UCS B-Series Servers: Replacing a RAID Controller with older firmware may cause the data-store mount to failure in ESXi hosts

Contents

Problem Statement:

After the replacement of the RAID Controller the NAA id of the VD was changed during foreign configuration import and that caused the datastore mount to fail.

Affected Hardware:

UCSB-MRAID12G

UCSC-MRAID12G

Servers with UCSB-MRAID12G RAID Controllers:

UCS B200 M4

UCS B200 M5
UCS B480 M5
UCS B420 M4

UCS C220 M4

UCS C240 M4

Affected Firmware:

RAID Controller Firmware : 24.5.x.x and 24.6.x.x

Example #
***mrsasctlr.**24.5.0-0043_6.19.05.0_NA.bin**

24.5.x.x controller firmware is seen in all the UCSM releases prior to 3.2.*
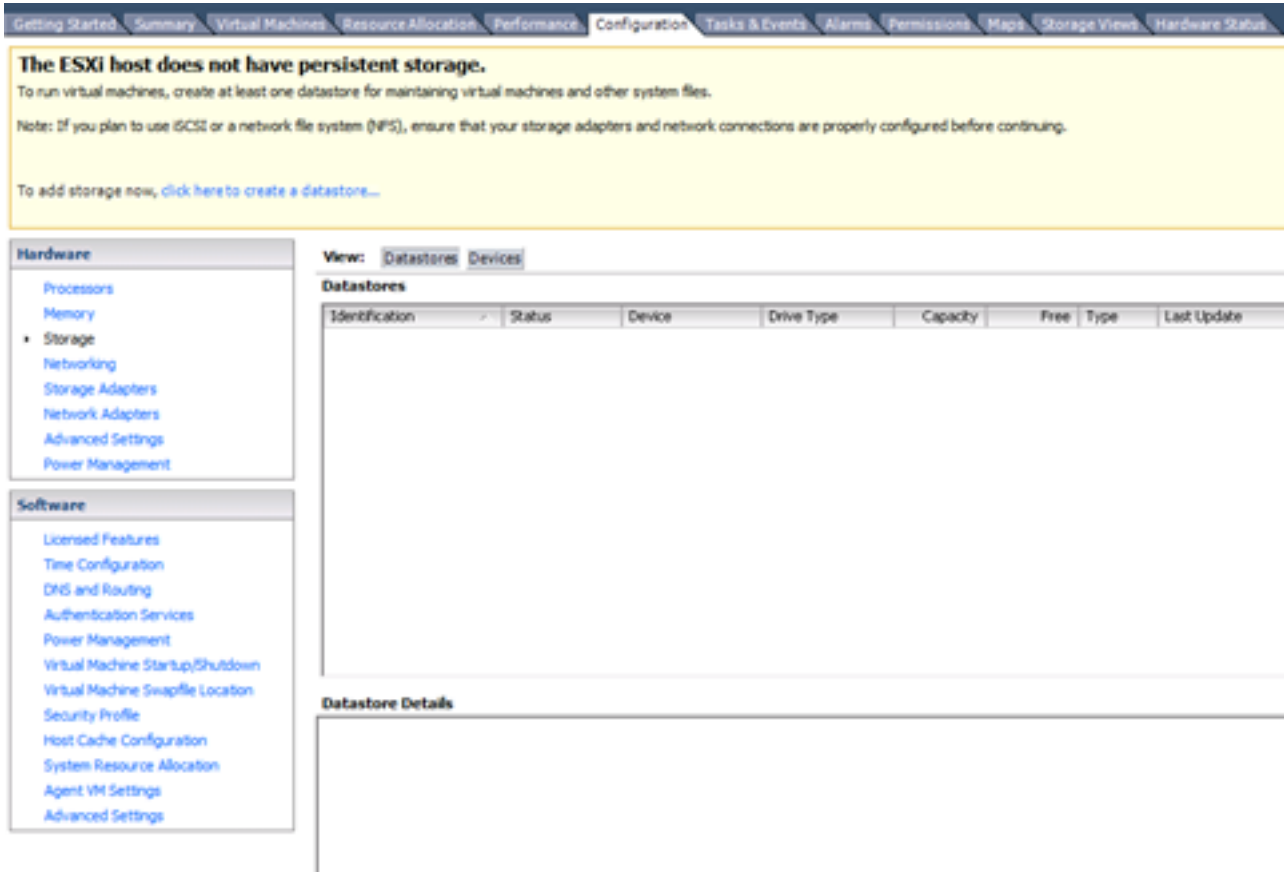
Release notes from 3.1 #

