

IBM System Storage N series



# OnCommand Unified Manager Core Package 5.2 Release Notes



---

# Contents

<b>OnCommand Unified Manager Core Package 5.2 Release Notes Overview . . . . .</b>	<b>1</b>
<b>OnCommand Unified Manager Core Package new features . . . . .</b>	<b>3</b>
<b>OnCommand Unified Manager Core Package changes . . . . .</b>	<b>5</b>
<b>System requirements . . . . .</b>	<b>7</b>
<b>OnCommand Unified Manager Core Package 5.2 fixed issues . . . . .</b>	<b>9</b>
The DataFabric Manager server communicates on the first available interface when the preferred interface is not running . . . . .	9
Cannot back up SnapVault source and SnapVault destination volumes in the same aggregate . . . . .	9
The dfmmonitor service stops responding . . . . .	9
Report share functionality unavailable through OnCommand console for users who do not have Read permission for the Global group . . . . .	9
Restore of database backup fails . . . . .	9
<b>OnCommand Unified Manager Core Package limitations . . . . .</b>	<b>11</b>
<b>OnCommand Unified Manager Core Package 5.2 known issues . . . . .</b>	<b>13</b>
Installer issues . . . . .	13
Cannot upgrade to Core Package 5.x from DataFabric Manager server 3.5 . . . . .	13
Provisioning capability issues . . . . .	13
Cannot provision a volume with the same size as its aggregate. . . . .	13
Provisioning operations are not supported on systems running Data ONTAP 7.3.5.1P5 . . . . .	14
Data migration rollback and cutover jobs sometimes time out in N series Management Console . . . . .	14
vFiler unit migration operation might fail . . . . .	14
High CPU usage on storage systems during a Data Motion migration cutover causes storage system panic . . . . .	15
Provisioning policies and online migration are not supported on SAN datasets with multiple-space names . . . . .	15
During online migration, SnapManager-initiated backups to secondary storage fail . . . . .	15
Sometimes the Cancel button does not cancel a failed migration . . . . .	16
Protection capability issues . . . . .	16
Creation of an Open Systems SnapVault relationship fails on Linux . . . . .	16
Cannot include an NFS datastore in a renamed volume in a dataset . . . . .	17
Creating a dataset fails if the dataset has storage service with provision only policy . . . . .	18
Users cannot modify values in custom fields . . . . .	18
Provisioning failure despite sufficient space . . . . .	19
Restore operations performed from most recent backup on the tertiary (mirror destination) node might fail. . . . .	19
Improperly specified SnapVault source to target access list causes restore operation failure . . . . .	19
Backup scripts that refer to directories or files require full path names . . . . .	20
Compression option for imported SnapMirror connections not preserved . . . . .	21

First update of a mirror connection after the source has shrunk and then grown might generate a false error message . . . . .	21
Open Systems SnapVault relationship baseline backup version is not registered . . . . .	22
Adding a directory from an unsupported file system to a dataset fails with an error. . . . .	22
Enabling on-demand backup operations on SnapDrive-generated datasets requires the snapdrive dataset backup_add command. . . . .	22
SnapMirror job fails with "process was aborted" message. . . . .	23
Temporary conformance impact of deleting volumes or qtrees . . . . .	23
vFiler units cannot provision secondary storage of virtual objects . . . . .	24
Qtree is not the source for the snapmirror destination . . . . .	24
Number of bytes of data transferred during backup is inaccurate . . . . .	24
Faulty export file entry might prevent inclusion of NFS-exported virtual machine or datastore objects in a dataset of VMware objects . . . . .	24
Mixed LUN connections in a virtual machine can cause backups to fail . . . . .	25
Performance Advisor issues . . . . .	25
Changing the Performance Advisor perfTopnchartinterval setting might require you to restart N series Management Console . . . . .	25
Delay in updating entries in the Top Network Interface view . . . . .	25
Client statistics data not collected. . . . .	25
Issues related to Data ONTAP . . . . .	26
Data Motion for vFiler fails if volumes are in a SnapMirror relationship . . . . .	26
Cutover fails during Data Motion for vFiler . . . . .	26
OnCommand console issues . . . . .	26
Groups tab displays incorrect values for total capacity and used capacity . . . . .	26
Memory leak on Internet Explorer 9 . . . . .	27
The OnCommand console repeatedly prompts for login credentials . . . . .	27
Charts are not displayed after disabling HTTP and enabling HTTPS . . . . .	27
Internet Explorer does not display content in the OnCommand console if HTTPS is configured . . . . .	27
Edit Group Membership page displayed incorrectly in the Operations Manager console . . . . .	27
Breadcrumb trail not updated . . . . .	28
Storage tab displays blank page . . . . .	28
Setting or modifying the storage efficiency script schedule fails with an error message . . . . .	28
Report issues . . . . .	29
No data in about volume Performance Advisor after a volume move. . . . .	29
The dfm report view command fails with an error message . . . . .	29
Events Current report shows wrong event name. . . . .	29
Cannot create alarms with an SNMP trap host . . . . .	30
Email alerts generated for deleted objects . . . . .	30
DataFabric Manager server service issues . . . . .	30
The dfm service start http command fails with an error message . . . . .	30
Windows server stops responding when the dfm service stop -k command is run . . . . .	30
Others issues . . . . .	31
No sample scripts for generating volume and Snapshot names . . . . .	31
Performance issues in backup-related operations after upgrade to OnCommand Unified Manager 5.2 . . . . .	31
DataFabric Manager server discovers obsolete storage objects after Data ONTAP reversion . . . . .	31
DataFabric Manager server configured on a Microsoft cluster and configured as cluster services fails . . . . .	32
Terminating a DataFabric Manager server restore operation by using Ctrl-C causes subsequent restore operations to end in SQL database failure . . . . .	32

Cannot use purely numeric host service-assigned object names or virtual object IDs . . . . .	32
Names generated for secondary qtree backup copies of root qtrees in datasets of virtual objects might be shortened . . . . .	33
Reregister host services after restoring an older version of a DataFabric Manager server . . . . .	33
Host service selects wrong DataFabric Manager server IP address . . . . .	34
Mount operation of remote backup fails with "write operations are not allowed" message . . . . .	34
Job status labeled as successful for incomplete scheduled backups. . . . .	34
Opening the Server tab by clicking a link from the Related Objects pane in the Datasets tab restricts the Server tab display. . . . .	34
<b>Documentation changes . . . . .</b>	<b>37</b>
Update to setting up your system . . . . .	37
Update to requirements for removing Open System SnapVault relationships from Operations Manager console . . . . .	37
Update to upgrading OnCommand Unified Manager Core Package and restoring the database . . . . .	37
Update to upgrading OnCommand Unified Manager Core Package and restoring the database . . . . .	38
Update to adding thresholds using Performance Advisor . . . . .	38
Incorrect syntax for dfpm migrate start command in the man pages . . . . .	38
Update to report scheduling . . . . .	38
Updated description of Volume Quota Overcommitted Threshold (%) . . . . .	39
Update to What Storage Capacity reports are topic . . . . .	39
Correction to the Configuration requirements for OnCommand Core Package in VCS topic . . . . .	39
Web search engines can access DataFabric Manager server web pages . . . . .	39
<b>Where to find product documentation and other information . . . . .</b>	<b>41</b>
<b>Websites . . . . .</b>	<b>43</b>
<b>Copyright and trademark information . . . . .</b>	<b>45</b>
Trademark information . . . . .	45
<b>Notices . . . . .</b>	<b>47</b>



---

# OnCommand Unified Manager Core Package 5.2 Release Notes Overview

The release notes describe new features, fixed issues, and known issues for OnCommand Unified Manager Core Package 5.2.

