

# **InMage Scout Standard Release Notes**

Version – 8.0.1 Update 7

**Table: Document History**

<b>Document Version</b>	<b>Document Date</b>	<b>Remarks</b>
1.0	March 1, 2015	Standard version
1.1	April 7, 2015	Minor update
1.2	Nov 20 ,2015	Updated CX and vContinuum MT known issues and limitations sections.
1.3	May 2, 2017	Scout Update 5 <ul style="list-style-type: none"><li>- Updated known limitation for SLES 11</li><li>- Added steps to upgrade protected P2V Windows Cluster to Scout Update 5</li></ul>
14	Oct 6,2017	Updated P2V Windows cluster upgrade section as per Update 6 and future update release
15	Dec 31,2018	Added link for steps to upgrade MySQL and PHP libraries in CX and RX servers.

## Contents

<b>1</b>	<b>Disclaimer of Warranty</b> .....	<b>4</b>
<b>2</b>	<b>About this document</b> .....	<b>4</b>
<b>3</b>	<b>Overview</b> .....	<b>5</b>
3.1	Scout .....	5
3.2	ScoutCloud RX.....	5
<b>4</b>	<b>What's new in this release?</b> .....	<b>5</b>
<b>5</b>	<b>Upgrade Path</b> .....	<b>5</b>
5.1	Agents .....	5
5.2	RX.....	6
5.3	CX Server .....	6
5.4	vContinuum .....	6
<b>6</b>	<b>Upgrade sequence</b> .....	<b>7</b>
<b>7</b>	<b>Installation/Upgrade Instructions</b> .....	<b>7</b>
<b>8</b>	<b>Supported Platforms</b> .....	<b>7</b>
<b>9</b>	<b>Issues Addressed</b> .....	<b>7</b>
9.1	Scout vContinuum.....	7
9.2	CX Server .....	8
9.3	VX Agent.....	8
9.4	RX.....	8
9.5	FX Agent .....	9
<b>10</b>	<b>Known Issues and Limitations</b> .....	<b>9</b>
10.1	CX .....	9
10.2	Agent .....	9
10.3	Fx.....	10
10.4	Scout vContinuum MT .....	11
10.5	ScoutCloud RX.....	14
<b>11</b>	<b>Sizing guideline for CS/PS/MT</b> .....	<b>15</b>
<b>12</b>	<b>Upgrade protected P2V Windows Cluster to Scout Update5 or above</b> .....	<b>15</b>
12.1	When do you need to use these manual steps? .....	15
12.2	Manual Steps .....	15

## 1 Disclaimer of Warranty

InMage Systems makes no representations or warranties, either express or implied, by or with respect to anything in this document, and shall not be liable for any implied warranties of merchantability or fitness for a particular purpose or for any indirect, special or consequential damages.

**Copyright (c) 2008-2014 InMage Systems Inc. All rights reserved.**

"**InMage Systems**" and InMage's products are trademarks of InMage Systems Inc. References to other companies and their products use trademarks owned by the respective companies and are for reference purpose only.

## 2 About this document

The purpose of the Release Notes is to communicate major changes as well as new features in this release of the Scout. It also documents known issues and their workarounds.

The 8.0.1 GA release of Scout introduces new features, enhancements to existing features, bug fixes and support for new applications and new platforms.

## 3 Overview

### 3.1 Scout

Scout Software is a continuous data protection product that allows users to asynchronously replicate their servers (physical or virtual) over LAN and/or WAN for local and/or remote recovery with near zero data loss.

The software is used by enterprises for disaster recovery (DR) and by Managed Service Providers (MSPs) to offer DR as a service (DRaaS) to their customers.

### 3.2 ScoutCloud RX

ScoutCloud RX is a multi-tenant portal intended for use by MSPs and end-customers with large Scout deployments to monitor the health of protected servers in a centralized fashion.

## 4 What's new in this release?

- Integration with Azure billing. Scout now sends a count of the number of instances protected by Scout to Azure Site Recovery, thereby allowing your usage of Scout to be tracked against your Azure bill.
- Simplified MSCS Physical servers or VMs protection and recovery through vContinuum wizard.
- Improved Configuration Server (CS) performance by optimizing memory usage.
- Bug fixes.

## 5 Upgrade Path

Scout 8.0.1 is a full build. User can directly upgrade from 7.x to 8.0.1 version.

### 5.1 Agents

Agent upgrade is supported from Scout 7.0 GA/7.1 GA/8.0 GA to Scout 8.0.1 GA. Agent upgrade is supported irrespective of updates installed.

Table 1

Base Release	Release Update	Upgrade to 8.0.1 GA Supported?
Scout 7.0 GA	Any Update	Yes
Scout 7.1 GA	Any Update	Yes
Scout 8.0. GA	Any hotfix	Yes

## 5.2 RX

RX upgrade is supported from Scout 7.0 GA/7.1 GA/8.0 GA to Scout 8.0.1 GA. RX upgrade is supported irrespective of updates installed.

**Table 1**

Base Release	Release Update	Upgrade to 8.0.1 GA Supported?
Scout 7.0 GA	Any Update	Yes
Scout 7.1 GA	Any Update	Yes
Scout 8.0 GA	Any hotfix	Yes

## 5.3 CX Server

CX Server upgrade is supported from 7.0 GA /7.1 GA/8.0 GA to Scout 8.0.1 GA. CX Server upgrade is supported irrespective of updates installed.

**Table 2**

Base Release	Release Update	Upgrade to 8.0.1 GA Supported?
Scout 7.0 GA	Any Update	Yes
Scout 7.1 GA	Any Update	Yes
Scout 8.0 GA	Any hotfix	Yes

## 5.4 vContinuum

vContinuum upgrade support is available from vContinuum previous releases 4.0(7.0 GA)/4.1(7.1GA)/8.0 GA to Scout 8.0.1GA. vContinuum upgrade is supported irrespective of updates installed.

Table 3

Base Release	Release Update	Upgrade to 8.0.1 GA Supported?
vContinuum 4.0(Scout 7.0 GA)	Any Update	Yes
vContinuum 4.1 (Scout 7.1 GA)	Any Update	Yes
vContinuum 8.0 GA	Any hotfix	Yes

## 6 Upgrade sequence

If earlier version of Scout already present, upgrade sequence should be as below.

1. Upgrade Rx (if exists)
2. Upgrade CS +PS
3. Upgrade PS (if exists separately)
4. Upgrade unified agents on secondary servers (target side)
5. Upgrade vContinuum wizard (if exists)
6. Upgrade unified agents on primary servers (source side)

## 7 Installation/Upgrade Instructions

**Fresh deployment:** Refer to InMage\_Scout\_Quick\_Install\_Guide.pdf for instructions to deploy and configure Scout components (Configuration Server(CS), Process Server (PS), Unified Agent (UA), and vContinuum) and ScoutCloud RX.

**Upgrade:** Refer to the upgrade section in Scout User guide.

## 8 Supported Platforms

Refer to Scout\_Compatibility\_Matrix.pdf for information on the operating system platforms on which Scout and ScoutCloud RX software can be deployed.

## 9 Issues Addressed

### 9.1 Scout vContinuum

1. For P2V scenario, when source server disk order gets changed due to reboot, for that protected server DR-Drill or Recovery operation will fail.
2. Retention logs are not getting removed when “remove VM” operation is used.
3. Change of IP fails in recovered VM when protected VM has VMXNet3.
4. vContinuum discovery takes more time in large ESX environment where lots of guest OS and ESX Clusters are present.

5. For P2V scenario, when hardware based services are present (HP services), recovered VM takes longer time to come up. It needed manual intervention to make the service start type as manual
6. On recovered Linux VM, configuration files permissions are not preserved for files like /etc/hosts, /etc/nsswitch.conf..etc
7. If source Linux server has virtual IPs, then in recovered VM, virtual IPs are not set properly.
8. Sometime in P2V scenario, recovered VM is failing to boot with the kernel panic error because vContinuum was cleaning the root data.
9. Resume protection fails in case of Windows server having EFI partition.
10. vCenter discovery fails for non-English language.
11. Remove and Recovery operation fails when vCenter credential get changed.

## 9.2 CX Server

1. For DR-Drill or Recovery operation of P2V, Fx job fails on copy and cause recovery to stuck at powering ON the recovered VM.
2. When one of the servers in the plan is decommissioned then it removes consistency jobs for all other servers protected in the same plan.
3. vContinuum wizard fails to show the plan details as it is unable to load the page.
4. In case of 1 to N, deletion of one pair sometimes causing Out of order for differential files.
5. Policy violation events table filling up without pruning and causes to out of memory.
6. Fx traps are not functioning as expected
7. CX does not retain the value of "Use process server NATIP address" check box.
8. Duplicate volsync processes are causing out of order differential files.
9. Addressed the Poodle Vulnerability(tagged as CVE-2014-3566 in the NVD) in Configuration Server.
10. When more than twenty replication pairs are assigned to same PS, then values of health report and bandwidth are not getting updated on Linux CS.

## 9.3 VX Agent

1. Data inconsistent issue on SUSE non-root volume.
2. In P2V case, recovered RHEL6UP5 64bit VM fails to boot after recovery.
3. MSCS Cluster storage validation tests fail when UA is installed on cluster nodes
4. On MSCS node sometime vxagent get crashed while booting the server.
5. When secure mode is enabled from source to PS, after sometime, replication pairs stop progressing and vxagent heartbeats will not get updated in CX.
6. When involflt driver is loaded after vxagent service starts, replication will not progress.

## 9.4 RX

1. Added fix to display authorized info for the respective customer.
2. Fixed XSS security issue.
3. Addressed the Poodle Vulnerability( tagged as CVE-2014-3566 in the NVD) in RX.



## 9.5 FX Agent

1. Fx copies non modified file to target server. With the fix Fx now copy only those files that are modified or failed to copy since last fx job run.

## 10 Known Issues and Limitations

### 10.1 CX

1. A replication pair cannot be moved from one protection plan to another.
2. On detecting a resize of a protected volume, CX sets the volume replication pair's *Resync Required* flag to *Yes* and pauses replication for that volume. This is to allow the Scout administrator to also resize the corresponding target volume. However, CX allows the administrator to resume replication without validating the target volume size. If replication is resumed without resizing the target volume appropriately, it would encounter disk write errors when attempting to write to unavailable blocks on that volume.
3. If force unregister is used on the Configuration Server, the unregister operation will be triggered but will not be validated for successful completion by the Configuration Server. If for some reason the action should fail to effect the unregister operation with Azure Site Recovery, stale registrations may continue to exist in the Azure Site Recovery vault. The user will need to contact support in order to delete the stale entries that may continue to persist in Azure Site Recovery.
4. Restore of the CS to a different server from CS backup requires to copy the certificates of the original CS server to new CS server.
5. After upgrading CX, the "Auto Refresh" check box is getting unchecked for "Protection Health", "CS/PS Health" and "Alerts and Notifications" dock of CX Dashboard. User need to set the dashboard preferences freshly.
6. "Reboot required" option is in disabled state while doing the push upgrade. User need to manually reboot the upgraded machines.
7. Standby CX is not supported on Windows CX.
8. CS and PS IP should be static. Change in CS/PS IP will break the communication between agent and CS/PS and replication pair will not progress.
9. Process Servers are not listed on Network traffic page.

