



IBM Systems Group

CSE For High Availability and System Management

Jay Brenneman
rjbrenn@us.ibm.com

Trademarks

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

DB2*	Resource Link
DB2 Connect	S/390*
DB2 Universal Database	Tivoli*
e-business logo*	Tivoli Storage Manager
e-business on demand	TotalStorage*
HiperSockets	WebSphere*
IBM*	z/OS*
IBM eServer	z/VM*
IBM logo*	zSeries*
IMS	

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

UNIX is a registered trademark of The Open Group in the United States and other countries.

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

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

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.


Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Agenda

- What does CSE do?
- What does CSE not do?
- How do I turn it on?
- How does this get me closer to High Availability?

z/VM without CSE



CSE allows you to:

- Manage only one VM Source Directory that is common to all nodes in Plex
- Extend the mini disk access control semantics across VM nodes in Plex
- Share Spool files across VM nodes in Plex
 - Practically impossible to retrofit onto existing VM systems
- Extend the query and messaging tools across VM nodes in Plex
- Share the RACF database

Ugly Bits: PVM and RSCS are not free...

- PVM is not included in the base cost of z/VM
 - Requires a Special Bid to get licensed on an IFL
 - needed for shared spooling
 - needed for cross system messaging, query
- RSCS is needed to support DirMaint communications if shared spool is not used
- DirMaint is required to support the single source directory

CSE does ****NOT**** allow you to:

- Share VM SysRes volumes
 - Each VM system must maintain its own object directory
- Get High Availability for free
 - Some infrastructure will have to be built
- Share SFS pools across VM systems
 - Need VTAM: TSAF, CS Collection, etc
- Virtual Reserve/ Release across VM systems

Enabling CSE – The SYSTEM CONFIG file:

- XLINK_System_Include Slot 1 LTICVM1
XLINK_System_Include Slot 2 RESRVED2
XLINK_System_Include Slot 3 RESRVED3
XLINK_System_Include Slot 4 LTICVM4
XLINK_System_Include Slot 5 LTICVM5
- XLINK_Volume_Include VMP*
- XSPPOOL_SYSTEM Slot 1 LTICVM1 Share_Spool NO
XSPPOOL_SYSTEM Slot 2 RESERVED
XSPPOOL_SYSTEM Slot 3 RESERVED
XSPPOOL_SYSTEM Slot 4 LTICVM4 Share_Spool NO
XSPPOOL_SYSTEM Slot 5 LTICVM5 Share_Spool NO
- XSPPOOL_XLIST_OUTPUT, XSPPOOL_XLIST_INPUT
 - Users who can be logged on multiple systems at once: MAINT, OPERATOR...
- System_Identifier 2064 %0nnnn LTICVM1

RACF stuff

- Change the RACF database from minis to dedicated volumes that support Hardware Reserve/ Release
 - DDR from the minis to full volumes on VM1
 - Change the directory entry and recycle RACF on VM1
 - Add VM4 missing entries to the database
 - Change the directory entry and recycle RACF on VM4
- That's pretty much it – do the same for VM5

RSCS stuff

- RSCS used for DirMaint message passing if not using Shared Spool
- 'CP ATTACH C315 RSCS C31 ' /* Link to LTICVM4 */
 'CP ATTACH C215 RSCS C21 ' /* Link to LTICVM5 */
- LINKDEFINE LTICVM4 AST TYPE NJE LINE C31
 LINKDEFINE LTICVM5 AST TYPE NJE LINE C21
- 'RSCS START LTICVM4'
 'RSCS START LTICVM5'

PVM stuff

- Enables system to system messaging for
 - Indicate
 - Q Names
 - SMSG
- 'CP ATT C314 PVM C31' /* LINK TO LTICVM4 */
'CP ATT C214 PVM C21' /* LINK TO LTICVM5 */
- LOCAL LTICVM1
 LINK C31 LTICVM4 CTCA
 LINK C21 LTICVM5 CTCA
- START LINE C31
 START LINE C21
- SMSG PVM START CSECOM LTICVM4
 SMSG PVM START CSECOM LTICVM5

Format a volume to support XLINK

- Rolling IPL to pick up SYSTEM CONFIG changes
- Make sure “q n” shows you everyone logged on everywhere – indicates CSE messaging is up
- Attach the volume to your id
- XLINK FORMAT <vaddr> <volid>
 - Defaults to adding CSE track in CYL 0
- Attach the volume to SYSTEM on all VMs
- XLINK CHECK <volid>
 - Volume <volid> is controlled by CSE LINK.