## Overview of the Core Package

OnCommand Unified Manager Core Package is part of the OnCommand family of IBM N series management software. OnCommand Unified Manager is the product name for the combination of the Core Package and the Host Package.

OnCommand Unified Manager Core Package brings together multiple products, including Operations Manager, Protection Manager, and Provisioning Manager, into a single framework that provides an integrated, policy-based data and storage management solution for virtual and physical environments. Storage management solution for a virtual environment also requires the installation of OnCommand Unified Manager Host Package.

## OnCommand Unified Manager Core Package components

The following components are included in the OnCommand Unified Manager Core Package:

### **OnCommand console**

Provides the primary graphical interface to monitor, manage, and report in virtual and physical environments.

### **DataFabric Manager server**

Provides the services and database for your storage environment.

### **Operations Manager console**

Enables comprehensive monitoring and management of storage objects, with alerts and reports, in addition to performance and configuration tools.

### **N series Management Console**

Helps you simplify and automate common data protection and storage provisioning management tasks.

It also provides a single location from which you can view and analyze comprehensive information about storage system performance.

**OnCommand Windows PowerShell cmdlets**

Enables you to perform a subset of operations using the Windows PowerShell command line.

The cmdlets are supported only for datasets containing VMware virtual objects.



---

## OnCommand Unified Manager Core Package new features

OnCommand Unified Manager Core Package 5.2 includes new features, such as support for the licensing infrastructure in Data ONTAP 8.2, updated reports to identify Flash Pool enabled aggregates, and capability to remotely mount and restore dataset backups on a VMware ESX server.

### **Data ONTAP 8.2 support**

OnCommand Unified Manager Core Package 5.2 supports the licensing changes implemented in Data ONTAP 8.2.

To monitor and manage Data ONTAP 8.2 controllers, you must use OnCommand Unified Manager Core Package 5.2.

### **Validation of DataFabric Manager server database**

The upgrade operation is aborted if the DataFabric Manager server database is corrupted.

### **Flash Pool support**

The Aggregates view and Aggregates report include information about aggregates that are Flash Pool enabled.

### **Tool for debugging**

Data Collector, a debugging tool, is available with OnCommand Unified Manager Core Package. You can use debugging tool to collect diagnostic information to resolve issues in the DataFabric Manager server.

### **Sybase support**

The DataFabric Manager server includes Sybase server 11.0.1.

### **Apache support**

The DataFabric Manager server includes Apache 2.2.23.

### **OpenSSL support**

The DataFabric Manager server includes OpenSSL 1.0.1e.



---

## OnCommand Unified Manager Core Package changes

In this release, the changes and enhancements to various features are as follows:

### **Removal of 32-bit installer**

You can install OnCommand Unified Manager Core Package 5.2 on 64-bit servers only.

### **Performance enhancement**

The upgrade operation cleans the DataFabric Manager server database by removing the storage objects that are marked-deleted in the database and the associated entries in the history tables. The performance is also enhanced by removing events older than 180 days from the day of installation or upgrade, or as specified in the `eventsPurgeInterval` option.

Cleaning the marked-deleted objects improves the performance of the DataFabric Manager server database.

### **End of support**

The Business Continuance License has reached end of support. For more information, see the communication at the IBM N series support website.

### **Data transfer reports are disabled**

By default, the data transfer reports are disabled on a newly installed DataFabric Manager server.

This option can be turned on manually.

### **CPU monitoring is disabled**

By default, CPU monitoring is disabled on a newly installed DataFabric Manager server. The CPU Usage (%) graph data is not available.

By default, CPU monitoring is also disabled, if you have upgraded to OnCommand Unified Manager Core Package 5.2. The data in the CPU Usage (%) graph expires over a period of time.

The historical data is available in the CPU Usage (%) graph in both the OnCommand console and Operations Manager console until it reaches the expiry date.

### **DataFabric Manager server uses SSH protocol**

By default, the DataFabric Manager server uses the SSH protocol for storage system communications.

Earlier, if you set the RSH protocol as the default on your Windows Server 2008 system, now you should the default to SSH.


---

## System requirements

Before you begin the software installation process, you must make sure that your host system conforms to all supported platform requirements. Servers running OnCommand Unified Manager Core Package or OnCommand Unified Manager Host Package must meet specific software, hardware, and operating system requirements.

For specific and current software, hardware, and operating system requirements, see IBM N series interoperability matrix.

**Related information:**

 IBM N series interoperability matrix: [www.ibm.com/systems/storage/network/interophome.html](http://www.ibm.com/systems/storage/network/interophome.html)



---

## OnCommand Unified Manager Core Package 5.2 fixed issues

Some known issues are fixed in this release of OnCommand Unified Manager Core Package 5.2.

---

### **The DataFabric Manager server communicates on the first available interface when the preferred interface is not running**

The DataFabric Manager server is not aware of the preferred interface that is set on systems running Data ONTAP 7.3 or later. When the option `ndmpDataUseAllInterfaces` is set to **No**, the DataFabric Manager server communicates through the first available interface assuming it to be the preferred interface.

---

### **Cannot back up SnapVault source and SnapVault destination volumes in the same aggregate**

When you try to back up a SnapVault source volume and a SnapVault destination volume in the same aggregate of a storage controller, you might see the following error message: Backup from the source data to destination data in the same aggregate is not allowed.

---

### **The dfmmonitor service stops responding**

If you have installed OnCommand Unified Manager Core Package on a Linux system that does not have RSH installed, the `dfmmonitor` might stop responding. This issue might occur on the DataFabric Manager server that manages either 7-Mode objects or clustered Data ONTAP objects.

---

### **Report share functionality unavailable through OnCommand console for users who do not have Read permission for the Global group**

If you do not have permission to read the Global group, you cannot use OnCommand console to schedule and share any reports.

---

### **Restore of database backup fails**

When you upgrade DataFabric Manager server from a 7-Mode environment to a clustered environment, and from DataFabric Manager 3.5 or earlier, Performance Advisor fails to restore the DataFabric Manager server database backup.





---

## OnCommand Unified Manager Core Package limitations

OnCommand Unified Manager Core Package 5.2 includes limitations that you should consider before you configure your environment, such as lack of support for Data Motion for Volumes on IPv6 interfaces in the N series Management Console data protection capability.