### 10.2 Agent

1. After installing agents on all the nodes of the Microsoft cluster nodes, , if a new cluster disk is added to the nodes all the cluster nodes need reboot before protecting the disk. Without reboot, replication will be marked for "resync" on subsequent failover/failback of the newly added disk. The resync notification will be indicated on CX-UI. In this case, restart resync after rebooting all the cluster nodes.
2. Protection for dynamic disks on MSCS cluster (with Symantec Storage Foundation) is not supported.
3. When source system time jumps from point A to future time B and back to a point between A&B then, until the time on production system reaches the time > B, all recovery points will show same timestamp B. Recovery using timestamp is not supported till this event. However, you can perform recovery using bookmarks.

4. Caution is required during uninstallation of EMC PowerPath or InMage Linux agent, when both are installed on the server. The uninstallation of these should happen exactly in the reverse order in which they are loaded to avoid undesired issues.
5. Recovery ranges shown based on source machine's local time on CX UI do not have daylight saving time (DST) consideration.
6. For Windows, while replication on basic volume is in progress, converting source disk storage type from basic to dynamic would mark volume replication for resync on next reboot.
7. For Windows, while replication is in progress, format of source volume may lead to resync.
8. Source device protection of disk devices that are renamed using Linux udev rules is not supported.
9. Scout does not support protection of a volume that has character '\$' in its name.
10. Mount point stale entries of recycle bin protection causes the replication progress stuck.
11. S2 process crashed when protecting more than one name to the same device.
12. Scout does not support protection of "cryptsetup" encrypted devices.
13. Protection of dev-mapper logical volumes created from disks/partitions on Linux using dmsetup is not supported.
14. System level protection for SuSE 10 & 11 requires resynchronization (restart resync) after a system reboot.
15. When protected source volume is formatted then Resync flag will be set to yes for the protection.
16. Consistency job starts failing after deletion of some of the replication pairs when Scout filter driver still tracks the changes for the deleted pairs.
17. For MSCS, if a volume is active on node 1 with drive letter say 'D', then on node 2, a newly added local disk volume takes up drive letter as 'D' causing Resync required set to yes as Scout expects 'D' to be offline on node 2.  
Workaround: Always choose unused shared drive letter for local volume.
18. For MSCS, say volume D and E are shared volumes online on node 1 and 'F' is newly added shared volume from cluster node 1. If all the volumes are failedover to node 2, the volume 'F' for some seconds comes up with drive letter 'D' on node 2. This causes Resync required set to Yes.

### 10.3 Fx

1. When a new Fx job is created and run immediately while any other Fx job is running on the target then newly created job may fail until existing job completes. This happens only when Fx job is set for two different machines. If source and target machines are same then issue will not occur.
2. Only Push mode is supported.
3. Fx replication between 8.0.1 agent and earlier version(<8.0) is not supported. Both source and target host should have either 8.0.1 agent or both should have earlier version of agent.

For example Fx job between source( 8.0.1 agent) and target (7.1 agent) is not supported.

## 10.4 Scout vContinuum MT

1. If the device order on the protected Linux virtual machine changes either due to reboot or due to reconfiguration, replication will not progress. Workaround: Remove the VM from protection and re-protect the VM.
2. Unselect "Microsoft VHD disks" while protecting Physical Machines. Protection will fail if VHD disks are selected.
3. Reverting to a VMware snapshot will require resync operation.
4. Recently added NICs shows only after 15 minutes in vContinuum wizard.  
Workaround: User can restart "Svagent(InMage Vx replication)" service on production server.
5. At times, Windows 2008 recovered VM's RAID 5 and Mirror volumes show up as "Failed redundancy". However, data on the volume is accessible and applications will continue to start & function.

**Workaround:**

- a) Open disk management, select one of the dynamic disks, right click and click on "reactivate disk".
  - (or)
  - b) Right-click on volume which is showing "Failed redundancy" and reactivate volume.
6. In case of cloning protected virtual machines and bringing up cloned virtual machines on same production network, replication will not progress and replication pairs will be marked for *Resync Required* to "YES".
  7. Protection will fail if VM's disks are located at 0:0, 1:0, 2:0, 3:0 with same size. This would result into scsi-id mismatch issue as OS level and Host level disk comparison fails.
  8. In case of Linux, if system partitions are spread across multiple disks then all disks need to be protected with a fresh protection plan. Add disk feature for system disks or having system partition is not applicable.  
**Workaround:** Fresh protection needs to be setup considering all system devices.
  9. Linux physical-virtual or virtual-virtual workflow does not support security policies with automation.
    - a) SELinux workaround: Target machine boots into maintenance mode. After passing selinux=0 in kernel boot parameters, system comes up fine. Similar change needs to be done in /boot/grub/menu.lst file for permanently fixing the problem after system boots up.
    - (b) Active AppArmor security policy is not supported
  10. Linux virtual to virtual workflow does not support device name schemes (i) /dev/disk/by-path (ii) /dev/disk/by-id in /etc/fstab & /boot/grub/menu.lst.

- a. If system disk uses above naming scheme, system fails to boot. Workaround: Root device should be made as /dev/sdaX in GRUB kernel entry depending on the order on source machine, where X is corresponding partition number
- b. User disk or partition in /etc/fstab with above naming scheme drops system boot process into maintenance if filesystem is auto mounted. Above device name scheme needs to be changed to appropriate device name in recovered virtual machine  
Physical to virtual does not have this limitation.

11. In Linux P2V scenario, if system disk is in LV without VG, then the system disk will not be reported to CX and vContinuum cannot protect that disk.
12. Cluster node IP's are not updated for second physical NIC in the failback recovery with error "Object already Exists".

**Workaround:** Set the IP address manually.

13. With EISA partition on production, dynamic mirror volumes will be converted to simple volumes on secondary site. However, it will not have any issues for accessing the drive and all

applications will continue to work.

**Workaround:** Re-create mirror before trying for failback operation.

14. Resume protection failed after datastore vMotion for DR-VM as datastore vMotion caused vmdk paths and vm folder to change.
15. Import dynamic disk(s) manually if any disk(s) is/are in foreign state on recovered VM/DR-Drill VM.
16. During recovery, if recovery point is selected with specific time option from Scout vContinuum wizard and if the specified time matches with the exact tag time, recovery will happen to earlier tag.
17. Linux MT name should not contain "." else readiness check will fail while performing recovery though RX.
18. After PS failover, replication pairs will not progress in resume scenario if first process server is down.  
Workaround: If replication pairs don't progress because of PS server unavailability, user need to click on each pair-->click on settings-->choose new progress server.
19. For V2V scenario, protection of Linux VM will fail when it has LSI logic SAS controller having more than six disks.
20. Unprotected MSCS node can't be added to the existing protected cluster plan. To protect new cluster node to the existing cluster, user need to remove the existing protection and re-protect the complete cluster.  
For example one node is protected out of two node MSCS cluster and then adding second node to protection will fail. Workaround is to remove the existing protection and protect both the nodes at the same time.
21. For MSCS, after upgrading to 8.0.1, it is recommended to remove existing protection plan and create new protection plan to make avail of all the new features added in Scout 8.0.1 for MSCS. For existing MSCS protection plan after upgrading agent to 8.0.1 the limitation remain the same as that on earlier version like volume resize and add disk will have the same limitation until MSCS nodes are freshly protected.
22. Delete disk and remove volume will not be allowed for local disk/volumes of protected MSCS node.
23. Resume operation for MSCS servers may fail when recovered MSCS VMs are in the same network as that of CS.  
Workaround: Ensure recovered MSCS cluster nodes are in powered off state.  
On all source cluster nodes:
  1. Stop VX Agent service
  2. Delete these keys if they are present  
HKEY\_LOCAL\_MACHINE\SOFTWARE\SV Systems\ClusterTracking  
HKEY\_LOCAL\_MACHINE\SOFTWARE\Wow6432Node\SV Systems\ClusterTracking
  3. Start VX Agent service
24. For MSCS cluster, disk deletion operation may not remove replication pairs of the volume of the disk. Workaround: Perform "Remove volume" operation from vContinuum wizard for the volumes whose pairs were not removed as part of remove disk operation.
25. When two different plans are recovered from two different MTs at the same time then it is possible that one of the plan may not be listed for resume operation.
26. Re-sizing of local disks is not supported in MSCS cluster protected machines in vContinuum.

### **Limitations:**

1. Recovery will fail if vContinuum is installed on the same server where CX is installed. Do not install vContinuum on the server where CX is installed.

2. Hostname of the Master Target cannot exceed 15 characters.
3. Master Target should not be under a vApp pool.
4. If vContinuum wizard shows incorrect free space of datastore, close the wizard, refresh datastore from vSphere client and launch vContinuum wizard again.
5. vContinuum does not support protecting iSCSI disks directly exported to source VMs. The protection of these disks would require manual steps using Scout CS UI. However iSCSI disks directly exported to physical servers can be protected using vContinuum with P2V model.
6. vContinuum does not support protecting VMs configured to use pass-through devices (like RAW devices, tapes, etc.).
7. Target vSphere version should be either same or higher than the source vSphere server. If it is lower, it should at least support source VM OS version.  
Example: - Primary vSphere 5.1 should be protected to either version 5.1 or above
8. Windows 2012 virtual/physical machine can only be protected to ESX 5.0U1,ESX5.1 or later versions (due to VMware limitation).
9. Remote install/uninstall of InMage Scout Unified Agent on primary servers and Master Target using vContinuum is not supported for Linux.
10. For RHEL 5U3/CentOS5.3 and earlier versions, IP address on recovered VMs must be assigned manually
11. In case of 1-N protection, local backup to same production vSphere ESX is not supported.
12. InMage Offline sync folder should be in a single datastore. Offline sync folder on multiple datastores is not supported.
13. McAfee anti-virus version 8.8 should not be installed on Master Target. However, other versions of McAfee anti-virus can be installed on Master target.
14. EFI with dynamic disks cannot be protected using vContinuum.
15. DR-Drill with array based snapshot works only with storage arrays having Virtual LUN Copy feature. Check with your vendor whether your storage system supports this feature or not.
16. At times, change in network configuration for Windows server may fail in recovery. User needs to manually set network configuration (IP, subnet, DNS).
17. Linux workflow does not support UEFI disk partitioning layout.
18. At times, P2V of Linux shows "Operating system not found".  
Work around: Swap the disks to boot properly at vSphere level by using "Edit settings" option on vSphere Client UI.
19. When source is Windows 2012 R2 then MT should be Windows 2012 R2. If MT is running on windows 2012, it has incompatibility issue with loading system registry hive to make configuration changes on MT machine and recovery will fail with "error loading registry hive".  
Workaround: Manually change the IP of recovered VM.
20. In case of local backup (Both production and DR VM) on same ESX server, failback is not supported.
21. A partition without a driver letter or mount point cannot be protected using vContinuum.
22. For Linux V2V, recovered VM can potentially run into maintenance mode if all source disks are not protected.
23. When protected dynamic disk volume (example G: volume) is removed and new dynamic disk from the same machine is added for protection then vContinuum again protects G: volume.  
Workaround: After pair creation, perform remove volume( G: volume).
24. When CX, vContinuum and MT are upgraded to 8.0.1 and if source side agent is not upgraded, then addition of disk and volume resize operations will fail. Workaround: upgrade agent as well to 8.0.1.
25. DR-Drill operation will fail when source and MT VMs have the same disk signatures and it generally happens when both the VMs are prepared from the VM template.

