

Tech Note: Managing IBM z BladeCenter Extension (zBX) Storage Resources

Gary Fisher
ultimate@us.ibm.com

Abstract

A “dual-use” problem can occur on the zBX when removing storage resources. This document describes the problem, how it is caused, and how to avoid it.

Table of Contents

Abstract.....	1
Table of Contents.....	2
Overview.....	3
Background.....	3
Detailed Description of the Problem.....	4
Connectivity Environment.....	5
“Dual-use” storage problem – Error Case Described in Detail.....	6
Basic Add, Remove, and Add Storage Resource scenario.....	6
Error Case: Remove and Add Storage Resource scenario.....	12
Best Practice Procedures.....	17
Best Practices: Procedure to Re-boot a Virtual Server.....	17
Best Practices: Procedure to Migrate Virtual Server.....	17
Best Practices: Procedure to Remove Storage Resources.....	19
Best Practices: Procedure to Add Storage Resources.....	19
Procedural Recommendations.....	23
Summary.....	24
Appendix A: Storage Concepts and Terms:.....	25
Hierarchy of storage virtualization.....	25
References.....	27

Overview

It is possible to have an IBM z BladeCenter® Extension (zBX) storage resource accessed with two different identifiers, which can cause two or more virtual servers (e.g. Linux® or Microsoft® Windows® running on zBX blades) to access the same storage resource when they intend to have private and separate storage resources.

There are several possible symptoms depending on the order of removing and adding storage resources:

- Two virtual servers accessing the same storage resource
 - A virtual server's storage is overwritten by another virtual server, even though the storage resources are supposed to be different
 - A virtual server can access and read the storage resource of another virtual server
- A virtual server's storage resource is not assigned to a volume group but the ensemble continues to discover it

This document explains the circumstances where this problem can occur, describes procedures for adding and removing storage resources, and migrating virtual servers, none of which have been described before in any IBM publication.

Background

An ensemble is a collection of highly virtualized heterogeneous systems that can be managed as a single logical entity and where heterogeneous workloads can be deployed. Some characteristics of an ensemble are:

- A IBM z Systems® Central Processing Complex (CPC), with or without an attached zBX, is called a *node*.
- One ensemble can contain from one to eight nodes.
- A single node can be a member of only one ensemble.
- It includes dedicated integrated networks for management and data.

An ensemble is managed as a single logical virtualized system by the IBM z Unified Resource Manager, also called zManager, through the use of a Hardware Management Console (HMC). The HMC is used to create and manage the resources of the ensemble.

IBM z Systems is a hardware infrastructure which is managed as a single logical virtualized system by the zManager. It can consist of any of

- An IBM zEnterprise® 196 (z196) (2817), IBM zEnterprise 114 (z114) (2818), and an attached IBM z BladeCenter Extension Model 002 (2458-002)
- An IBM zEnterprise EC12 (zEC12) (2827), IBM zEnterprise BC12 (zBC12) (2828), and an attached IBM z BladeCenter Extension Model 003 (2458-003)
- An IBM z13™ (z13) or IBM z13s™ (z13s)
- AN IBM z BladeCenter Extension Model 004 (2458-004)

This technote describes a recently discovered issue with storage resources with a zBX (either a model 002, 003, or 004) using an IBM DS8000® class data storage device (DS8000) if “discover storage resources” is not performed after removing storage resources from both the ensemble and the DS8000 volume group. This

document describes how the problem is caused, and contains detailed procedures that keeps the problem from happening.

Detailed Description of the Problem

BladeCenter blades hold onto information, in their VIOS (Virtual I/O System) about storage resources even when all paths fail, on the assumption that the storage resource might be offline temporarily. Periodic automatic tests on failed paths will automatically recover any failed storage resources when they come back online. An external command, such as a “discover storage resources” procedure, will trigger the removal from the BladeCenter’s list of storage resources of any storage resources with all paths in a failed state. The zManager’s Storage Resource Manager running on the HMC gets its information about storage resources and paths from the BladeCenter, via the “Import” function. If a storage resource is removed from the HMC and the storage device, but no “discover” function is performed to remove it, the BladeCenter’s entry remains and can be re-imported.

Open system storage devices using SCSI or Fibre Channel Protocol (FCP) create a shortcut or “handle” to a specific volume identifier (VOL ID) in a Logical Subsystem (LSS) called a Logical Unit Number (LUN). The DS8000 class DASD devices, hereafter called DS8000, will re-use a LUN entry if a gap between successive entries is formed when a volume is removed from a volume group. If a LUN of a storage resource that was not deleted at the BladeCenter is re-used, the BladeCenter will consider this to be a re-activation of a temporarily offline resource, even though it is a different Logical Subsystem (LSS) and volume ID.

Thus, these three operations must be performed to remove a storage resource from an ensemble:

1. “Remove storage resource” from the HMC
2. Remove volume from a volume group at the DS8000
3. “Discover storage resources” from the HMC to delete the entry from the BladeCenter

If the discover storage resources function (step 3 above) is not performed after a storage resource is removed from the HMC (step 1 above) and the volume group (step 2 above), the same LUN is used for the next volume that is added to the volume group and will be used as the *actual new* storage resource. However, the BladeCenter and the HMC will use the old and incorrect information about the *previously deleted* storage resource. Depending on subsequent storage actions (e.g. add, remove, or assign storage resources), some types of storage problems can occur:

1. Using a storage resource that is no longer assigned to the volume group for the z Systems
2. Different “dual-use” problems, where two or more virtual servers running in the ensemble can access the same storage resource when they think it is two separate storage resources.
3. Attempting to perform I/O to a volume that is not defined.

Connectivity Environment




Figure 1: Ensemble Connectivity

Many common, but not all, types of ensemble connectivity are shown in Figure 1: Ensemble Connectivity:

- The ensemble is controlled by the Hardware Management Console (HMC) with functions provided by the zManager. An ensemble is comprised of 1 to 8 of:
 - An IBM zEnterprise 196 (z196) (2817) or IBM zEnterprise 114 (z114) (2818) which can attach to 1 to 8 zBX model 002, 2458-002
 - An IBM zEnterprise EC12 (zEC12) (2827) or IBM zEnterprise BC12 (zBC12) (2828) which can attach to 1 to 8 zBX model 003, 2458-003
 - An IBM z13 (z13) or IBM z13s (z13s), standalone (not in picture)
 - The zBX model 004, 2458-004, is stand-alone and not attached to a processor
- Each BladeCenter has at least two adapters attached to Fibre Channel Protocol (FCP) Storage Area Network (SAN) switches.
- In this example, storage used by virtual servers running on zBX blades is provided by a single IBM DS8800 (2107). It has four adapters used for the zBX SAN access. Other adapters can be attached to other SANs or other open systems.

“Dual-use” storage problem – Error Case Described in Detail

The following storage resource “Basic” and “Error Case” scenarios use the configuration shown in Figure 2: Basic Ensemble Connectivity




Figure 2: Basic Ensemble Connectivity

Basic Add, Remove, and Add Storage Resource scenario

Figure 3: Ensemble Storage Resources: Step 1 - Add through Figure 12: Ensemble Storage Resources: Step 9 - re-import graphically depict the processes at the DS8000, zBX, and HMC for adding and removing storage resources. The initial volume and its associated information are outlined with heavy solid black lines. The second volume and its associated information are outlined with heavy dotted black lines. New information or actions for the current step are shown in yellow.

Note that there are 3 separate tables for storage resource information, as identified in Figure 3: Ensemble Storage Resources: Step 1 - Add and used in subsequent figures:

1. DS8000 Volume Group Properties
2. Blade Storage Resources
3. Ensemble Storage Resources

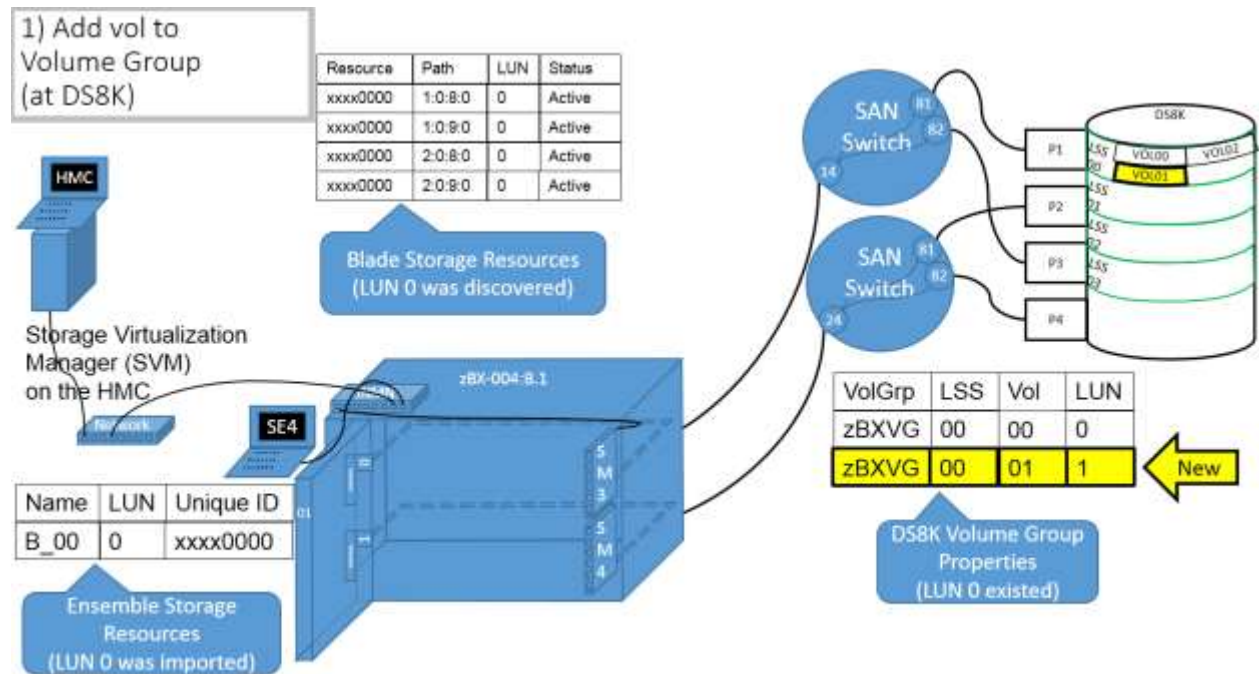


Figure 3: Ensemble Storage Resources: Step 1 - Add

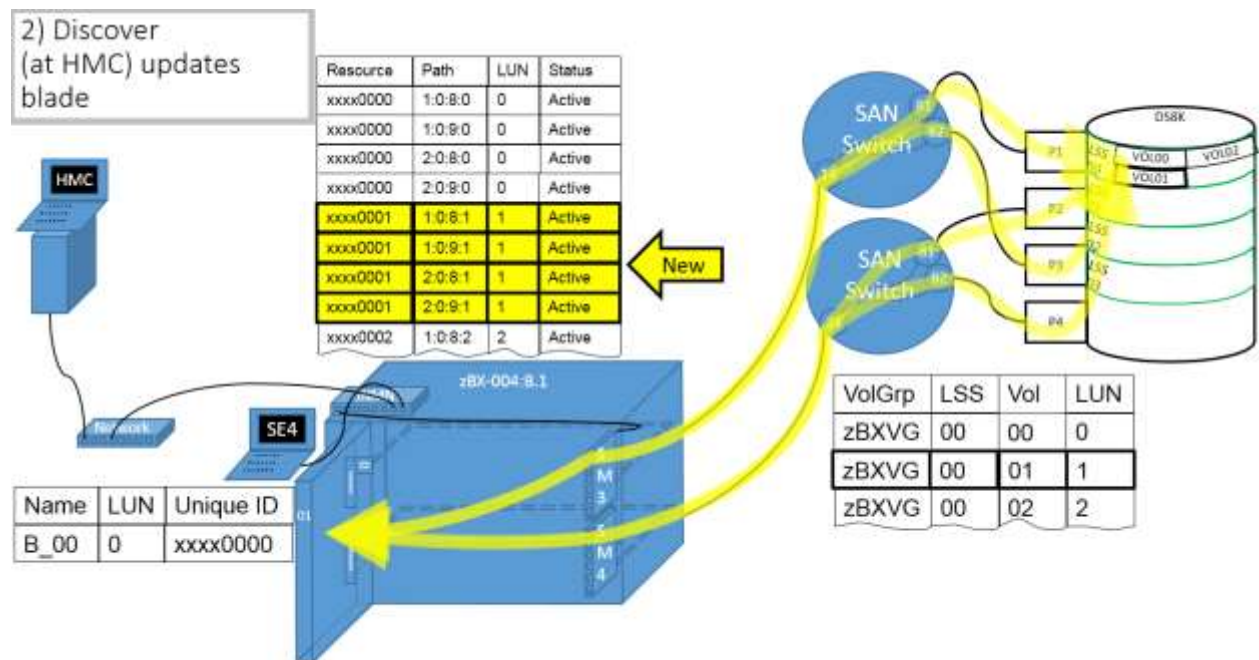


Figure 4: Ensemble Storage Resources: Step 2 - Discover

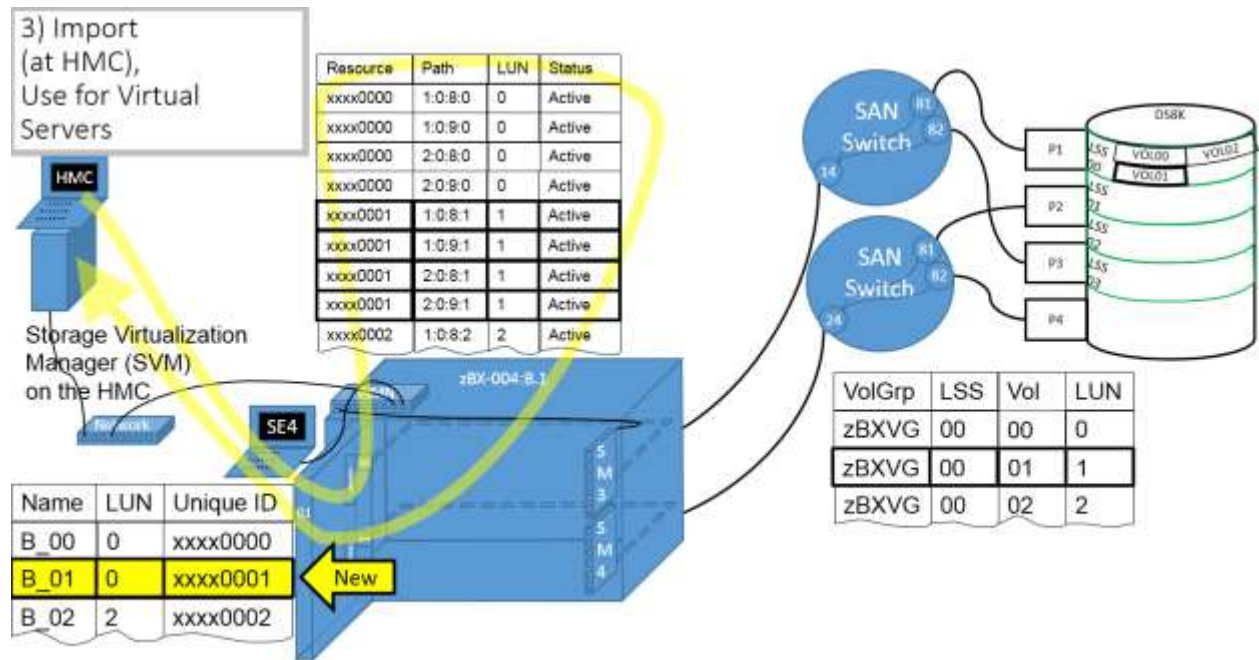


Figure 5: Ensemble Storage Resources: Step 3 - Import

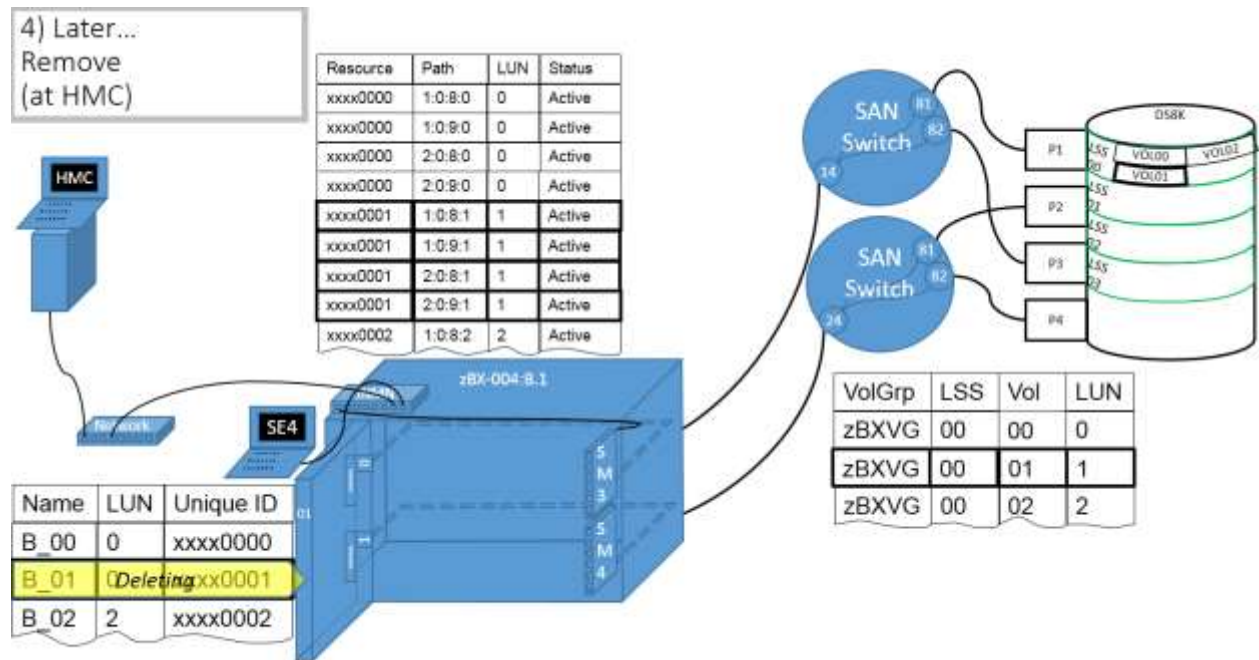


Figure 6: Ensemble Storage Resources: Step 4 - Remove at HMC

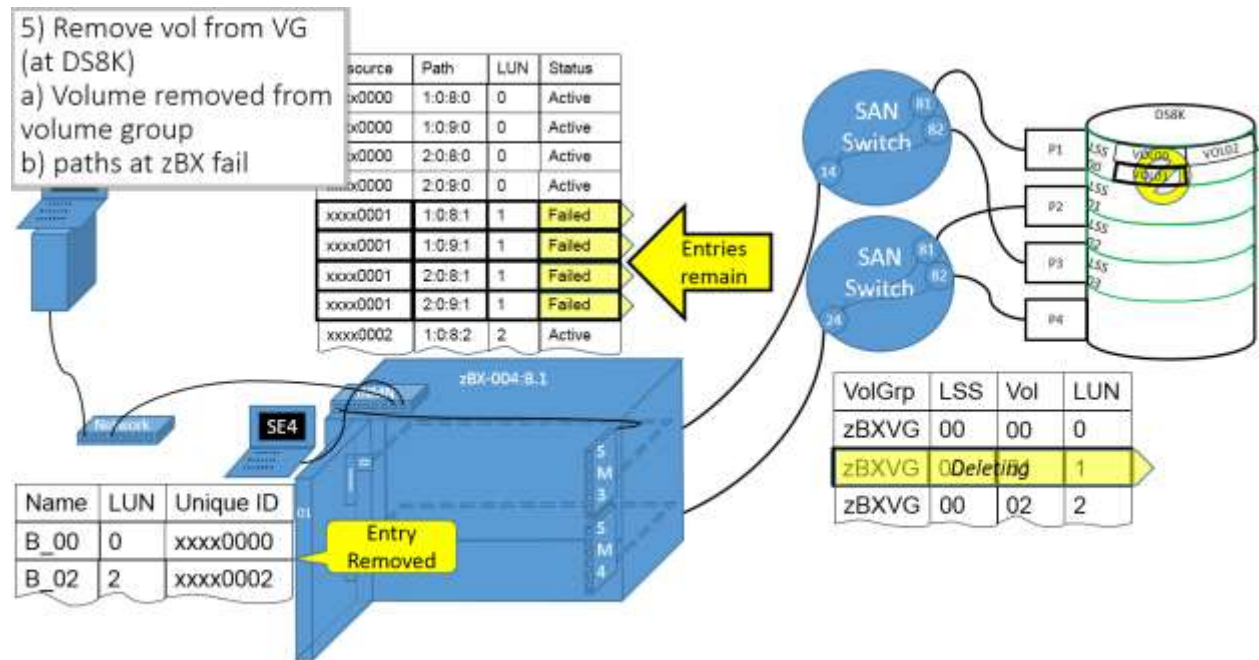


Figure 7: Ensemble Storage Resources: Step 5 - Remove at DS8000