- IPv6 interfaces are not supported by the N series Management Console data protection capability.
- Support for DS4486 disk shelves is limited to monitoring the disk shelves by the DataFabric Manager server.
- If you delete the configuration files from the DataFabric Manager server, the files cannot be retrieved because they are permanently deleted.
- Data Motion for vFiler is not supported on a Flash Pool aggregate.
- You cannot enable, disable, or schedule reports from the CLI.
- You cannot over-provision 64-bit volumes.

Provisioning is limited to the aggregate size. For example, if you have an aggregate of 10 TB, you cannot create a 64-bit volume greater than the size of the aggregate, which is 10 TB. However, you can create multiple thin-provisioned volumes of up to 10 TB each.

- You cannot push a user role with the *login-sp* capability to another storage system.

### Related information:

 IBM N series interoperability matrix: [www.ibm.com/systems/storage/network/interophome.html](http://www.ibm.com/systems/storage/network/interophome.html)



---

## OnCommand Unified Manager Core Package 5.2 known issues

The known issues are categorized based on the functionality available in the product.

- Installer issues
- Provisioning capability issues
- Protection capability issues
- Performance Advisor issues
- Console issues
- Report issues
- DataFabric Manager server service issues
- Other known issues

---

### Installer issues

You should be aware of the known issues that you might face while installing or upgrading OnCommand Unified Manager Core Package.

#### **Cannot upgrade to Core Package 5.x from DataFabric Manager server 3.5**

You cannot upgrade from DataFabric Manager server 3.5 to OnCommand Unified Manager Core Package 5.x.

#### **Workaround**

1. Upgrade to DataFabric Manager server 4.0 or later.
2. Upgrade to OnCommand Unified Manager Core Package 5.x.

---

### Provisioning capability issues

You should be aware of the known issues when using the provisioning capability of the Core Package.

#### **Cannot provision a volume with the same size as its aggregate**

Provisioning a secondary volume fails when you try to provision a volume that is the same size as the aggregate, and the volume size limit set by Data ONTAP is lesser than the aggregate size limit.

## Workaround

Provision the secondary volume by setting the `pmAutomaticSecondaryVolMaxSizeMb` option to less than the volume limit.

### **Provisioning operations are not supported on systems running Data ONTAP 7.3.5.1P5**

Storage systems running Data ONTAP 7.3.5.1P5 do not support provisioning operations executed by the DataFabric Manager server.

## Workaround

Either exclude resources on storage systems running Data ONTAP 7.3.5.1P5 from DataFabric Manager server provisioning operations, or upgrade those storage systems to Data ONTAP 7.3.6 or later.

### **Data migration rollback and cutover jobs sometimes time out in N series Management Console**

Sometimes a data migration rollback and cutover operation that is invoked from the N series Management Console GUI times out even though the source and target systems are carrying acceptable input output processing loads.

In such circumstances, the following message is displayed:

```
=== SEVERITY ===  
Error: Requested operation did not complete in  
60 seconds. This could be due to network latency  
or server problems. The operation is still  
running and may succeed or fail later.
```

## Workaround

Under these circumstances data migration rollback and cutover can still be completed by using the following CLI commands:

- **`dfpm migrate rollback`**
- **`dfpm migrate cutover`**

### **vFiler unit migration operation might fail**

When you start a migration operation for vFiler units (Data Motion for vFiler), the operation might fail if you have different versions of Data ONTAP running on the primary and secondary storage systems or if the versions do not support the migration of vFiler units. No error messages are displayed if the operation fails.

## Workaround

Before you start the migration operation, ensure that you are running the same version of Data ONTAP on the primary and the secondary storage

systems and that the versions support the migration of vFiler units. For more information about the Data ONTAP versions that support Data Motion for vFiler, see IBM N series interoperability matrix.

**Related information:**

➡ IBM N series interoperability matrix: [www.ibm.com/systems/storage/network/interophome.html](http://www.ibm.com/systems/storage/network/interophome.html)

### **High CPU usage on storage systems during a Data Motion migration cutover causes storage system panic**

If source or destination storage system CPU usage rises above 90 percent during a Data Motion migration cutover phase, those systems might panic when the source vFiler is destroyed.

Although the DataFabric Manager server employs mechanisms that usually postpone migration cutover attempts in cases of high CPU usage, sometimes these mechanisms do not.

#### **Workaround**

Before beginning the cutover phase of an online migration operation, manually verify that the CPU usage on the source and destination physical storage systems is below 50 percent.

### **Provisioning policies and online migration are not supported on SAN datasets with multiple-space names**

SAN provisioning policies and online migration are not supported for datasets with names that contain multiple spaces.

For example, creating a SAN dataset named, "My First Dataset" causes the provisioning capability to display the following error message: `igroup create failed`. If this dataset is created with the online migration capability enabled, that capability is disabled by this condition.

#### **Workaround**

Name the SAN datasets with names containing not more than a single space.

### **During online migration, SnapManager-initiated backups to secondary storage fail**

During online migration initiated through N series Management Console, backups of migrated objects to secondary storage locations initiated through a SnapManager product fail sometimes.

Because, during online migration, objects with identical names and paths exist at both migration source and destination locations, the DataFabric Manager server running a backup job cannot determine which object to back up to secondary storage.

## Workaround

On the DataFabric Manager server, create a text file named testpoints.txt.

- In the text file insert the following text:

```
testpoint -m check-datamotion-progress -n enforce -e 1 -r 1
```

- Locate the text file on the following path:

```
dfm_installed_path /log/testpoints.txt
```

If *dfm\_installed\_path* is not specified, the value of the environment variable NETAPP\_TESTPOINTS is the default path.

## Sometimes the Cancel button does not cancel a failed migration

Sometimes a data migration operation initiated from the N series Management Console vFiler Units window might fail, and even though the Conformance Results window instructs you to cancel the operation, clicking the **Cancel** button only displays a message that the Cancel operation failed.

## Workaround

Complete the following actions:

1. Go to N series Management Console/ Storage Systems Host.
2. Click **Refresh**.  
Wait for the Refresh operation to finish.
3. In the CLI, enter the command: `dfpm migrate fix`
4. Return to N series Management Console vFiler Units window and click **Cancel**.

After the migration operation is cancelled, you can attempt the migration again.

---

## Protection capability issues

You should be aware of the known issues when using the protection capability of the Core Package.

### Creation of an Open Systems SnapVault relationship fails on Linux

When you create an Open Systems SnapVault relationship from the DataFabric Manager server, the operation stops responding.

You should verify if running the **snapvault status** command displays the following error message: Transferring (0 KB done). You should also verify if the Open Systems SnapVault space estimator stops responding while processing a mount point. If you experience both these issues, then you can perform the following workaround.

### **Workaround: Disable AutoSupport and the space estimator**

**Note:** Ensure that you have access to the `snapvault.cfg` file.

1. Open the `snapvault.cfg` file.
2. Set Value to **FALSE** for `[QSM:EnableASUP]`.  
Result: AutoSupport is disabled.
3. Set Value to **FALSE** for `[Configuration:Run Estimator]`.  
Result: The space estimator is disabled.