Dirmaint Stuff

- Dirmaint will run on one node in the plex
- DirmSats will run on all other nodes in the plex
- Directory changes are made everywhere.
 - This can take some stern user re-education
- Dirmaint must be equal to or higher than the Dirmsats in code level

Directory Stuff

- Merging the directory is the hardest part of implementing CSE
- On VM1:
 - Enable dirmaint, make sure the lock disk (15D) is defined on XLINK controlled DASD
 - Add the system affinity information for both VM1 and VM4 to the DIRECTORY control statement
 - Use the CNSLDIR tool to sysaffin-ize the guests and profiles common to both VM1 and VM4
 - Check the guest and profile DIRECT files to make sure the tool worked
 - Check again... The tool makes lots of assumptions...
 - Once the sysaffin-ing is done correctly, commit the changes


More Directory Stuff

- Once VM1 has absorbed VM4's uniqueness into the collective, on VM1:
 - Using dirmsat as a template, create a new dirmsat user to run on VM4 – say dirmsat4
 - Add RACF privileges for dirmsat4
 - Run the BLDCSEDR tool to consolidate the dirmaint files into a CSE DIRECT file
 - send CSE DIRECT to VM4

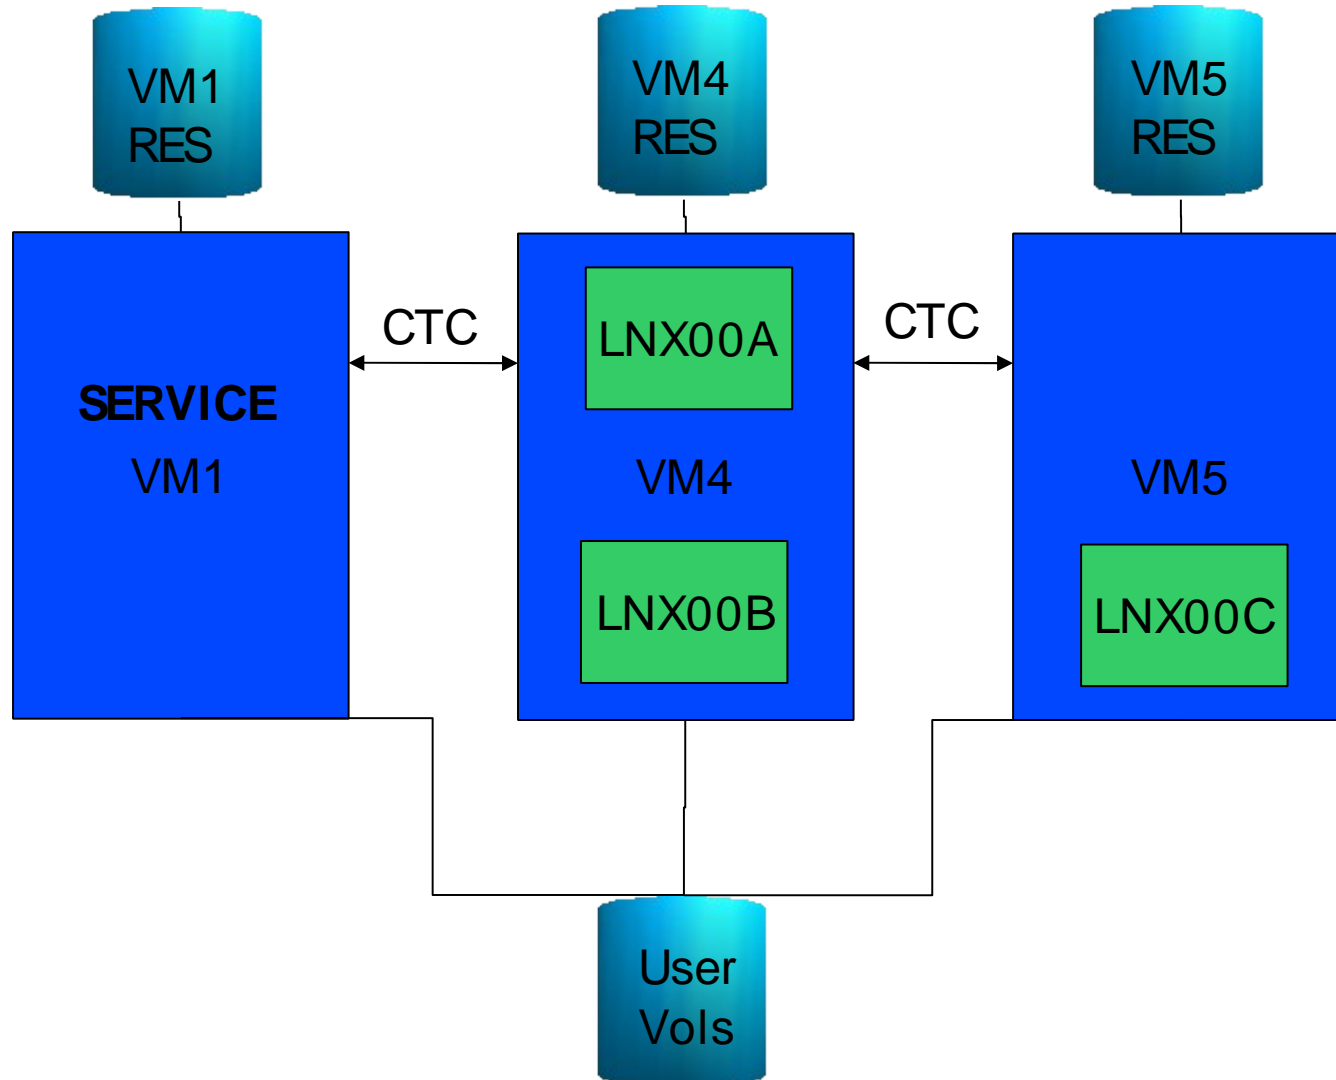
Yet More Directory Stuff

- On VM4:
 - Rebuild the directory with directxa using the CSE DIRECT file from VM1
 - Xautolog dirmsat4
 - Add the RACF privileges you forgot about when you created it
- Add FROM= TO= and SATELLITE_SERVER= statements to Dirmaint Config
- Force and restart dirmaint on VM1 and dirmsat4 on VM4
- Enjoy the utopia of CSE enabled VM.
- Repeat for VM5