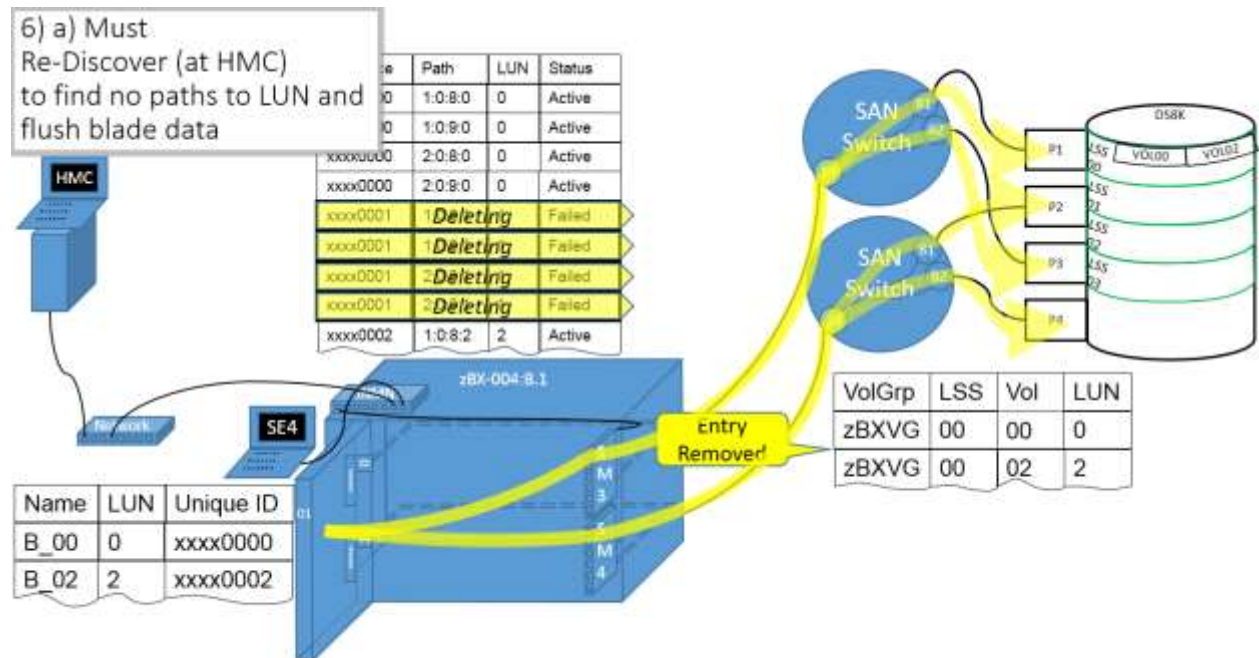


Figure 8: Ensemble Storage Resources: Step 6a - Re-Discover after Remove

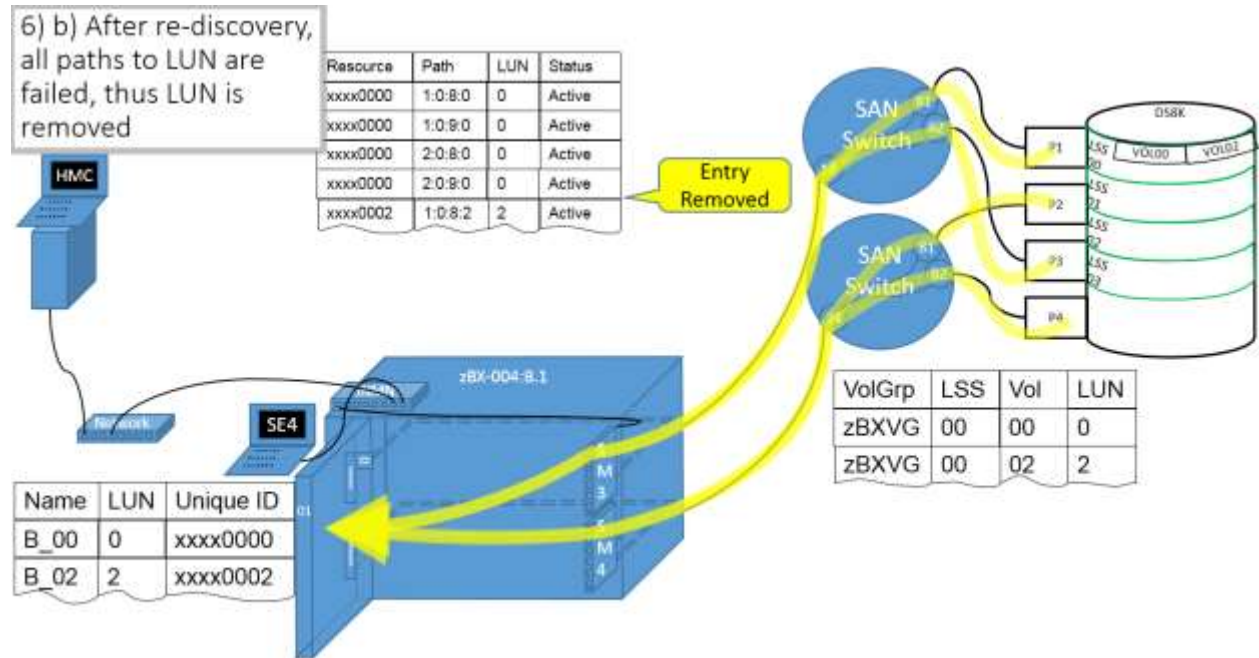


Figure 9: Ensemble Storage Resources: Step 6b - Finding all paths failed

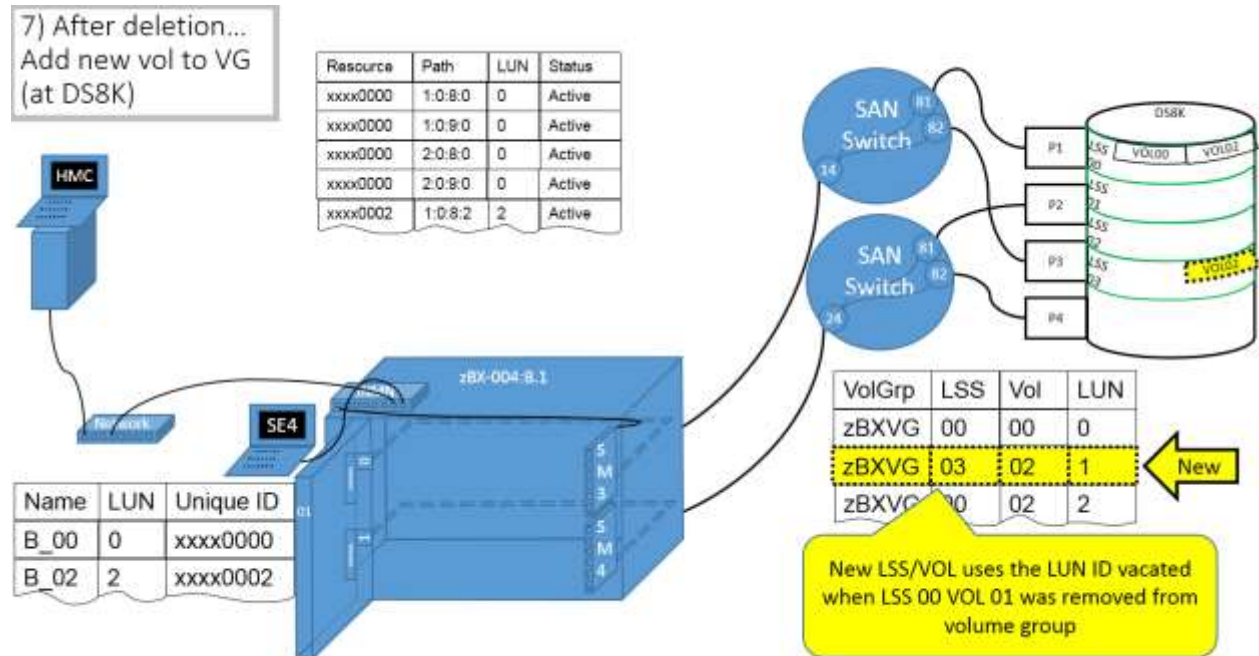


Figure 10: Ensemble Storage Resources: Step 7 - Add new

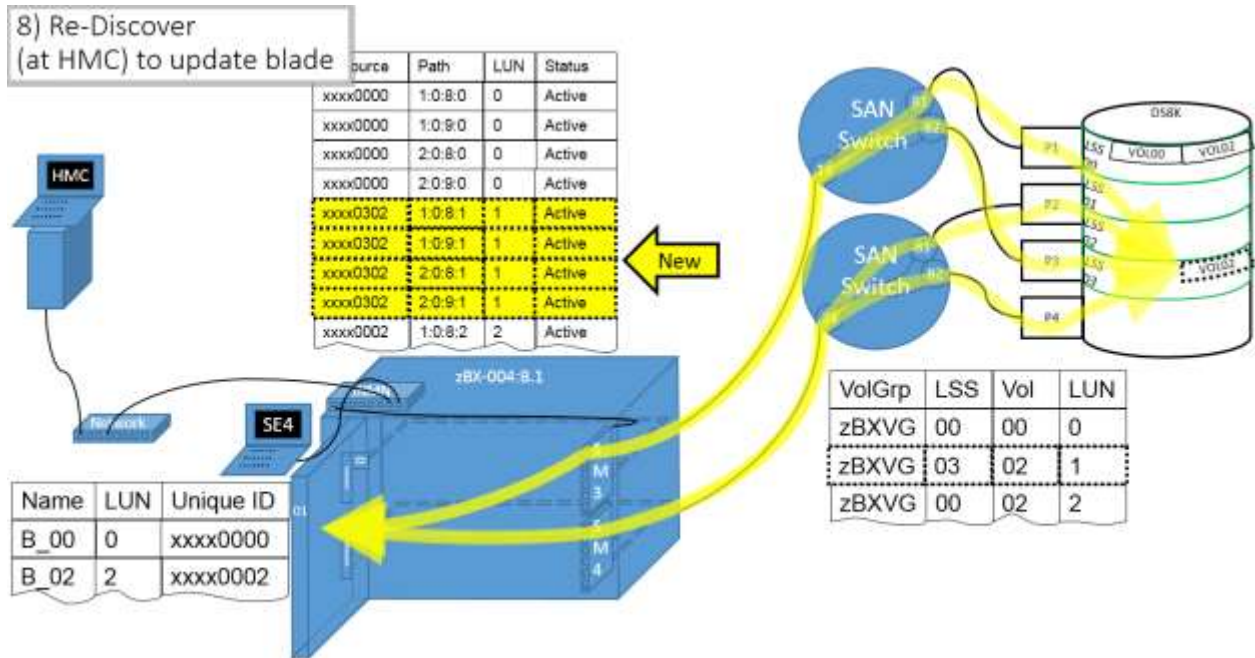


Figure 11: Ensemble Storage Resources: Step 8 - re-discover

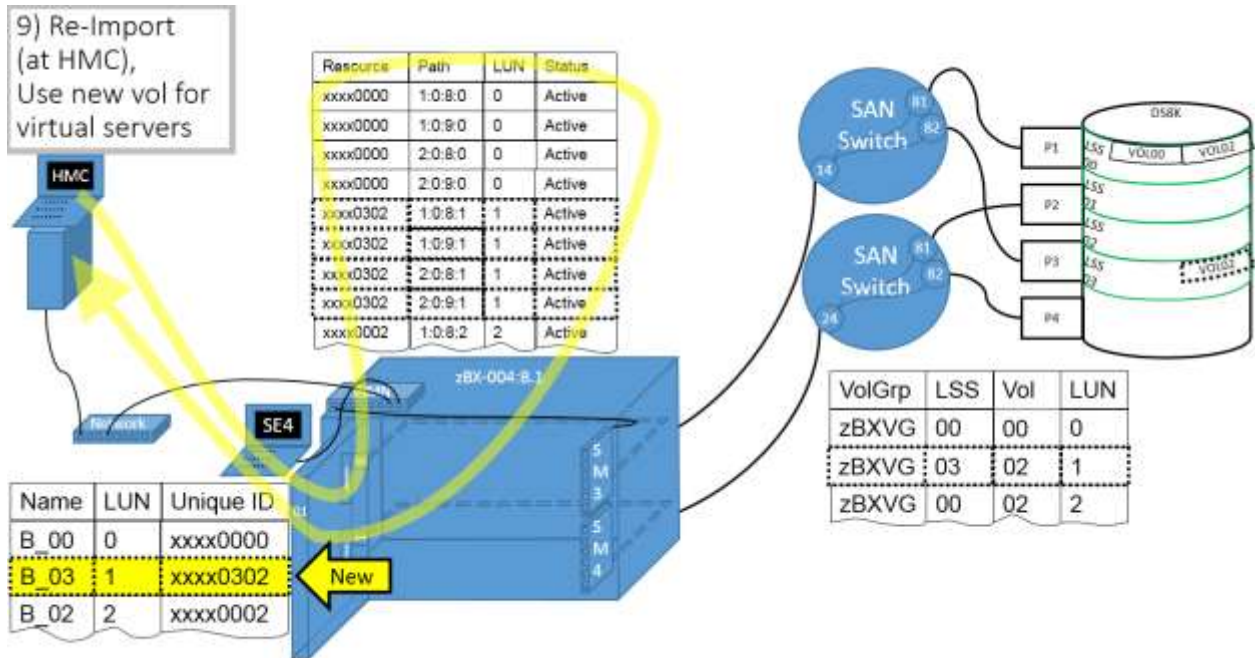


Figure 12: Ensemble Storage Resources: Step 9 - re-import

Error Case: Remove and Add Storage Resource scenario

Figure 13: Ensemble Storage Resources, Error Case: Step 1-5 through Figure 12: Ensemble Storage Resources: Step 9 - re-import graphically depict the processes at the DS8000, zBX, and HMC when step 6, re-discovery, is not performed, initiating the “dual-use” storage resource problem. The initial volume and its associated information are outlined with heavy solid black lines. The second volume and its associated information are outlined with heavy dotted black lines. New information for the current step is shown in yellow. Incorrect information is shown in red.