### **Workaround: Enable AutoSupport and the space estimator, and unmount the failed mount points**

**Note:** Ensure that you have access to the `svestimator.exe` file.

Locate the mount point at which the operation stopped responding by performing the following steps:

1. Run the Open Systems SnapVault utility `svestimator.exe`.  
Result: One or more messages are displayed on the screen: 1  
ASUPDATA:-136288576 18 0 045656 Skipping mount point  
'/t/99.888.777.666/root'.
2. Disable `svestimator.exe`, and set `EnableASUP` to **false**.  
Result: The Open Systems SnapVault relationship creation stops responding indefinitely.
3. Open the `qsmserver` log.
4. Locate the mount point at which the creation operation stopped responding.
5. Unmount that mount point from the storage system.
6. Verify that the Open Systems SnapVault relationship is being created.  
If the creation fails, repeat the above steps to locate the other mount point that is not working, and then unmount it.

### **Cannot include an NFS datastore in a renamed volume in a dataset**

Renaming the volume that contains an NFS datastore, and then attempting to include that datastore in a dataset can result in a `Could not find any storage mapping... error message`.

When you change the volume name, the datastore's NFS export path might be automatically changed on the storage system. However, the original NFS export path remains in effect on existing clients, such as the ESX server.

In such cases, you cannot add that datastore to a dataset because the OnCommand console cannot map the datastore to an existing NFS export.

## Workaround

None.

## Creating a dataset fails if the dataset has storage service with provision only policy

You can use the Add Dataset Using Storage Service wizard or the storage-service-dataset-provision API with the **provision-member** parameter to create a dataset with provisioning policy and provision a member to the dataset.

In both cases, provisioning the member to the dataset fails with the following error message: There is no node by name in the DP Policy No protection NM.

## Workaround

You can avoid this error message by using one of the following methods:

- Using the Add Dataset Using Storage Service wizard:
  - Create an empty dataset by using the wizard.
  - Provision members to the dataset after you have created the dataset.
- Using the storage-service-dataset-provision API:
  - Create an empty dataset by using the storage-service-dataset-provision API, and by not providing inputs to the **provision-member** parameter.
  - After you have created the dataset, use the dataset-provision-member API to provision members to the dataset.

## Users cannot modify values in custom fields

Even though you have the appropriate permissions to create and edit a custom field object that is configured for a dataset, protection policy, or provisioning policy, you must also have the DFM.Database.write permission to edit a value in a custom field using N series Management Console. Permission to create and edit a custom field object configuration does not automatically include the permission to edit the value in the field.

## Workaround

None.



## Provisioning failure despite sufficient space

An overcommitment threshold set for the DataFabric Manager server can impact the provisioning of flexible volumes in the resource pool.

The DataFabric Manager server enables you to set a `NearlyOvercommittedThreshold` parameter that reserves space on your system for projected future use. If this threshold is set too low, you can encounter problems provisioning volumes on destination storage systems. When such provisioning problems occur, you might see an inconsistency in which the dry run result reports insufficient space when the reported space available for the destination aggregate appears to be sufficient.

### Workaround

To resolve this problem, increase the Nearly Overcommitted Threshold in the DataFabric Manager server. By default, the `aggrNearlyOvercommittedThreshold` value is set to 95%. Increasing this value causes the DataFabric Manager server to allocate more storage in an aggregate before reporting it as full. Setting this value over 100% could result in an out-of-space condition. Note that other factors, such as Snapshot copy retention settings or resizing primary volumes, can also result in an out-of-space condition.

## Restore operations performed from most recent backup on the tertiary (mirror destination) node might fail

If you use a back up, then mirror topology to restore your most recent backup of the secondary node, the restore operation might fail. Secondary qtrees are present in the base Snapshot copy of the backup relationship, which can cause the restore operation to fail because LUNs exported from such qtrees are read-only.

### Workaround

You can retry the restore operation from the same backup version on the primary or the secondary node of the dataset, or you can restore from an older backup version.

## Improperly specified SnapVault source to target access list causes restore operation failure

A restore operation from a SnapVault secondary system to a SnapVault primary system might fail if the SnapVault access list settings that are specified on the primary system and secondary system by the **options snapvault.access host** command do not properly enable access from primary system to host and from host to primary system.

This type of failure results in empty qtrees at the restore destination volume and an uninitialized SnapVault relationship that you must clean up manually by issuing the **snapvault abort** command on the restored host.

Further complicating the cleanup, if the owner of the restore destination volume is a vFiler unit, the **snapvault status** command lists this uninitialized relationship under the name of the vFiler unit's physical storage system. The **snapvault abort** command, however, must use the name of the vFiler unit.

## Workaround

To avoid these restore failure and cleanup problems, ensure that both the SnapVault primary system and SnapVault secondary system enable SnapVault access to one another in their **options snapvault.access host** settings before you start a restore operation from the OnCommand console.

For example, if the SnapVault primary system is primary1 and the SnapVault secondary system is secondary2, then the following settings are the minimum required to ensure a successful restore:

- On the SnapVault primary system primary1, enter the following command:  
options snapvault.access host=secondary2
- On the SnapVault secondary system secondary2, enter the following command: options snapvault.access host=primary1

## Backup scripts that refer to directories or files require full path names

Backup scripts that refer to directories or files for creation or output redirection might cause these directories or files to be created in unexpected locations if the script uses a relative path rather than the full path to refer to these directories and files.

## Workaround

To ensure that directories and files are created where you intend them to be created, include the full path name to your intended location on the DataFabric Manager server or the host services system.

For example, the following full paths guarantee that the specified directory abc and the specified output redirection file abc.txt are created in the root directory of the DataFabric Manager server or host services system:

- mkdir c:\abc
- dir >> c:\abc.txt

However, the following relative paths might cause the specified directories or files to be created in a subdirectory that is not obvious:

- mkdir abc

- `dir >> abc.txt`

## **Compression option for imported SnapMirror connections not preserved**

When you use N series Management Console to import a SnapMirror relationship that has compression enabled and if the DataFabric Manager server option `hostPreferredAddr1` is set for both the source and destination storage systems, the compression option for SnapMirror connections is not preserved.

After the SnapMirror relationship is imported, the next transfer update job for the dataset creates a new connection with the preferred host addresses and modifies the `snapmirror.conf` entry. This entry does not have the `compression=enable` option.

If you manually edit the `snapmirror.conf` file and add the `compression=enable` option for any existing connections, and if `hostPreferredAddr1` is set for both the source and destination storage systems, the next transfer update job for the dataset overwrites the `snapmirror.conf` entry and deletes the compression option.

## **Workaround**

Ensure that the `hostPreferredAddr1` option is not set for the source and destination storage systems. If the option is not set, N series Management Console leaves the `compression=enable` option intact in the `snapmirror.conf` file.

## **First update of a mirror connection after the source has shrunk and then grown might generate a false error message**

When Dynamic Secondary Sizing is enabled for a mirror relationship and the source volume has shrunk for one update, the next update might generate a secondary resizing error message that can be ignored.

If the source volume has shrunk for one update, the next update attempts to shrink the destination volume. The update job might display an error for the resize operation, but the mirror job is successful nevertheless.