- Workaround: Change the disk signature of MT local disks.
26. MSCS cluster IP change is not supported by vContinuum wizard. User need to manually update cluster network setting on recovered VMs.
  27. During recovery, newly assigned IP may not be set to one of the NIC of recovered MSCS VM. In this case user has to manually assign the IP to the NIC.
  28. Array based snapshot DR-Drill is not supported at ESX level when ESX is connected to vCenter.
  29. ASR Scout does not support non-English OS.
  30. VMware VM with Paravirtual scsi controller is not supported.
  31. Enabling of UUIDs is a pre-requisite for enabling V2V protection. UUID property created inside the VM should be enabled from vSphere/vCenter.
  32. At the time of MSCS Cluster recovery, all the cluster nodes need to be selected for recovery. Partial recovery of cluster nodes is not supported.
  33. When new disk is added to a MSCS Cluster, then ensure that drive letter assigned to it is not present in any of the cluster nodes. Otherwise, volume re-size will be triggered by vContinuum as the newly assigned drive letter matches with the existing protected volumes but with different volume size.

## 10.5 ScoutCloud RX

1. The protection status information at RX may vary with that of CX because of periodic data synchronization.
2. When the vContinuum setup is having multiple IPs, the recovery readiness check might fail if the vContinuum IP address is provided instead of hostname. Hostname can be provided to resolve this.
3. The information in the RPO/Data change graphs is limited to last 7 days.
4. The RPO graph for a protected server will be empty on the first day of the protection
5. Shared folders are not being shown in shared documents tab and only the files are getting listed.
6. Only three health factors "RPO Health", "Retention Health" and "Resync Health" are considered while showing the Protection status at RX. Other health factors are not considered.
7. After upgrade, if CX synchronization mode is changed from PULL to PUSH, synchronization may not happen. Workaround: unregister and register CX back with Rx.
8. When only RX is upgraded to 8.0.1 but CX is not upgraded, bandwidth report value shows as zero. Workaround: Upgrade CX to 8.0.1
9. When RX is upgraded and CX and MT are not upgraded then recovery through RX will fail. Workaround: upgrade CX and MT to 8.0.1
10. Specify tag option is not available in RX while doing recovery. User has to use vContinuum wizard for this option.
11. If force unregister is used on Scout Cloud RX Server, the unregister operation will be triggered but will not be validated for successful completion by the RX Server. If for some reason the action should fail to affect the unregister operation with Azure Site Recovery, stale registrations may continue to exist in the Azure Site Recovery vault. The user will need to contact support in order to delete the stale entries that may continue to persist in Azure Site Recovery.
12. Protection and recovery plan should not contain any space. If space is present, recovery through RX will fail.

## 11 Sizing guideline for CS/PS/MT

The recommended system requirements for CS, PS and MT VMs vary depending on several factors including the rate of data change rates of primary server, bandwidth, retention requirement, hardware make and model, etc. Refer to section 7 'Sizing Guideline' of [ASR Scout Quick Install Guide](#) to learn more about the CS, PS, MT VM size.

## 12 Upgrade protected P2V Windows Cluster to Scout Update5 or above

Manual steps required for upgrading Windows Cluster protected in P2V mode using Scout 8.0.1 Update 4 or earlier version to Scout 8.0.1 Update5 or above version.

Since all the updates installer are cumulative, it contains all the previous update fixes. This upgrade steps for P2V Windows Cluster is applicable if you are upgrading your source from Scout Update4 or earlier to Update5 or above. For example, latest Scout update version is Update6. The below steps are still applied when you upgrade your servers from Scout Update4 or earlier version to Scout Update6. If you have protected your physical Windows Cluster with Scout update 5 then you don't need these upgrade steps.

In Scout 8.0.1 Update5, the disk names format is changed for P2V clusters. Because of this the P2V Cluster protections that are already enabled using Scout 8.0.1 Update 4 or earlier version should follow the following sequence of steps to avail the cluster fixes that are added in Scout Update5 or above.

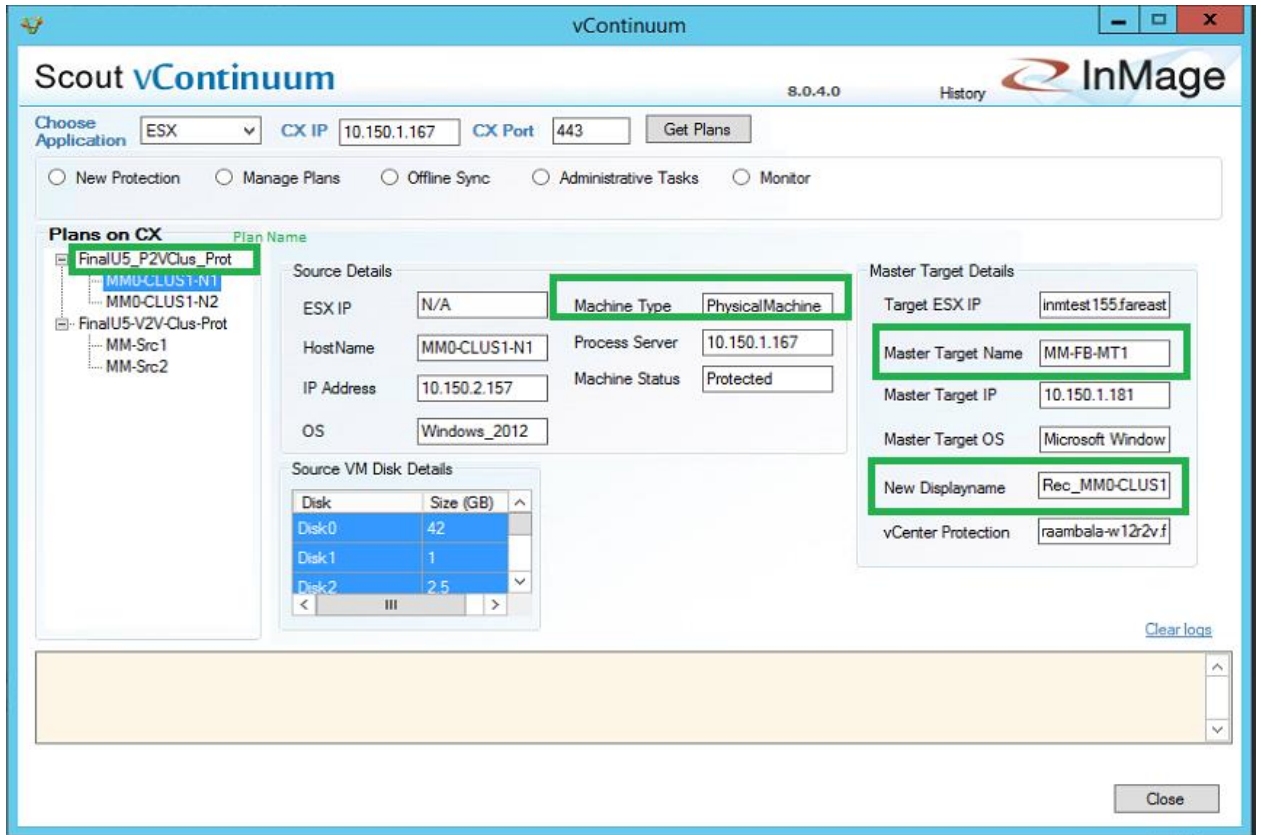
### 12.1 When do you need to use these manual steps?

If any of the following is true, you need to perform the manual steps before upgrading to Update5.

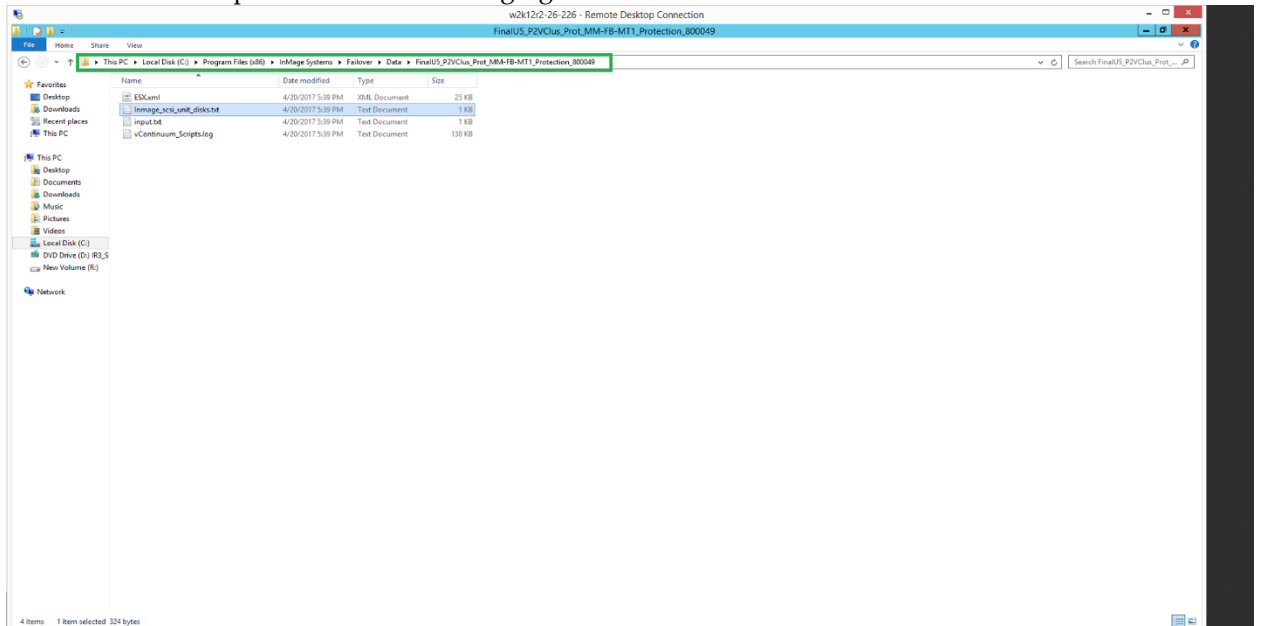
- a) If you want to re-protect physical cluster and want to re-use the existing target disk.
- b) If you want to add new cluster disk to the existing protection (protected in P2V mode)
- c) If you plan to remove a protected disk from the plan

### 12.2 Manual Steps

1. From the vContinuum user interface, note the plan names which have protected MSCS clusters in P2V mode. You will know that a plan is protected in P2V mode if the "Machine Type" is "Physical Machine". Also, note the Master Target Name from the "Master Target Name" field as shown in the below figure.

