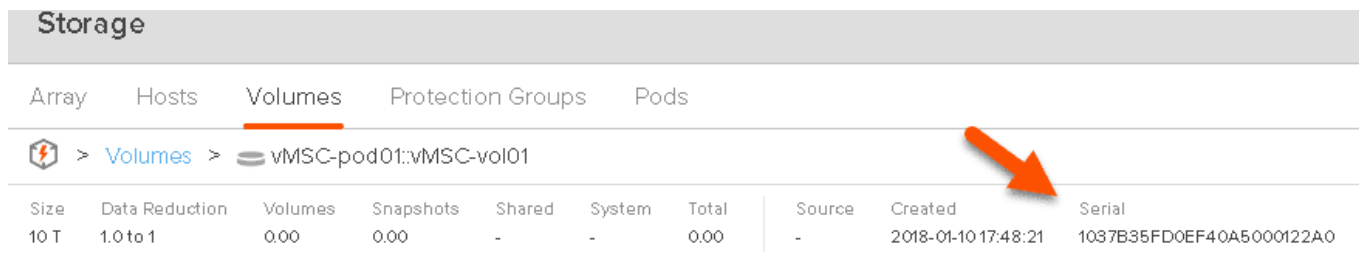


Issue with Consistent LUN ID in ESXi 6.5



Storage

Array Hosts Volumes Protection Groups Pods

> Volumes > vMSC-pod01:vMSC-voi01

Size	Data Reduction	Volumes	Snapshots	Shared	System	Total	Source	Created	Serial
10 T	1.0 to 1	0.00	0.00	-	-	0.00	-	2018-01-10 17:48:21	1037B35FD0EF40A5000122A0

Having consistent LUN IDs for volumes in ESXi has historically been a gotcha-though over time this requirement went away.

For example:

- 2012 <https://cormachogan.com/2012/09/26/do-rdms-still-rely-on-lun-id/>
- 2013 <https://cormachogan.com/2013/12/11/vsphere-5-5-storage-enhancement-part-7-lun-idrdm-restriction-lifted/>

These days, the need for consistent LUN IDs is mainly gone. The lingering use case for this is Microsoft Clustering and how persistent SCSI reservations are handled. The below KB doesn't mention 6.5, but I believe it to still be relevant to 6.5:

<https://kb.vmware.com/s/article/2054897>

There are two places having consistent LUN IDs have mattered. Same LUN ID for a volume within a host (so each path shows the same LUN ID for a volume) and also having the same LUN IDs for a volume on different hosts.

Today, device recognition is used via information in the page 83 of the device VPD. This has things like the device serial number (sometimes called a WWN). This, in combination with the vendor Organizational Unique Identifier OUI, creates what is called a Network Address Authority. This does not use the LUN ID in calculation.

naa.624a93701037b35fd0ef40a5000122a0

The OUI for Pure Storage is 2439a70

<https://www.wireshark.org/tools/oui-lookup.html> can show you this. My volume serial number is 1037b35fd0ef40a5000122a0

Storage

Array Hosts Volumes Protection Groups Pods

 > Volumes >  vMSC-pod01:vMSC-vol01



Size	Data Reduction	Volumes	Snapshots	Shared	System	Total	Source	Created	Serial
10 T	1.0 to 1	0.00	0.00	-	-	0.00	-	2018-01-10 17:48:21	1037B35FD0EF40A5000122A0

So this allows ESXi to identify volume uniquely without using the LUN ID anymore. Though they still do keep some old identifiers that do use the LUN ID to generate the ID. Namely the VML, or VMware Legacy identifier.

vml.0200fd0000624a93701037b35fd0ef40a5000122a0466c61736841

This is a much longer ID and depending on the ESXi version, which use the LUN ID in a few places. This KB article explains:

<https://kb.vmware.com/s/article/2078730>

In short, the LUN ID is used in the early part of the ID and in version prior to 5.5 used as part of a hash for the tail portion. In 5.5 and later, ESXi just zeroed out the LUN portion to the hash wouldn't change due to LUN ID. But the earlier part of the ID that used the ID does change. If you had different LUN IDs, multiple VMLs would be generated for a volume.

Now on the FlashArray, we ensure LUN ID consistency when presenting a volume to a host or a group of hosts. But with our release of [ActiveCluster](#), you could present the same volume from two arrays at once (using active/active replication) when configured in a [Uniform configuration](#). This introduced the possibility to use one LUN ID for a device's paths on one FlashArray and a different LUN ID for the paths to the other