Affected OS:

VMware ESXi

Cause:

With older firmware releases, if there is a DDF(Device Data Format) workspace version mismatch found, the controller FW is not able to restore the NAA ID from DDF during foreign import.

MR 6.4 has DDF_WORK_SPACE version 1, whereas MR 6.10 has DDF_WORK_SPACE version 3. Later releases of FW post-MR 6.4, fixes were made that allow the controller FW to restore NAA IDD from DDF even if a DDF workspace mismatch is found. NAA ID cannot be analyzed properly when the replacement controller Firmware is old(Example: 24.5.x and 24.6.x). However 24.12.x version can properly parse NAA ID.

Before Replacement:
Server 2/2:
   Equipped Product Name: Cisco UCS B200 M5 2 Socket Blade Server
   Equipped PID: UCSB-B200-M5
   Equipped VID: V06
   Equipped Serial (SN): FCH222973K5
   Slot Status: Equipped
   Acknowledged Product Name: Cisco UCS B200 M5 2 Socket Blade Server
   Acknowledged PID: UCSB-B200-M5
   Acknowledged VID: V06
   Acknowledged Serial (SN): FCH222973K5
   Acknowledged Memory (MB): 524288
   Acknowledged Effective Memory (MB): 524288
   Acknowledged Cores: 28
   Acknowledged Adapters: 1
      Virtual Drive 0:
         Type: RAID 1 Mirrored
         Block Size: 512
         Blocks: 1560545280
         Operability: Operable
         Presence: Equipped
         Size: 761985
         Lifecycle: Allocated
         Drive State: Optimal
         Strip Size (KB): 64
         Access Policy: Read Write
         Read Policy: Normal
         Configured Write Cache Policy: Write Through
         Actual Write Cache Policy: Write Through
         IO Policy: Direct
         Drive Cache: No Change
         Bootable: True
         Unique Identifier: bcc0dd21-2006-4189-86c1-132017ad0958
         Vendor Unique Identifier: 618e7283-72eb-6460-240f-d02c0bbd9310   <<<<<<<<<<<<<<

After Replacement:
Server 2/2:

Equipped Product Name: Cisco UCS B200 M5 2 Socket Blade Server
Equipped PID: UCSB-B200-M5
Equipped VID: V06
Equipped Serial (SN): FCH222973K5
Slot Status: Equipped
Acknowledged Product Name: Cisco UCS B200 M5 2 Socket Blade Server
Acknowledged PID: UCSB-B200-M5
Acknowledged VID: V06
Acknowledged Serial (SN): FCH222973K5
Acknowledged Memory (MB): 524288
Acknowledged Effective Memory (MB): 524288
Acknowledged Cores: 28
Acknowledged Adapters: 1
    Virtual Drive 0:
    Type: RAID 1 Mirrored
    Block Size: 512
    Blocks: 1560545280
    Operability: Operable
    Presence: Equipped
    Size: 761985
    Lifecycle: Allocated
    Drive State: Optimal
    Strip Size (KB): 64
    Access Policy: Read Write
    Read Policy: Normal
    Configured Write Cache Policy: Write Through
    Actual Write Cache Policy: Write Through
    IO Policy: Direct
    Drive Cache: No Change
    Bootable: True
    Unique Identifier: 7a894b44-721a-41ae-a3bf-380102b9e64e
    Vendor Unique Identifier: 618e7283-72ea-3f20-ff00-005a0574b04b   <<<<<<<<<<<<<<<

In this case [Vendor Unique Identifier] id of server 2/2 got changed from [618e7283-72eb-6460-240f-d02c0bbd9310] to [618e7283-72ea-3f20-ff00-005a0574b04b]

How to prevent hitting the issue?