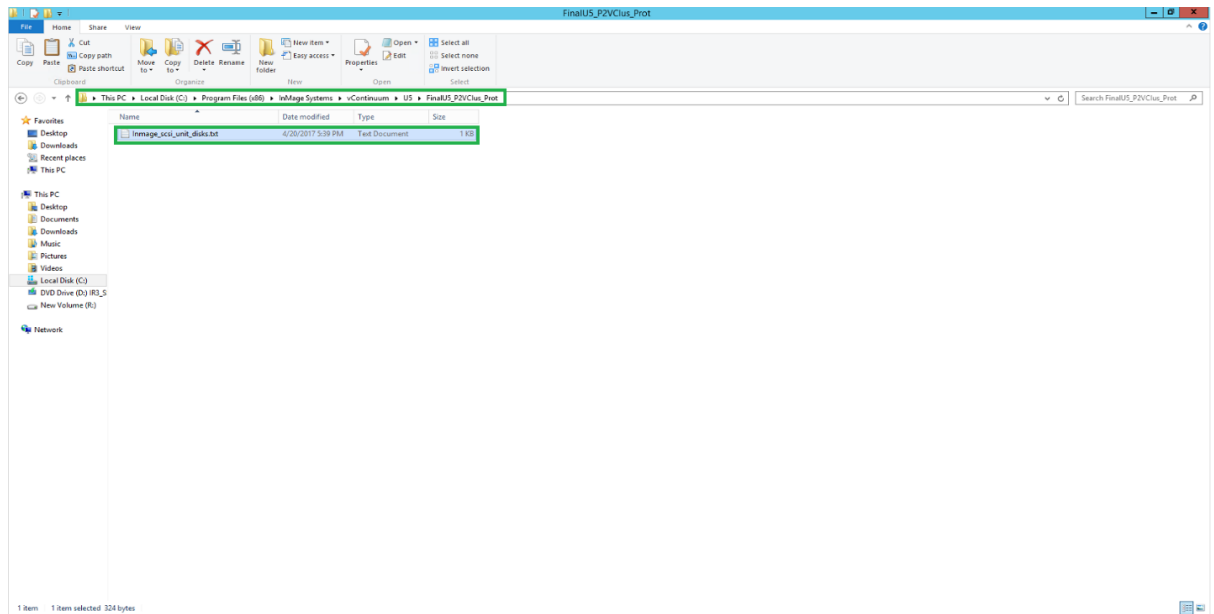


- For each MSCS Cluster that is protected in P2V mode, note the plan names as stated in above step and create a folder called U5\See the Plan Name path from the following figure.

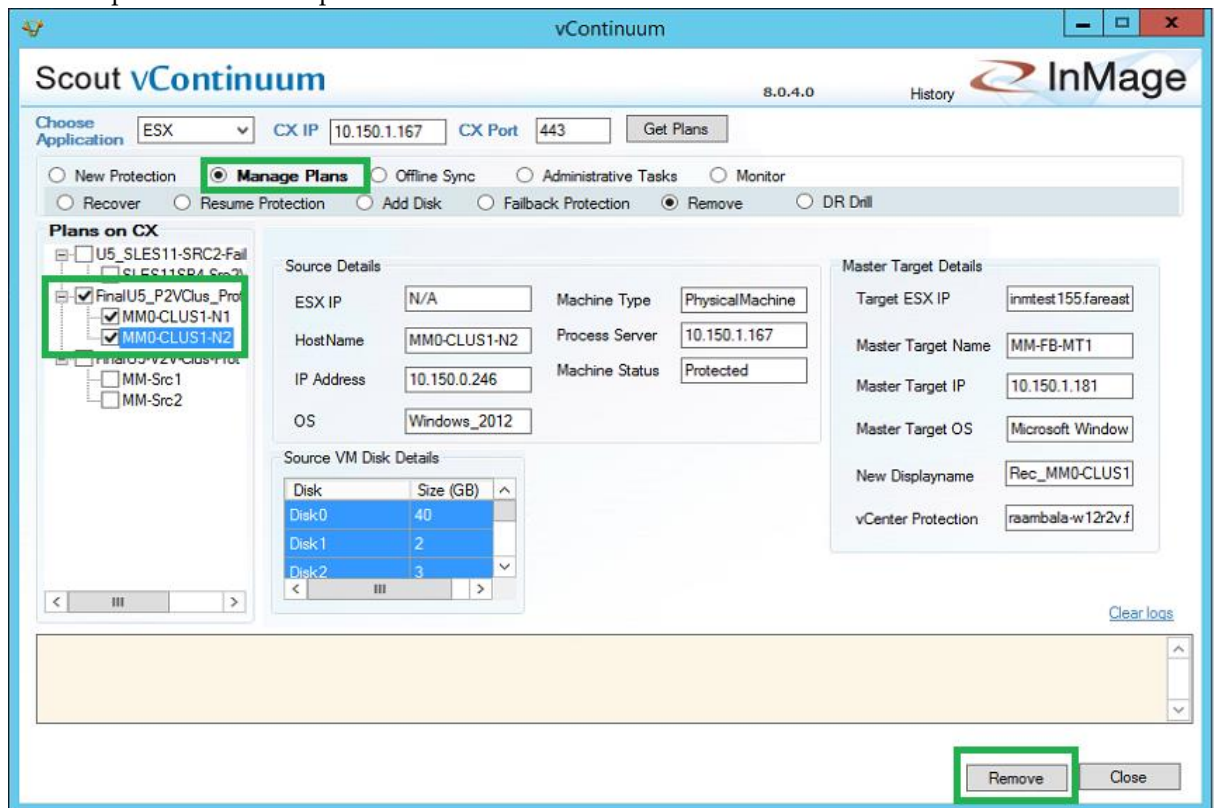




- Now copy the <InMage Installation Folder>\Failover\Data\<plan\_name\_folder>\Inmage\_scsi\_unit\_disks.txt file to <InMage Installation Folder>\vContinuum\U5\<plan\_name\_folder> as show in the figure below.



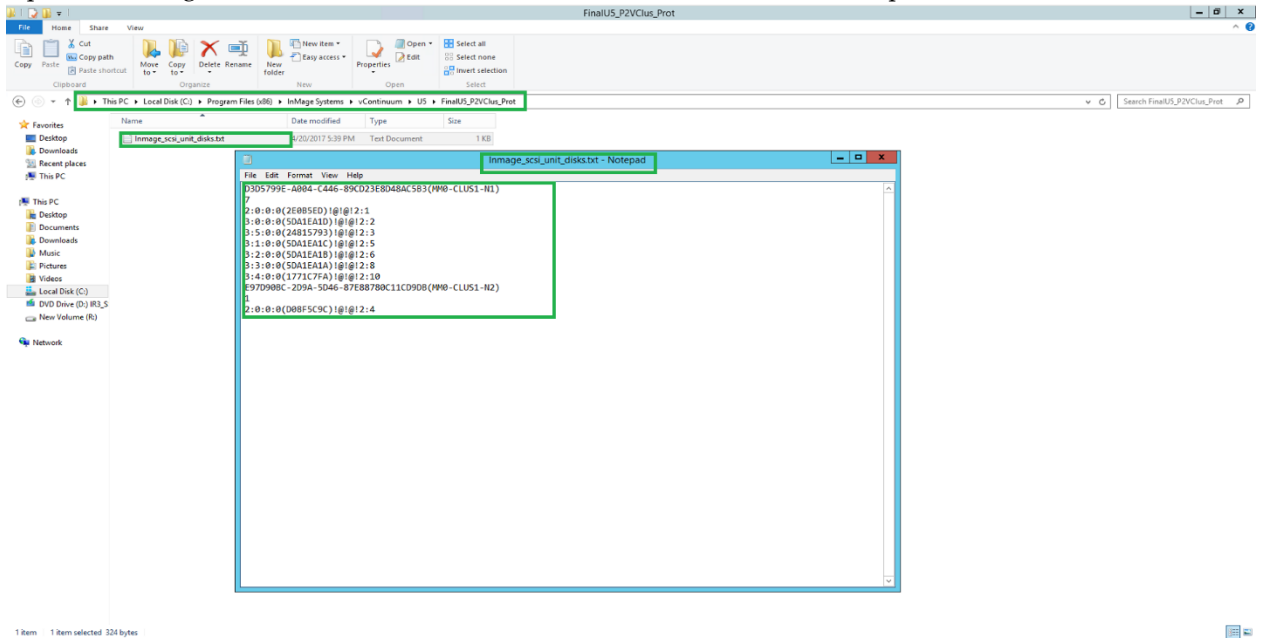
- Remove protection of this plan from vContinuum UI as shown below.



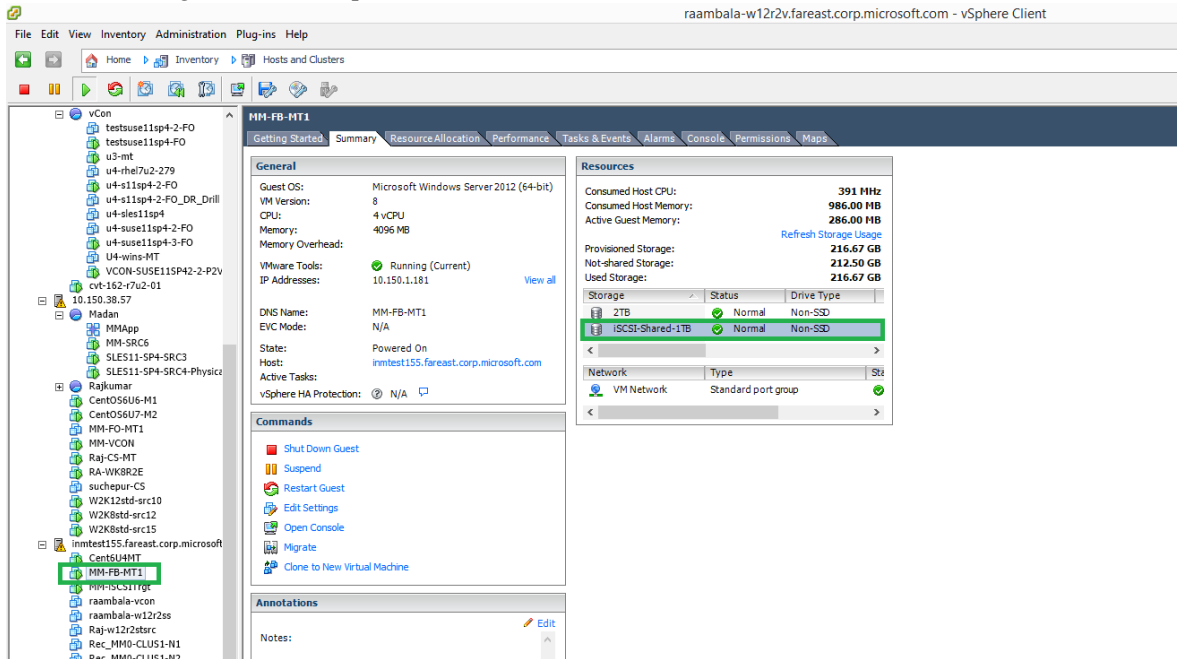
- Observe that the protections of disks under this plan are being removed from the CX UI.