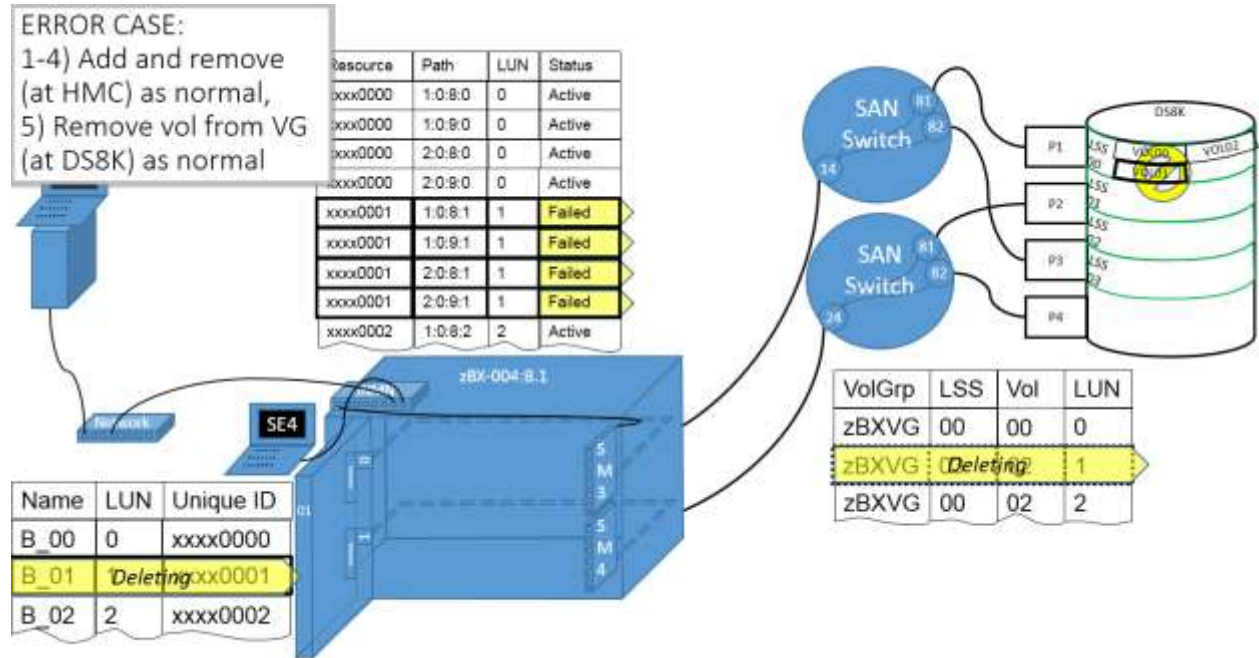


Figure 13: Ensemble Storage Resources, Error Case: Step 1-5

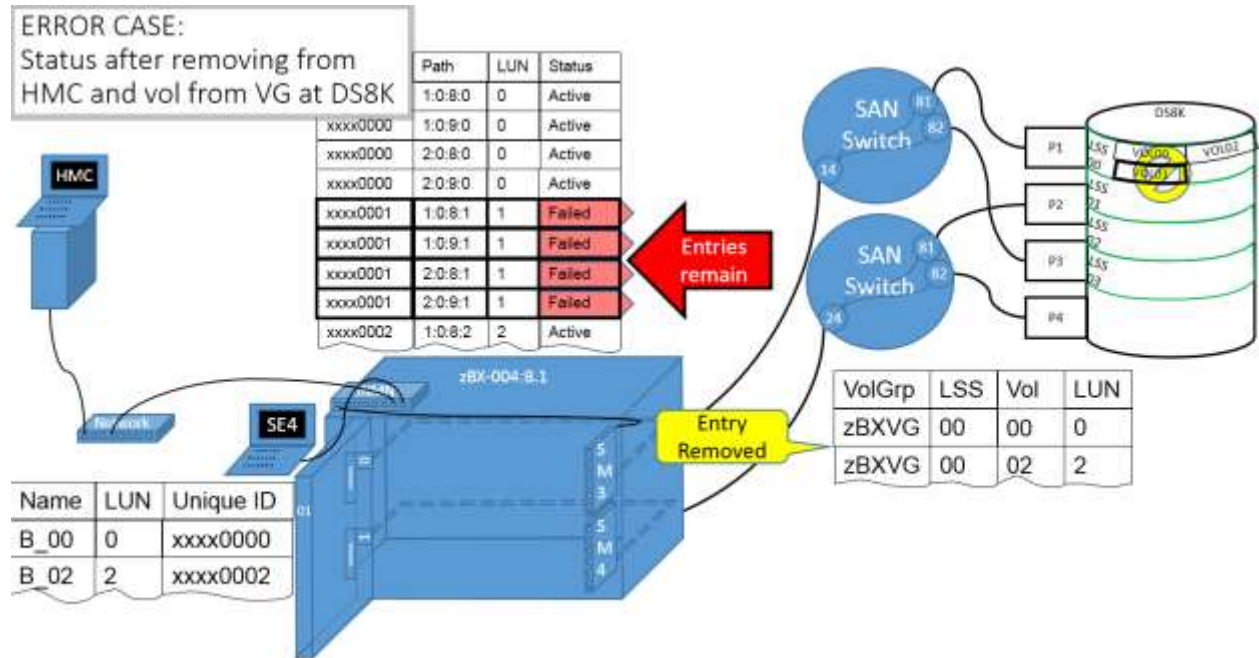


Figure 14: Ensemble Storage Resources, Error Case: Status after step 5

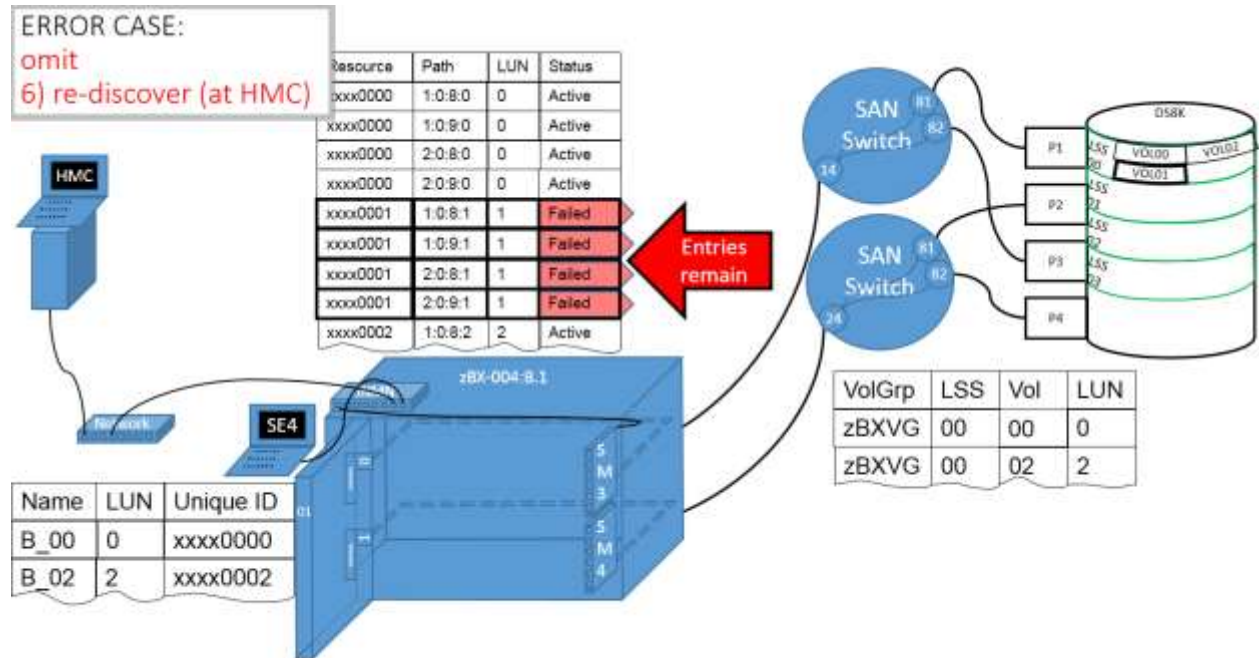


Figure 15: Ensemble Storage Resources, Error Case: Omitting step 6 (re-discovery)

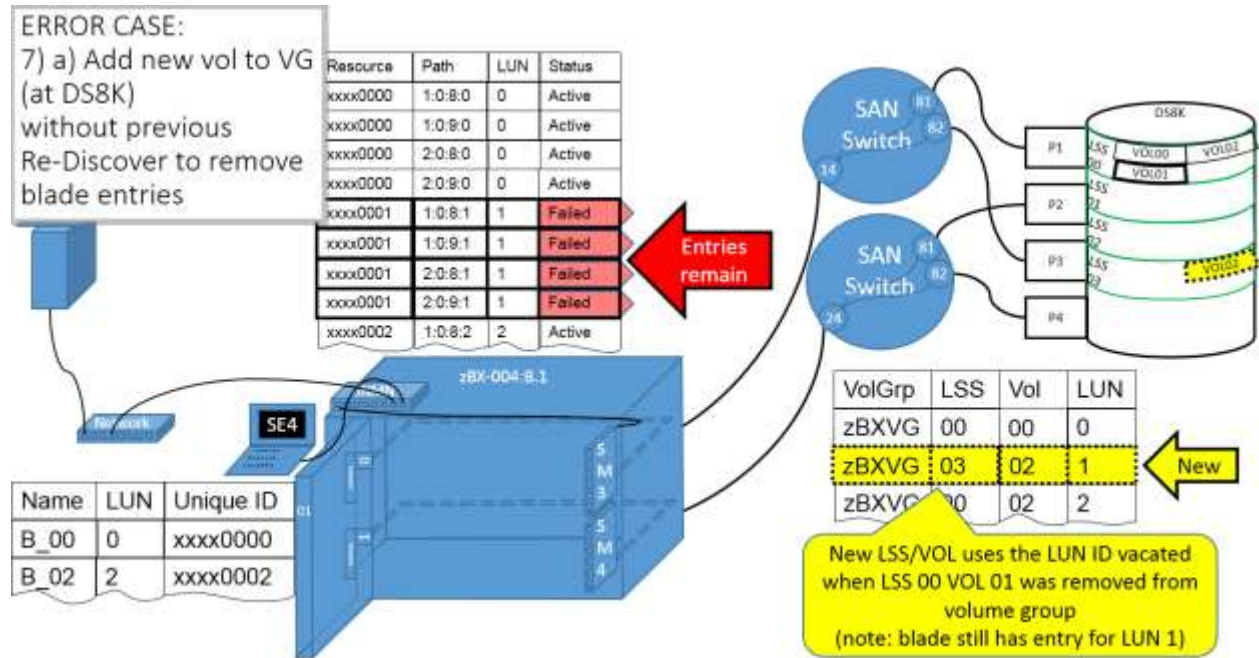


Figure 16: Ensemble Storage Resources, Error Case: Step 7a – Add new (after omitting step 6, re-discovery)

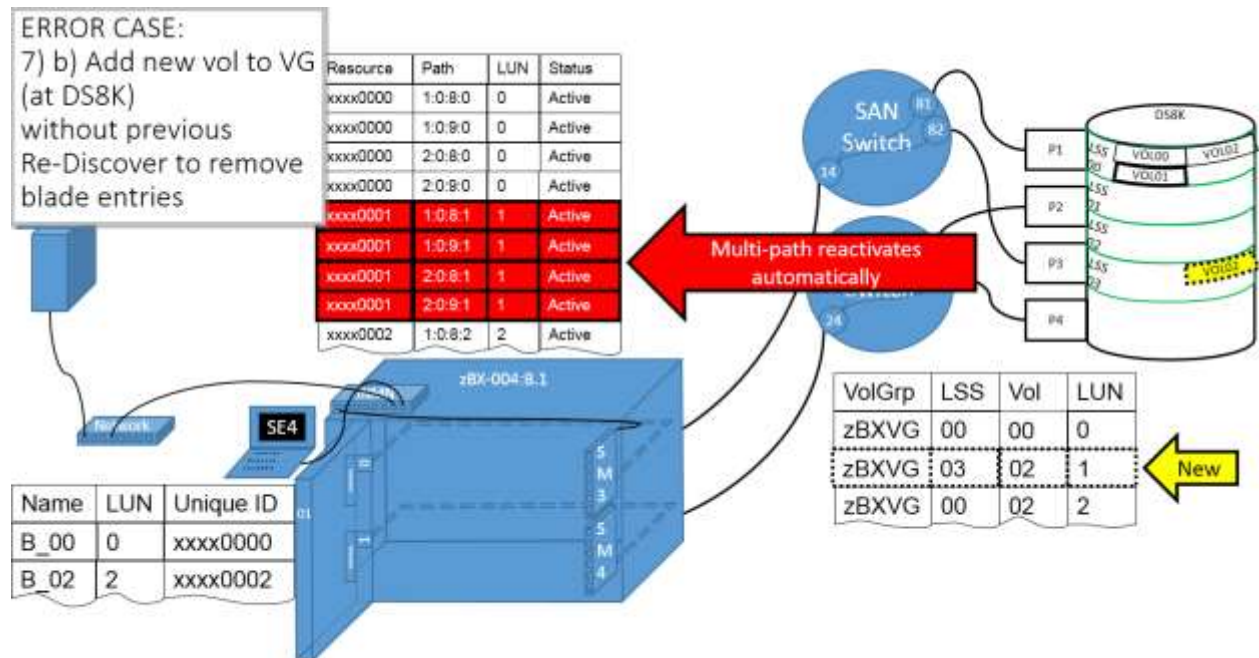


Figure 17: Ensemble Storage Resources, Error Case: Step 7b – Multi-path reactivate (after omitting step 6, re-discovery)

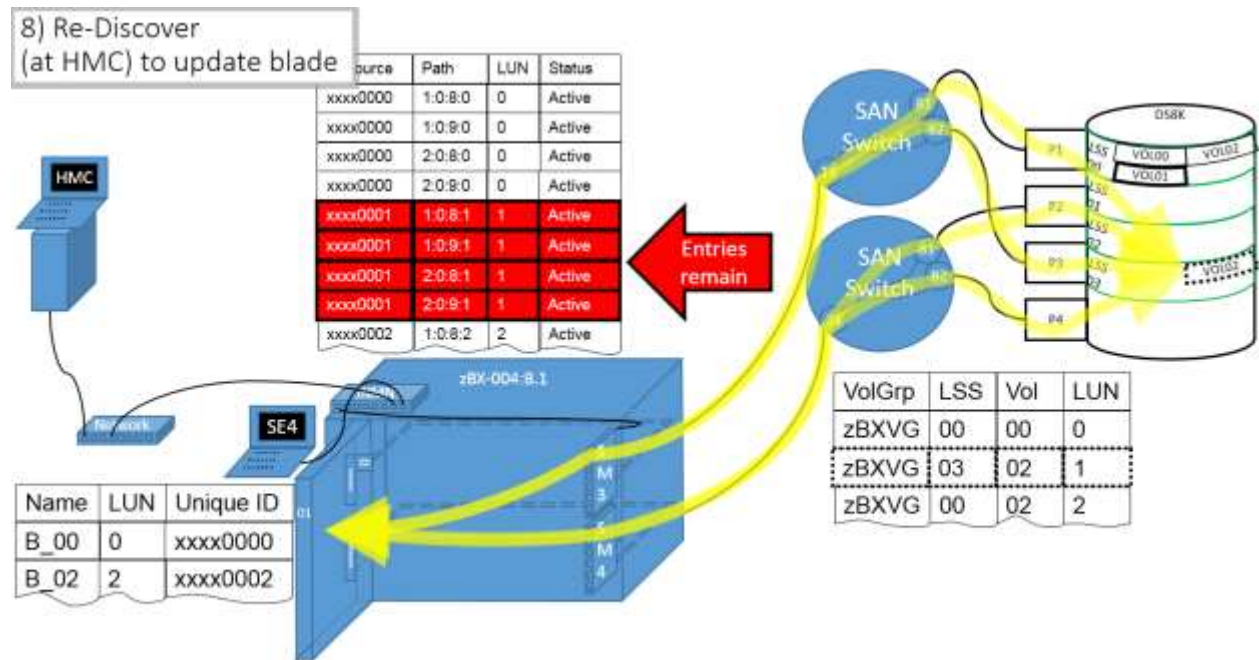


Figure 18: Ensemble Storage Resources, Error Case: Step 8 - Discover to update blade (after omitting step 6, re-discovery)

