HPE 3PAR OS Upgrade Preparation Guide

September-2017

Abstract

This document is intended for use as a planning tool for HPE customers. It contains the HPE recommendations for HPE 3PAR OS upgrade process.

This document is divided into 3 sections. Refer to each section and follow the checklist given below and mark the task status as complete and NA if the recommendation is not applicable:

Pre-upgrade Validation
Checks to be done during the upgrade
Post upgrade tasks

NOTE:

HPE recommends updating host drivers to a supported level of the targeted HPE 3PAR OS upgrade version before attempting an online update. This information is available on the HPE Single Point of Connectivity (HPE SPOCK) website:
https://www.hpe.com/Storage/Spock/

HPE recommends checking the compatibility information for HPE Cloud System Matrix versions. See
The HPE Cloud System Matrix Compatibility website:
HP CloudSystem Matrix Compatibility

These documents are provided to assist you in setting up host servers with HPE-supported configuration information. In addition to this planning document, the implementation guides are available on the HPE Business Support Center (BSC) website:
www.hpe.com/support/hpesc

Pre-upgrade validation

Applies to	Host OS/Array	Task (Click on hyperlink)	Impact
All 3PAR OS Versions	End of Support Host OS	Take the EOSL hosts Offline during the upgrades	EOSL hosts can lose access to 3PAR if remained Online during the upgrade
HPE 3PAR OS versions 3.2.2-GA-MU3.	HPE 3PAR Array	Advisory	Running Concurrent updatevv Sessions May Have Unpredictable Results
All 3PAR OS Versions	Windows 2008/12	Update MS 2008/2012 MPIO hotfix & persona check, non-PP VLUNs check	Absence of MPIO hotfix may lead to loss of access to 3PAR LUN
All 3PAR OS Versions	Windows 2008/12	Windows Server OS Guidelines	LUNs are marked offline post OS 3PAR upgrade upon reboot of host
All 3PAR OS Versions	Windows 2008/12/16	HPE QMH2572 8Gb and QMH2672 16Gb Fibre Channel HBA	Path loss experienced in windows system running with driver versions 9.1.15.21,9.1.16.21 and 9.1.17.22
All 3PAR OS Versions	CLX and 3PAR CLI recommendations	CLX and CLI software versions must be upgraded prior & post the 3PAR OS upgrade	CLX failover will fail
All 3PAR OS Versions	Red Hat, Oracle, Linux & SUSE OS	Verify multipath.conf as per IG	Loss of access to 3PAR LUN after 3PAR Node reboot
All 3PAR OS Versions	VMware	MPIO Path Policy set to Round-Robin	Loss of access to 3PAR Luns on ESX server during the upgrade
Upgrade to 3PAR OS 3.1.3.x	VMware	Set auto_failover for PP RCOPY groups	Loss of access to 3PAR Lun from the ESX hosts(both primary and secondary)
Upgrades to 3.2.2.GA or later	VMware	New LUNs can no longer be added to ESX 6.0 Environment	3.2.2 upgrade triggered this issue with new device discovery that the symptom of persistent checksum feature introduced in 3.2.2.GA.
All 3PAR OS versions	NetApp	Dedicated ports for NetApp Controllers	NetApp hosts needs the host port to be dedicated for Ontap only
All 3PAR OS Versions	HPE insight Control	Advisory	Supportability
All 3PAR OS Versions	HPE Cloud System	Validate compatibility with 3PAR OS	Cloud system would break if 3PAR OS is upgraded to an non-compatible version
All 3PAR OS Versions	VERITAS DMP	Set VERITAS DMP Restore Interval	Loss of access to 3PAR LUNs during the OS upgrade
All 3PAR OS Versions	Oracle RAC	Verify timeout values	Loss of access to 3PAR LUN after 3PAR Node reboot
All 3PAR OS Versions	Solaris	Verify minimum patch levels	Loss of access to 3PAR LUNs during the OS upgrade

All 3PAR OS Versions	OpenVMS	Validate supported patch levels	Loss of access to 3PAR LUNs during the OS upgrade
312.MU5 and higher	Host based applications	Update host based applications	Host Based applications might cease to work if not updated
All 3PAR OS Versions	HPE 3PAR Array	Remote Copy support	Remote Copy failure
3.1.2 and higher	HPE 3PAR Array	Persistent ports	Avoids host outage
3.1.3 and higher	HPE 3PAR Array	Stop and start Quorum witness	Host outage if VV involved is part of PP Rcopy groups
All 3PAR OS Versions	HPE 3PAR Array	Validate 3PAR Array Health	If a pre-check fails, upgrade cannot be completed
Upgrades from 3.1.2.x and higher	HPE 3PAR Array	HP Proliant Model Service Processor	Post reboot, HP Proliant Model SP might not come back
Upgrade to 3PAR OS 3.1.3.GA and higher	HPE 3PAR Array	Set speed to auto on admin and RCIP ports	MC and CLI sessions will stop working causing array to be unmanageable
Upgrade from 3.1.x. to 3.1.2.x and above	HPE 3PAR Array	Stop AO, DO and active tasks	Upgrade will not proceed if any active tasks are detected
All 3PAR OS Versions	HPE 3PAR Array	Capture existing policies	Existing AO policies will be lost

Checks to be done during the upgrade

Applies to	Host OS/Array	Task (Click on hyperlink)	Impact
All 3PAR OS Versions	AIX	Monitor paths in AIX	Risk of losing access to 3par LUN
All 3PAR OS Versions	Solaris	Monitor paths in Solaris	Risk of losing access to 3par LUN
3.1.2.x and Higher	HP-UX	HP-UX Path rescan	Host path shows "No_Hw"

Post upgrade tasks

Applies to	Host OS/Array	Task (Click on hyperlink)	Impact
All 3PAR OS Versions	HPE 3PAR Array	Change persona	Unexpected behavior of LUNs exported from 3PAR
All 3PAR OS Versions	HPE 3PAR Array	Update GUI and CLI clients	Applications might not work with older 3PAR OS

HPE 3PAR OS 3.3.1.X specific guidelines

Applies to	Host OS/Array/Fabric Switch	Task (Click on hyperlink)	Impact
SP 5.0 and 3.3.1.X	Password changes	General guidelines	NA
3.3.1.GA/EGA	Tunesys/TuneVV	Do not run tunesys/tunevv on a TDVV3 DDS volume	The new SA space LDs created as destination space for the tune are not marked as exclusive to the DDS.

General Recommendations

Applies to	Host OS/Array/Fabric Switch	Task (Click on hyperlink)	Impact
All 3PAR OS Versions	HPE 3PAR Array / Brocade switch	Persistent Ports Failback	Loss of Path between the Array and the Host

End of Support Life Host OS

HPE recommends that all Host Operating Systems that are deemed End Of Support Life (EOSL) by the OS vendor remain offline during the 3PAR OS upgrade to avoid any unforeseen issues.
As a best practice, HP recommends that the customer deploy configurations based on currently supported configuration.

Examples of EOSL Host OS:
ESX 4.1 and lower, Windows 2003 and lower, RHEL 4 and lower, HP-UX 11.11,Citrix XenServer 5 and lower, Solaris 9 & lower, SUSE Linux 11 and below, IBM AIX 5.3, 6.1-Ext (TL00 - TL05) & 7.1-Ext(TL00-TL02)

Please refer to HPE SPOCK Support Matrix for complete EOSL Host OS details –
https://h20272.www2.hpe.com/SPOCK/index.aspx

HPE recommends all EOSL HPE 3PAR OS with below details to be performed in an offline manner.

HPE 3PAR OS 3.1.1 EOSL as of 5th May 2014:
http://h20564.www2.hpe.com/hpsc/doc/public/display?docId=emr_na-c04274026&sp4ts.oid=5047531

HPE 3PAR OS 3.1.2 EOSL as of 6th Feb 2015:
http://h20564.www2.hpe.com/hpsc/doc/public/display?docId=c04577713

HPE 3PAR OS 3.1.3 EOSL as of 29th July 2017
http://h20564.www2.hpe.com/hpsc/doc/public/display?docId=c05059466

HPE 3PAR StoreServ Storage - Running Concurrent updatevv Sessions May Have Unpredictable Results

DESCRIPTION
The updatevv command for HPE 3PAR OS is used to update the snapshot Virtual Volume (VV) with new snapshot content. For each snapshot VV name specified, the command puts the snapshot in standby mode when updating the VV maps, and hosts may see disk resets while the VV is in standby mode. Normally, the VV will be in standby mode for a very short duration. Running concurrent updatevv sessions for VV sets or a list of VVs which have common VVs can have unpredictable results, including an unexpected node restart.

SCOPE
This issue affects all HPE 3PAR StoreServ Storage arrays running HPE 3PAR OS versions 3.2.2.x.

RESOLUTION
Fixed in (3.3.1 MU1).

Workaround
Do not use the updatevv command with multiple VVs listed or a VVSET. Use updatevv only on one VV at a time (for example, updatevv VV1).

TOP

Windows Server 2008/2012/2016

NOTE: Failure to follow the below recommendations from Microsoft may lead to host losing access to storage.

1. Recommended Hotfixes by Microsoft for Windows 2008 & 2012 servers

- For Windows Server 2008 using MS MPIO or Windows 2008 VM running on ESX using a HBA in pass-through mode or an iSCSI LUN, following Microsoft hotfixes must be applied before updating the HPE 3PAR Storage System:

- Windows Server 2008, SP1, SP2, R2 and R2 SP1 require KB2754704 installed to resolve issue with MPIO path failover. Read MS Article: http://support.microsoft.com/kb/2754704 for more details.