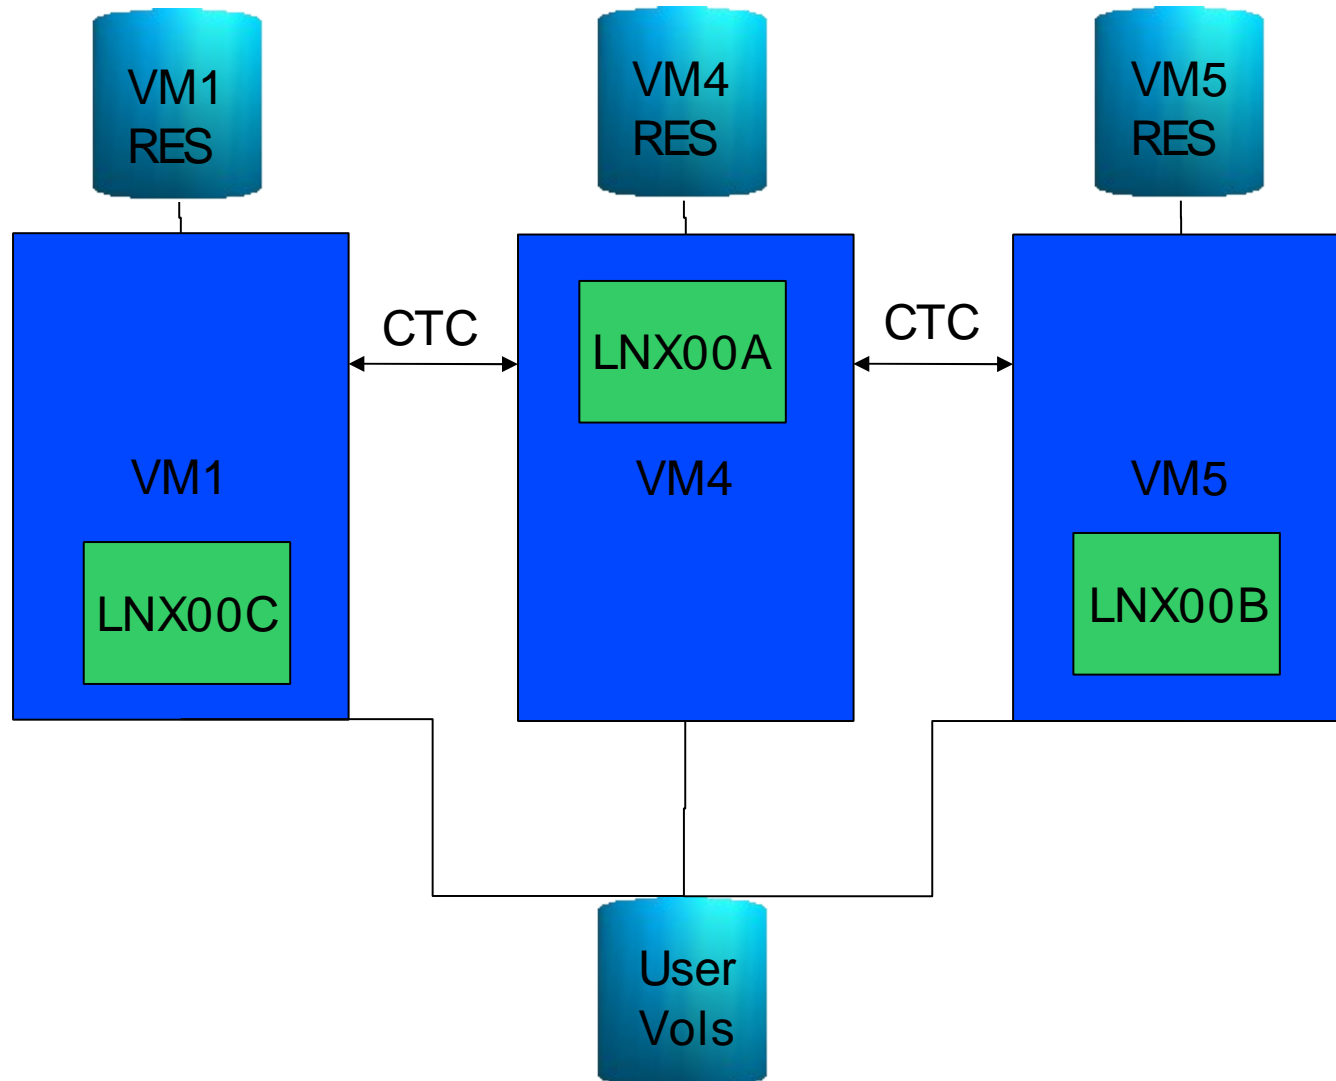
z/VM with CSE



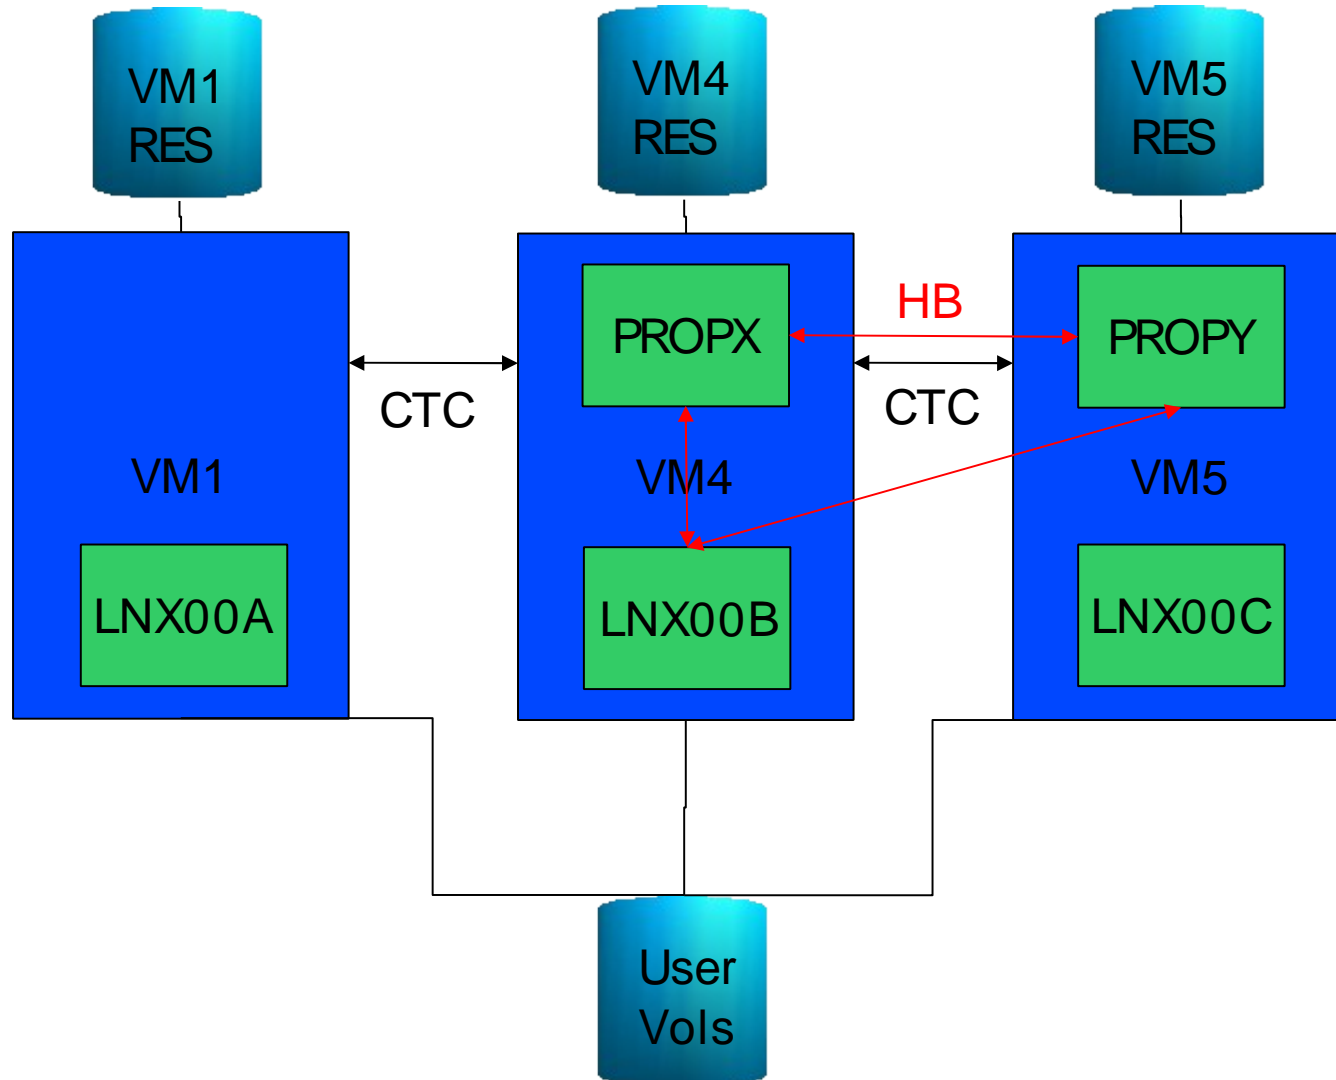