Server	VX Agent Pair	RPO	Resync progress	Status	Resync Required	Resync in Transit		Differential Data in Transit (MB)				Action
						Step1	Step2	On Primary Server	On CX-PS	On Secondary Server		
MM0-CLUS1-N1->MM-FB-MT1	C -> C:\ESX\0005799E-A004-C446-89C023E8D48AC5B3_C	1.17 min	N/A	Differential Sync [Detection pending]	NO	0	0	0	0.5	0	Summary	
MM0-CLUS1-N2->MM-FB-MT1	C -> C:\ESX\97D908C-209A-5D46-87E88780C11C0908_C	0.87 min	N/A	Differential Sync [Detection pending]	NO	0	0	0	0.42	0	Summary	
MM0-CLUS1-N1->MM-FB-MT1	C\SRV [ System Reserved ] -> C:\ESX\0005799E-A004-C446-89C023E8D48AC5B3_C_SRV	0.65 min	N/A	Differential Sync [Detection pending]	NO	0	0	0	0.01	0	Summary	
MM0-CLUS1-N2->MM-FB-MT1	C\SRV [ System Reserved ] -> C:\ESX\97D908C-209A-5D46-87E88780C11C0908_C_SRV	1.48 min	N/A	Differential Sync [Detection pending]	NO	0	0	0	0.01	0	Summary	
Cluster:MM0-Clus1, Group:Cluster Group Servers: MM0-CLUS1-N1,MM0-CLUS1-N2->MM-FB-MT1	D [ New Volume ] -> C:\ESX\0005799E-A004-C446-89C023E8D48AC5B3_D	0.65 min	N/A	Differential Sync [Detection pending]	NO	0	0	0	0.01	0	Summary	
Cluster:MM0-Clus1, Group:PFS2 Servers: MM0-CLUS1-N1,MM0-CLUS1-N2->MM-FB-MT1	I [ New Volume ] -> C:\ESX\0005799E-A004-C446-89C023E8D48AC5B3_I	0.62 min	N/A	Differential Sync [Detection pending]	NO	0	0	0	0.01	0	Summary	

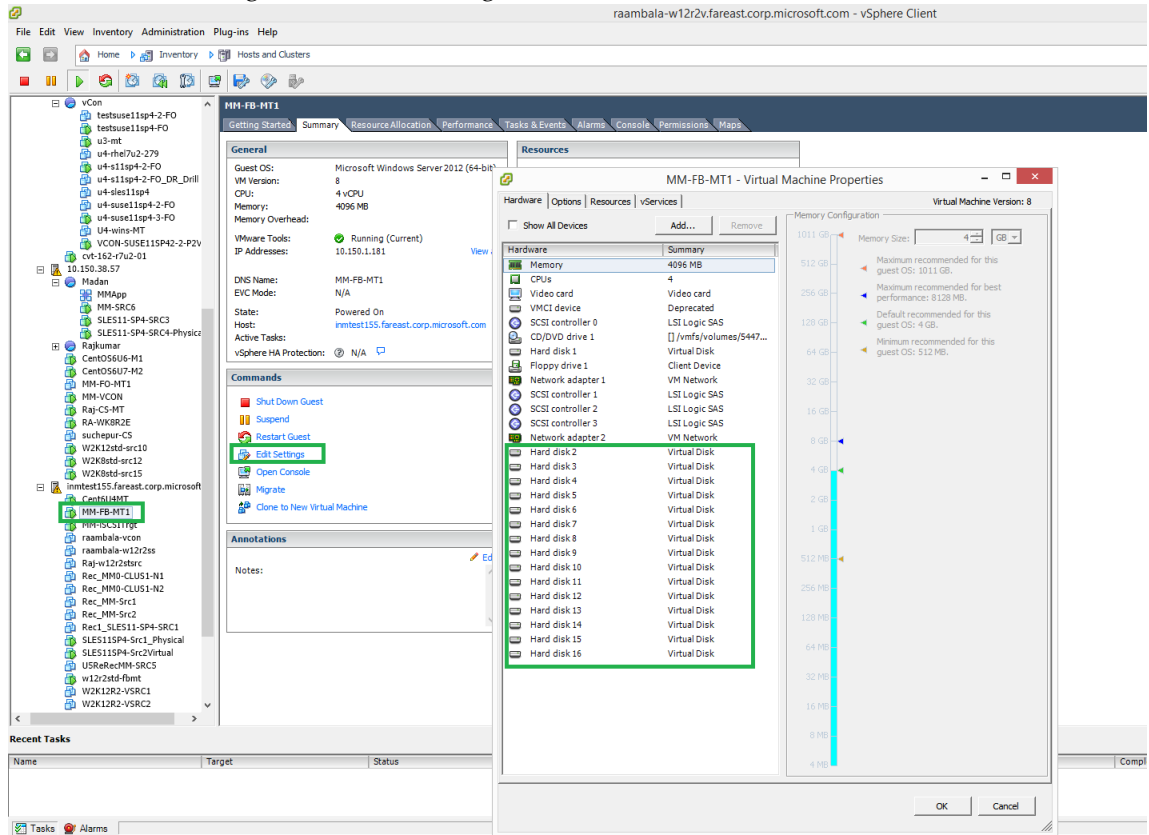
- Open the Inmage\_scsi\_unit\_disks.txt file from the new folder created in step 3 above.



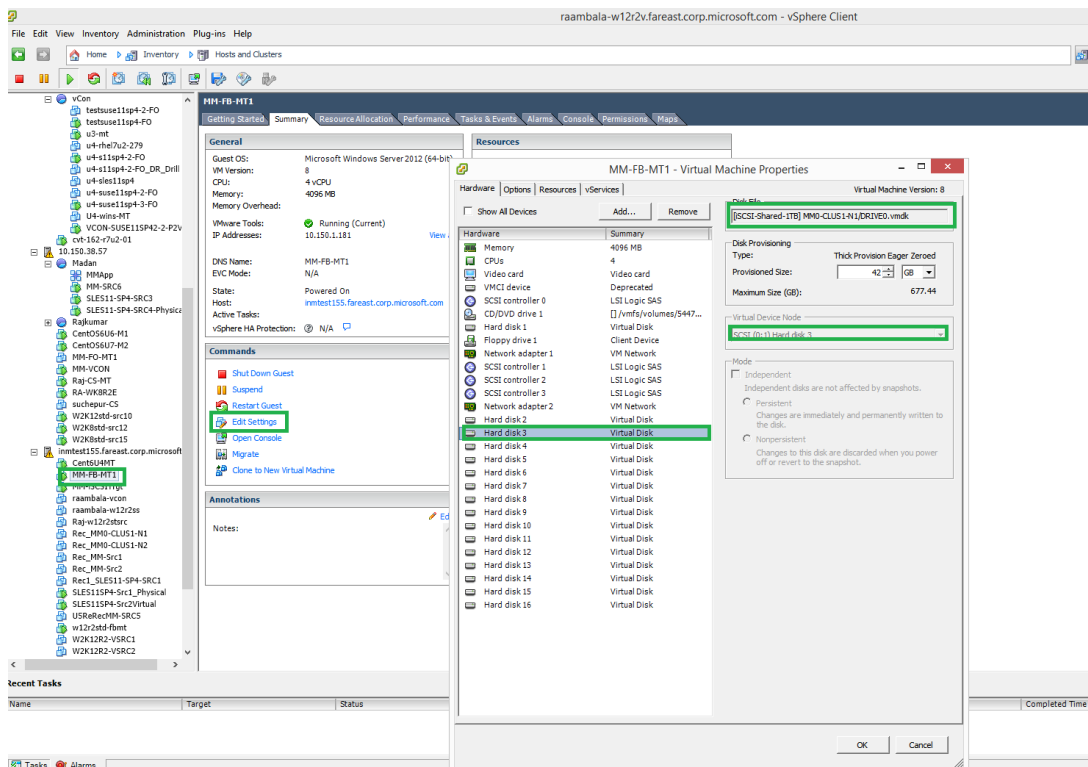
- Go to the target vCenter where the Master Target is present and identify the datastore where the Master Target's disks are protected.



8. Note the Master Target's disks in the target vCenter.

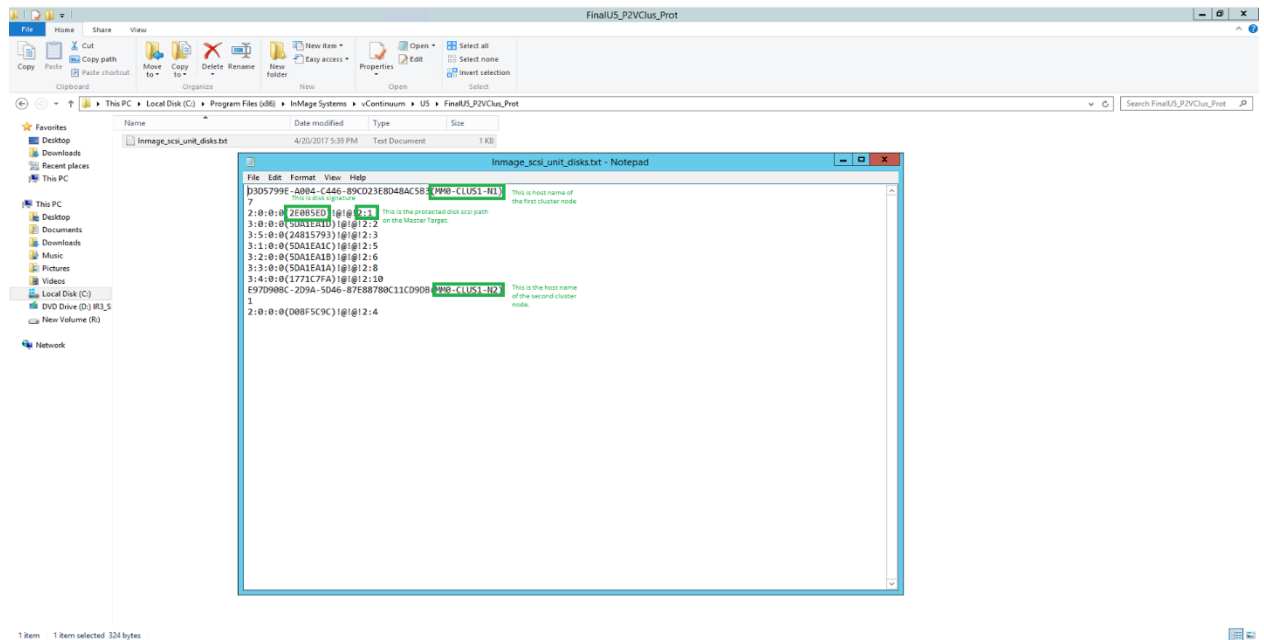


9. Note the path of the disks that are protected under MT as shown below.

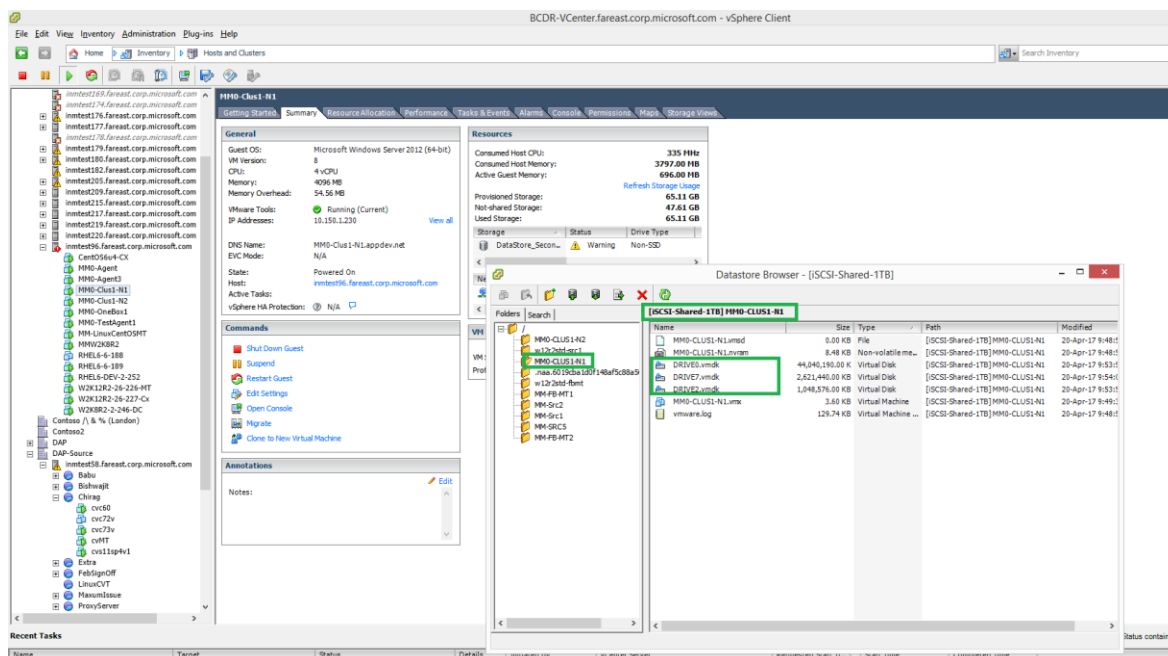


- Note the SCSI path of the protected disk in Master Target which is 0:1 in the figure shown above. These two values represent the SCSI controller and SCSI Target in the target ESX Host. Increment the SCSI Controller value by 2. Currently the SCSI Controller value is 0 and if you increment it by 2, it will be come  $0 + 2 = 2$ . Now the SCSI path of this Master Target disk becomes 2:1. Now open the Inmage\_scsi\_unit\_disks.txt file as shown in step 6 above and identify the row where the scsi path 2:1 is present. If you see the example SCSI unit disks text file that have chosen in step 6, it will match with the row **2:0:0:0(2E0B5ED)!@!@!2:1**.