- Windows Server 2012 and Windows Server 2008 R2, R2 SP1requires KB2821052 installed to resolve issue with MPIO path failover. Read: http://support.microsoft.com/kb/2821052 for more details.

- Windows Server 2008 SP2 requires KB2878031 installed to resolve an issue with MPIO path failover

- Windows Server 2012 and 2012 R2 requires KB3046101 installed to resolve issue with MPIO path failover

2. Load Balancing Policy should be Round-Robin

For Windows 2008 that are using the Microsoft native MPIO driver, it is required the load-balancing policy be set to "Round-Robin".

The procedure for checking or changing load-balancing policy can be found at
http://technet.microsoft.com/en-us/library/ee619752%28WS.10%29.aspx

See section "To configure the load-balancing policy setting for a LUN"

You must restart the computer after applying the above listed hotfix.

3. Remote Copy Groups Using Non-Peer Persistent VLUNs and Windows 2012 Host Clusters

DESCRIPTION

After upgrading from HPE 3PAR OS 3.1.2.x to 3.1.3 or higher, the passive node (seeing the read-only copy of the RC group) will hang during a reboot. Furthermore, if only the boot disk is presented (no clustered disks are presented), the system boots correctly. If one of the clustered disks is added, the passive node will hang during the reboot. Then when adding the clustered disks while the passive node is up and running, both Device Manager and Disk Manager hang during a re-scan.
The results to the active node are assumed to be the same if it is accessing a read-only copy of an RC group.

This situation can occur in the following environment:

- Cluster nodes are Windows 2012 R1 Standard Edition
- Windows Failover Cluster with CLX 4.0
- HPE 3PAR arrays involved in RC Groups are running HPE 3PAR OS v3.1.3
- Boot from SAN
- Clustered nodes are configured as Persona 15
- Peer persistent VLUNs are not active/presented to clustered nodes

RESOLUTION

Make sure that the Windows hosts are accessing non-peer persistent VLUNs! DO NOT use the following steps on peer persistent VLUNs!

These issues may be avoided or corrected by performing the following steps. Follow these ONLY on Non-Peer Persistent VLUNs as Peer Persistent VLUNs need auto_failover to be set prior to the upgrade.

- ✓ To prevent this issue from occurring prior to upgrading to 3.1.3 or higher:

Perform the following on primary groups on both arrays:

setrcopygroup pol no_auto_failover <group_name>

The policy, path_management, does not exist prior to 3.1.3.

- ✓ To correct this issue if it occurs after upgrading to 3.1.3 or higher:

Perform the following on primary groups on both arrays:

For the non-peer persistent group, the following policies need to be removed:

auto_failover

path_management

Remove the policies using:

setrcopygroup pol no_auto_failover <group_name>

setrcopygroup pol no_path_management <group_name>

4. Windows Server OS Guidelines

Before upgrading to 3PAR OS 3.1.1, 3PAR OS 3.1.2, or associated MUs, Hewlett Packard Enterprise recommends the execution of Microsoft KB2849097 on every Windows Server 2008/2012 host connected to a 3PAR array prior to performing an initial array firmware upgrade. Subsequently, the script contained in KB2849097 will have to be rerun on a host each time new 3PAR LUNs are exported to that host.
KB2849097 is a Microsoft PowerShell script designed to modify the Partner Attributes registry value that is located at
`HKLM\System\CurrentControlSet\Enum\SCSI\<device>\<instance>\Device Parameters \Partmgr`
The value is responsible for the state of 3PAR LUNs following an array firmware upgrade. The script sets the value to "0" essentially changing its policy to "online."
Please refer to ***HPE 3PAR Operating System Upgrade Planning Guide*** for complete/detailed steps on the below link
http://h20565.www2.hpe.com/hpsc/doc/public/display?docId=c02660486 **Pg.**33-34

Caution!! HPE guideline is to mandatorily update to the latest Compatible version of Host Explorer before changing the host persona to 15. If the Host Explorer application is in use, please check compatibility of Host Explorer with HPE 3PAR OS using link below & update the Host Explorer before Host persona is changed".
http://h20566.www2.hpe.com/hpsc/doc/public/display?sp4ts.oid=5044390&docLocale=en_US&docId=emr_na-c04790896

TOP

5. Path Loss Experienced from Windows Servers with Qlogic based HBA following an HPE StoreFabric B-Series switch FOS Upgrade

(1) DESCRIPTION

Some customers have observed a path loss from Windows Servers with HPE QMH2572 and QMH2672 Host Bus Adapters (HBAs) with driver versions 9.1.15.21, 9.1.16.21 and 9.1.17.22 following a Hewlett Packard Enterprise StoreFabric B-series (or Brocade OEM equivalent) Fabric OS (FOS) upgrade. The reason for path loss is that Host Bus Adapters log out of the Fabric Name Server due to the lengthy HAreboot times associated with single CPU switches, and do not log back in. The HBAs will continue to process I/O following the HAreboot until they receive a Registered State Change Notification (RSCN) at which point they log out all of their target ports.
This path loss event could also be triggered by an HAreboot command on an HPE StoreFabric B-series fixed-port (single CP/"pizza box") switch with attached QLogic QMH2572 and QMH2672 Host Bus Adapters.
The affected Host Bus Adapters were connected to non-director, fixed-port (single CP/"pizza box") class switches, which could include Access Gateway switches in a blade server enclosure.
This issue has not been observed when the Host Bus Adapters were connected to director class switches due to the much shorter timeframe associated with the HA Failover restart used by these switches compared to the HAreboot restart used by the fixed port switches.
The path failure may not occur until the FOS upgrade or significantly later because the Name Server logout is prolonged. The name server logout will persist until the HBA performs a new fabric login (FLOGI).
When path loss was observed during the FOS upgrade, it was determined that it was an RSCN event associated with the 8 Gb 3PAR StoreServ 7000 series or 3PAR T400 that caused the HBA to log out of its attached target ports.
In cases where the path failure occurred sometime after the FOS upgrade, wherein no logical connection was made to the FOS upgrade, the trigger event may be any RSCN that occurs for any reason.

(2) SCOPE

This issue is seen on Windows servers running Windows 2008, 2012, and 2016.

(3) RESOLUTION

Driver version 9.1.17.25 is targeted to be included in the 2017 April/May Service Pack for ProLiant (SPP) once generally available.
Driver version 9.1.17.25 is available by Controlled Release .

A resolution for systems containing the SN1600Q adapter will be made available in a future release. Until this driver patch is available, the configuration using the 9.2.2.20 driver will be vulnerable to this issue.
NOTE: The workaround should be used until a driver is available.

Original Driver	Patched Driver	O/S
9.1.15.21	9.1.17.25	WIN 2008/2012+R2
9.1.16.21	9.1.17.25	WIN 2008/2012+R2
9.1.17.22	9.1.17.25	WIN 2008/2012+R2
9.2.2.20	9.1.17.25	WIN 2008/2012+R2*
9.2.2.20	Future Release WIN 2016	

TOP

6. CLX & 3PAR CLI Recommendations for 3PAR OS upgrades:

CLX depends on the 3PAR Remote Copy functionality in the backend. Hence it is always recommended to have the same version of the HPE 3PAR OS on the Primary and Secondary 3PAR Arrays.

Pre-Upgrade recommendations:

1. CLX and CLI software versions must be upgraded prior or post the 3PAR OS upgrade.
2. Check the CLX Config matrix for the supportability details under HPE SPOCK Matrix.
   https://h20272.www2.hpe.com/SPOCK/Pages/spock2Html.aspx?htmlFile=sw_array_3par.html
3. When one array is upgraded and the other array upgrade is pending, it is not recommended to do any CLX failovers, as there will be a mismatch of the 3PAR CLI versions and hence the CLX failover may fail.
4. Hence, please plan to upgrade both the arrays within a minimal gap between the upgrades.
5. Prior to the upgrade, it is recommended to failover to the other Cluster Node which will not be affected by the 3PAR Upgrade.

Post Upgrade Recommendations:

1. Post the 3PAR OS upgrade. The 3PAR CLI must be upgraded to the same version as of the 3PAR OS version.
2. After Installing the new 3PAR CLI, the new Array certificate must be accepted by logging into the 3PAR CLI. This is required for the CLX to communicate with the 3PAR Array.
   NOTE: Unless the array certificate is accepted in the 3PAR CLI, the CLX connection test will not be successful.
3. Also the CLX Application must be upgraded to the supported version. Ensure that the CLX dependency is removed on the all Storage disks configured in the MS Failover Cluster, prior to the CLX upgrade.
4. Care must be taken to pause the Failover of the Roles in the MS failover cluster, when the CLX software is being upgraded. Failing to do so, may cause an unexpected Cluster failover to be triggered causing Host outages.
5. The CLX Software must be upgraded on all the Nodes of the Cluster. All nodes should have the same version of the CLX software installed, otherwise, the CLX failover will fail.
6. Please raise a case with the 3PAR Host Applications team (3PAR 3SC Breakfix), in case of any queries related to the CLX upgrade procedure.

7. HPE 3PAR Peer Persistence

Path Verify Enabled must be set in the MPIO settings on Windows Server 2016/2012/2008 hosts configured for an HPE 3PAR Peer Persistence implementation.

Please refer to ***HPE 3PAR Windows Server 2008/2012/2016 Implementation Guide***
http://h20566.www2.hpe.com/hpsc/doc/public/display?docId=c04448812 Pg.11

For All host and Array based application (RMC, RMS, SRM, SRA) please refer to HPE SPOCK Compatibility Matrix:
https://h20272.www2.hpe.com/SPOCK/Pages/spock2Html.aspx?htmlFile=sw_array_3par.html

TOP

Red Hat, Oracle, Linux, and SUSE OS

NOTE: Failure to follow the below recommendations may lead to host losing access to storage.