This error message can be ignored if seen only as part of a first update after a source volume has grown. The displayed message is similar to the following:

```
Error Message: myDfmStation: Could not
resize secondary volume myFiler:/myVolumeName
(10565) to 10.0 GB. The new volume size would
be less than that of the replica file system.
```

This issue fixes itself and no backups are lost, so no workaround is necessary.

## Open Systems SnapVault relationship baseline backup version is not registered

When N series Management Console creates a new Open Systems SnapVault relationship that uses extended ASCII characters in the source directory path, the relationship gets created but a baseline backup version is not registered.

This occurs because the create job includes a name from the storage system that is encoded differently from what it expects, which causes the job to fail when trying to match the directory with the newly created relationship.

### Workaround

Either manually run the backup job immediately after the relationship create job ends, or wait for the next scheduled job to run. These options work because neither on-demand or scheduled jobs require extra information from the storage system to create a backup version.

## Adding a directory from an unsupported file system to a dataset fails with an error

An attempt to add a directory that resides on a host file system that is not supported by Open Systems SnapVault causes N series Management Console to issue an error indicating that the directory does not exist or is not suitable for backup.

You can use Open Systems SnapVault to back up supported file system types only. For example, Open Systems SnapVault supports only NTFS on systems running Windows.

### Workaround

To determine which file system types are supported by the version of Open Systems SnapVault that you are using, you can run the **svinstallcheck** utility on the Open Systems SnapVault host. See the Open Systems SnapVault documentation for details.

## Enabling on-demand backup operations on SnapDrive-generated datasets requires the `snapdrive dataset backup_add` command

For on-demand backup operations to apply to SnapDrive-generated datasets, you must first invoke the **snapdrive dataset backup\_add** command.

Clicking the **Protect Now** button in the N series Management Console Datasets tab, does not necessarily create a backup copy in secondary storage. To enable datasets generated by SnapDrive for Unix or SnapDrive for Windows for on-demand backup operations in N series Management Console,

the SnapDrive database administrator must invoke the **snapdrive dataset backup\_add** command in the SnapDrive for Windows or SnapDrive for UNIX interface.

### **SnapMirror job fails with "process was aborted" message**

Problems on the data source node can prevent N series Management Console from creating a SnapMirror relationship, even though a preview of your dataset configuration detected no configuration problems.

When the problem is with the data source node itself, the job that creates the SnapMirror relationship fails and the protection application displays the following message: process was aborted.

Problems with a data source node that can prevent the creation of a SnapMirror relationship include the following:

- The maximum number of Snapshot copies that can be created on the source node has been reached; therefore, no new Snapshot copies can be created.
- The Snapshot copy on the source node is locked, possibly by another operation.

You can find additional information about problems that can prevent the creation of a SnapMirror relationship by checking the following files on the SnapMirror destination node:

- /vol/vol0/etc/logs/snapmirror
- /vol/vol0/etc/logs/messages

### **Temporary conformance impact of deleting volumes or qtrees**

If you delete a volume or qtree that is part of a dataset without deleting the volume or qtree from the DataFabric Manager server database, the deleted object is temporarily flagged in the database as "disappeared" for up to an hour before it is marked deleted.

N series Management Console references the DataFabric Manager server database regarding objects used in protection implementations. Until a deleted object is marked as deleted in the database, N series Management Console attempts to create relationships for it, as specified in the policy applied to its dataset, generating conformance error messages. After a deleted object is marked as deleted in the database, N series Management Console no longer attempts to create relationships for it, and the conformance error messages are no longer generated.

### **Workaround**

You can perform one of the following actions:

- To stop the conformance error messages, manually delete the volume or qtree from the DataFabric Manager server database.
- Wait up to an hour for the DataFabric Manager server to mark the “disappeared” object as deleted, and the conformance error messages are no longer generated.

### **vFiler units cannot provision secondary storage of virtual objects**

You cannot use vFiler units to provision secondary storage nodes in a dataset of virtual objects.

### **Qtree is not the source for the snapmirror destination**

If N series Management Console generates the error message the qtree is not the source for the snapmirror destination, there are three possible causes:

1. The SnapVault relationship was released manually from the source node.
2. The base Snapshot copy was deleted manually from the source node.
3. The source qtree for which the relationship was originally created was deleted and later re-created with the same name.

### **Number of bytes of data transferred during backup is inaccurate**

Monitoring a backup job reveals that the number of bytes of data transferred during the job does not match the number of bytes displayed in the N series Management Console Jobs window.

This is expected behavior. The number of bytes transferred is an approximation and does not reflect an exact count; it is always less than the actual number of bytes transferred. For jobs that take a short time to complete, N series Management Console might report a data transfer size of zero.

### **Faulty export file entry might prevent inclusion of NFS-exported virtual machine or datastore objects in a dataset of VMware objects**

The DataFabric Manager server cannot read exports file entries whose path names do not begin with **vol** if the path name for any NFS export virtual machine or datastore entry in a storage system's etc/exports file does not start with **/vol**.

The objects that those unread entries specify remain undiscovered by the DataFabric Manager server, and any attempt to include those objects in a dataset results in a storage mapping error.

### **Workaround**

If a storage mapping error is caused by an incorrect path name, you must delete the entries that do not start with **/vol** from the etc/exports file, and then allow the DataFabric Manager server monitor to run again. After you run

host service discovery (**dfm hs discover**) one more time, the DataFabric Manager server allows inclusion of the newly discovered objects in a dataset.

### **Mixed LUN connections in a virtual machine can cause backups to fail**

Virtual machines containing mixed environments for example, an ESX-enabled environment with a direct-connect iSCSI environment can cause a backup to fail.

#### **Workaround**

You can avoid this by not mixing LUN connections in a virtual machine.

---

## **Performance Advisor issues**

You should be aware of the known issues when using Performance Advisor in your environment.

### **Changing the Performance Advisor perfTopnchartinterval setting might require you to restart N series Management Console**

If you modify the Performance Advisor perfTopnchartinterval setting to change your data sampling rate, the chart data in Performance Advisor Monitor Dashboard panels might not be consistently displayed until you restart your N series Management Console session.

#### **Workaround**

After you use the CLI command **dfm options set perfTopnchartinterval** to change the data sampling and update rate for the panels in the Performance Advisor's Monitor Dashboard window, you must login again to N series Management Console to reliably view those updated panels.

### **Delay in updating entries in the Top Network Interface view**

The Top Network Interface view does not display updated information about global groups because there is a delay in updating the objectMap table.

#### **Workaround**

None.

### **Client statistics data not collected**

Performance Advisor is unable to collect client statistics data for NFS and CIFS objects in storage systems running Data ONTAP 8.1 operating in 7-Mode.

#### **Workaround**

None.

---

## Issues related to Data ONTAP

You must be aware of the known issues related to Data ONTAP.

### Data Motion for vFiler fails if volumes are in a SnapMirror relationship

You cannot perform online migration of vFiler units (also called Data Motion for vFiler) if the volumes in the vFiler unit are in a SnapMirror relationship. The operation fails with an error message.

#### Workaround

1. Remove the SnapMirror relationship by running the following command on the source volume: `snapmirror release source_volume destination_system:destination_volume`
2. Delete the old Snapshot copies.