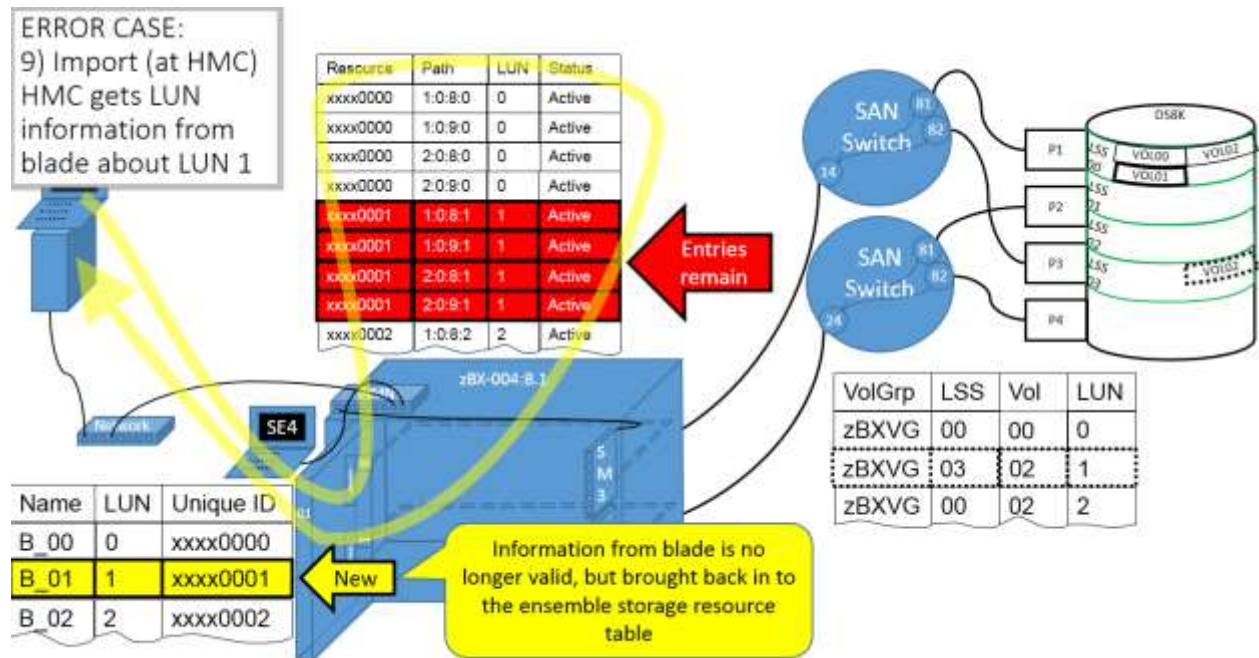


Figure 19: Ensemble Storage Resources, Error Case: Step 9 - Import (after omitting step 6, re-discovery)

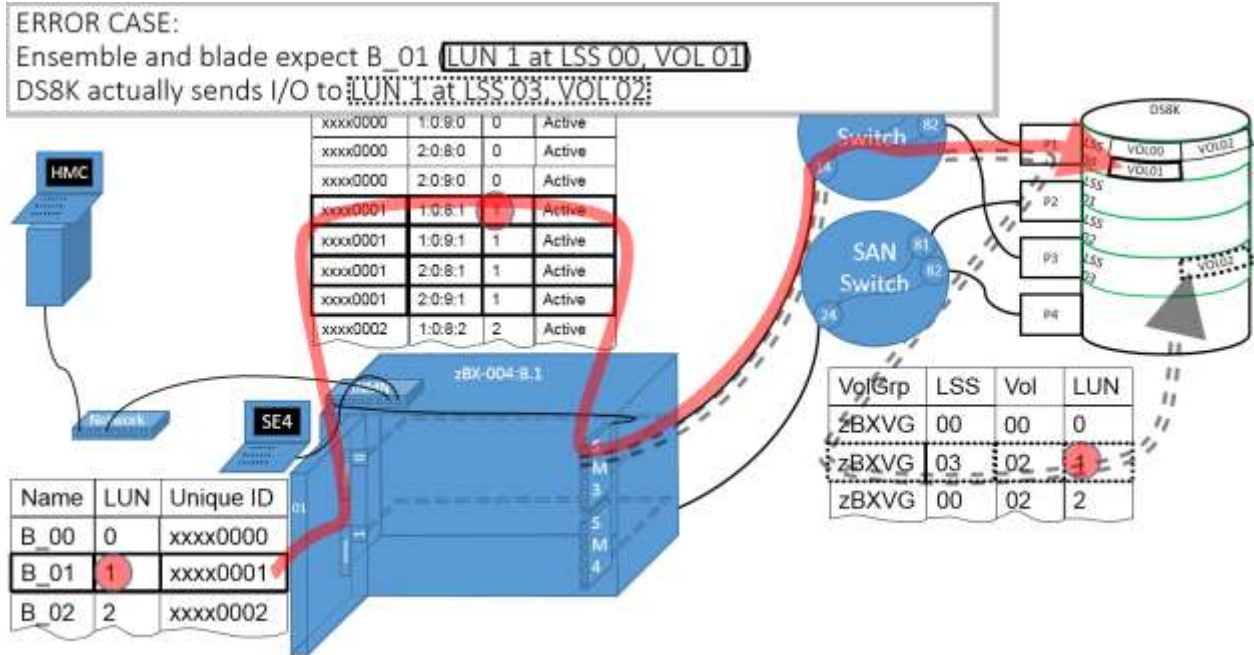


Figure 20: Ensemble Storage Resources, Status of Storage Resource Manager, blade data, and DS8000 Volume Group information

Best Practice Procedures

Best Practices: Procedure to Re-boot a Virtual Server

Some installations re-boot virtual servers periodically for maintenance that requires a re-boot, or to clean up various software leaks. One method is a software-only re-boot by using Linux's "reboot" or "shutdown -r" commands, or Windows Server's "shutdown /r" commands.

The suggested best practice is to also "bounce" the virtual server image with "deactivate" and "activate", which re-verifies the virtual server configuration including storage:

1. Shutdown the virtual server operating system without the restart option, which "hardens" any journaled I/O's that might be pending, performs a clean shutdown of all software servers and applications, etc. On Linux, this is "shutdown" or "shutdown -h" without the "-r". On Windows server, this is "shutdown /s" without the "/r" option.
2. Virtual servers on the HMC that remain in a "Stopped" state following the operating system shutdown will also require a deactivate of the virtual server from the HMC. Some operating system shutdowns bring the virtual server status to the "Not Operating" state so a "deactivate" is not required. Deactivate from the HMC verifies that the operating system is not running. Note that deactivate alone (without a software shutdown) is supposed to send a signal to the OS to do a software shutdown, but an independent first-step software shutdown is generally cleaner because the operator can see any messages or status at the OS screen.
3. Activate the virtual server from the HMC, which starts the OS.

Best Practices: Procedure to Migrate Virtual Server

zBX storage experts specify this procedure as a "best practice" for migrating virtual servers between different blades. The goal of this procedure is to help prevent creating out-of-sync storage resources/LUNs between the zManager and the customer's storage solution.

Use this procedure to migrate a virtual server defined to run on one blade to any other blade in the ensemble; the blades can be on different blade centers (e.g. B.1 to C.2) or different zBX's (e.g. 002 to 004).

1. Stop old (source) server
 - a. Deactivate the virtual server to be migrated
 - i. A software shutdown may be performed first, e.g. Linux's "shutdown -h" or Windows Server's "shutdown /s"
 - ii. An HMC deactivate may be done with or without a prior software shutdown, e.g. "deactivate server". An HMC deactivate is required if the virtual server status is "Stopped" following the software shutdown.
 - b. Wait until server shows "Not Operating" in the HMC
2. Add storage to new (target) server if target is in different volume group
 - a. On the DS8000 add the volume to the volume group assigned to the target blade.
The volume will now be a member of both volume groups assigned to the source and target blades.

Note: Implementation is of course storage dependent, however on the DS8000 you should be able to map the volume into two volume groups concurrently

- b. Use the HMC Manage Storage Resources task to discover storage resources for the target blade only to discover the new paths to the target blade's volume group
 - c. Add the new LUNs to the target blade
 - i. Select only the newly assigned LUN paths from the discovered list of paths
 - ii. Import them for the target blade
 - iii. Ensure the target has access to the shared volume
3. Migrate source to target
- a. Prepare the Virtual Server migration to the target blade using the HMC task
 - i. Select the virtual server
 - ii. Select Configuration
 - iii. Select Migrate Virtual Server
 1. New window to select destination from list of possible destination blades
 2. Select destination blade

Note: The target blade **will not appear** in the candidate list of potential blades to migrate to **until the storage resource(s) of the source are available to the target** (previous "Add Storage" step).

- b. Select Migrate
 - c. Wait until the migrate task is complete
4. Remove storage from old (source) server
- a. While the Virtual Server remains deactivated, remove the Storage Resource definition from the source blade using the HMC Manage Storage Resources task.

Perform this "**remove from HMC**" step **prior to deleting access** to the volume on the DS8000.

- b. On the DS8000 remove the volume from the source blade's volume group, which removes access to the volume from the blade where the old server ran
- c. Use the HMC Manage Storage Resources panel to discover storage resources for the source blade only

Note well: failure to perform this "re-discover" step will cause the "dual-use" problem

- d. Ensure the shared storage resource is no longer accessible to the source blade
 - i. Examine storage resources for the blade
5. Start new (target) server
- a. Activate the Virtual Server on the target blade
 - i. Activate the (new) virtual server on the new platform (i.e. the zBX-004)
 - ii. Watch it boot, either:
 - Open a GUI with the HMC
 - Use Remote Desktop
 - iii. Verify storage access
 1. Log into the server
 2. Verify storage

Best Practices: Procedure to Remove Storage Resources

HMC Discover Storage Resources should always be performed during removal of storage. The steps are as follows:

1. On the HMC, ensure the Storage Resource to be removed is not currently being used by any Virtual Server.
2. Use Manage Storage Resource task to remove the Storage Resource and all its defined paths from the blade(s).
3. On the DS8000, remove the Volume from the Volume Group(s) associated to the blade(s).
6. On the HMC, execute Discover Storage Resources and confirm the old LUN/UDID paths are no longer being discovered by the blade(s).