Observe the following host configuration guidelines when planning an online HPE 3PAR OS upgrade for Red Hat, Oracle, Linux®, or SUSE OS:

1. We have seen a few instances where the **multipath.conf** file is not configured per the guide. This may cause incorrect detection of path removal during a node reboot or incorrect path failover.

For the Red Hat hosts, please ensure that the system is configured per the current 3PAR/HP Implementation Guide. It can be downloaded from:

Documents are available at HPE Information Library.

http://www.hpe.com/info/storage/docs

Select:

- Products and Solutions: HPE 3PAR StoreServ Storage
- Information Type: Configuration
- Language: English

2. The SCSI timeout value for RHEL 4 and SUSE 9 must be changed from the default of 30 seconds to 60 seconds. To change the timeout value, see the HPE 3PAR host OS implementation guide specific to the OS. This needs to be performed for both FC- and iSCSI-connected hosts.

This change is not required in RHEL 5 and 6.

WARNING: If not set to 60 seconds, the SCSI timeout will result in host disks being taken offline during InServ Storage Server rolling upgrades. Furthermore, Remote Copy requires the SCSI timeout value of 60 seconds, otherwise remote copy operations will become stale with a node reboot.

3. QLogic drivers for SUSE 9, 10 and Red Hat 4, 5 should be running with driver parameter PORT-DOWN retry count (qlport_down_retry) set to 1. To set the value, see the HPE 3PAR host OS implementation guide specific to the OS.

NOTE: If the HPE 3PAR StoreServ Storage array is running HPE 3PAR OS 3.1.1 or later you must modify the HBA parameter by setting qlport_down_retry to 10. If the HPE 3PAR StoreServ Storage array is running an HPE 3PAR OS version earlier than 3.1.1, set qlport_down_retry to 1.

5. For hosts using Oracle Clusterware, make the following required changes or adjustments before performing an upgrade:

∘ For 10gR1 10.1.0.x and 10gR2 10.2.0.1 (without bug 4896338 patched), set css miscount to 195:

#set css miscount=195

◦ For 10gR2 10.2.0.1 (with bug 4896338 patched) and later, 11gR1, and 11gR2, css disktimeout must not be lowered from its default value of 200.

◦ For all Linux hosts with Oracle Clusterware, install and adjust Linux Hangcheck tick parameters to 30 seconds.

Note: This does not apply to oracle 11g R2 because Hangcheck process is removed.

TOP

VMware

1. Path Policy Recommendation

VMware ESX/ESXi server - Active/Active, with a path policy of Most Recently Used (MRU), does not maintain or re-instate balancing of I/O load after a failover/failback multipath event.
This could leave I/O in an unbalanced state, which may yield significant I/O performance issues.

Implementation of an MRU path policy is not recommended by HP. As of ESX 4.0, a third path policy choice of Round Robin is available. This is the recommended path selection policy on ESX systems where it is available.
The VMware vCenter Site Recovery Manager and the HPE 3PAR Adapter for VMware Site Replicator require newer versions to be compatible with HPE 3PAR OS 3.1.1 or later. For details, see VMware's Site Recovery Manager

For All host and Array based application (RMC, RMS, SRM, SRA) please refer to HPE SPOCK Compatibility Matrix:
https://h20272.www2.hpe.com/SPOCK/Pages/spock2Html.aspx?htmlFile=sw_array_3par.html

2. Upgrades to HPE 3PAR OS 3.1.3 & Higher when Persona 11 is used with Peer Persistent Remote copy Groups:

Peer Persistence is ONLY supported using ESX hosts. When Remote Copy is configured with RC groups (Sync Mode) that are in 3.1.2 MU2 and higher system, VVs are exported from both Primary/Secondary RC systems to the same ESX host. Volume paths for a given volume are "Active" only on the array where the "Primary" copy of the volume resides. Other volume paths are marked "Standby." Primary volume is in read/write (RW) mode, secondary volume is in read/only (RO) mode. This can be viewed on the ESX host path management.

Peer Persistence has an existing setting called "auto_failover" to support Automatic Transparent Failover (ATF) in high availability (HA) configuration that are in synchronous replication. If this setting is NOT turned on before

Upgrading to 3.1.3.x and higher, the paths for the PP RC Groups will change from "standby" to "active". ESX host will no longer be able to write to LUNs as these paths on the secondary array are marked as read-only (RO).

1. Check if the system is licensed for Peer Persistence by issuing the Cli Command: Cli%showlicense. If Peer Persistence license is not found, this recommendation can be ignored.

2. To check if this is a Peer Persistent configuration, the same WWN is exported from both primary/secondary site and exported LUN numbers are the same for those LUNs. NOTE: If you **know** this is a Peer Persistent configuration and the WWN and LUN numbers are NOT the same, they need to be configured as such or Peer Persistence will not function when a failover is initiated and the LUNs will not be accessible after the upgrade without considerable support work

3. To check if auto_failover is set for all RC Groups:

a) Use the "showrcopy -d" output to verify. You will find "auto_failover" under the options column in the Output:

```
Group Information
Name                ID   Target     Domain    Status   Role       Mode     LocalUserCpg
LocalSnapCpg RmUserCpg  RmSnapCpg  Options
BCK-CLU-01-DP_VMGRE  284 3PAR01-DC2 -         Started  Primary    Sync
auto_failover,path_management
  LocalVV             ID   RemoteVV              ID   SyncStatus    Resync_ss Sync_ss
VV_iter R_iter S_iter LastSyncTime
  BCK-CLU-01-DP_VMGRE 1038 BCK-CLU-01-DP_VMGRE_DR  980 Synced        none      none
NA      NA     NA     NA
```

b) Setting `auto_failover` can be done using the following command:

```
Cli%setrcopygroup pol auto_failover <Group Name>
```

This command can be issued only on a group that is in the source role. The policy is also assigned to the target group when issued on the source.

3. Verify Hosts in this Peer Persistent configuration are only ESX hosts. Windows hosts are not supported.
4. Verify the ESX host's persona is set to 11

3. HPE 3PAR StoreServ Storage - Array Host Disconnects While Running VMware vSphere 5.5 Update 2 and Later

DESCRIPTION

With arrays running HPE 3PAR OS 3.2.2 or 3.2.1, hosts using VMware vSphere 5.5 Update 2 and later may experience host disconnects due to a low timeout value on the host for the heartbeat of Atomic Test and Set (ATS). A delay of 8 seconds or more for an individual ATS heartbeat I/O anywhere in the SAN infrastructure can result in a host disconnect from the storage.

VMware has indicated that the following observations are an indication this issue has been encountered:

- An ESXi 5.5 Update 2 or ESXi 6.0 host loses connectivity to a VMFS5 datastore
- In the **/var/run/log/vobd.log** file and Virtual Center Events, you see the following VOB message:
  Lost access to volume <uuid><volume name> due to connectivity issues. Recovery attempt is in progress and the outcome will be reported shortly
- In the **/var/run/log/vmkernel.log** file, you see the following message:
  ATS Miscompare detected beween test and set HB images at offset XXX on vol YYY
- You see error messages indicating an ATS miscompare similar to this in **/var/log/vmkernel.log**:
  2015-11-20T22:12:47.194Z cpu13:33467)ScsiDeviceIO: 2645: Cmd(0x439dd0d7c400) 0x89, CmdSN 0x2f3dd6 from world 3937473 to dev "naa.50002ac0049412fa" failed H:0x0 D:0x2 P:0x0 Valid sense data: 0xe 0x1d 0x0

SCOPE

This issue affects VMware hosts running VMware 5.5 Update 2 or later connecting to HPE 3PAR StoreServ arrays running all versions of HPE 3PAR OS 3.2.2 or 3.2.1.

RESOLUTION

HPE has released HPE 3PAR OS 3.2.2 MU3 and 3.2.2 EMU2 P33 with enhancements to improve ATS processing to minimize the likelihood of experiencing this issue.
The following workaround must be applied on the host side, until the upgrade or patch can be installed:
To prevent the host disconnects due to the ATS timeout, disable the ATS heartbeat feature and revert to the pre-5.5 Update 2 behavior on the VMware host. To revert the heartbeat to non-ATS mechanisms, disable this feature on ALL hosts sharing the datastore where these errors are seen (refer to the Vmware KB article at http://kb.vmware.com/kb/2113956).

Customer Advisory

http://h20564.www2.hpe.com/hpsc/doc/public/display?docId=c05228694

TOP

4. HPE 3PAR StoreServ Storage – ESX 6.0x servers freeze and lose paths to the array when doing rescans

Description

ESX servers running 6.0x connected to a 3PAR array could experience path issues leading to freezing VMs when a lun rescan has been initiated.
When a VMware vSphere ESXi 6.0 host requests SMART data from the storage array, a response may be received from the storage array that can trigger a Permanent Device Loss (PDL) condition.
VMware versions may send unsupported commands to the PE LUN (LUN 256). Such commands are currently returned with sense data 0x5 0x25 0x0. This response causes VMware to encounter an unexpected PDL alert on the LUN. The response is now changed to 0x5 0x20 0x0 to avoid this issue.

This issue occurs because in this specific scenario, an ESXi host has sent a request for SMART data to a storage array, and the array has responded with an unexpected illegal request error. The response received by the host triggers a Permanent Device Loss (PDL) detection, and the kernel performs a path evaluation to determine if there is need to fail the link in question.