### Cutover fails during Data Motion for vFiler

On a Data ONTAP 8.2 system, cutover fails during online migration of vFiler units (also called Data Motion for vFiler).

#### Workaround

1. Create a file `testpoints.txt` at the following location:  
`installation_dir/log/.`
2. Add the following to the `testpoints.txt` file: `testpoint -m datamotion -n disable-load-checks -e 1 -r 1`

Result: The CPU and disk load check module is disabled.

---

## OnCommand console issues

You should be aware of the known issues when using the OnCommand console.

### Groups tab displays incorrect values for total capacity and used capacity

When storage controllers, aggregates, volumes, qtrees, and clusters are members of a group, the values displayed for total capacity and used capacity are incorrect.

#### Workaround

View the correct values for total capacity and used capacity from either the appropriate storage inventories ( Storage Controllers view, Aggregates view, Volumes view, Qtrees view, and Clusters view) or the storage capacity reports.



## **Memory leak on Internet Explorer 9**

If you have logged in to the OnCommand console and it is left unattended for 5 days or more on Internet Explorer, a memory leak occurs. The memory leak might also occur if the DataFabric Manager server is loaded with too many controllers.

### **Workaround**

Close the Internet Explorer browser and launch OnCommand console in a new browser window.

## **The OnCommand console repeatedly prompts for login credentials**

After the Core Package installation is complete, the OnCommand console repeatedly prompts for the user name and password.

### **Workaround**

Clear the browser cache, and restart the browser.

## **Charts are not displayed after disabling HTTP and enabling HTTPS**

If you disable HTTP and enable HTTPS, and then access the OnCommand console by using the IP address, charts and icons are not displayed.

### **Workaround**

Use an Fully qualified domain name (FQDN) the first time you access the OnCommand console to view the charts and icons.

## **Internet Explorer does not display content in the OnCommand console if HTTPS is configured**

If the DataFabric Manager server is configured for HTTPS, Internet Explorer displays the following message: This Web page contains content that will not be delivered using a secure HTTPS connection, which could compromise the security of the entire Web page.

After the error message is displayed, when you choose to view the content of the page either securely or not securely, the OnCommand console does not display any content.

## **Edit Group Membership page displayed incorrectly in the Operations Manager console**

The Edit Group Membership page is not displayed correctly when the Operations Manager console is opened in the Internet Explorer 8 browser. The **Groups Member** text box is displayed in a separate pane at the bottom of the Edit Group Membership page.

## Workaround

1. Launch the OnCommand console.
2. Click the **Administration** menu.
3. Click the **Groups** option.
4. From the **Groups** tab, select the group you want to modify.
5. Click **Edit**.

## Breadcrumb trail not updated

When you navigate from one group to another in the Groups list, the breadcrumb trail in the Storage tab and Server tab does not display the updated trail information.

For example, in the Storage Controller view, when you double-click an object O1, which belongs to the group G1, the breadcrumb trail is modified to display the list view of O1. However, if you select another group, G2, from the Groups list to which O1 does not belong, the breadcrumb trail is not updated.

## Workaround

Click the parent breadcrumb object to update the trail information.

## Storage tab displays blank page

When you perform a particular task from the Storage tab, then navigate away from the tab and return to it, a blank page is displayed. This occurs when you add an object to a group from any of the inventory views, navigate to another tab, change from the Global group to the group you added the object to, and then return to the Storage tab.

## Workaround

To update the page, click any inventory view and return to your previous view.

## Setting or modifying the storage efficiency script schedule fails with an error message

While you are setting the storage efficiency script schedule or modifying an existing schedule in the Operations Manager console, an Internal Server Error message is displayed.

The following error message is displayed:

Internal Server Error

The server encountered an internal error or misconfiguration and was unable to complete your request.

Please contact the server administrator, root@localhost and inform them of the time the error occurred, and anything you might have done that may

have caused the error.

More information about this error may be available in the server error log.  
Apache/2.2.23 (Unix) mod\_ssl/2.2.23 OpenSSL/1.0.1e Server at 127.0.0.1  
Port 80

## Workaround

Add or modify the script schedule by entering the appropriate command in the CLI:

- `dfm script schedule add`
- `dfm script schedule modify`

---

## Report issues

You should be aware of the known issues related to report generation of the DataFabric Manager server database, and issues related to the display of reports.

### **No data in about volume Performance Advisor after a volume move**

When you move a volume from one aggregate to another aggregate within the same node, or between aggregates of different nodes, Performance Advisor does not retain any historical data about the volume.

## Workaround

None.

### **The dfm report view command fails with an error message**

When you run the **dfm report view** command, the following error message is displayed when the Web UI service is down: `http_open_url_socket` failed with return value -2. Please validate `dfm http` and `webui` services are running correctly.

When you log in to view the DataFabric Manager server reports from the OnCommand console, the session expires.

## Workaround

None.

### **Events Current report shows wrong event name**

The events Volume Quota Overcommitted and Volume Quota Not Overcommitted are named incorrectly in the Events Current report. Information provided as Volume Quota Not Overcommitted should be read as Volume Quota Overcommitted. Information provided as Volume Quota Overcommitted should be read as Volume Quota Not Overcommitted.

## Workaround

None.

### Cannot create alarms with an SNMP traphost

If you create an alarm with an SNMP traphost in the *community@<ip>:port* format, the alarm fails with an error message indicating that the traphost address is invalid.

## Workaround

Use the `dfm alarm create -A` command with full Lightweight Directory Access Protocol name to create an alarm with an SNMP traphost.

### Email alerts generated for deleted objects

The DataFabric Manager server sends email alerts for deleted storage objects. The email provides a link to the event details on the OnCommand console; however, no details are available because the object is deleted.

## Workaround

None.

---

## DataFabric Manager server service issues

You should be aware of the known issues in the DataFabric Manager server services.

### The `dfm service start http` command fails with an error message

If you cancel the `dfm ssl server setup` command, and then stop and restart the HTTP service, the `dfm service start http` command fails with the following error message: `SSLCertificateKeyFile: file '/opt/IBMdsm/conf/server.key' does not exist or is empty.`

The `dfm service start http` command fails because the `server.key` file is deleted when you cancel the `dfm ssl server setup` command.

## Workaround

Run the `dfm ssl server setup` command before running the `dfm service start http` command.

### Windows server stops responding when the `dfm service stop -k` command is run

After a restore operation is complete, when you run the `dfm service stop -k` command, the Windows server might stop responding and reboot.

## Workaround

1. Stop all services except the SQL service.
2. Take a new backup.
3. Restore the new backup.

---

## Others issues

You should be aware of the known issues related to the DataFabric Manager server, such as email alert generated for deleted vFiler units, discovery of obsolete storage objects, failure to mount remote backups.

### **No sample scripts for generating volume and Snapshot names**

There are no sample scripts that you can use as examples to create custom scripts for the DataFabric Manager server to generate volume and Snapshot names.

However, the DataFabric Manager server generates some environment variables that can be used in custom scripts to generate volume and Snapshot names for the volumes that are managed by datasets.