This problem can be avoided by updating the firmware of the replacement controller before inserting the VD / disk.

Detailed steps:

1. Shutdown the server
2. Remove all the disk one by one & leave disks the same slot not fully inserted so that their placement order isn't disturbed(If fully removing out of the slot, please keep a **note of the slot** as drives have to be placed back in the same slot)
3. Install a new RAID controller for replacement without inserting a disk.
4. The server will recognize the new RAID controller
5. Update the firmware of the Raid controller.
6. Post successful firmware upgrade, power off the server & insert the disk into the server.
7. Now power on the server

How to recover if the server is hit with this issue?

Detailed steps:

```
===================
Procedure to restore datastore
===================
```
1 Log in to the vSphere Client and select the server from the inventory panel.

2 Click the Configuration tab and click Storage in the Hardware panel.

3 Click Add Storage.

4 Select the Disk/LUN storage type and click Next.
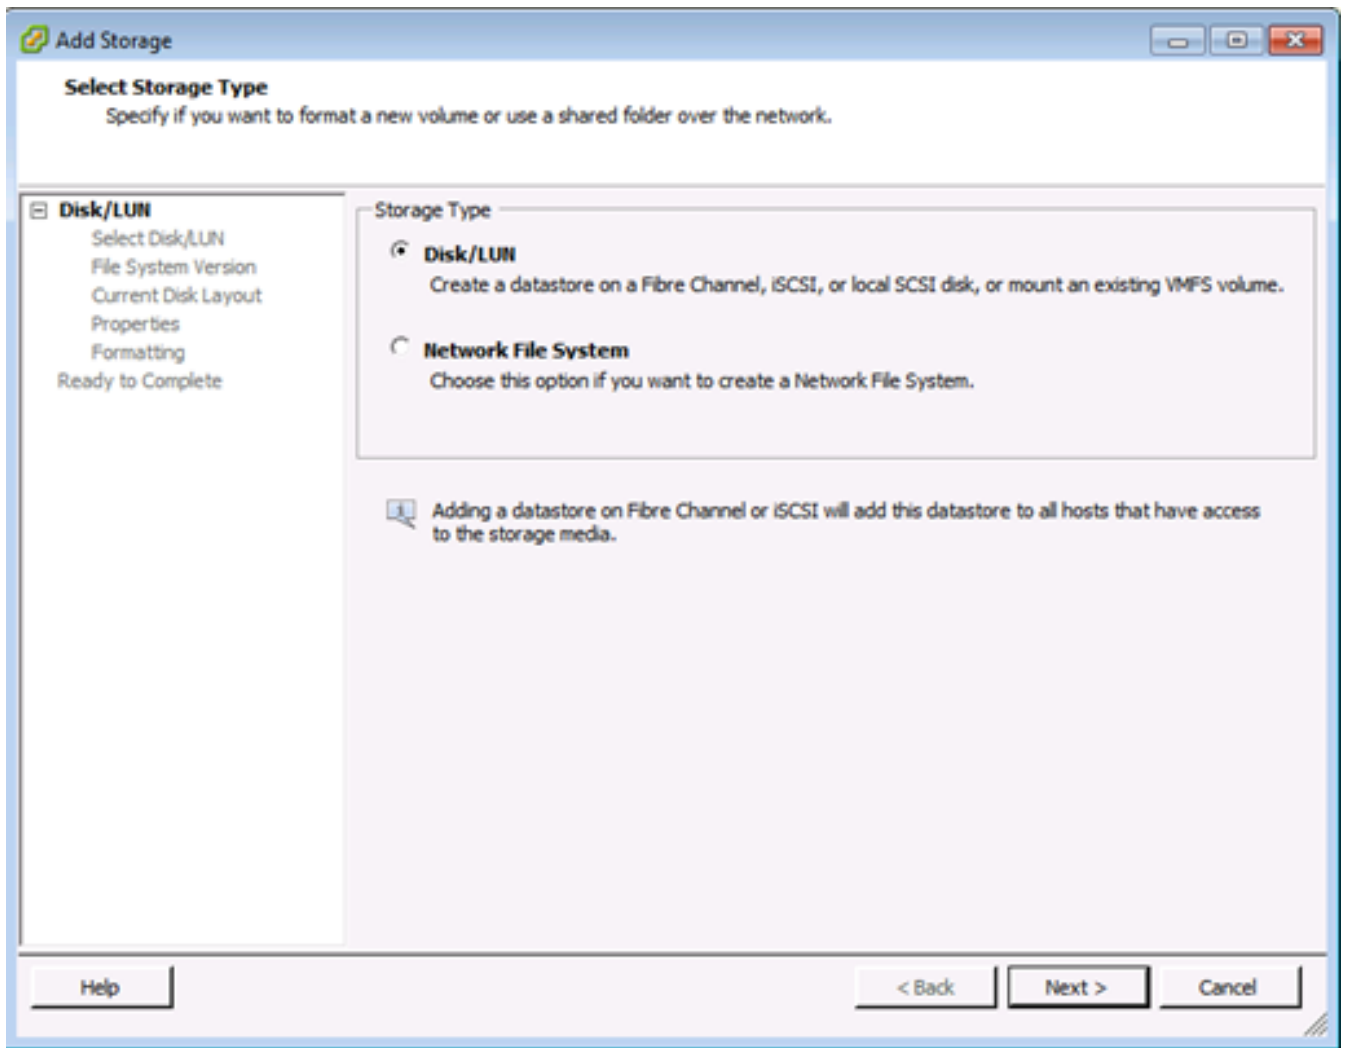
5 From the list of LUNs, select the LUN that has a datastore name displayed in the VMFS Label column and click Next.