This file represents the mapping of the SCSI tuple of protected physical source disk with that of the protected target disk on the Master Target. The source and target SCSI tuple is separated by the symbols **!@!@!**. The value on the left side of this symbol represents the scsi tuple of the protected source disk and that on the right side represents the scsi tuple of the protected target disk on Master Target. The one within the parenthesis is the disk signature. In the current example it is **2E0B5ED**. Observe the following diagram.



- Now from the above diagram go to the protected target disk path in the target data store. This path is obtained from the step 9 shown above.



12. Now rename vmdk file in the format as given below....

<Source Host Name>\_<Source Disk Signature>.vmdk

So the disk [iSCSI-Shared-1TB] MM0-CLUS1-N1\DRIVE0.vmdk should be renamed as [iSCSI-Shared-1TB] MM0-CLUS1-N1\MM0-CLUS1-N1\_2E0B5ED.vmdk.

This renaming can be done by opening the Target ESX Shell and executing the following command from the path /vmfs/volumes/<datastore name>/<Host Folder Name>.

The command to rename the vmdk file as described above is ...

**vmkfstools -E <old vmdk file with name> <new vmdk file name>**

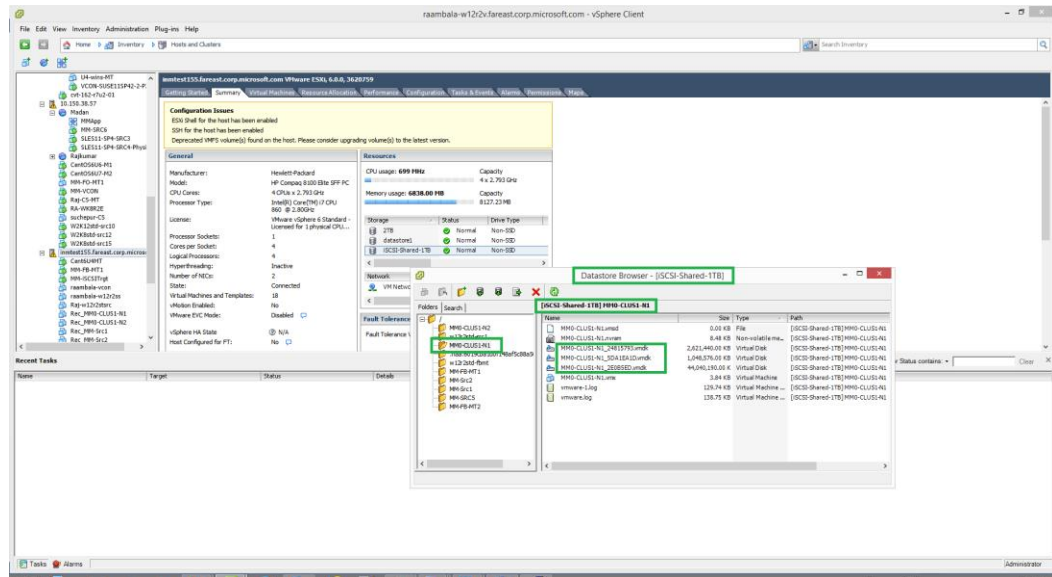
**Example:**

1. Go to the path /vmfs/volumes/iSCSI-Shared-1TB/MM0-CLUS1-N1
2. Rename the vmdk file DRIVE0.vmdk like this...

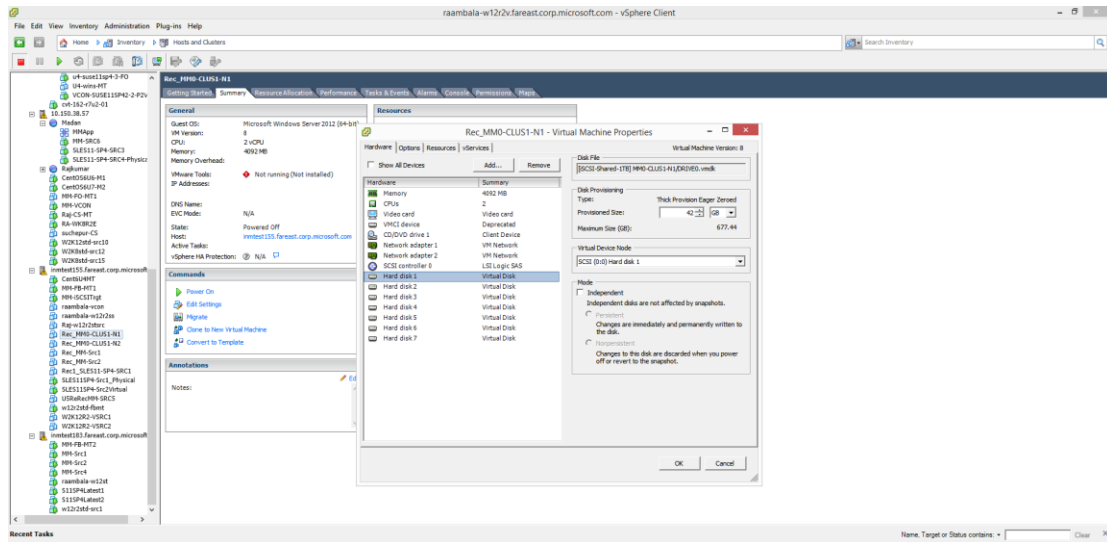
**vmkfstools -E DRIVE0.vmdk MM0-CLUS1-N1\_2E0B5ED.vmdk**

Like this rename for all the protected disks of MSCS cluster protected in P2V mode.

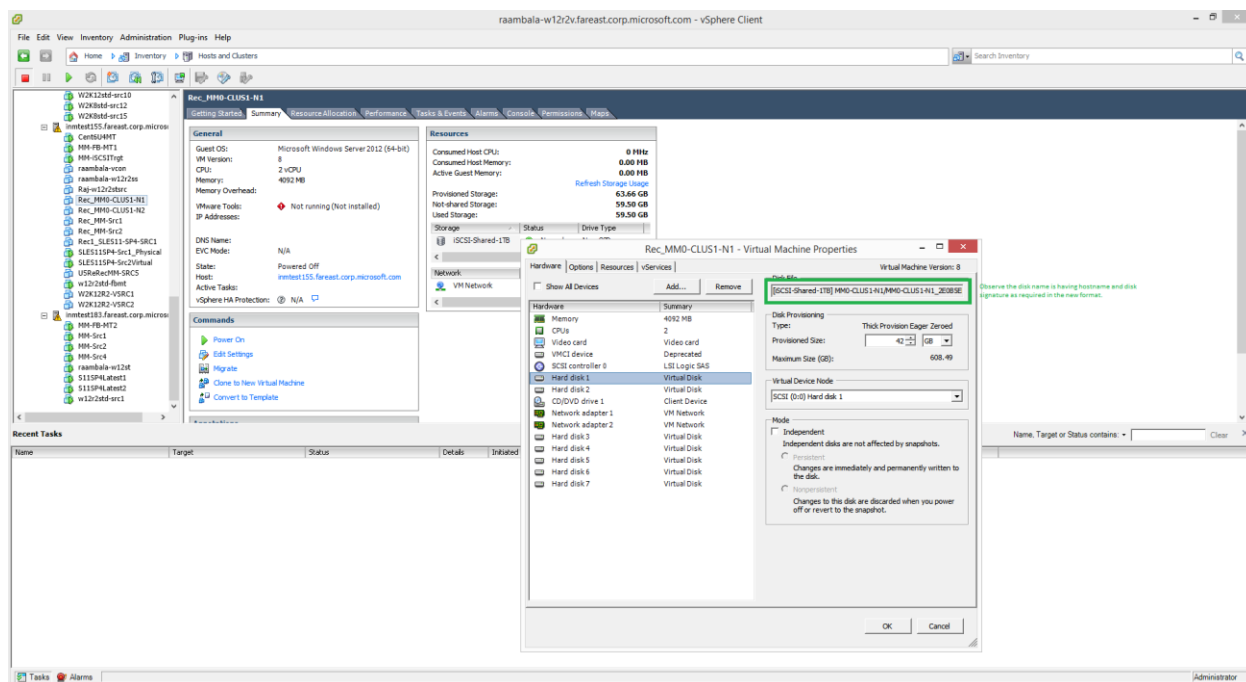
After renaming the target ESX datastore .vmdk files name in the new format and it would show up as shown below.



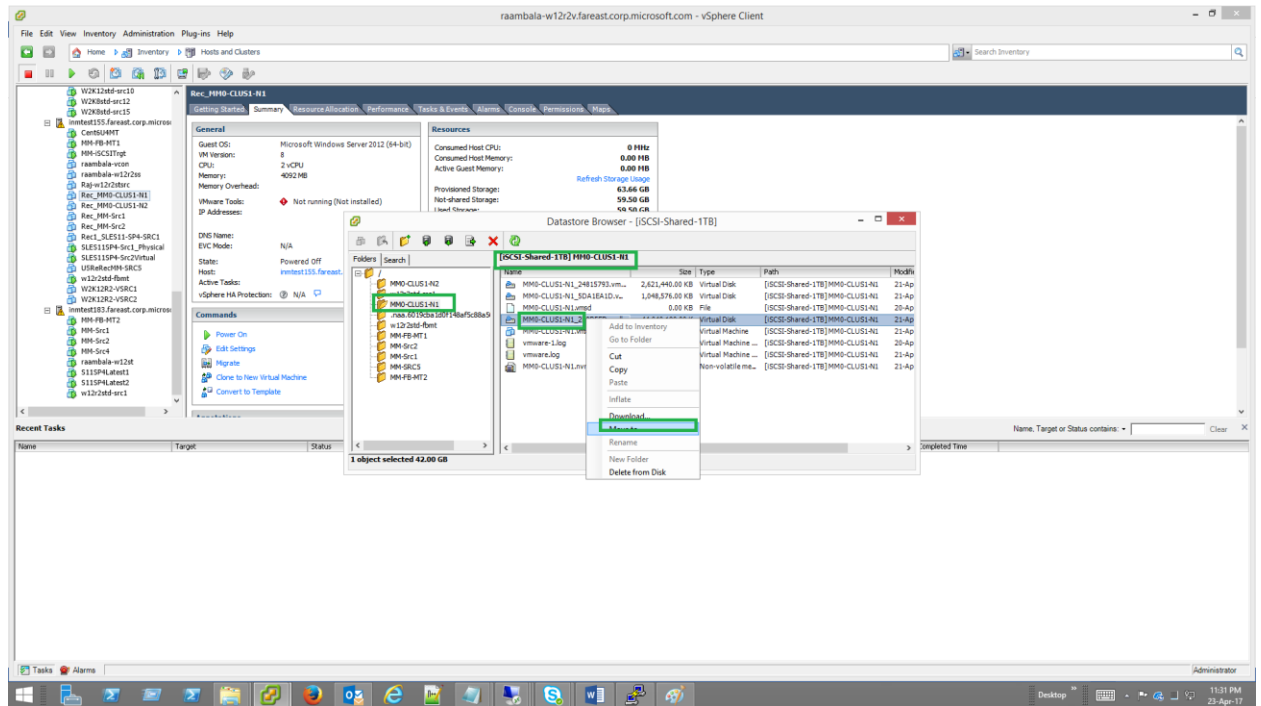
13. Once the vmdks are renamed in the new format, go to the DR VMs in the target ESX and detach the old disks (old .vmdk files) and attach the newly renamed disks (new .vmdk files). For this first identify the DR VM. This DR VM is obtained from the diagram shown in Step 1 under the field "New Display Name". Open the vSphere client and select the target ESX Host and from there identify the DR VM with the help of "New Display Name" field. Click on edit settings and remove all the disks that are attached. After that attach the newly renamed vmdk files to each DR VM available in the P2V cluster protection plan. The following diagram shows the DR VM and its disks in the vSphere client before detaching the older disks...