Note well: failure to perform this “re-discover” step will cause the “dual-use” problem

Suggestion: **remove only 1 or very few LUNs at any single time**, because determining that a path is unresponsive may take several seconds per path, and then associating all paths to a single LUN may take several more seconds per LUN. If many LUNs are removed at once, this operation might take a considerable amount of time (i.e. many minutes). If the discovery is performed immediately after removing the LUN, it might need to be repeated to allow the Blade(s) multipath daemon to process the path state changes and ensure a removed LUN is not re-discovered. If a removed LUN is re-discovered, there is the potential for “dual-use”.

Best Practices: Procedure to Add Storage Resources

HMC Discover Storage Resources should always be performed during addition of storage. The steps are as follows:

1. Add the Volume to the Volume Group(s) associated to the target blade(s).
2. On the HMC, execute Discover Storage Resources and confirm the new LUN/UDID paths are being discovered for the correct target blade(s).
3. Ensure the LSS and Volume ID on the DS8000, shown in Figure 21: DS8000's Volume Group Properties showing LSS and Volume ID, matches the last four digits of the Unique Device Identifier on the Discover Storage Resources panel, shown in Figure 22: Unified Resource Manager's Discover Storage Resources showing Unique Device Identifier.




Figure 21: DS8000's Volume Group Properties showing LSS and Volume ID

ypervisor	Host WWPN	Target WWPN	LUN	Defined	Unique Device Identifier
ATCH15K.C.2.09	21000024ff2759a0	500507630648d067	0004000000000000	Defined	scsi-36005076306ffd0670000000000001204
ATCH15K.C.2.09	21000024ff2759a1	500507630653d067	0004000000000000	Defined	scsi-36005076306ffd0670000000000001204
ATCH15K.C.2.09	21000024ff2759a1	500507630643d067	0004000000000000	Defined	scsi-36005076306ffd0670000000000001204
ATCH15K.C.2.09	21000024ff2759a0	500507630658d067	0007000000000000	Defined	scsi-36005076306ffd0670000000000001207
ATCH15K.C.2.09	21000024ff2759a1	500507630643d067	0007000000000000	Defined	scsi-36005076306ffd0670000000000001207
ATCH15K.C.2.09	21000024ff2759a0	500507630648d067	0007000000000000	Defined	scsi-36005076306ffd0670000000000001207
ATCH15K.C.2.09	21000024ff2759a1	500507630653d067	0007000000000000	Defined	scsi-36005076306ffd0670000000000001207

Figure 22: Unified Resource Manager's Discover Storage Resources showing Unique Device Identifier

- Ensure the LUN ID on the DS8000, shown in Figure 23: DS8000's Volume Group Properties showing LUN ID, matches the LUN value on the Discover Storage Resources panel, shown in Figure 24: Unified Resource Manager's Discover Storage Resources showing LUN. It is important to note on most, if not all, storage devices (including the DS8000) the LUN ID is *decimal*, while the Unified Resource Manager's LUN ID is *hexadecimal*. A conversion from decimal to hexadecimal is required (e.g. if the DS8000 LUN ID is "10" it will be "0A" at the HMC,

DS8000 "16" will be "10" at the HMC). Refer to Table 1: Hex to Decimal Conversion.

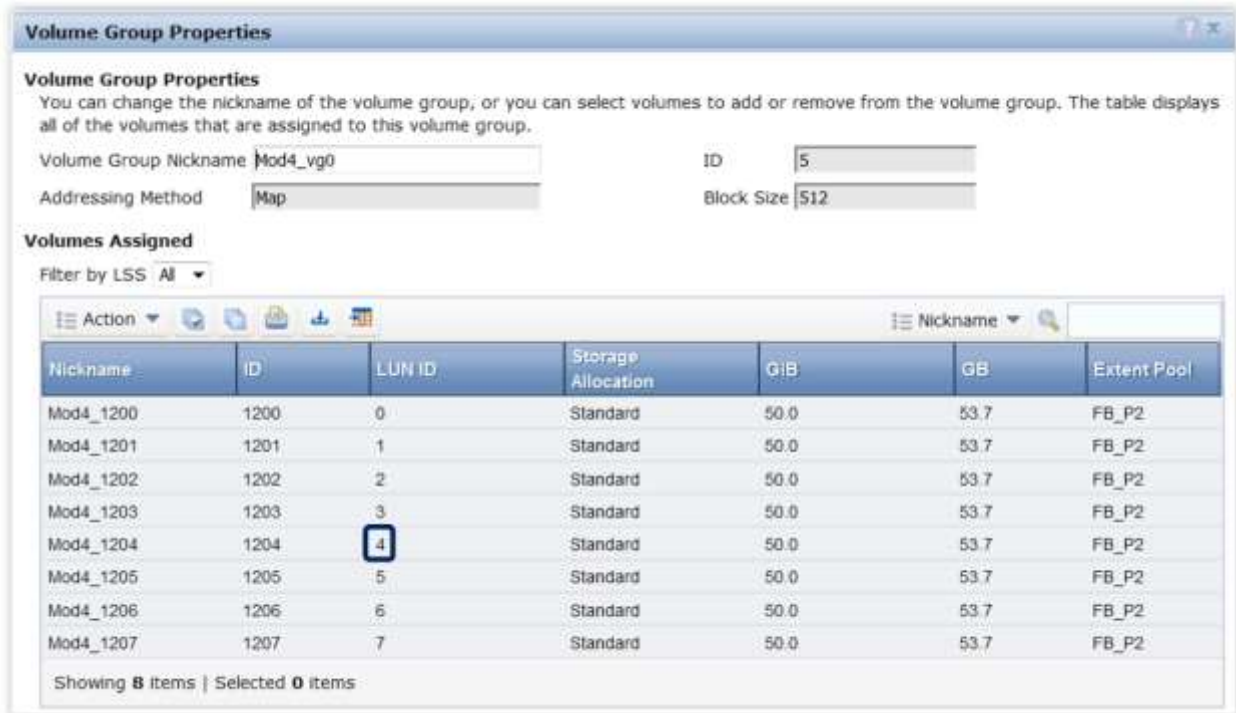


Figure 23: DS8000's Volume Group Properties showing LUN ID

ypervisor	Host WWPN	Target WWPN	LUN	Defined	Unique Device Identifier
ATCH15K.C.2.09	21000024ff2759a0	500507630648d067	0004000000000000	Defined	scsi-36005076306ffd0670000000000001204
ATCH15K.C.2.09	21000024ff2759a1	500507630653d067	0004000000000000	Defined	scsi-36005076306ffd0670000000000001204
ATCH15K.C.2.09	21000024ff2759a1	500507630643d067	0004000000000000	Defined	scsi-36005076306ffd0670000000000001204
ATCH15K.C.2.09	21000024ff2759a0	500507630658d067	0004000000000000	Defined	scsi-36005076306ffd0670000000000001204
ATCH15K.C.2.09	21000024ff2759a0	500507630658d067	0007000000000000	Defined	scsi-36005076306ffd0670000000000001207
ATCH15K.C.2.09	21000024ff2759a1	500507630643d067	0007000000000000	Defined	scsi-36005076306ffd0670000000000001207
ATCH15K.C.2.09	21000024ff2759a0	500507630648d067	0007000000000000	Defined	scsi-36005076306ffd0670000000000001207
ATCH15K.C.2.09	21000024ff2759a1	500507630653d067	0007000000000000	Defined	scsi-36005076306ffd0670000000000001207

Figure 24: Unified Resource Manager's Discover Storage Resources showing LUN

Table 1: Hex to Decimal Conversion

		2nd Hex															
		0x0	0x1	0x2	0x3	0x4	0x5	0x6	0x7	0x8	0x9	0xA	0xB	0xC	0xD	0xE	0xF
1st Hex																	
	_0	_1	_2	_3	_4	_5	_6	_7	_8	_9	_A	_B	_C	_D	_E	_F	
0x0	0_	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0x1	1_	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0x2	2_	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
0x3	3_	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
0x4	4_	64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
0x5	5_	80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
0x6	6_	96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
0x7	7_	112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
0x8	8_	128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
0x9	9_	144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
0xA	A_	160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
0xB	B_	176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
0xC	C_	192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
0xD	D_	208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
0xE	E_	224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
0xF	F_	240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255

5. If the Unique Device Identifier and LUN ID on the HMC match the DS8000, import the newly defined Storage Resource to the target blade(s).

Figure 25: Example showing an incorrect Unique Device Identifier shows a Unique Device Identifier that is imported from a deleted LUN whose information was retained at the blade because no re-discovery was performed. Note the Unique Device Identifier shows 1307, the LUN shows 0008, but the Volume Group information at the DS8000 shows that LUN 8 is ID 1306. Thus the Unified Resource Manager expects 1307 but I/O will actually go to 1306.