In ESXi 6.0 Update 2, a change to the PDL response behavior can result in this condition blocking additional IO operations, resulting in the aborts and timeouts described in the Symptoms section. For more information, see *General Storage Issues* section in the ESXi 6.0 Update 2 Release Notes.

SCOPE

This issue affects VMware hosts running VMware 6.0 or later connecting to HPE 3PAR StoreServ arrays running all versions of HPE 3PAR OS 3.2.2 or 3.2.1.

Resolution

HPE has released HPE 3PAR OS 3.2.2 MU3 with enhancements to address this behavior of the hosts to minimize the likelihood of experiencing this issue
(Refer to KB article: http://kb.vmware.com/kb/2133286)

TOP

5. ESXi 5.5/6.0 with Emulex 10.7.110.4 FC/FCoE driver could run in to PSOD during an online OS upgrade or patch install

Description

Possible PSOD on ESXi 5.5/6.0 with Emulex 10.7.110.4 FC/FCoE driver, The PSOD stack trace shows a summary line similar to the following:

WARNING: Heartbeat: 785: PCPU 22 didn't have a heartbeat
for 21 seconds; *may* be locked up

Cause:

Root cause of this issue has been identified as the lpfc driver failing to unlock a spinlock when the lock-owning function exits execution. When the next driver thread attempts to lock the same spinlock, it can't because the ESXi kernel book-keeping shows the lock held. This causes the new kernel thread to "spin" waiting for the lock and eventually causes the PSOD.

Technical Information:

The PSOD failure frequency is random because it depends on what is happening in the SAN and how those events impact driver processing. In several cases, there were one or more FCIDs that were unresponsive to FC-LS

commands even though they were zone members. The driver eventually terminated those attempted connections and this could trigger the lock issue.
Solution:

Broadcom ECD recommends the 10.7.170.0 (or later) revision FC/FCoE driver as a minimum to correct this issue; but customers should consult their OEM partner to ensure that they install the supported OEM kit containing the fix.. All drivers are available on VMware's website

Ref: Vmware KB Published: 2146526
https://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&docType=kc&externalId=2146526&sliceId=1&docTypeID=DT_KB_1_1&dialogID=183144167&stateId=1%200%20183146466

Customer Advisory
https://h20565.www2.hpe.com/hpsc/doc/public/display?sp4ts.oid=316558&docLocale=en_US&docId=mmr_sf-EN_US000013314

6. Advisory: Host Bus Adapters - Fibre Channel Driver lpfc 11.1 Might Encounter a Purple Diagnostic Screen Displaying Page Fault #PF Exception 14 on VMware ESXi 5.5, 6.0 and 6.5.

Description
Fibre Channel (FC) and Fibre Channel over Ethernet (FCoE) driver lpfc might occasionally display a purple diagnostic screen on VMware ESXi 5.5 and 6.0. The diagnostic output includes #PF Exception 14 in world, addr 0x0, and lpfc_scsi_cmd_iocb_cmpl.

Resolution
Please refer to the below **Customer Advisory**:
For VMware 5.5 and 6.0 http://h20566.www2.hpe.com/hpsc/doc/public/display?docId=emr_na-a00020440en_us
For VMware 6.5 http://h20565.www2.hpe.com/hpsc/doc/public/display?docId=emr_na-a00022166en_us

Note: http://vibsdepot.hpe.com/hpq/recipes/HPE-VMware-Recipe.pdf -> Please use this link to check the latest VMWARE FW and Software Recipe & update the driver version as per the recommendations here.

TOP

7. New LUNs can no longer be added to ESX 6.0 environment after 3PAR OS upgrade to 3.2.2.GA or higher

Description
The 3.2.2 upgrade on EOS did trigger this issue with new device discovery that the symptom of persistent checksum feature (DIF protection mode) that was introduced in 3.2.2GA.

The HPE 3PAR Persistent Checksum feature, available on HPE StoreServ 20000 Storage and HPE StoreServ 8000 Storage systems (arrays with GEN 5 ASIC), provides end-to-end data integrity protection from the host initiator HBA through the data network to the backend drives on the StoreServ storage system. This feature is based on the DIF protection model defined by the SCSI T-10 committee. It provides data protection from silent data corruption of any media and transmission errors caused by any component in the I/O stack across the data network. Detection and recovery features have been built into the HPE 3PAR OS to correct issues discovered through this feature.
Minimum host HBA driver versions for support of HPE 3PAR Persistent Checksum over FC are required. With ESXi 5.5 and later, supporting drivers are available. With respect to these drivers, the HPE 3PAR Persistent Checksum on HPE

3PAR StoreServ Storage systems is enabled by default, and there is no need to configure anything on the array or the host HBA driver.

Workaround

If there are VM hosts that will attach to 8k/20k arrays, the customer would need to upgrade their HBA firmware/driver.This T-10 DIF error requires an update to HBA firmware/driver for VM hosts. There is a white paper ("VMWare FW and Software Recipe"), published April 2017, that describes this. The white paper is located here:

http://vibsdepot.hpe.com/hpq/recipes/HPE-VMware-Recipe.pdf .

Additional information about the recipe that is customer-consumable can be found in this FAQ document:
http://h20564.www2.hpe.com/hpsc/doc/public/display?docId=c04849476

For 7k and 10k arrays (such as the array that underwent the upgrade), the VMWare host driver has this turned on by default, In these cases, the customer should turn this off.

- To disable 3PAR Persistent Checksum on the host, use the one of the following commands according to the HBA type in use.
    - `QLogic-based HBAs: On the host command line, perform esxcfg-module -s 'ql2xtlodifvendor.0' qlnativefc followed by a host reboot.`
    - `Emulex-based HBAs: On the host command line, perform esxcfg–module -s 'lpfc_external_dif-0' lpf c followed by a host reboot.`
      For more information, please see page 21 in the Vmware ESX/ESXi Impementation Guide, located at:

http://h20566.www2.hpe.com/hpsc/doc/public/display?docId=emr_na-c03290624&lang=en-us&cc=us

TOP

Ontap Hosts shares port with Non-Ontap Host(s)

Environment

HPE 3PAR StoreServ 10000 Storage
HPE 3PAR StoreServ 7000 Storage
HPE 3PAR StoreServ 7450 Storage

1. NetApp hosts needs the host port to be dedicated for Ontap only, sharing the port with a non-ONTAP host is not supported.

Validate if you are using a Netapp V-series controller in your environment .In case of a Netapp V-series controller, kindly note that storage array target ports cannot be shared with other hosts or even with other V-series controllers. Ensure that storage array target ports should be dedicated to the V-Series controller.

2. Check for any ONTAP hosts in the host configuration using Cli command showhost –d

host.xxx.xxxxxxx - ONTAP-legacy 500xxxxxxxxx1x 4:2:3
host.xxx.xxxxxxx - ONTAP-legacy 210xxxxxxxxx6x 5:2:4
host.xxx.xxxxxxx - ONTAP-legacy 500xxxxxxxxx1x 5:2:3
host.xxx.xxxxxxx - ONTAP-legacy 210xxxxxxxxx6x 4:2:4

If the 3PAR OS is below 3.1.2, shared RCFC ports and host ports on the same HBA are not supported.

Solution:

As of HPE 3PAR OS 3.1.2, the HBAs on each HPE 3PAR StoreServ 7000 Storage system can use shared RCFC and host cable connections to connect the systems through the FC SAN. **This configuration is only supported on HPE 3PAR StoreServ 7000 Storage systems running HPE 3PAR OS 3.1.2 or later.**

For other 3PAR Storage models (V, F & T class), NetApp hosts needs the host port to be dedicated for Ontap only, sharing the port with a non-ONTAP host is not supported.

The command controlport config can be used to set the specified connection mode and type on FC ports. It also allows setting the unique node WWN option for the port. When unique_nwwn is enabled, the port presents a unique node name on the connection and this is needed by certain initiators such as ONTAP. Please refer to the HPE 3PAR Command Line Interface Reference for more details for perspective 3PAR OS level.

HPE Insight Control for VMware vCenter

HP Insight Control for VMware vCenter Server does not support cluster provisioning operations for HPE 3PAR Storage arrays running firmware 3.1.2 MU23PAR CLI Client Recommendation. Customer Advisory below:

HP Insight Control for VMware vCenter – Advisory

HPE Cloud systems/Converged systems 3PAR OS

If a cloud system matrix is deployed in the environment, it is important that the HPE 3PAR OS to which the 3PAR Array will be upgraded is compatible with the CSM version.

To validate the compatibility, refer link below:

http://www.hpe.com/go/insightmanagement/docs

Under the "Getting started" section, look for the document called "HP Insight Management Support Matrix (including HP Insight Control and HP Matrix OE support matrix)". Once that is opened, go to the section Supported Storage Products and you will see a chart listing the supported versions of 3PAR OS with MOE and Insight Control.

TOP

VERITAS DMP

NOTE: The following recommendations are from Symantec. Please contact the vendor for further support on this.
The default value of the VxDMP link restore (dmp_restore_interval tunable) interval must temporarily be changed to a value of 1 second before performing an online update of the HPE 3PAR OS for the following VxDMP versions:

- All Linux/Solaris/AIX VxDMP versions older than 5.0MP1
- All HP-UX VxDMP versions older than 5.0MP3

Changing the link restore interval causes the VxDMP software to react to the return of missing paths more quickly than it normally would with default settings. Failure to perform the change to the shorter restore interval for the noted VxDMP versions above, may cause a loss of host connectivity to HPE 3PAR volumes during the upgrade.

For instructions, see "Changing the VERITAS DMP Link Restore Interval" (below).