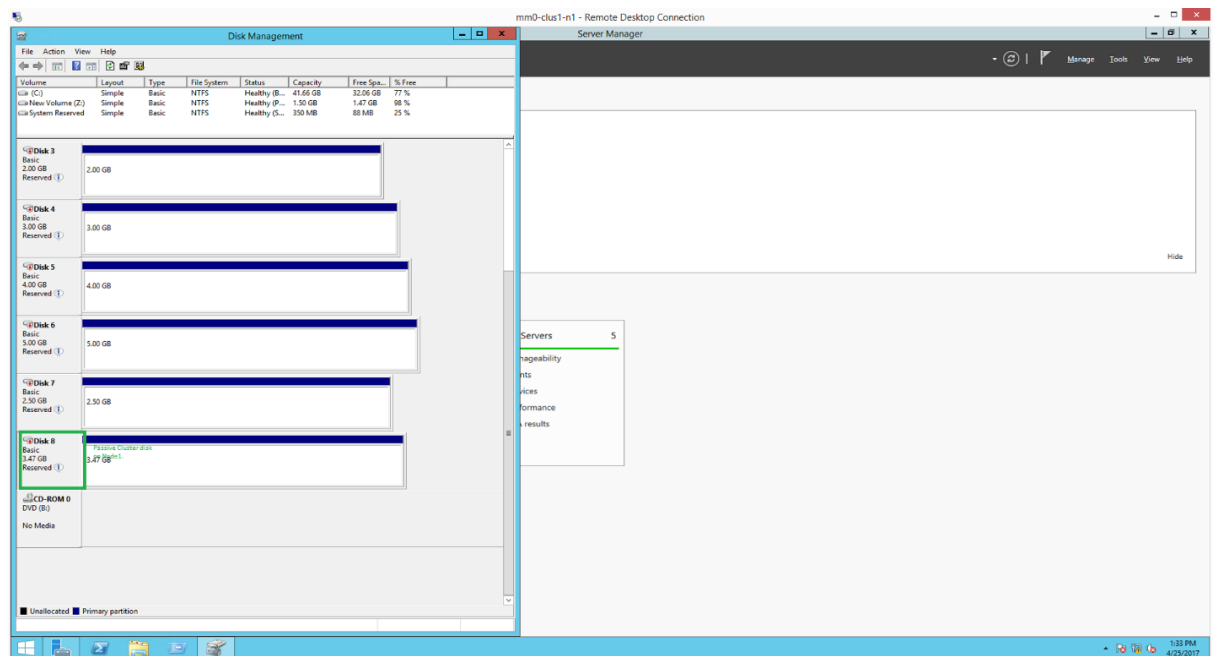
The following figure shows the disks path in the DR VM after attaching the new vmdk files



- The next step is to protect the MSCS cluster again in P2V mode as the protection is removed in Step 4. The cluster disks can move from one node to the other and as they move from one cluster node to the other cluster node, the disks become active on the new node and passive on the old node. So, at the time of protection of this cluster earlier, the active – passive disks on each node could be different than what is the current active – passive disks on each cluster node. So, if a cluster disk is found to be passive now and active earlier, then you have to move the vmdk file from the older path to the new path in the target ESX datastore. This can be done quite easily as now the vmdk files are having the disk signatures in their names. So, the end user should check if a disk is passive on that node, then she should note the disk signature and check for the availability of the vmdk file with that disk signature on the Target ESX Data Store folder. If the vmdk file is indeed there, then it should be moved to the new Target ESX Data Store folder by right clicking that vmdk file and choosing the new path in the Target ESX Data Store folder. By doing this step, vContinuum can re-use the disk during protection and the initial seeding time of protection will get reduced and the protection will complete faster. If this step is not done, then a new disk will be created and the older disk which is available in a different path in the ESX data store will not be re-used. Thus, this is only an optimizing step and not a mandatory step.



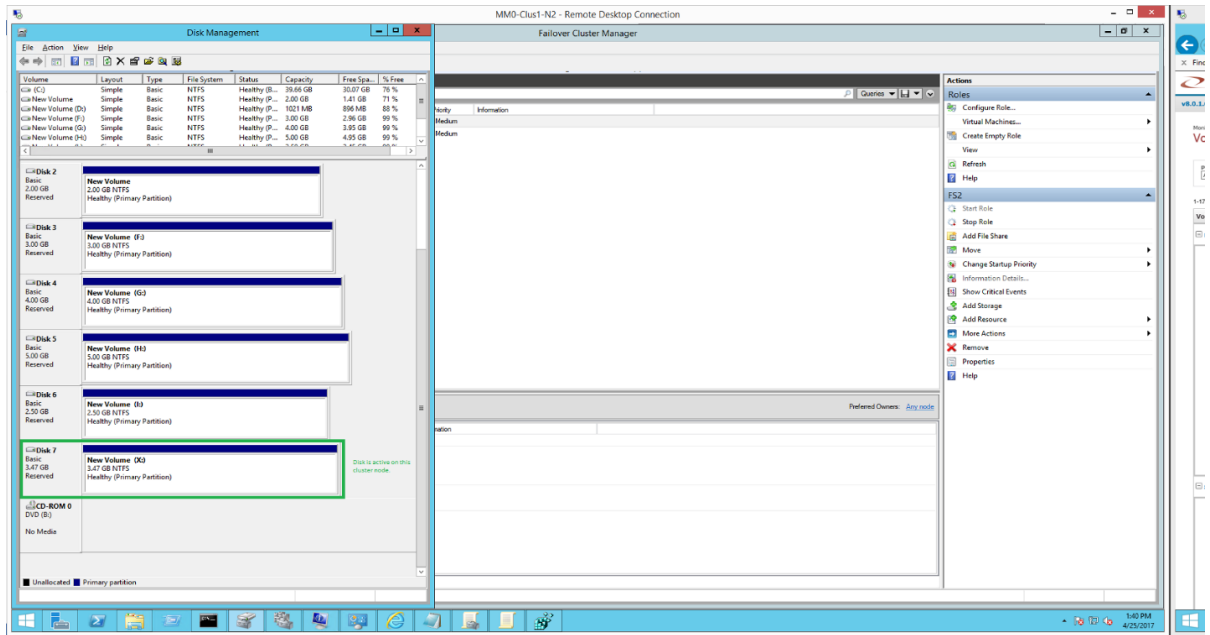
To explain this, further, let's see the disk management console on one of the source node MM0-CLUS1-N1.



Observe the **Disk 8** is in **Reserved** and in-active on MM0-Clus1-N1.

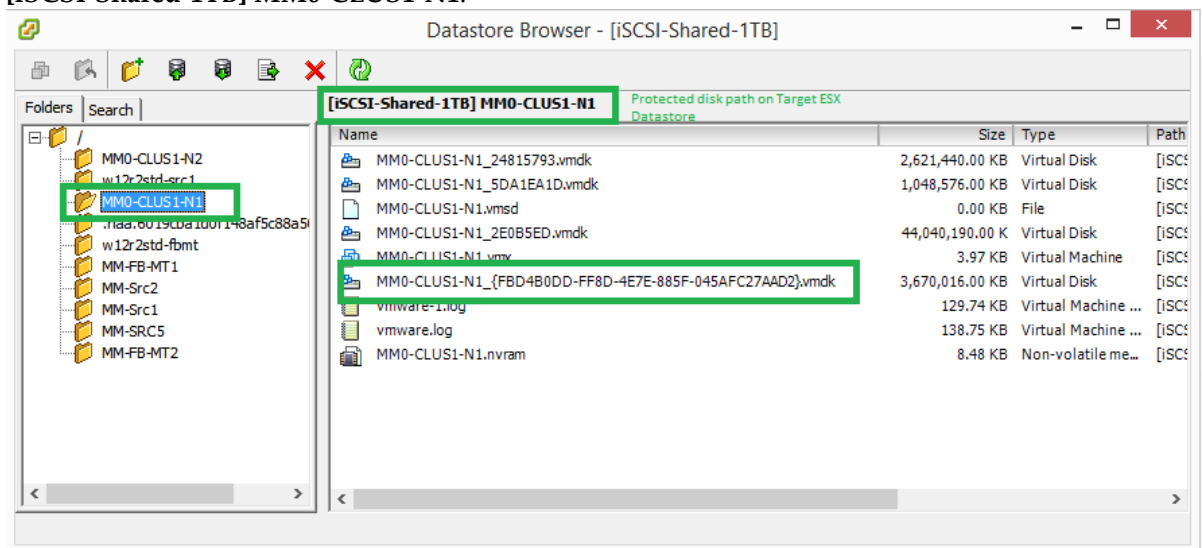
Now, check the state of this disk on the second cluster Node MM0-CLUS1-N2 in the disk management console.



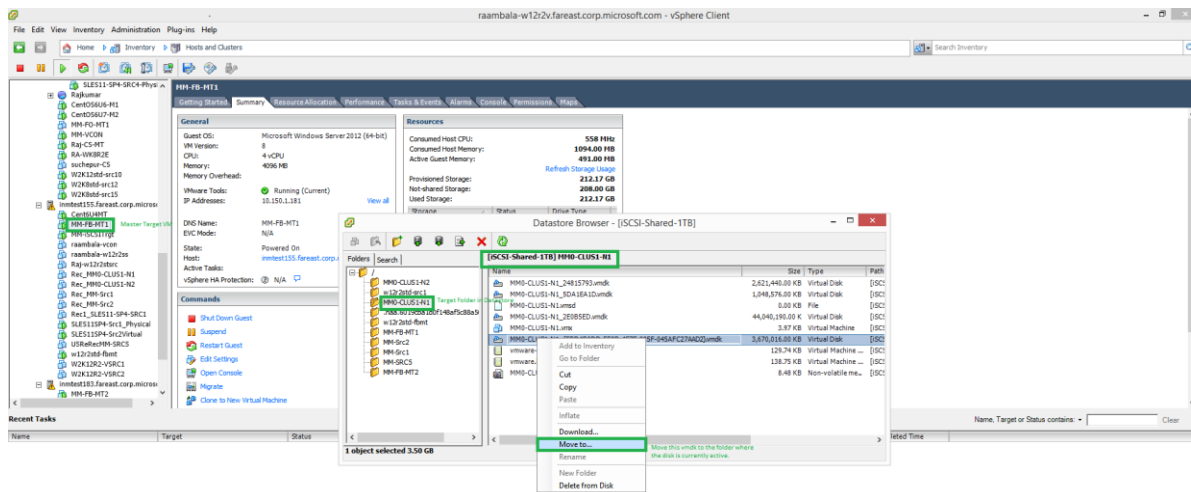


So, from the above figure, it is clear that the disk is passive on MM0-CLUS1-N1 and active on MM0-CLUS1-N2.