Note: The name present in the VMFS Label column indicates that the LUN is a copy that contains a copy of an existing VMFS datastore.

Add Storage

Select Disk/LUN
Select a LUN to create a datastore or expand the current one

Disk/LUN
  Select Disk/LUN
  File System Version
  Current Disk Layout
  Properties
  Formatting
Ready to Complete

Name, Identifier, Path ID, LUN, Capacity, Expandable or VMFS Label contains: ▾ [   ] Clear

Name	Path ID	LUN	Drive Type	Capacity	VMFS Label	Hardware Acceleration
Local LSI Disk (naa.600605b006b49...	vmhba2:C2:T0:L0	0	Non-SSD	278.88 GB	datastore1 (...	Not supported

Help   < Back   Next >   Cancel

6 Under Mount Options, these options are displayed:

Keep Existing Signature: Persistently mount the LUN (for example, mount LUN across reboots)

Assign a New Signature: Resignature the LUN

Format the disk: Reformat the LUN

Notes: **Format the disk** option deletes any existing data on the LUN. Before attempting to resignature, ensure that there are no virtual machines running off that VMFS volume on any other host, as those virtual machines become invalid in the vCenter Server inventory and they are to be registered again on their respective hosts.

select Assign a New Signature and click Next.
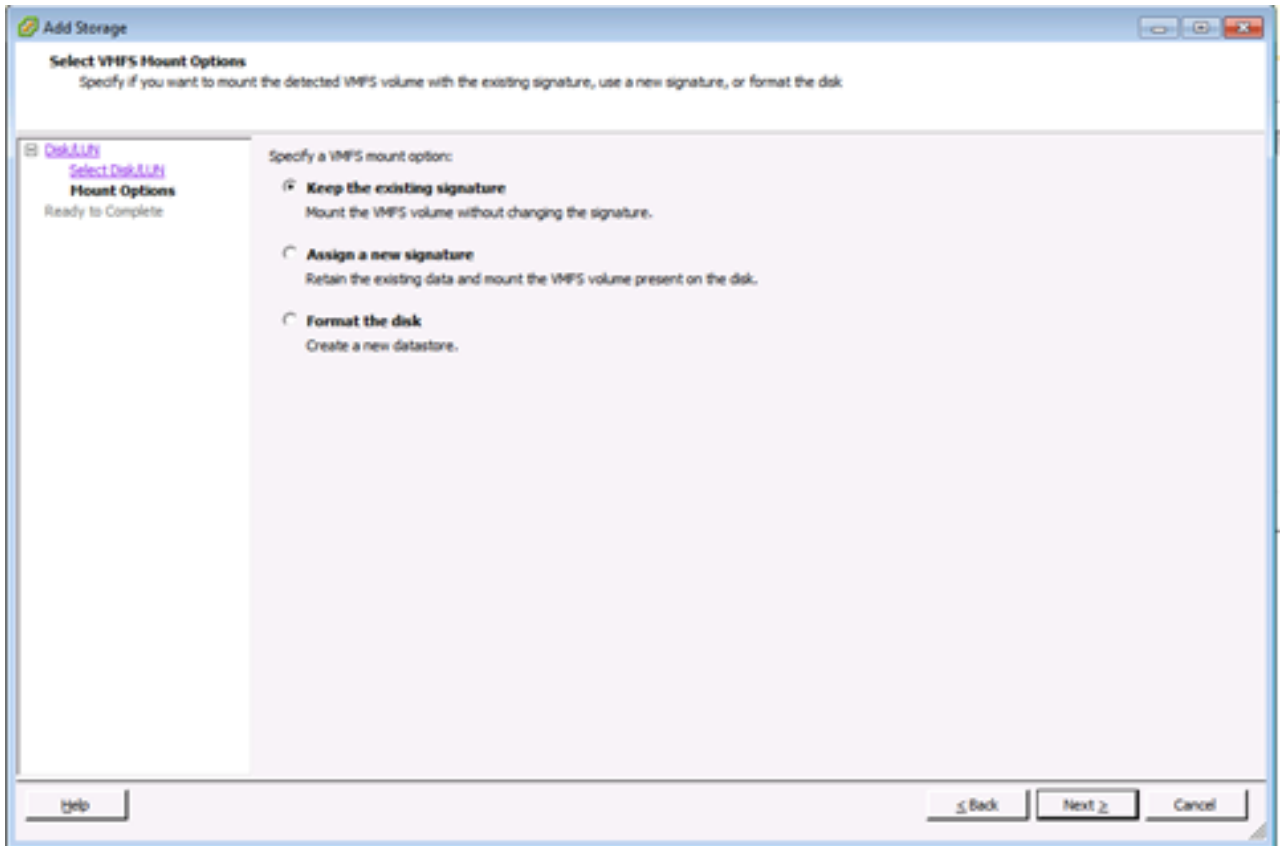
7 Select the desired option for your volume