After the online update, reset the link restore interval to the previously configured value. For instructions, see "Restoring the VERITAS DMP Link Restore Interval" (below).
As of VxDMP 5.0MP1 for Linux/Solaris/AIX and 5.0MP3 HP-UX, Veritas software was modified to circumvent the path loss issue related to the link restore interval during an HPE 3PAR OS upgrade and no changes to the restore interval are necessary.

Changing the VERITAS DMP Link Restore Interval:

The following procedure may be required if you are performing an online update to the HPE 3PAR OS and VERITAS DMP is being used for the multipathing solution on the host. The procedure must be repeated for each host server that is connected to the HPE 3PAR StoreServ Storage and using VERITAS DMP.

Use the following procedure to gather the current settings and change the link restore interval:
Display the current setting of the link restore interval as follows. Record the interval so that it can be restored later.
1. Display the current setting of the link restore interval as follows. Record the interval so that it can be restored later.

 # vxdmpadm stat restored (note the current setting)

2. Stop the daemon.
vxdmpadm stop restore

3. Restart the daemon with the interval set to 1.
vxdmpadm start restore interval=1

4. Verify that the daemon is running with the interval set to 1.
vxdmpadm stat restored

NOTE: The restore daemon reverts to the default value of 300 if or when a host reboot occurs.
Restoring the VERITAS DMP Link Restore Interval

TOP

If the VERITAS DMP link restore interval has been changed to 1 second for HPE 3PAR OS online update, use the following commands to change the link restore interval to the previous setting on storage server hosts after completing an online update. The following procedure must be repeated for each host server that is connected to HPE 3PAR Storage and using VERITAS DMP.

Use the following procedure to restore the link restore interval to its original setting:
1. Stop the daemon.

```
# vxdmpadm stop restore
```

2. Restart the daemon with the interval set to what it was previously.

```
# vxdmpadm start restore interval=<prior setting>
```

3. Verify the daemon is running with the interval set to the prior setting.

```
# vxdmpadm stat restored
```

Sun Solaris

1. Please ensure that Solaris host matches the minimum patch requirements for various versions of Solaris and other associated drivers as listed in the Solaris Implementation guide available at:

Solaris Implementation guide is available at HPE Information Library.

http://www.hpe.com/info/storage/docs

Select:

- Products and Solutions: HPE 3PAR StoreServ Storage
- Information Type: Configuration
- Language: English

2. Solaris using VxDMP 5.0MP3 RP1HF3: dmp_fast_recovery turned off as a minimum.
However, 5.0MP3 RP2HF1 does not require this tunable turned off.

Please contact Oracle for further Support on this.

3. Solaris 10 MU3 with an iSCSI configuration is not supported for online updates.

HPE Open VMS

Recommendations for OpenVMS hosts:

It is important to have all Open VMS hosts running with “recommended” Patch levels prior to performing a 3PAR OS upgrade. HP recommends all hosts be configured as per the Implementation guides. Click on link below and look for HPE 3PAR OpenVMS Implementation Guide:

http://h17007.www1.hpe.com/us/en/storage/info-library/index.aspx
http://h20564.www2.hpe.com/hpsc/doc/public/display?docId=c04739526

TOP

HPE 3PAR Host Based Software

Installing **3.1.2.MU5**, **3.1.3.x** and **3.2.x** version of HPE 3PAR OS requires that all affected software components in the storage environment be upgraded or installed in a specific order as detailed in the following list.

Any listed host applications accessing the HPE 3PAR storage system must be upgraded to the specified software levels in order for the storage environment to take full advantage of the new SSL certificate management functionality.
Depending on the number of host applications in the storage environment, the full installation process might take a considerable amount of time to complete.

You will be required to create a HPE Passport login account to validate the compatibility for HPE 3PAR Host based applications through SPOCK matrix and to download the software provide the contractual support agreement (SAID).

For All host and Array based application (RMC, RMS, SRM, SRA) please refer to HPE SPOCK Compatibility Matrix:
https://h20272.www2.hpe.com/SPOCK/Pages/spock2Html.aspx?htmlFile=sw_array_3par.html

These software can be obtained from HPE USB Portal:
https://h20392.www2.hpe.com/portal/swdepot/displayProductsList.do?category=3PAR

Below is the extract from SPOCK

HPE StoreOnce Recovery Manager Software

- HPE StoreOnce Recovery Manager Central Software
  (326 KB PDF, 2017_02_21)

HPE 3PAR Recovery Manager Software

- 3PAR Recovery Manager Software for Microsoft Exchange
  (19 KB PDF, 2016_10_13)
- 3PAR Recovery Manager Software for Microsoft Hyper-V
  (22 KB PDF, 2016_08_18)
- 3PAR Recovery Manager Software for Microsoft SQL Server
  (78 KB PDF, 2016_12_13)
- 3PAR Recovery Manager Software for Oracle DB
  (200 KB PDF, 2016_08_12)
- 3PAR Recovery Manager Software for VMware vSphere
  (91 KB PDF, 2015_10_30)

Note: -
For the below RMx products, the OS support is up to 322.MU3 as per SPOCK Matrix.
RMH (Recovery Manager for Hyper-V)
RME (Recovery Manager for Exchange)
RMS (Recovery Manager for SQL)
RMO (Recovery Manager for Oracle)
RMV (Recovery Manager for Vmware)

Recovery Manager Central - RMCO/RMCV/RMCS (Recovery Manager Central For Oracle/ Recovery Manager Central For VMware/ Recovery Manager Central For SQL) are supported with 322.MU4 & RMC must be upgraded as per SPOCK matrix (Refer to this link provided above)

The host application upgrades listed in step 1 below can be completed weeks ahead of the remaining upgrade/install steps, if needed.

TOP

1. Upgrade the following applications, in any order desired, to the revision level indicated:
    - Recovery Manager for Exchange (RME)
    - Recovery Manager for SQL (RMS)
    - Recovery Manager for HyperV (RMH)
    - Recovery Manager for VMware (RMV)
    - Recovery Manager for Oracle (RMO)
    - VSS Provider
    - HPE 3PAR Cluster Extension for Windows (CLX)
    - Host Explorer
    - HP Storage Plug-in for SAP Landscape
    - HP StoreFront Analytics Pack for VMware vCOPS™
    - MetroCluster/Serviceguard
2. Upgrade the HPE 3PAR Management Console (IMC /SSMC)
3. Upgrade the CLI and SNMP to respective upgraded OS version
4. Upgrade the HPE 3PAR StoreServ 7000 Storage SmartStart Software (SS)

5. If Recovery Manager Exchange in use, Post OS upgrade please validate the array certificate using the HP 3PAR VSS Provider **provcfg.exe** Please refer to RME user guide for more information- http://h20565.www2.hpe.com/hpsc/doc/public/display?docId=a00008993en_us

For compatible version of RME and VSS provider, kindly refer to HPE SPOCK matrix for more details - https://h20272.www2.hpe.com/spock/utility/document.aspx?docurl=Shared%20Documents/sw/array_3par/3par_rme.2017_05_03.pdf

HPE 3PAR Remote Copy Support – Updated!!

If you have an HPE 3PAR array in remote copy relationship with another 3PAR, the allowed delta between major versions of the HPE 3PAR OS release is limited to adjacent major release levels—Please refer to the below link for details on RC hardware platform compatibility.

http://h20565.www2.hpe.com/hpsc/doc/public/display?docId=c02660486 **Pg.** 13 -21

TOP

Persistent Ports

If SAN/Fabric does not support NPIV, "Persistent ports" functionality on 3PAR needs to be disabled to avoid any path loss issues during the 3PAR OS upgrade. This is applicable for OS upgrades from 3.1.2.x and above.

This can be done using the following command:
setsys PortFailoverEnabled no

Persistent port failback

Issue: During a persistent ports failback, a WWN tried to register on a switch port while that same WWN was still registered on a different switch port. This causes the switch to disable a switch port due to duplicate WWN detection. The switch port has to be manually re-enabled in this scenario.

Brocade FOS 7 switches have a "F-Port login parameter" named "Enforce FLOGI/FDISC login" set to 0. Which means it will NOT allow a duplicate pWWN to be advertised on multiple ports. And due to that it has a tendency to break the 3PAR persistent port feature at times since at failover or failback time it is not possible to get the pWWN moved. When this is determined by the switch it will place the target port into a persistent disabled state.
By changing this parameter value to 2 persistent port will not be affected by this issue and work well. Unfortunately you will have to disable the switch before you can adjust this parameter setting.

More information regarding this FOS parameter settings can be found at:
https://www.brocade.com/content/html/en/administration-guide/fos-740-admin/GUID-53169C82-7080-446F-B4B0-7E282BFEE19C.html - Duplicate PWWN handling during device login
https://www.brocade.com/content/html/en/administration-guide/fos-740-admin/GUID-EDC003D0-470B-42C8-8D24-CEF11DC57CF6.html - Setting the behavior for handling duplicate PWWNs

Switch setting to change the behavior into a port persistent failover friendly state:

```
admin> switchdisable
admin> configshow |grep enforce
switch.login.enforce_login:0 <==
admin> configure

Fabric parameters (yes, y, no, n): [no]
Virtual Channel parameters (yes, y, no, n): [no]
F-Port login parameters (yes, y, no, n): [no] y <==
...
Enforce FLOGI/FDISC login: (0..2) [0] 2 <==
...
D-Port Parameters (yes, y, no, n): [no]
RDP Polling Cycle(hours)[0 = Disable Polling]: (0..24) [0]
..
webtools attributes (yes, y, no, n): [no]

admin> configshow |grep enforce
switch.login.enforce_login:2 <==
admin> switchenable
```

Above configure steps will add (if it does not exist) or change the switch switch.login.enforce_login parameter value from 0 (default) to 2.