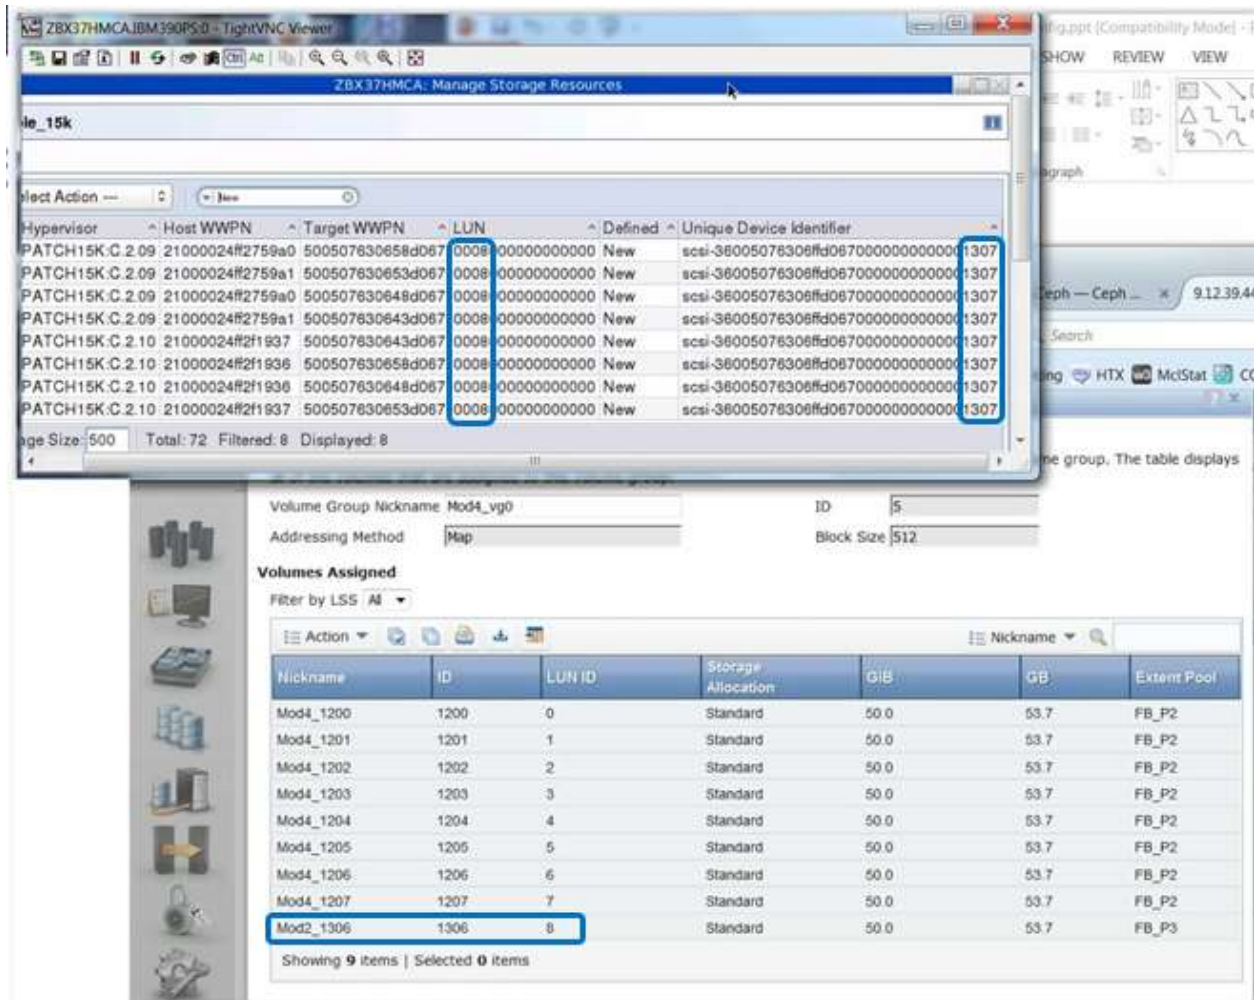


Figure 25: Example showing an incorrect Unique Device Identifier

Procedural Recommendations

Priority (H/M/L)	Effort (H/M/L)	Description
H	L	<p>Re-discover storage resources after deleting any storage resource from zBX access or removing any storage from a volume group</p> <p>Re-discovery (and finding the storage resource is removed from all paths) deletes the storage resource from the blade.</p> <p>Failure to re-discover will cause the “dual-use” storage resource problem.</p>

Summary

This has described the cause of the “dual use” problem found on the zBX. It also describes procedures from zBX storage experts for how to add and remove storage resources, how to reboot, and migrate virtual servers in such a way as to avoid the problem.

Appendix A: Storage Concepts and Terms:

Copied or paraphrased from IBM System Storage® DS8000: Architecture and Implementation:

Hierarchy of storage virtualization

Going through the virtualization hierarchy, we start with “just a bunch of disks” that are grouped in *array sites*. The array sites are created automatically when the disks are installed.

The next steps are initiated by a user:

- An *array site* is transformed into an *array*, with spare disks.
- The *array* is further transformed into a *rank* with formatted *extents*.
- The *extents* from selected ranks are added to an *Extent Pool*. The combined extents from those ranks in the Extent Pool are used for subsequent allocation to one or more *logical volumes*.
 - One or more ranks can be aggregated into an Extent Pool and logical volumes configured in that Extent Pool are not bound to any specific rank.
- Next, we create *logical volumes* within the *Extent Pools* (by default, striping the volumes), assigning them a *logical volume number* that determines which *logical subsystem* they would be associated with and which server would manage them.
 - A logical volume composed of fixed block extents is called a *LUN*. A *Logical subsystem* groups logical volumes (LUNs) into groups of up to 256 logical volumes.
 - Logical volumes (LUNs) can be defined in one or more volume groups. It allows a logical volume (LUN) to be shared by host HBAs configured to separate volume groups.
 - Different logical volumes (LUNs) on the same logical subsystem can be configured in separate Extent Pools.
- The *logical volumes (LUNs)* are then assigned to one or more *volume groups*. A *volume group* is a named construct that defines a set of logical volumes (LUNs).
 - You define the volume group by indicating which logical volumes (LUNs) are to be placed in the volume group. Logical volumes (LUNs) can be added to or removed from any volume group dynamically.
 - For each logical volume (LUN), you can choose an LSS. You can have up to 256 logical volumes (LUNs) in one LSS.
One restriction:
- Logical volumes (LUNs) are formed from a number of extents from an Extent Pool, which belong to one server (CEC), server 0 or server 1, respectively.
- LSSs also have an affinity to the servers. All even-numbered LSSs (X'00', X'02', X'04', up to X'FE') are handled by server 0 and all odd-numbered LSSs (X'01', X'03', X'05', up to X'FD') are handled by server 1.
- Finally, the *Host bus adapters (HBAs)* are configured into a *host attachment* that is associated with a *volume group*.
 - HBAs are identified to the DS8000 in a host attachment construct that specifies the HBAs' *World Wide Port Names (WWPNs)*.
 - A set of host ports can be associated through a port group attribute that allows a set of HBAs to be managed collectively. This port group is referred to as a *host attachment* within the GUI.
 - When used in conjunction with open systems hosts, a host attachment object that identifies the HBA is linked to a specific volume group.
 - Each host attachment can be associated with a volume group to define which LUNs that HBA is allowed to access.
 - Multiple host attachments can share the same volume group.

- The host attachment can also specify a port mask that controls which DS8000 I/O ports the HBA is allowed to log in to. Whichever ports the HBA logs in on, it sees the same volume group that is defined on the host attachment associated with this HBA.

References

- IBM Corporation. (2011). *zEnterprise 196 Service Guide* (GC28-6892-02 ed.). Poughkeepsie, NY: IBM.
- IBM Corporation. (2011). *zEnterprise System Ensemble Performance Management Guide* (GC27-2607-02 ed.). Poughkeepsie, NY: IBM.
- IBM Corporation. (2011). *zEnterprise System Ensemble Planning and Configuring Guide* (GC27-2608-02 ed.). Poughkeepsie, NY: IBM.
- IBM Corporation. (2011). *zEnterprise System Introduction to Ensembles* (GC27-2609-02 ed.). Poughkeepsie, NY: IBM.
- IBM Corporation. (2012). *Building an Ensemble Using IBM zEnterprise Unified Resource Manager* (SG24-7921-00 ed.). Poughkeepsie, NY: IBM.
- IBM Corporation. (2012). *IBM System Storage DS8000: Architecture and Implementation* (SG24-8886-02 ed.). San Jose: IBM.
- IBM Corporation. (2012). *System z Hardware Management Console Operations Guide Version 2.12.0* (SC28-6919-00 ed.). Poughkeepsie, NY: IBM.
- IBM Corporation. (2012). *z/Architecture® Principles of Operation* (SA22-7832-09 ed.). Poughkeepsie, NY: IBM.
- IBM Corporation. (2012). *zEnterprise System Support Element Operations Guide Version 2.12.0* (SC28-6920-00 ed.). Poughkeepsie, NY: IBM.
- IBM Corporation. (2013, 6 23). *IBM United States Withdrawal Announcement*. Retrieved from Hardware withdrawal: IBM zEnterprise 196, IBM zEnterprise 114, IBM zEnterprise BladeCenter Extension Model 002, and selected IBM zEnterprise EC12 features:
http://www.ibm.com/common/ssi/rep_ca/4/877/ENUSZG13-0204/ENUSZG13-0204.PDF
- IBM Corporation. (2013). *IBM zEnterprise BladeCenter Extension: Network Connectivity Options* (REDP-5036-00 ed.). Poughkeepsie, NY: IBM.
- IBM Corporation. (2014). *zEnterprise System Readme First Guide for zBX Move* (P/N 41T9404 ed.). Poughkeepsie, NY: IBM.
- IBM Corporation. (2015, 1 14). *IBM United States Hardware Announcement: The IBM z13*. Retrieved from The IBM z13:
http://www.ibm.com/common/ssi/rep_ca/1/897/ENUS115-001/ENUS115-001.PDF
- IBM Corporation. (2015, 1 14). *IBM United States Withdrawal Announcement*. Retrieved from Hardware withdrawal: Select IBM z Systems products:
http://www.ibm.com/common/ssi/rep_ca/5/897/ENUS915-005/ENUS915-005.PDF

Technote: Managing zBX Storage Resources

IBM Corporation. (2015). *z BladeCenter Extension Installation Manual for Physical Planning 2458-004* (GC27-2630-00 ed.). Poughkeepsie, NY: IBM.

IBM Corporation. (2015). *z BladeCenter Extension Model 004 Service Guide* (GC28-6950-00 ed.). Poughkeepsie, NY: IBM.

IBM Corporation. (2015). *z Systems Ensemble Planning Guide* (GC27-2631-00 ed.). Poughkeepsie, NY: IBM.

IBM Corporation. (2015). *z Systems Ensemble Workload Resource Group Management Guide* (GC27-2633-00 ed.). Poughkeepsie, NY, NY: IBM.

Morse, K. (2012). *zBX SAN Cabling and Zoning Recommendations*. IBM Washington Systems Center. Gaithersburg, MD: IBM Corporation. Retrieved from ibm.com/support/techdocs



©Copyright IBM Corporation 2016
IBM Systems
Route 100
Somers, New York 10589
U.S.A.
10/2016

IBM, IBM logo, BladeCenter, DS8000, System Storage, z13, z13s, zEnterprise, and z Systems are trademarks or registered trademarks of the International Business Machines Corporation.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

InfiniBand and InfiniBand Trade Association are registered trademarks of the InfiniBand Trade Association.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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.

OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the [OpenStack website](#).

TEALEAF is a registered trademark of Tealeaf, an IBM Company.

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

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

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates. It is the user's responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.