## Workaround

None.

### **Performance issues in backup-related operations after upgrade to OnCommand Unified Manager 5.2**

When you upgrade to OnCommand Unified Manager 5.2 from DataFabric Manager server 4.0.2 or earlier, you might experience performance issues related to backup operations, even if the option `purgeJobs0lderThan` is set.

This performance issue might occur due to the retention of a large amount of conformance and job history data. Conformance and job history data older than the value specified in the `purgeJobs0lderThan` option is cleared on a nightly basis. You can contact technical support if you require further assistance.

## Workaround

None.

### **DataFabric Manager server discovers obsolete storage objects after Data ONTAP reversion**

When you revert the Data ONTAP version from 8.1.1 to 8.0.3, the DataFabric Manager server continues to discover objects from the earlier version.

## Workaround

1. Delete the systems running Data ONTAP 8.1.1.
2. Add the systems running Data ONTAP 8.0.3.

### DataFabric Manager server configured on a Microsoft cluster and configured as cluster services fails

If you configure DataFabric Manager server on a Microsoft cluster, and perform the **dfm datastore setup** operation, the DataFabric Manager server pauses to move the directories, and failover occurs. This causes the operation to fail.

## Workaround

1. Stop DataFabric Manager server cluster services.
2. Using the CLI, start the DataFabric Manager server services.
3. Run the **dfm datastore setup** command.
4. Start DataFabric Manager server cluster services.

### Terminating a DataFabric Manager server restore operation by using Ctrl-C causes subsequent restore operations to end in SQL database failure

If you terminate a **dfm backup restore** operation by using Ctrl-C, subsequent attempts to restore end in SQL database failure, even if rollback of the terminated original restore operation succeeds.

A message similar to the following is displayed:

```
***** SQL error: Item
'/opt/IBMdfm/data/monitordb.dbR' already exists
Error: Unable to upgrade the database to the latest file
       format.
Error: Service start failed.
```

## Workaround

After you start **dfm backup restore**, observe the warning message and avoid terminating the operation by using Ctrl-C. If you cannot avoid this action, contact technical support for help in restoring database operations. Do not attempt to delete any database file without first contacting technical support.

### Cannot use purely numeric host service-assigned object names or virtual object IDs

DataFabric Manager server commands do not accept purely numeric host service-assigned object names or virtual object IDs.

## Workaround

Use alternate IDs for both host service-assigned object names and virtual objects.

### **Names generated for secondary qtree backup copies of root qtrees in datasets of virtual objects might be shortened**

The names of secondary qtrees that are generated by OnCommand console backups of primary root qtrees might be renamed and limited to 30 characters.

If the custom naming format for secondary qtrees is %Q (which specifies using the primary qtree name for the secondary qtree name also), and if in a dataset of virtual objects a primary qtree is a root qtree with the name "-" or "etc", then an OnCommand console protection job substitutes the %Q format with a %L\_%S\_%V (*custom-label\_primary-storage-system\_primary-volume*) format when generating a name for the secondary qtree. The resulting qtree name is limited to 30 characters. Any characters beyond that number are truncated.

If, in the previous case, the naming format for secondary qtrees specifies the primary qtree name plus extension characters (for example, %Q%1, %Q%2, or %Q%3), the qtree name is limited to 30 characters plus the extension characters.

Limiting the length of the virtual object name accommodates restoration of the virtual object if restoration is necessary.

## Workaround

If you specify a custom label to be associated with a dataset, consider the possibility that it might be combined with the storage system name and the primary volume name to generate the name for a secondary qtree backup copy of a root qtree, which might in turn be subject to truncation if the name is over 30 characters.

### **Reregister host services after restoring an older version of a DataFabric Manager server**

You must reregister a host service after you restore an older version of a DataFabric Manager server. If you restore an older version of a DataFabric Manager server and a host service is not reregistered with the DataFabric Manager server, any operations on that host service fails until the host service is reregistered.

## Workaround

None.

## Host service selects wrong DataFabric Manager server IP address

While registering a host service in MSCS, VCS, or a multihomed environment, the DataFabric Manager server does not select the right IP address. This might cause server notifications to fail.

### Workaround

- Use the following command to register the host service and explicitly mention the right IP address: `dfm hs register -i dfm_server_ip_address hostservice_dns_or_ip_address`
- If the host service is registered with the Cluster IP address, set the new IP address by entering the following command: `dfm hs configure -i dfm_server_ip_address hostservice_dns_or_ip_address`
- View the configured DataFabric Manager server IP address by entering the following command: `dfm hs diag hostservice_dns_or_ip_address`

## Mount operation of remote backup fails with "write operations are not allowed" message

When mounting a backup that contains the latest Snapshot copy on a SnapVault destination, the backup is mounted as read-only and write operations are not allowed on the mounted datastore.

### Workaround

To recover your data, you must copy the VMDK to a read/write datastore.

## Job status labeled as successful for incomplete scheduled backups

If a dataset consists of virtual machines from multiple host services and one of the host services goes offline, the associated scheduled backup job does not fail or display an error. Dataset backup failure alarms are not triggered under these circumstances.

### Workaround

Set up alarms that monitor the host service status. This way, you will know when a host service goes offline.

**Note:** This issue does not affect on-demand backup jobs, which display partial success if the host service goes down during the operation.

## Opening the Server tab by clicking a link from the Related Objects pane in the Datasets tab restricts the Server tab display

If you click the link of a VMware object that is listed in the Related Objects pane of the Datasets tab, the OnCommand console opens the Server tab to display the details for the clicked object; however, only the data for the clicked object is displayed.



The other objects of that type that are normally displayed in the Server tab are not listed when it is displayed in this way. After you access the Server tab in this way, you cannot refresh it or modify its filtering to list or display details for any other virtual objects.

## **Workaround**

Click the **All** breadcrumb button in the top left corner to restore the display of all the objects of the same type.

You can also select another object type in the navigation pane, and then reselect the original object type to restore the display of all objects of that type.



---

## Documentation changes

There are corrections or changes to the documentation provided with OnCommand Unified Manager Core Package 5.2. This information should be added to the documentation.

---

### Update to setting up your system

You must configure Lightweight Directory Access Protocol (LDAP) to enable LDAP authentication to the DataFabric Manager server.

LDAP authentication is supported on a DataFabric Manager server that is installed on a Linux server.

For information about the LDAP options and its configuration steps, see *OnCommand Unified Manager Online Help*.

---

### Update to requirements for removing Open System SnapVault relationships from Operations Manager console

The `hostLogin` and `hostPassword` options must be set for any secondary storage system for data cleanup after removal of the relationship. If these options are not set or are not set correctly, the following error message is displayed when you try to delete the relationship: Unable to obtain filer from host stab.

---

### Update to upgrading OnCommand Unified Manager Core Package and restoring the database

You can view the `sybase.log` file, which is located at `install_directory/log/`, to monitor the progress of the database delete operation for the mode that is being removed.

This information should be included in the *OnCommand Unified Manager Installation and Setup Guide*.