Further, the specific value to use to correct this ('2'), requires at least FOS 7.2.x being installed on the FC Switch. FOS 7.0.x and 7.1.x only allows the values 0 and 1, where '1' is not sufficient enough.

On the switch side you can confirm you have hit this particular issue by not only reviewing the state of the port as being flagged as disabled but also by reviewing the errdump output which should contain a message for the affected

port:
2016/06/06-12:44:55, [FCPH-1003], 275, FID 128, WARNING, STESSDFS01, New port 18 has same Port WWN as old port 19 as part of duplicate Port WWN detection policy.

Within example below <n:s:p> refers to the actual port on the array to be changed to normal (e.g. 3:2:1)
1. To assure IO continues on the partner host port, use the "cli% controlport failover <n:s:p>" command.
This command may return an error that there is no partner host port. If so, ignore the error and move on to the next step.

If no error is returned, verify if the host port failed over by giving the "cli% showport <n:s:p>" command, which should show the port successfully failed over. One can also verify that the port is idle by using the "cli% statport –host –iter 5" command and verify that IO traffic indeed moved to the partner port.

2. Set the port offline using the "`cli% controlport offline <n:s:p>`" command.

3. Re-enable the SAN switch port for the affected 3PAR <n:s:p> at the switch level

4. Restart the port using the "`cli% controlport rst <n:s:p>`".
Pause 10 seconds and then issue a "`cli% showport <n:s:p>`" command and verify that the port is as target port and in the "ready" state.

5. If the port shows in the "failed-over" state, so step1) was successful, use the "cli% `controlport failback <n:s:p>`" command to failback IO to the hostport.

Issue a "`cli% showport <n:s:p>`" command to assure the port successfully failed back. The "`cli% statport -host -iter 5`" command should show that the port is handling IO-traffic again.

Note: above steps require 3paradm level of access.
Resolution
This issue is addressed in Inform OS 3.2.2.MU4.

Start-and-stop Quorum witness steps

Before performing an online upgrade to HPE 3PAR OS 3.1.3 MU1 or later, stop HPE 3PAR Quorum Witness, and restart it upon completion of the upgrade.
To stop and start a quorum witness, please follow the steps from the below link –
http://h20565.www2.hpe.com/hpsc/doc/public/display?docId=c02660486 **Pg.** 22 & 23

TOP

Validate 3PAR Array Health

Run the following commands from InServ CLI to validate the status of the array.

Note: Cli% symbol before a command means that it needs to be executed from 3PAR CLI

- `Cli%showversion –a –b` to verify the current InForm OS level on the array
- `Cli%checkupgrade` to verify if the system is ready to undergo an ONLINE upgrade. Sample output below:

```
  == checkupgrade ==
     The current system configuration is valid for proceeding with an online
upgrade.
     --> The system can undergo an online upgrade.
     Return Code = 0
```

- `Cli%checkhealth -svc -detail` performs a full Healthcheck on the system, if there are any issues noticed (such as failed drive), please notify HP support to fix it before scheduling the upgrade.
- `Cli%statcpu -t -d 15 -iter 1` (Applicable only for ONLINE upgrades)

CPU idle time should be more than 50 % during the InForm OS upgrade. As one node in a node pair reboots at a time during the upgrade, the other node has to handle the entire load during that time. If a system is already running low on resources, we do not recommend performing an OS upgrade. We therefore strongly recommend the InformOS upgrade is scheduled during off-peak hours.

- `Cli%statport -d 15 -iter 1 -ni` (Applicable only for ONLINE upgrades)

The I/O load limit recommended for host ports during OS upgrade is as shown below. We therefore strongly recommend the InformOS upgrade is scheduled during off-peak hours. Port Connection parameters can be found using the below commands from the StoreServ CLI

Check the host and "Partner" port connections with showport:
`Cli%showport`
Then, get the configured rates of the primary and partner host ports both with showport –par. We will use the configured rate for the subsequent port I/O load review:
`Cli%showport -par`
Average I/O and KB/s on each port can be found using the below command from the StoreServ CLI:
`Cli%statport –d 15 –host –iter 1 –ni`
For each host port, the average I/O - "KBytes per sec" - should not be greater than:

```
200,000 KBytes per sec for 4 GB connections
400,000 KBytes per sec for 8 GB connections
800,000 Kbytes per sec for 16 GB connections
```

`Cli%statvlun -ni –hostsum`
Shows I/O coming from each host connected to the HPE 3PAR.

TOP

HPE Proliant Model Service Processor

The below-mentioned workaround is applicable only to HP Proliant Service Processors, where a reboot is required (for example applying patch P001 on SP OS 4.3.0, P003 on SP OS 4.2.0, and P011 on SP OS 4.1.0.GA-97 or any other equivalent patches released for addressing the Bash ShellShock vulnerability) or where the SP has to be rebooted as a part of troubleshooting. Please contact the 3PAR support team for assistance with procedure below.

Follow the below-mentioned steps before applying the above-mentioned patches or before issuing the reboot command on HP Proliant Service Processors (Engage the assigned deployment engineer to follow steps below)

Step 1: Login as root
Step 2: Create a copy of the grub configuration file by using the command –
`cp /etc/grub.conf /root/grub.conf.orig`
Step 3: Create a copy of the inittab file by using the command **–**
`cp /etc/inittab /root/inittab.orig`
Step 4: Edit the current grub.conf file by using the command – **`vim /etc/grub.conf`**

Step 5: Delete the entries for console within the grub.conf file as shown in the picture below and save the grub.conf file by using – **`EscKey`** followed by **`:wq`** and exit from the **`grub.conf`** file

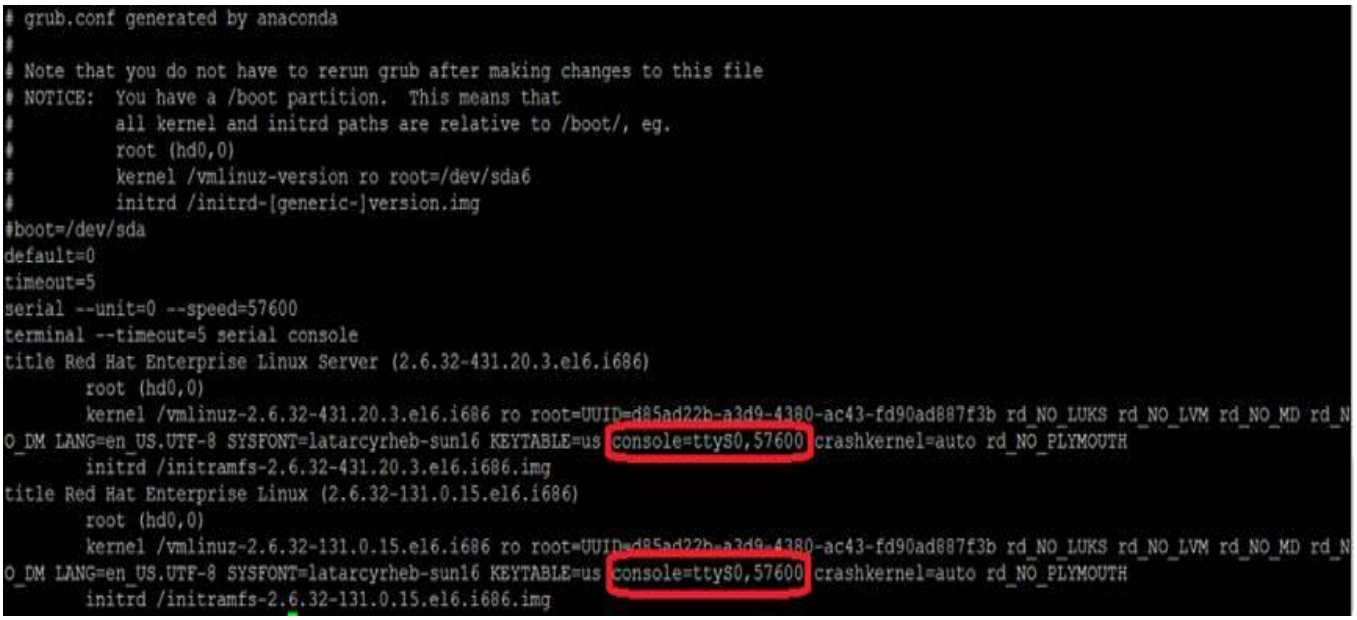

Step 6: Edit the current inittab file by using the command **– `vim /etc/inittab`**

Step 7: Comment out the entry for agetty serial redirection by placing a # symbol at the beginning of the line shown below so that the line should read as –