8 In the Ready to Complete page, review the datastore configuration information and click Finish.
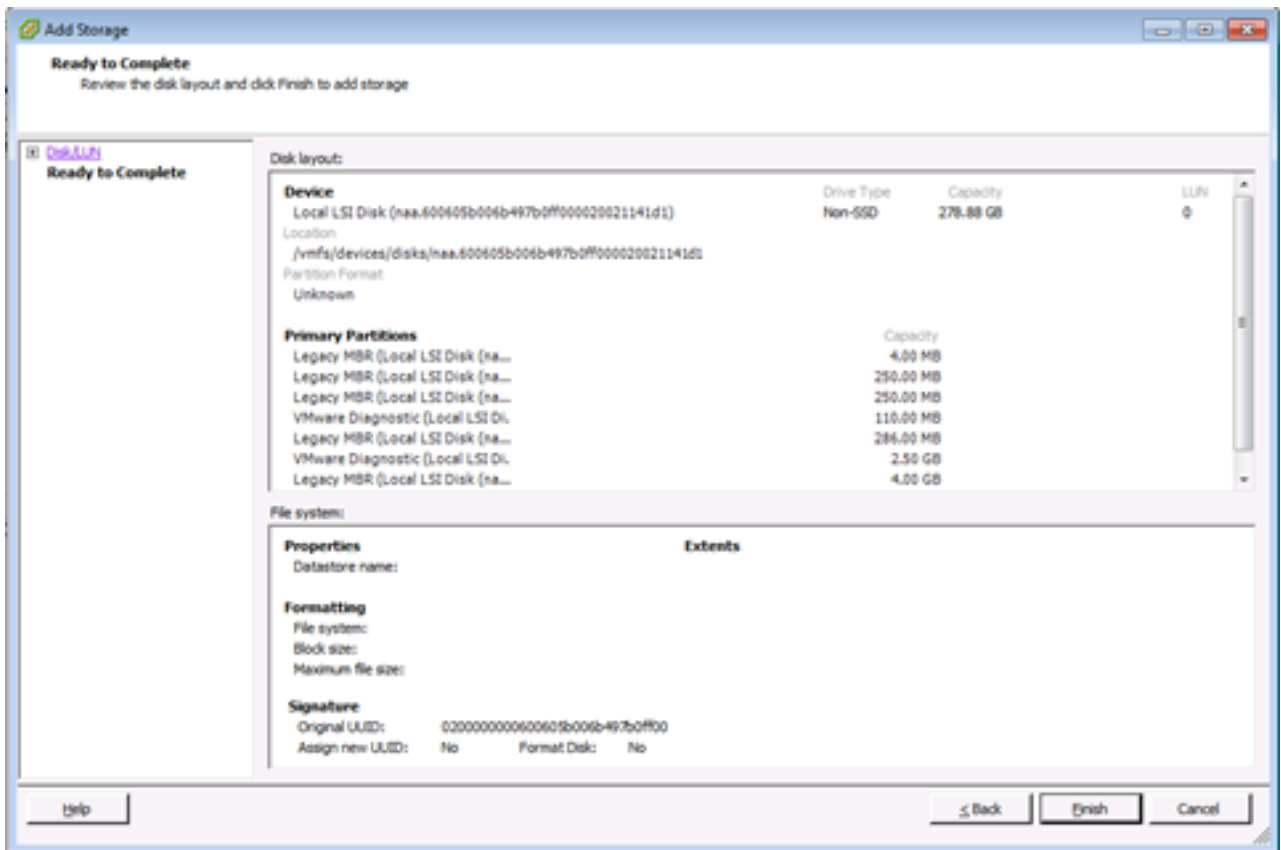
====================
What to do next

=========================
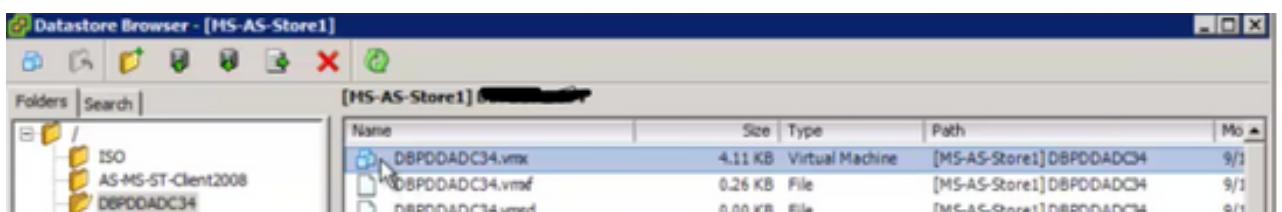After resignaturing, you might have to do the following:

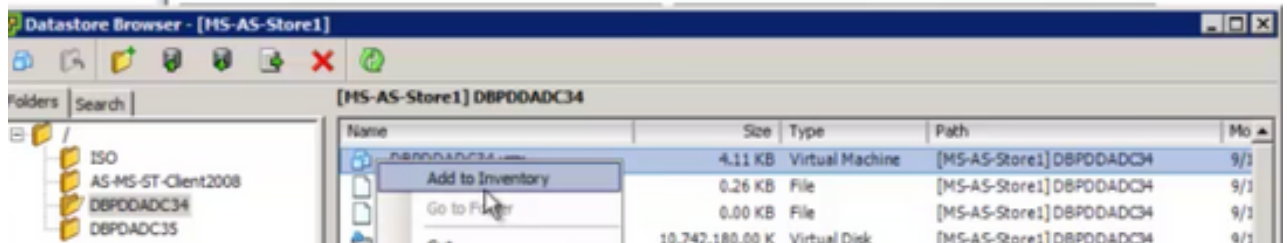1 Log in to the vSphere Client ,Under Inventory List > Click Datastore
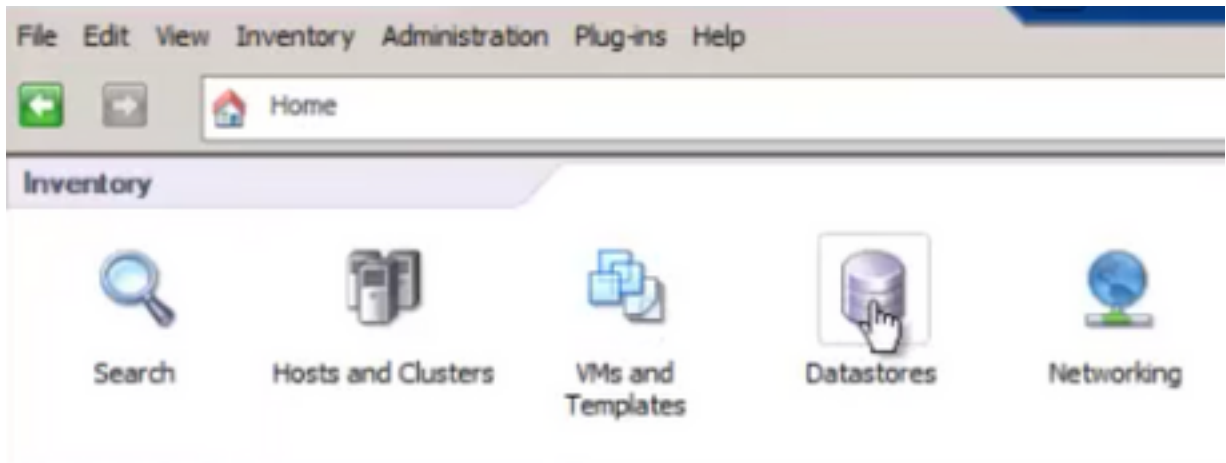
