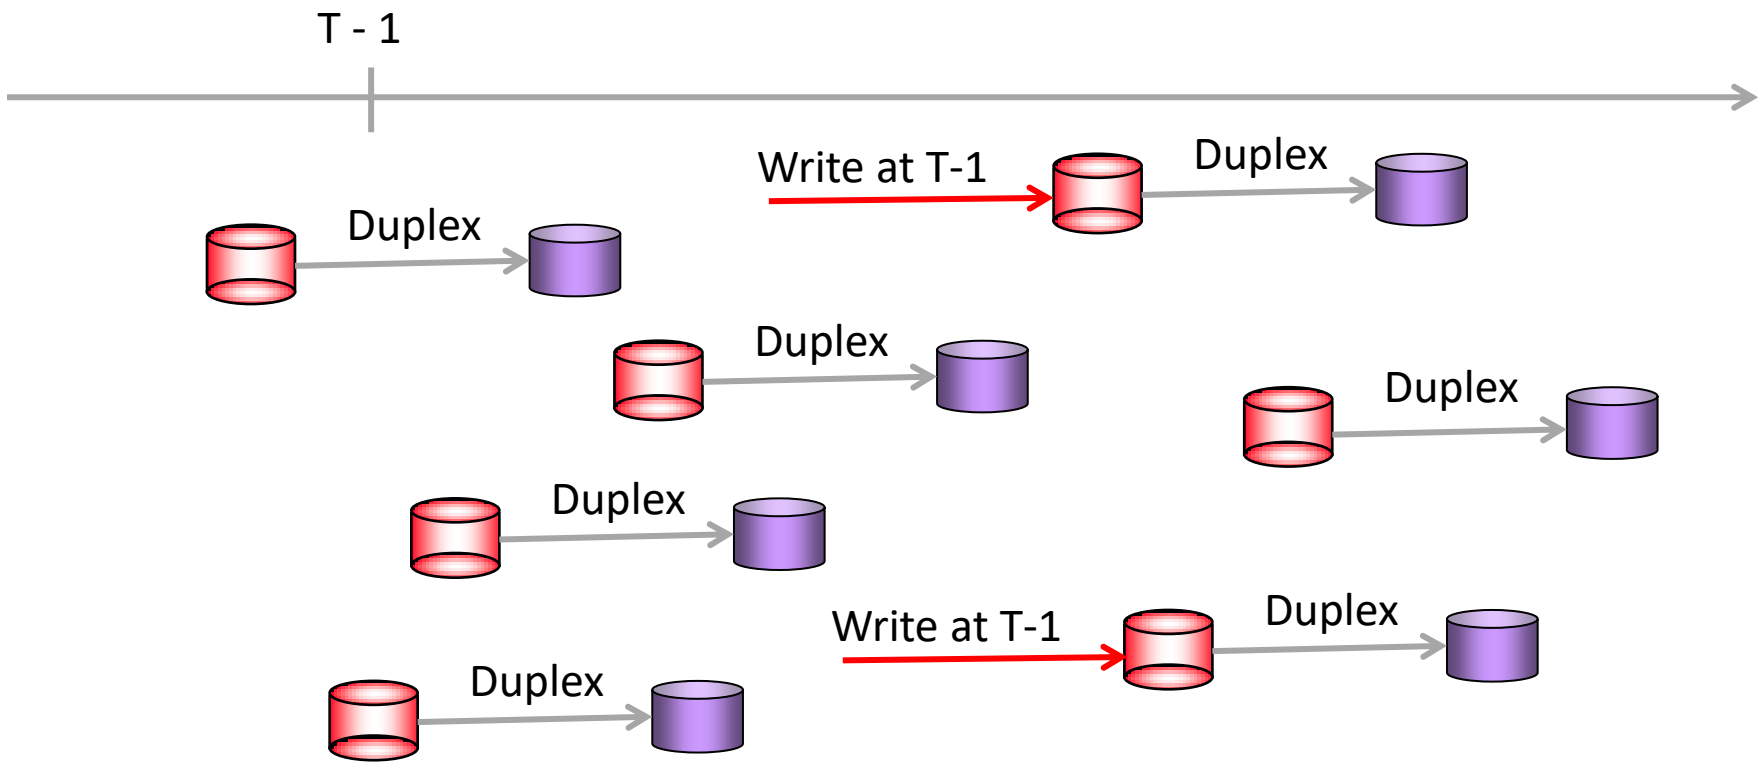

Introduction : ShadowImage

- ShadowImage At Time Split :
  - Suspend of each SI pair based on timestamp
  - Allow multi-controller consistency