z/VM with CSE



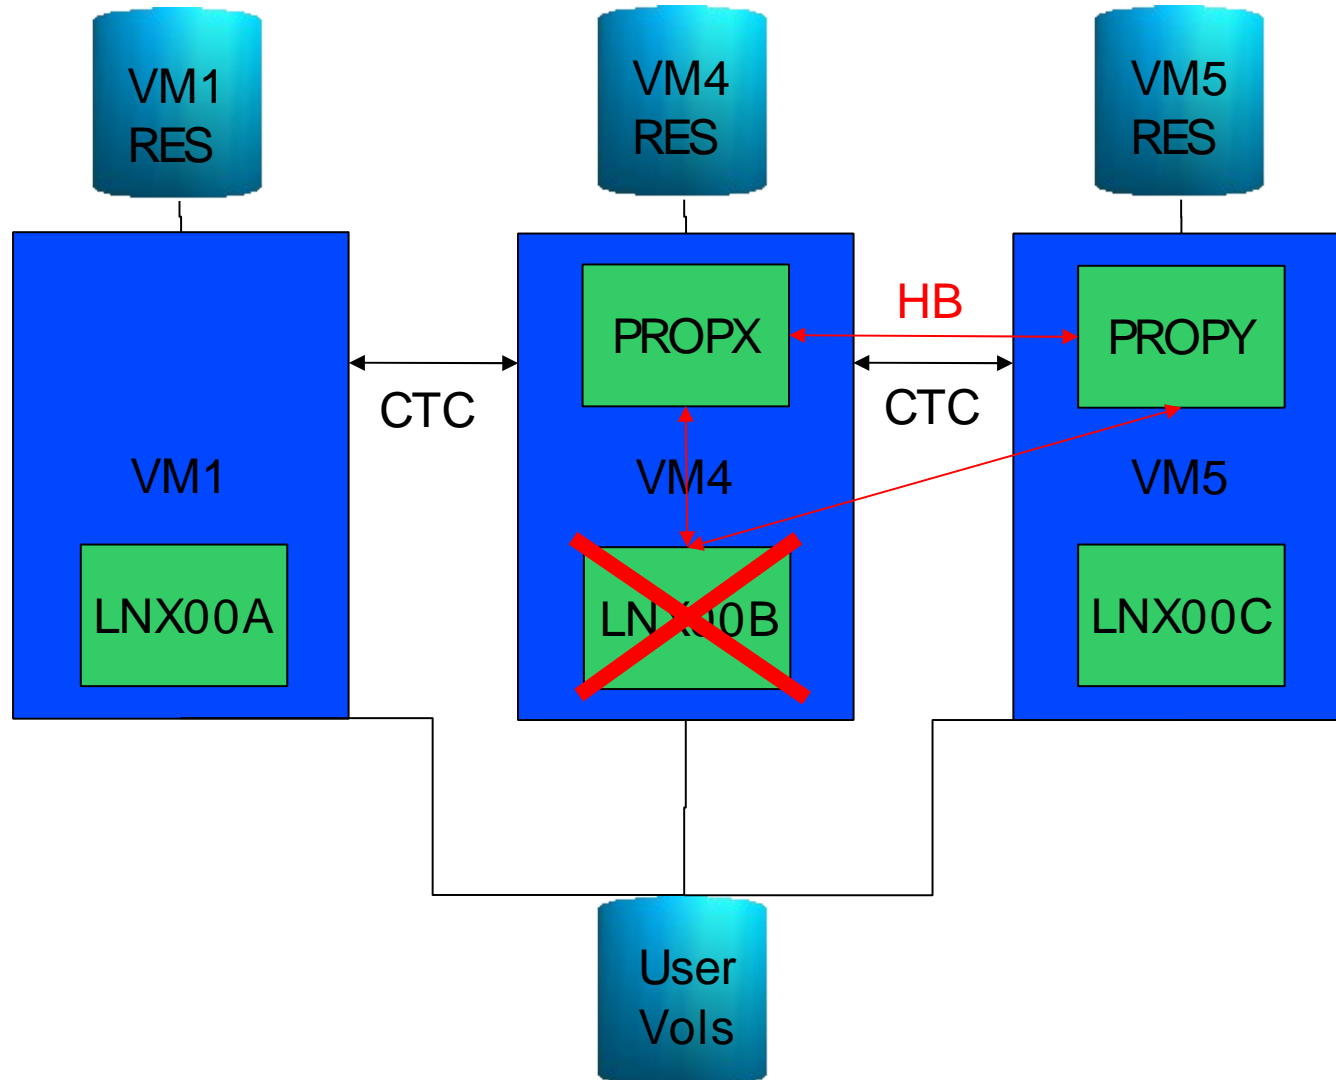
z/VM with CSE