2 Right-click the datastore and click "Browse Datastore "
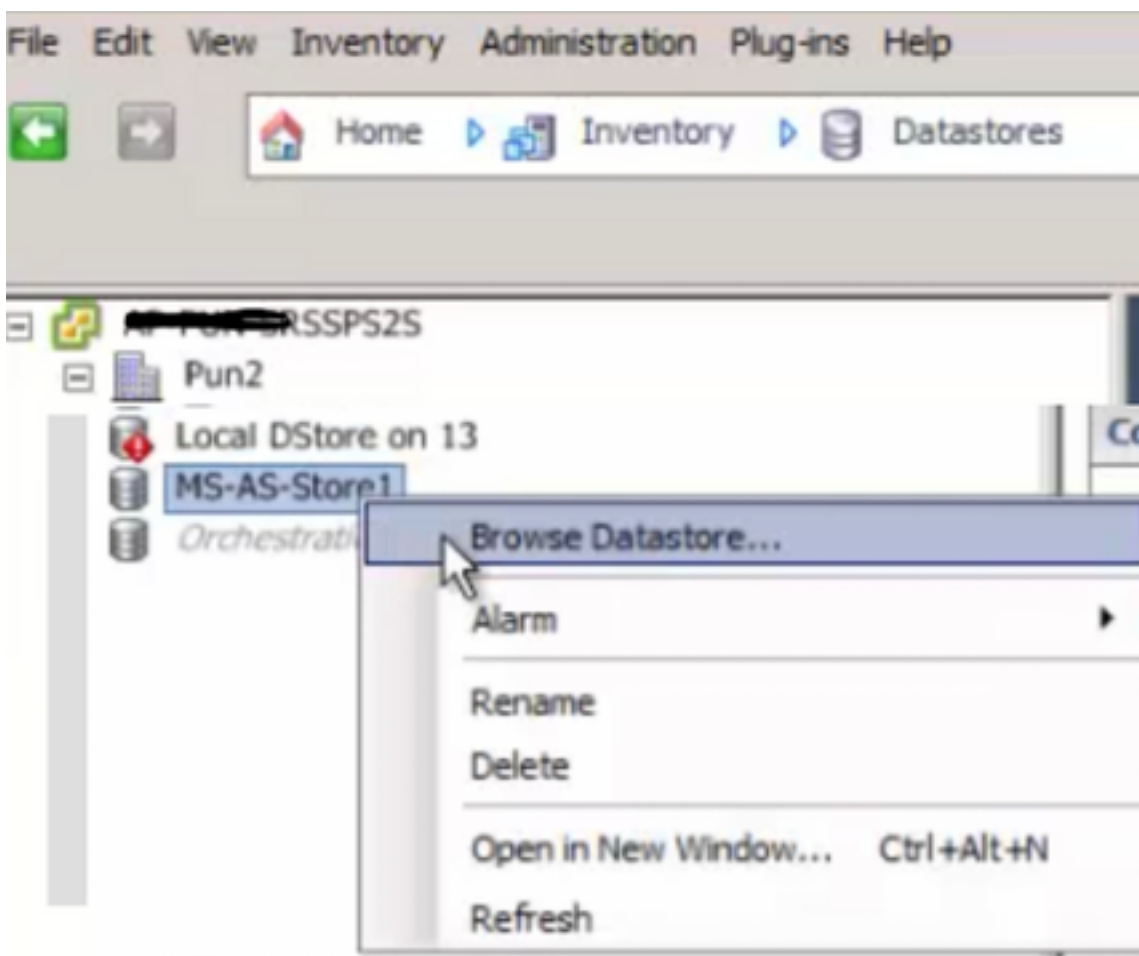
3 On the left pane, click a VM folder to display the contents on the right pane