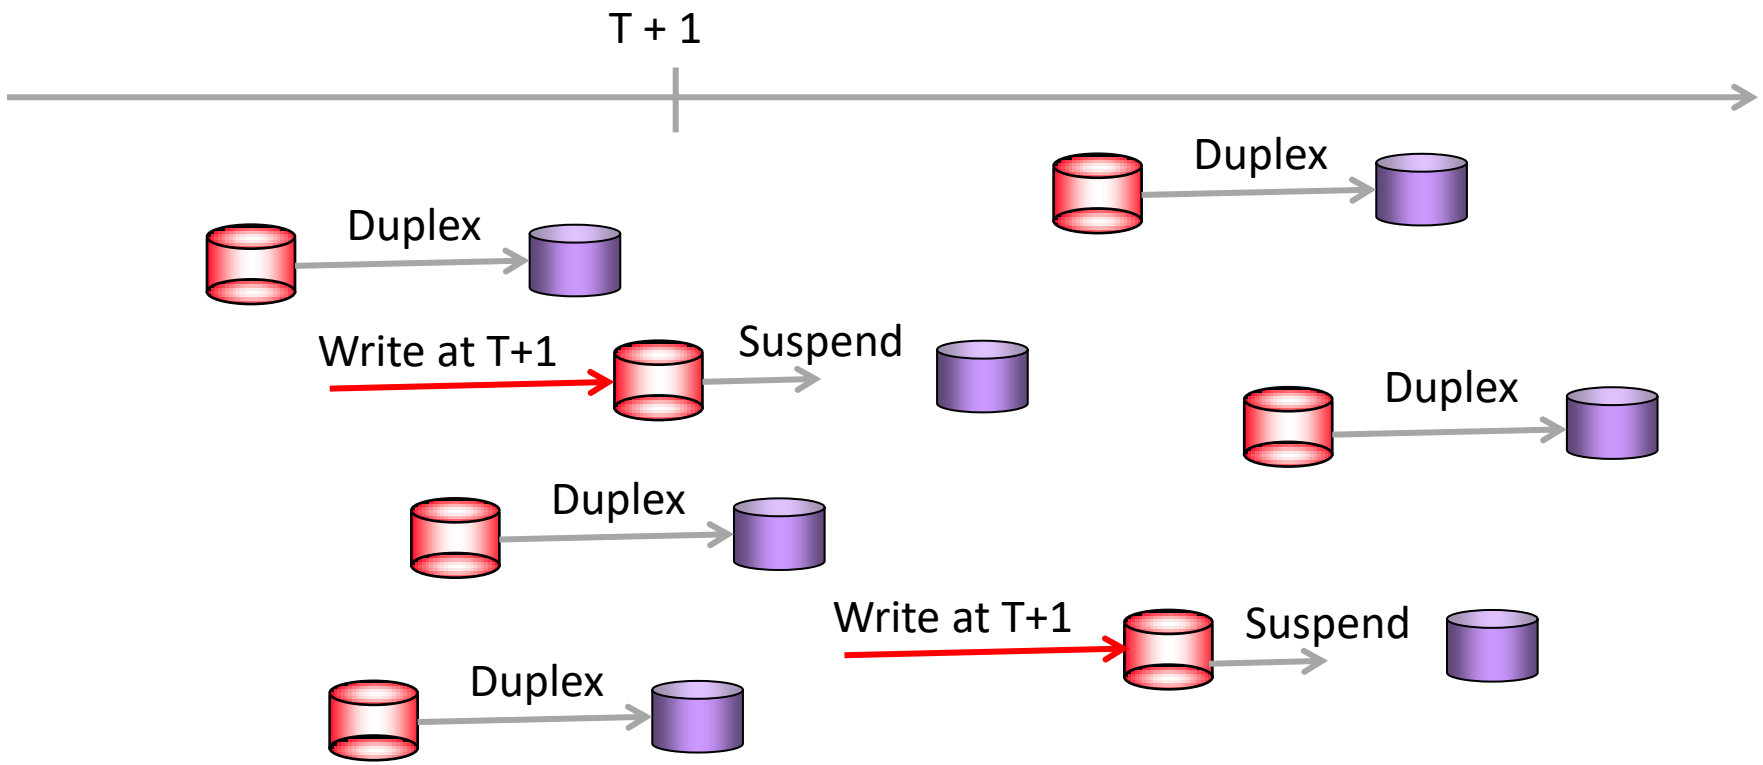

Introduction : ShadowImage

- ShadowImage At Time Split :
  - Suspend of each SI pair based on timestamp
  - Allow multi-controller consistency