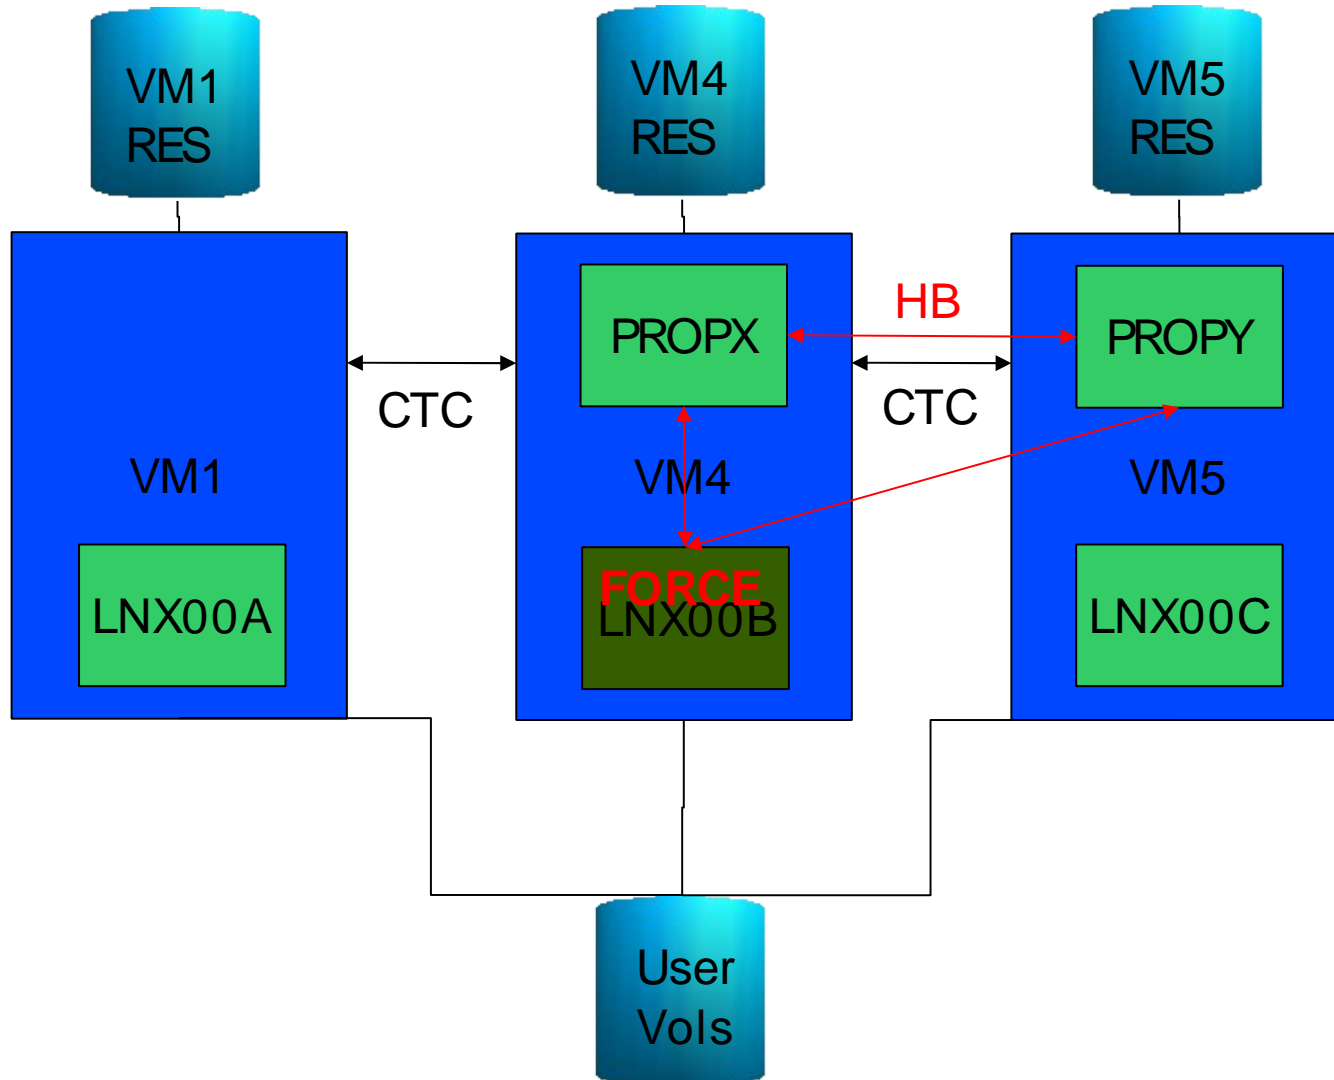
High Availability



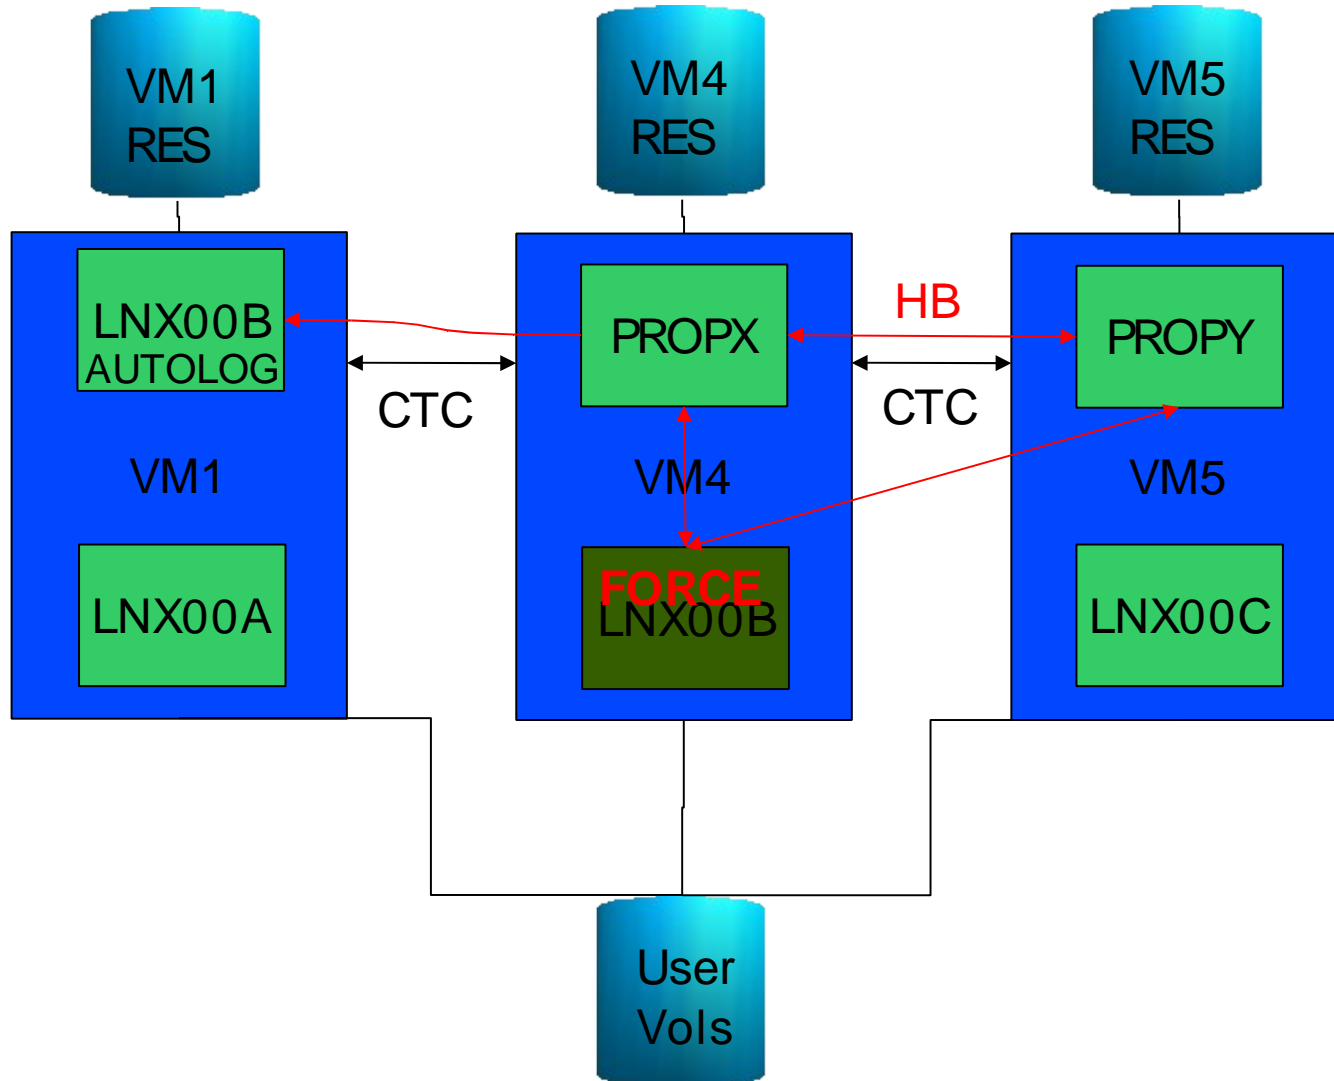
High Availability



High Availability



High Availability



Summary

- Maintaining Multiple VM systems is easier with CSE
- CSE allows greater flexibility in choosing where to run a workload
- CSE provides the infrastructure needed for HA