Device Details

Properties Paths

Runtime Name	Status	Device	Target	Name	1
vmhba1:C0:T10:...	Active	PURE Fibre Channel Disk (naa....	52:4a:93:72:e3:85:21:00 52:4a:...	vmhba1:C0:T10:L252	
vmhba1:C0:T11:L...	Active	PURE Fibre Channel Disk (naa....	52:4a:93:72:e3:85:21:10 52:4a:...	vmhba1:C0:T11:L252	
vmhba1:C0:T12:...	Active	PURE Fibre Channel Disk (naa....	52:4a:93:72:e3:85:21:02 52:4a:...	vmhba1:C0:T12:L252	
vmhba1:C0:T13:...	Active	PURE Fibre Channel Disk (naa....	52:4a:93:72:e3:85:21:12 52:4a:...	vmhba1:C0:T13:L252	
vmhba1:C0:T14:...	Active	PURE Fibre Channel Disk (naa....	52:4a:93:72:e3:85:21:06 52:4a:...	vmhba1:C0:T14:L252	
vmhba1:C0:T15:...	Active	PURE Fibre Channel Disk (naa....	52:4a:93:72:e3:85:21:16 52:4a:...	vmhba1:C0:T15:L252	
vmhba1:C0:T4:L...	Active (I/O)	PURE Fibre Channel Disk (naa....	52:4a:93:75:15:74:75:12 52:4a:...	vmhba1:C0:T4:L253	
vmhba1:C0:T5:L...	Active (I/O)	PURE Fibre Channel Disk (naa....	52:4a:93:75:15:74:75:10 52:4a:...	vmhba1:C0:T5:L253	
vmhba1:C0:T6:L...	Active (I/O)	PURE Fibre Channel Disk (naa....	52:4a:93:75:15:74:75:16 52:4a:...	vmhba1:C0:T6:L253	
vmhba1:C0:T7:L...	Active (I/O)	PURE Fibre Channel Disk (naa....	52:4a:93:75:15:74:75:02 52:4a:...	vmhba1:C0:T7:L253	
vmhba1:C0:T8:L...	Active (I/O)	PURE Fibre Channel Disk (naa....	52:4a:93:75:15:74:75:06 52:4a:...	vmhba1:C0:T8:L253	
vmhba1:C0:T9:L...	Active (I/O)	PURE Fibre Channel Disk (naa....	52:4a:93:75:15:74:75:00 52:4a:...	vmhba1:C0:T9:L253	
vmhba2:C0:T10:...	Active	PURE Fibre Channel Disk (naa....	52:4a:93:72:e3:85:21:11 52:4a:9...	vmhba2:C0:T10:L252	

FlashArray.

You can see in the above screen shots, from my preferred FlashArray (the one local to my host) uses LUN ID 253, and the paths to my remote FlashArray use LUN ID 253.

Pure Storage UI - Storage Hosts view for ac-esxi-a-07. Summary table:

Size	Data Reduction	Volumes	Snapshots	Shared	System	Total
18336 G	15.8 to 1	208.29 G	0.00	-	-	208.29 G

Connected Volumes table:

Name	Shared	LUN
vmsc	True	253

Pure Storage UI - Storage Hosts view for ac-esxi-a-07. Summary table:

Size	Data Reduction	Volumes	Snapshots	Shared	System	Total
14337 G	7.4 to 1	13.86 M	0.00	-	-	13.86 M

Connected Volumes table:

Name	Shared	LUN
vmsd	True	252

In this case, my host will only use the remote paths if the local one goes away.

Anyways, the point is that since I have two different LUN IDs for this device on this host, ESXi generates two VMLs:

```
[root@sn1-r720-b05-13:~] esxcfg-scsidevs -l -d naa.624a93701037b35fd0ef40a5000122a0
naa.624a93701037b35fd0ef40a5000122a0
  Device Type: Direct-Access
  Size: 10485760 MB
  Display Name: PURE Fibre Channel Disk (naa.624a93701037b35fd0ef40a5000122a0)
  Multipath Plugin: NMP
  Console Device: /vmfs/devices/disks/naa.624a93701037b35fd0ef40a5000122a0
  Devfs Path: /vmfs/devices/disks/naa.624a93701037b35fd0ef40a5000122a0
  Vendor: PURE      Model: FlashArray      Revis: 8888
  SCSI Level: 6  Is Pseudo: false Status: on
  Is RDM Capable: true  Is Removable: false
  Is Local: false Is SSD: true
  Other Names:
    vml.0200fd0000624a93701037b35fd0ef40a5000122a0466c61736841
    vml.0200fc0000624a93701037b35fd0ef40a5000122a0466c61736841
  VAAI Status: supported
```

You can see the LUN ID in it if you look at the start of it:

vml.0200fd
vml.0200fc

Hexidecimal FC is 252 in decimal and FD is 253.

So why am I bringing this up if it doesnt matter any more? Well in ESXi 6.5 support for multiple VMLs were dropped. So when a device was discovered and it had more than one LUN ID on that host, ESXi would not claim it and would not use it. You'd see messages like in this KB article:

<https://kb.vmware.com/s/article/2148265>

Messages like:

[crayon-6427904b43da8354340351/]

The insidious thing about this, was that let's say you added a device on a certain number of paths. In my case, I added it from only my first array and all paths had the same LUN IDs. ESXi would be great with that and it could be used. But then later on, I enabled active-active replication and then presented paths to it from the second array. But these paths had a different LUN ID. ESXi 6.5 would actually be fine with that. This is because the device has already been claimed, and it seems VML addresses are only added when the device is first claimed, but not when just new paths are added. Since basically nothing uses VMLs, things would work fine.

But then you reboot. This causes all devices to be reclaimed. Oh, ESXi sees different LUN IDs, doesn't like multiple ones, so it doesnt claim the device. So that device becomes unavailable after a reboot even though it worked before.

So the above KB article claims this is an expected result. This is not true, it is a regression bug. And it has been fixed in ESXi 6.5 U1:

<https://docs.vmware.com/en/VMware-vSphere/6.5/rn/vsphere-esxi-651-release-notes.html>

[crayon-6427904b43db1337997471/]

So, in short, upgrade to ESXi 6.5 U1. Especially if you are using uniform active/active replication.