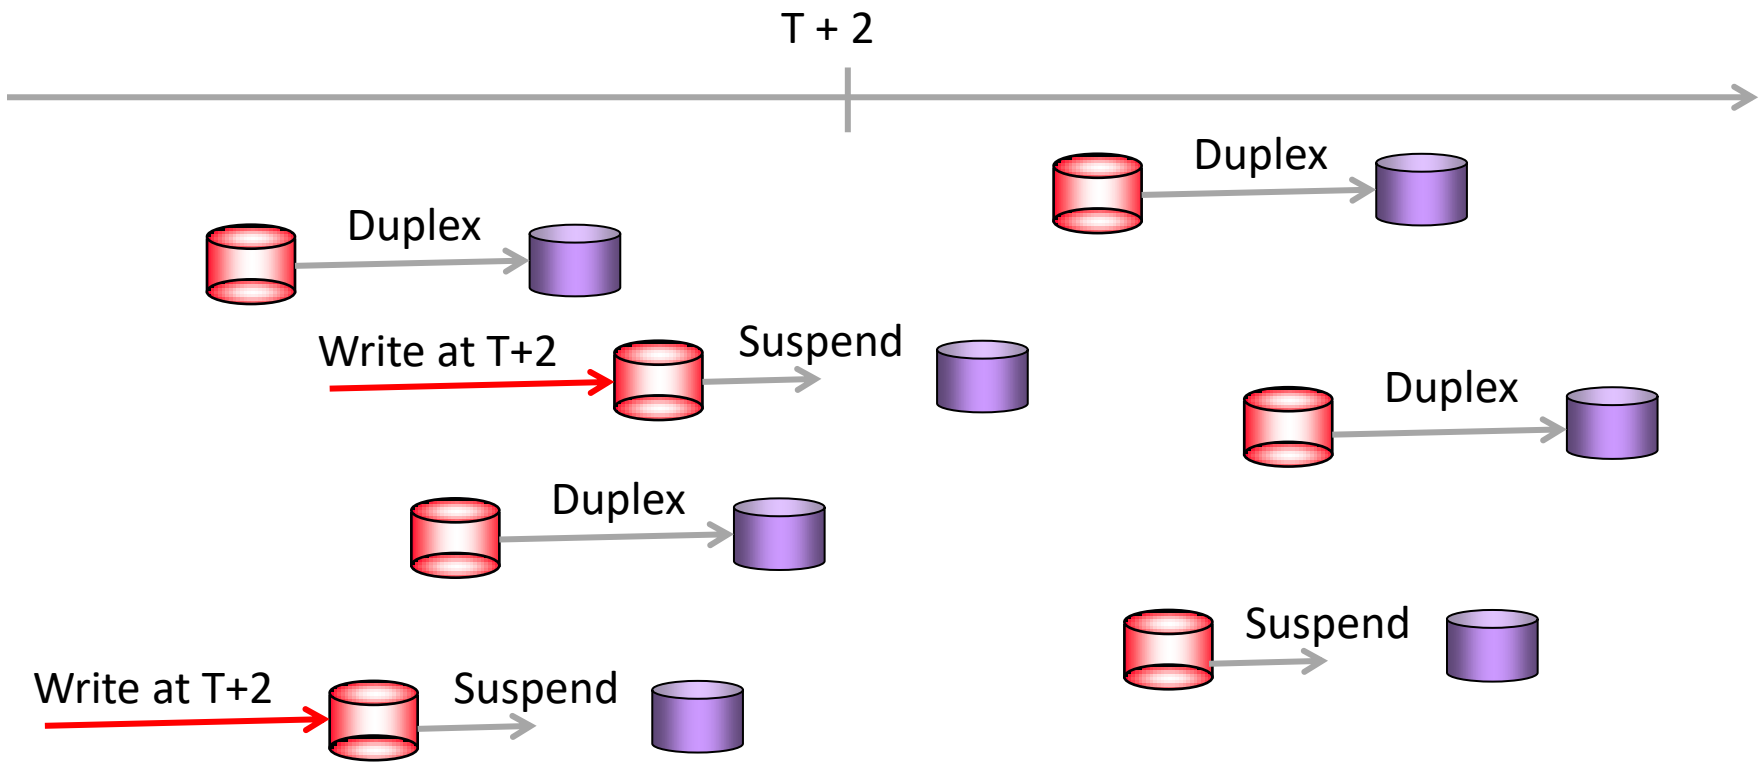

Introduction : ShadowImage

- ShadowImage At Time Split :
  - Suspend of each SI pair based on timestamp
  - Allow multi-controller consistency