`#co:2345:respawn:/sbin/agetty 57600 /dev/ttyS0 vt100`

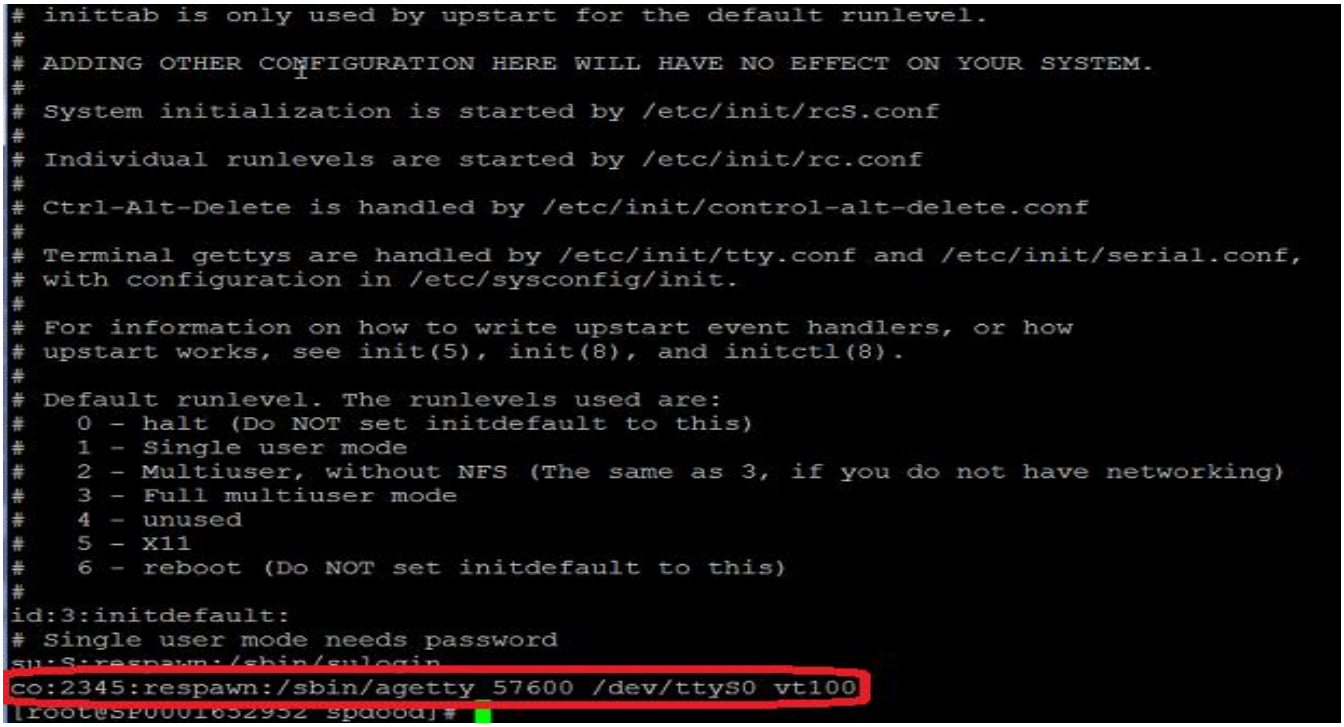

After performing the task involving the SP reboot, follow the steps below to restore the original configuration

Step 1: Once the SP is rebooted, login as root again

Step 2: Copy the original grub.conf file back using the command –

cp /root/grub.conf.orig /etc/grub.conf

Step 3: Verify the console entries are back in the grub.conf file using the command –

cat /etc/grub.conf

Step 4: Copy the original inittab file back using the command –

cp /root/inittab.orig /etc/inittab

Step 5: Verify the agetty serial redirection entries are back in the inittab file using the command –

cat /etc/inittab

TOP

Set network speed to auto

Prior to upgrade to any HPE 3PAR OS, network speed on the admin ports on the 3PAR array needs to be set to auto negotiate. This is to avoid the "netc" process on the 3PAR nodes to hang. This may prevent successful completion of the upgrade or after upgrade prevent netc from starting on some nodes and subsequently cause frequent netc dumps until the issue is addressed.

Follow the workaround to set the speed on admin ports and RCIP ports to "auto negotiate"

A) Run command below to set auto negotiate before upgrade

1) `cli% shownet`

```
IP Address   Netmask/PrefixLen    Nodes    Active Speed Duplex  AutoNeg
Status
10.0.121.221  255.255.248.0        01       0      100    Full   No
Active
Default route: 10.0.120.1
NTP server: None
DNS server: None
```

2) `cli% setnet speep <IP Address of the array> auto`

Note: On successfully setting auto speed it returns: Interface speed set. If not, keep retrying few times and/or consult HP-3PAR.

3) `cli% shownet`

```
IP Address  Netmask/PrefixLen    Nodes    Active Speed Duplex  AutoNeg Status
10.0.121.221  255.255.248.0        01       0      100    Full   Yes
Active

Default route: 10.0.120.1
NTP server: None
DNS server: None
```

B) Setting auto speed negotiate on rcip ports

1. Get list of rcip ports:

```
cli% showrctransport -rcip
N:S:P   State       HwAddr   IPAddress PeerIPAddress       Netmask    Gateway  MTU  Rate Duplex
0:9:1   ready 0002AC80041D 10.97.13.32   10.97.13.33 255.255.255.0 10.97.13.1 1500 1Gbps   Full
1:9:1   ready 0002AC802ADF 10.98.13.32   10.98.13.33 255.255.255.0 10.98.13.1 1500 1Gbps   Full
```

2. Setting auto speed for rcip

Usage: `controlport rcip speed auto <node:slot:port>`

Example:

```
cli% controlport rcip speed auto 0:9:1
```
Remote Copy interface change successful

```
cli% controlport rcip speed auto 1:9:1
```
Remote Copy interface change successful

```
cli% controlport rcip speed auto 2:9:1
```

2:9:1 is not configured for remote copy. <-- Ignore it as port is not configured as rcip

```
cli% controlport rcip speed auto 3:9:1
```
3:9:1 is not configured for remote copy. <-- Ignore it as port is not configured as rcip

Note: For valid ports if you find any errors keep retrying the command few times and/or consult HP-3PAR GDC Team.

AO, DO recommendations

Recommendations for AO & DO

A) Stop AO, DO, and RC tasks before the upgrade

All system administration activities must be stopped during an InForm OS online update. All active tasks must complete or be stopped before proceeding. This includes any provisioning activity, physical/virtual copy activity, and dynamic optimization. Be sure to stop any automated administration scripts that may be running as well (Recovery Manager and VSS Provider for Microsoft Windows).

If this system is licensed for Adaptive Optimization (AO), please check if the Adaptive Optimization (AO) cycle is going to overlap with the OS upgrade window. If it will, please deactivate AO before OS upgrade window and reactivate once OS upgrade is complete.
To stop AO - up to InForm OS 3.1.1, please launch System Reporter, click on "Policy Settings" on top menu bar-->a new window pops-up-->Click on "Adaptive Optimization" tab-->clcik "Change" next to the InServ to be upgraded (you need to do it for all Polcy Ids) --> Change "Configuration Active " to False

B) On Node AO:

The data for AO is now maintained on the StoreServ Storage in an internal database. AO configured from external System Reporter will not work for 3.1.2 systems, region mover log will generate the following error: No license for Adaptive Optimization.
All existing system reporter based AO configurations will need to be manually recreated through CLI or IMC. Please refer to the attached procedure -On Node SR & AO - Inform OS 3.1.2 Upgrade Process- to port the existing System Reporter based AO configuration to 3.1.2 node based AO configuration. Please let us know in advance if any help is needed from us in migrating the AO configuration after the upgrade

To stop Remote Copy on both primary and DR sites:

stoprcopygroup [option] <group_name>

Stopping a Remote Copy group while the group is currently syncing will delay a return of the command.

To stop DO - please cancel all tune tasks (e.g. tunevv)

TOP

During HPE 3PAR OS upgrade

IBM AIX

HPE 3PAR OS online automatic upgrades with AIX Direct Connect attached hosts are not supported by HP for any HPE 3PAR OS upgrades. Online maintenance upgrades can be successfully performed by means of a node-by-node (or advanced) upgrade method. Using this method will cause the system to wait for user input between each node reboot needed during an online HPE 3PAR OS upgrade.

This pause time in the upgrade process is used to allow all of the I/O paths to the AIX direct-connect hosts, for all of the presented LUNs, to recover (fail back).

To display the paths that must be monitored for recovery, use the AIX lspath –s failed command on the attached host:

```
# lspath –s failed
Failed hdisk4 fscsi0
Failed hdisk5 fscsi0
...
#
```

After all the paths have recovered, continue with the next node of the HPE 3PAR OS.

```
# lspath –s failed
```

Sun Solaris

As of HPE 3PAR OS 3.1.1 GA, online updates can be successfully performed by means of a node-by-node (aka "Advanced") update and pausing between HPE 3PAR StoreServ Storage node updates. Wait for all I/O paths for all Sun Cluster nodes to recover during each pause between HPE 3PAR StoreServ Storage node updates. Refer to "I/O Path Monitoring for Solaris" page in the Solaris section of the "***HPE 3PAR Operating System Upgrade Planning Guide***". Link below:

http://h20565.www2.hpe.com/hpsc/doc/public/display?docId=c02660486 **Pg.** 26-29

TOP

HP-UX Troubleshooting

The following HPE 3PAR OS online upgrades have been determined to cause HP-UX servers to report device discovery problems until the upgrade is complete:

- HPE 3PAR OS 3.1.1 (including all MUs) to 3.1.2 MU1, 3.1.2 MU2, or 3.1.2 MU3
- HPE 3PAR OS 3.1.2 GA to 3.1.2 MU1, 3.1.2 MU2, or 3.1.2 MU3

Issue

If a rescan of the HP-UX I/O system hardware is run during the HPE 3PAR OS online upgrade process, ioscan might report device paths with a Software State of NO_HW. This is due to a defect that causes an unexpected SCSI inquiry response to LUN 0 for HP-UX hosts during the HPE 3PAR OS online upgrade. It was verified that I/O is not affected and paths remain accessible as long as the server is not restarted during the online upgrade process. When the entire online upgrade process is complete, the SCSI inquiry response to LUN 0 is normal and a rescan reports device paths with a Software State of CLAIMED

```
#ioscan -fnC disk
Class | H/W Path         Driver S/W State H/W Type  Description
============================================================================
disk 67 0/0/0/5/0/0/0.20.131.97.0.0.1  sdisk  NO_HW   DEVICE  3PARdataVV
     /dev/dsk/c15t0d1 /dev/rdsk/c15t0d1
disk 69 0/0/0/5/0/0/0.20.131.97.0.0.2 sdisk  NO_HW   DEVICE  3PARdataVV
     /dev/dsk/c15t0d2 /dev/rdsk/c15t0d2
```

If an HP-UX server is restarted during an HPE 3PAR OS online upgrade, 1 or more previous known LUN paths might not be discovered by the HP-UX server. The HP-UX operating systems will not continue device discovery when a LUN 0 SCSI Inquiry response returns an unexpected value. Once the upgrade is complete, manual intervention by a system administrator is required for the restarted HP-UX server to rediscover HPE 3PAR devices as described in the next section

Resolution

NOTE: This issue was resolved with HPE 3PAR OS 3.1.3.GA. If you are upgrading to any of the affected HPE 3PAR OS versions listed above, follow the steps listed below.