4 On the right pane, right-click the .vmx file and select "Add to Inventory"

5 Walkthrough the "Add to Inventory" wizard to complete adding the VM to the ESXi host

6 Repeat steps for all remaining VMs

7 Once all VMs have been re-registered, remove all Inaccessible VMs from inventory by right-clicking on each one and selecting "Remove from Inventory"
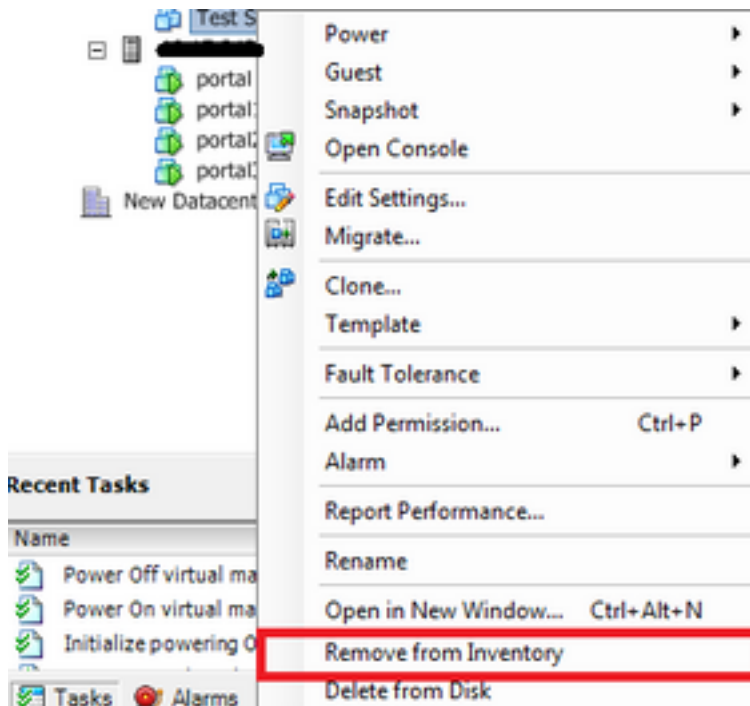
8 Power on each VM and verify it's operational and accessible

Note: Before powering on the VM, reboot the ESXi host and after it's come back online and accessible via vSphere client, confirm VMs are still visible and have not gone to "Inaccessible" state

Related BUG: CSCvr11972

CSCvr11972 Vendor Unique Identifier changed after replacing MRAID12G

https://bst.cloudapps.cisco.com/bugsearch/bug/CSCvr11972