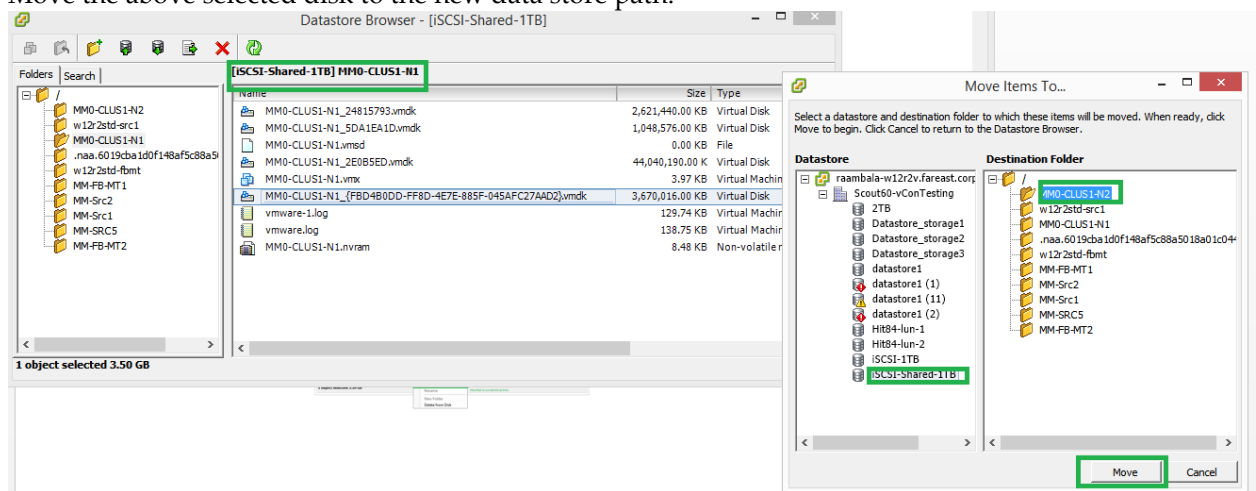
Now, go to the data store of MM0-Clus1-N1 on the Target ESX host and check the availability of this disks corresponding vmdk file by going to the target ESX data store path **[iSCSI-Shared-1TB] MM0-CLUS1-N1**.



But this disk is passive on this node, so before you start protection of this Cluster (MM0-CLUS1), you must move this vmdk from this path to the path where the second cluster Node's disks are present...i.e. you should move this disk to the Target ESX datastore path **[iSCSI-Shared-1TB] MM0-CLUS1-N2**.

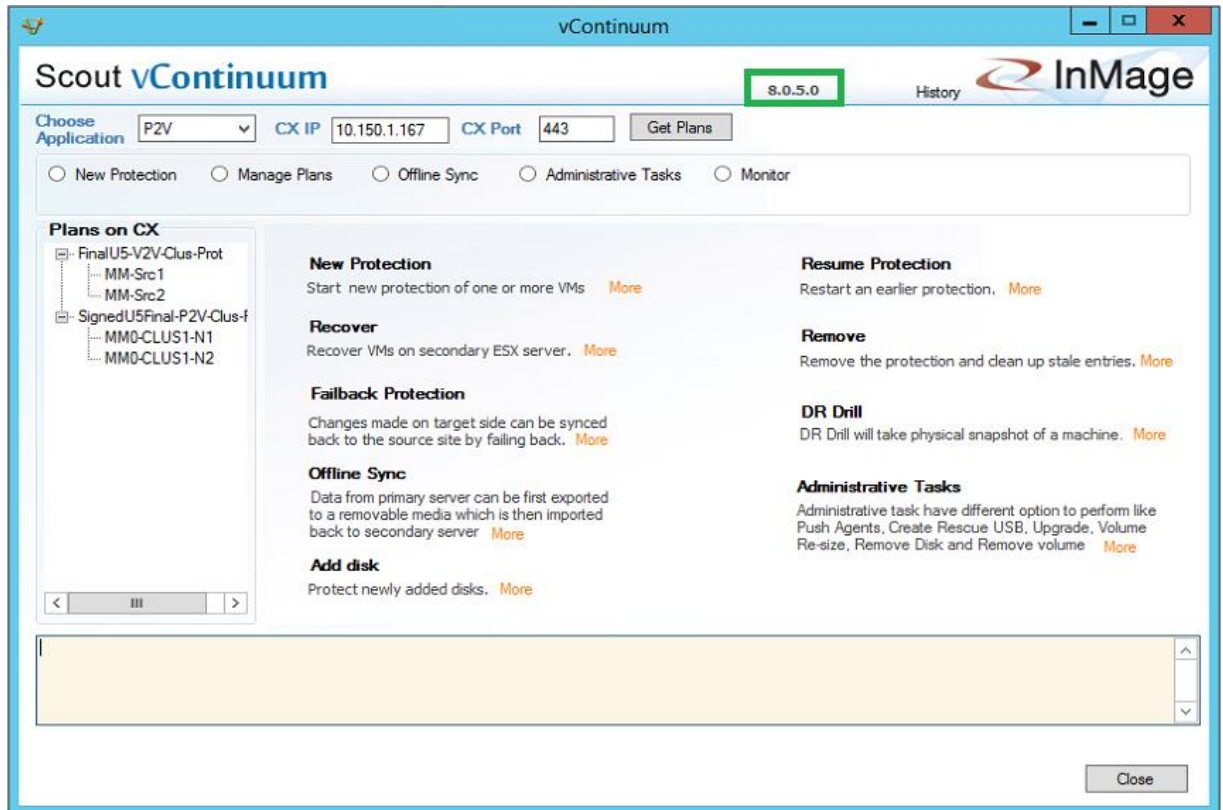


Move the above selected disk to the new data store path.

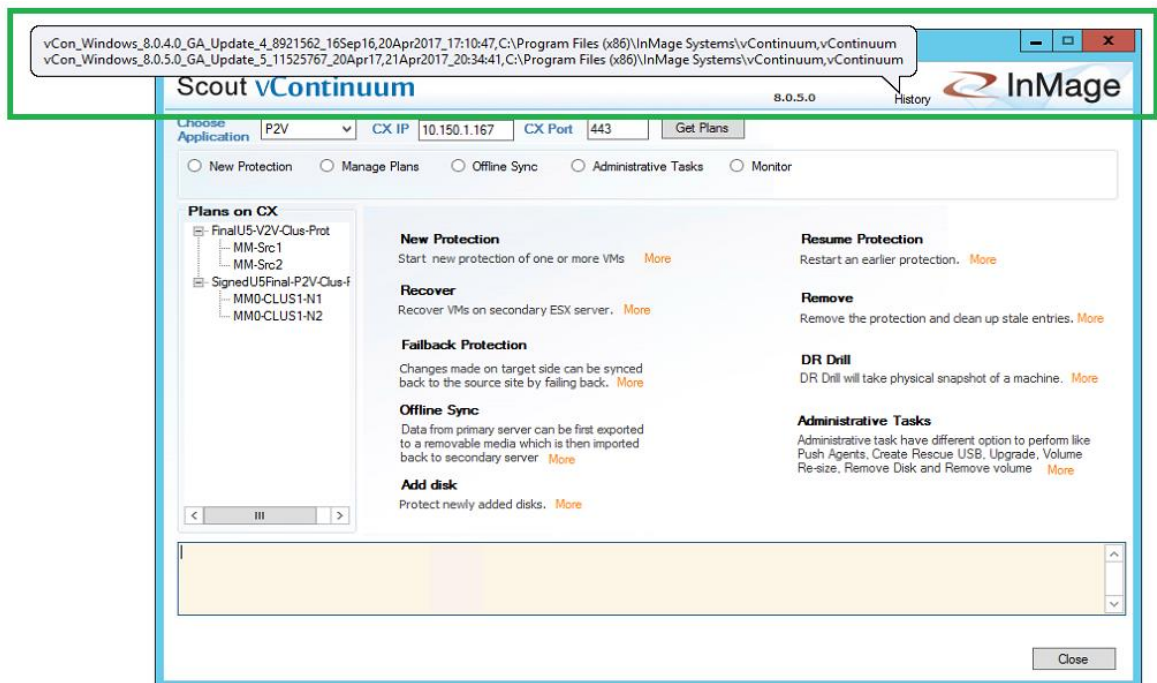


This step is not mandatory as described above. If this step is done, then the target side ESX disks will be re-used during re-protection and otherwise, new disks will get created during re-protection.

15. Now, upgrade vContinuum to Scout 8.0.1 Update 5 by running the **vCon\_Windows\_8.0.5.0\_GA\_Update\_5\_11525767\_20Apr17.exe** (Note: If there is any new update after update5, use the latest vContinuum update installer. For example, for update 6, vContinuum installer is **vCon\_Windows\_8.0.6.0\_GA\_Update\_6\_11525767\_21Sep17.exe**. Run the installer to upgrade to latest version.)
16. After vContinuum upgrade is over, launch vContinuum from the Start menu and you can check the updated version as below...



17. Observe the Upgrade history by hovering the mouse on “History” label on vContinuum user interface like this...



18. Now protect the same Cluster which you have earlier protected in P2V mode.

With this the upgrade to latest Scout update version completes.

## 13 Scout Update 7 enables TLS v1.2 support

The PHP Scripting Platform needs to be upgraded to v7.2.10 32 bit and the MySQL database needs to be upgraded to v5.7.23 on CX and RX Servers. Please refer to the steps given in the **section 7.4** of [Scout\\_Standard\\_Quick\\_Install\\_Guide.pdf](#).

## 14 Please ignore the Version mismatch is found for the VX agent alerts in the CS UI after upgrading to Scout 8.0.7.

After upgrading the agents to 8.0.7, you may find alerts in the CS UI as shown below.

```
■ Protection Alerts : 2019-01-07 14:57:52  
Version mismatch is found for the VX agent U4-RHEL7U2-1 ipaddress 10.150.1.18  
Version mismatch is found for the VX agent V2V-W2K12R2-2 ipaddress 10.150.38.94  
Version mismatch is found for the FX agent V2V-W2K12R2-2 ipaddress 10.150.38.94  
Version mismatch is found for the VX agent V2V-VCON-MT2 ipaddress 10.150.1.65  
Version mismatch is found for the FX agent V2V-VCON-MT2 ipaddress 10.150.1.65  
Version mismatch is found for the VX agent V2V-W12KR2-1 ipaddress 10.150.172.108  
Version mismatch is found for the FX agent V2V-W12KR2-1 ipaddress 10.150.172.108
```

You may ignore these alerts.