During the HPE 3PAR OS Online Upgrade process, a rescan of the HP-UX11.31 I/O system hardware reports device paths with a Software State of NO_HW. Perform the following steps to check the health of the data LUN paths and confirm they are reported as online.

1. Ensure that the HP-UX 11.31 server was not restarted.
2. Scan the HP-UX system hardware and identify the hardware paths for any devices that report a Software State as NO_HW.

```
#ioscan -fnC disk
Class | H/W Path         Driver S/W State H/W Type Description
============================================================================
disk  67 0/0/0/5/0/0/0.20.131.97.0.0.1 sdisk  NO_HW   DEVICE 3PARdataVV
   /dev/dsk/c15t0d1   /dev/rdsk/c15t0d1
disk  69 0/0/0/5/0/0/0.20.131.97.0.0.2 sdisk  NO_HW   DEVICE 3PARdataVV
   /dev/dsk/c15t0d2   /dev/rdsk/c15t0d2
```

3. Verify the health of the data LUN path using the hardware paths from the previous step.

TOP

```
#ioscan -P health -H <hardware_path>
#ioscan -P health -H 0/0/0/5/0/0/0

Class | H/W Path                  health
================================================================================
lunpath 9 0/0/0/5/0/0/0.0x23240002ac000e2c.0x0       offline
lunpath 68 0/0/0/5/0/0/0.0x23240002ac000e2c.0x4001000000000000  online
lunpath 77 0/0/0/5/0/0/0.0x23240002ac000e2c.0x4002000000000000  online
```

If the LUN paths display as online, I/O is not affected and paths remain accessible as long as the server is not restarted during the online upgrade process.

NOTE: HP-UX OS version 11i v2 (11.23) and earlier do not provide the health check utility, however it has been verified that even though the devices are in the NO_HW state, the online upgrade will resume without I/O interruption. If an HP-UX host is restarted during the online upgrade operation, the following steps must be performed to re-discover 3PAR devices:

1. Ensure all InServ nodes successfully completed the online upgrade and show the correct HPE 3PAR OS version as follows:

```
# cli upgradesys -status
System is not currently undergoing an online upgrade.
First node to be rebooted is node 3
   Which will be followed by node 0
   Which will be followed by node 2
   Which will be followed by node 1
```

```
# showversion –b
Release version 3.1.2.322 (MU1)
Patches: P05, P13

Component Name      Version
CLI Server          3.1.2.322 (MU1)
CLI Client          3.1.2.322 (MU1)
System Manager      3.1.2.428 (P05)
Kernel              3.1.2.322 (MU1)
TPD Kernel Code     3.1.2.322 (MU1)
TPD Kernel Patch    3.1.2.466 (MU1)
```

2. Scan the HP-UX system hardware for usable I/O system devices:

```
#ioscan –fn
```

3. If you are using LVM/PVLinks, scan the physical volumes for LVM volume groups and activate them:

```
#vgscan -v –a
```

Run the following command for each volume group using 3PAR LUNs:

```
#vgchange -a y <VG Name>
```

4. If you are using VxVM, rebuild the volume device nodes and start the volumes:

#vxdctl enable

#vxvol –g <Disk Group> start <Volume Name>

Verify that the volumes are online:

#vxprint –ht –g <Disk Group>

TOP

Post HPE 3PAR OS upgrade

Host Persona – Updated!!

ONLY for 3.1.3.x and higher Upgrades - Persona 15 for Windows 2008 and Windows 2012 Server:
New Host Persona 15 for Windows Hosts Needs Set after upgrading to HPE 3PAR OS 3.1.3.x and higher to facilitate appropriate LUN behavior.

Details: HPE 3PAR Windows Server 2012 and Windows Server 2008 Implementation Guide, http://h20628.www2.hpe.com/km-ext/kmcsdirect/emr_na-c03290621-14.pdf , See page 8.

Previously with HP3PAR OS 3.1.1 and 3.1.2, KB2849097 was required to present the LUNs appropriately. The new Persona 15 presents LUNs in the manner Microsoft Windows expects.

Action: After upgrading to HPE 3PAR OS 3.1.3.x or higher , HP requires host persona 15 to be used for Windows Server 2008, 2008 R2, 2012, or 2012 R2 to prevent this issue from occurring. Changing to host persona 15 is an online procedure.

"HP guideline is to mandatorily update to the latest Compatible version of Host Explorer before changing the host persona to15. If the Host Explorer application is in use, please check compatibility of Host Explorer with HPE 3PAR OS using link below & update the Host Explorer before Host persona is changed"

Compatibility information for HPE 3PAR including Host Explorer is available at:
http://h20272.www2.hpe.com/Pages/spock2Html.aspx?htmlFile=sw_array_3par.html

Download latest version of host explorer, please visit the following HP URL:
https://h20392.www2.hpe.com/portal/swdepot/displayProductInfo.do?productNumber=HP3PARHE

Kindly refer to ***HPE 3PAR Operating System Upgrade Planning Guide*** for all persona related Information
http://h20565.www2.hpe.com/hpsc/doc/public/display?docId=c02660486 **Pg.** 9-12

TOP

Updating GUI/CLI Clients – Updated!!

1. If you are running HP/3PAR software that uses 3PAR CLI, (i.e.: System Reporter, Cluster Extension Software), please ensure that you upgrade the CLI of the hosts to latest CLI client.

To download 3par software (e.g. System Reporter), please visit the following HP URL:
https://h20392.www2.hpe.com/portal/swdepot/displayProductsList.do?category=3PAR

You will be required to create a HP Passport login account and provide the contractual support agreement (SAID).

2. Post OS upgrade to **312.MU5,313.x and 321.x**, upgrade the following applications to the revision level indicated in any order desired:

- EVA to 3PAR Online Support
- Storage Replication Adapter for VMware depending on your environment. See "Storage Replication Adapter for VMware Support Matrix" (page 9) for more information.
- SCVMM UI add-in
- System Reporter (SR)
- HP Metrocluster / Serviceguard

TOP

HPE 3PAR OS 3.3.1.X specific updates

General Guidelines

✓ Virtual service processor are supported only on ESXi Version 5.5/6.0/6.5 or Hyper-V 2012/2012R2/2016 & reconfigure VM to meet the SP 5.0 requirements (Memory >= 4GB, Virtual CPU cores >= 4)
Note: For Hyper-V, do not reconfigure virtual memory until after the upgrade to SP 4.5 is complete.

✓ SP 5.0 User credentials have changed as below :-

Service Processor	Prior to Upgrade User name	Post SP OS 5.0 upgrade User name	HPE standard password
TUI /CLI /3PAR Service Console	3parcust	admin	3parInServ

Note: The passwords for the accounts has been preserved during the upgrade. Use the current password with the new user name to login into SP 5.0

✓ The network port the Service Processor's webserver listens on has changed from 443 to 8443. Ensure that you will be able to connect from your web browser to the Service Processor (Service Console) on port 8443, which may require firewall rule updates in your environment.

✓ Use the following URL in your web browser to connect to the 3PAR Service Console: https://<SP IP Address>:8443

✓ Ensure to validate the '3parcust' credentials for service processor prior to the SP upgrade & in case it is not working please do engage support to reset the same.
Note: The only way to recover the credentials post the SP 5.0 upgrade is to rebuild it.

HPE 3PAR OS 3.3.1 GA/EGA DO NOT run tunesys/tunevv on a TDVV3 DDS volume

Issue : When the TDVV3 DDS is rebalanced to include new nodes (tuned to additional nodes by, for example, tunesys), the new SA space LDs created as destination space for the tune are not marked as exclusive to the DDS - they do not have the for_dds_sa flag set. This means that the SA LDs can be used by any other VV in the system for growth and may cause issues to the different internal structure of these TDVV3 SA space LDs

This can also be seen with a direct tunevv (conversion) of the DDS, e.g. DO tune to a different RAID type. This is because the problem is in the DDS SA space where the code to tune DDS SA space to new nodes has not been implemented

Workaround

A workaround is provided in 3.3.1 MU1:

- With 3.3.1 MU1 DDS volumes will be skipped during tunesys operations.
- All tunes, including direct tune operation, such as tunevv or conversions via the SSMC UI, will be disallowed on dedup DDS volumes.

Update

- A limitation exists whereby the Dedup data store (DDS) volume for TDVV volumes cannot be tuned in 3.3.1.MU1. This restriction will be removed in a later release

CLI Error Message:

"Dedup Data Store volume <DDS VV> cannot be tuned in this release."

Fix

This issue is addressed with 3.3.1 MU2, and the restriction will be removed.

TOP