When you upgrade a DataFabric Manager server that was earlier managing both 7-Mode and clustered Data ONTAP objects to a 7-Mode environment, a message is displayed after you select the mode: Purging objects of undesired mode from the database. During a restore operation, you are prompted to delete the objects.

---

## Update to upgrading OnCommand Unified Manager Core Package and restoring the database

You can view the `sybase.log` file, which is located at `install_directory/log/`, to monitor the progress of the database delete operation for the mode that is being removed.

This information should be included in the *OnCommand Unified Manager Installation and Setup Guide*.

When you upgrade a DataFabric Manager server that was earlier managing both 7-Mode and clustered Data ONTAP objects to a 7-Mode environment, a message is displayed after you select the mode: Purging objects of undesired mode from the database. During a restore operation, you are prompted to delete the objects.

---

## Update to adding thresholds using Performance Advisor

For storage systems running Data ONTAP 8.1.1 or later, you cannot use Performance Advisor to add thresholds for the `nfsv3:nfsv3_ops` counter.

---

## Incorrect syntax for `dfpm migrate start` command in the man pages

The manual (`man`) pages incorrectly states the syntax for the `dfpm migrate start` command for volume and aggregate mapping using the `-f` option as:

```
vfiler-volume1=destination-aggregate1,vfiler-volume2=destination-  
aggregate2,....
```

The correct syntax is:

```
vfiler:/volume1=dest-filer:aggregate1,vfiler:/volume2=dest-  
filer:aggregate2,....
```

---

## Update to report scheduling

When you retain saved reports for a long period of time, you must delete the old reports or move the report archive directory `reportsArchiveDir` to a new location that has sufficient space.

You can delete the reports by using the `dfm report output delete` command.

---

## Updated description of Volume Quota Overcommitted Threshold (%)

Volume Quota Overcommitted Threshold (%) specifies the percentage at which a volume is considered to have consumed the whole of the overcommitted quota space for that volume. A volume crosses this threshold if the sum of the tree quotas of all the qtrees in that volume is more than the volume quota overcommitted threshold.

This updated information should be included in the description for this field in the *OnCommand Unified Manager Online Help*.

---

## Update to What Storage Capacity reports are topic

Space efficiency reports are available for 7-Mode environments only. This information is missing in the list item and section related to space efficiency reports in the “What Storage Capacity reports are” topic in the OnCommand console Help.

---

## Correction to the Configuration requirements for OnCommand Core Package in VCS topic

DataFabric Manager server supports Native ext3 File System and Logical Volume Manager (LVM) only. The following statement in the Installation and Setup Guide is incorrect: Veritas File System and Volume Manager on Native ext3 File System and Logical Volume Manager (LVM) must be used.

The statement should read as follows: Native ext3 File System and Logical Volume Manager (LVM) must be used.

---

## Web search engines can access DataFabric Manager server web pages

The DataFabric Manager server enables search engines to search, find, and index the web pages in the Operations Manager console; however, you can configure a text file to block search engine access, if required.

To block access by search engines to the web pages, create a file `robots.txt` in the directory `installpath/DFM/web`.

If the file contains only the following text, search engines do not search or index Operations Manager web pages:

```
User-agent:*  
Disallow:/
```



## Where to find product documentation and other information

There are additional sources of information for this release of the OnCommand Unified Manager Core Package.

The following table describes where to find additional information about OnCommand Unified Manager Core Package functionality.

To learn about...	Go to this information source...
How to monitor and manage multiple N series systems, storage resources, user quotas, content distribution, and appliance configuration	<i>OnCommand Unified Manager Operations Manager Administration Guide</i> and the <i>Operations Manager Help</i>
How to download and install N series Management Console	<i>OnCommand Unified Manager Installation and Setup Guide</i>
How to use the Performance Advisor application to view performance data and customize the tools; also, how to use the statistical counters for storage system performance that are contained in Data ONTAP	<i>OnCommand Unified Manager Performance Advisor Administration Guide</i> and the <i>N series Management Console Help</i>
The workflows you might perform by using the N series Management Console provisioning and data protection capabilities to provision space and protect data by using real-life scenarios	The workflow examples in the <i>OnCommand Unified Manager Guide to Common Provisioning and Data Protection Workflows for 7-Mode</i>
The tasks involved when you use the N series Management Console data protection disaster recovery capabilities to set up and manage data protection relationships and disaster recovery, in addition to related concepts and reference information	<i>Protection Manager and Disaster Recovery Help</i> in N series Management Console
The tasks involved when you use the N series Management Console provisioning capability to provision space for your data, in addition to monitoring and managing space utilization	<i>Provisioning Manager Help</i> in N series Management Console

<p>How to use the CLI commands for the DataFabric Manager server, Backup Manager, Content Management System, and N series Management Console protection, provisioning, and disaster recovery capabilities</p>	<p>The help pages accessible from the command-line interface</p> <p>Manual (man) pages are also included in the <i>OnCommand Unified Manager Operations Manager Administration Guide</i></p>
---	--

## Documentation on the N series support website

The OnCommand Core Package documentation set is available on the N series support website (accessed and navigated as described in Websites) at [www.ibm.com/storage/support/nseries/](http://www.ibm.com/storage/support/nseries/).

## Online information

There are several ways to access online documentation and technical resources.

### Online documentation

Included with your DataFabric Manager software is a Web-based administration tool, Operations Manager, and its Help.

Included with your N series Management Console software is Help for protection, provisioning, and disaster recovery capabilities, and for Performance Advisor.

You can also obtain online documentation and answers to frequently asked questions from the N series support website (accessed and navigated as described in Websites) at [www.ibm.com/storage/support/nseries/](http://www.ibm.com/storage/support/nseries/).

## Where to find third-party source material

Operations Manager, Provisioning Manager, and Protection Manager contain LGPL-licensed packages, which are available at <ftp://ftp.netapp.com/frmntap/opensource>.



---

## Websites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. The following web pages provide N series information:

- A listing of currently available N series products and features can be found at the following web page:  
[www.ibm.com/storage/nas](http://www.ibm.com/storage/nas)
- The IBM System Storage N series support website requires users to register in order to obtain access to N series support content on the web. To understand how the N series support web content is organized and navigated, and to access the N series support website, refer to the following publicly accessible web page:  
[www.ibm.com/storage/support/nseries/](http://www.ibm.com/storage/support/nseries/)  
This web page also provides links to AutoSupport information as well as other important N series product resources.
- IBM System Storage N series products attach to a variety of servers and operating systems. To determine the latest supported attachments, go to the IBM N series interoperability matrix at the following web page:  
[www.ibm.com/systems/storage/network/interophome.html](http://www.ibm.com/systems/storage/network/interophome.html)
- For the latest N series hardware product documentation, including planning, installation and setup, and hardware monitoring, service and diagnostics, see the IBM N series Information Center at the following web page:  
<http://publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp>



---

## Copyright and trademark information

This section includes copyright and trademark information, and important notices.

### Copyright information

Copyright ©1994 - 2013 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2013 IBM Corporation. All rights reserved.

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Network Appliance reserves the right to change any products described herein at any time, and without notice. Network Appliance assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Network Appliance. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of Network Appliance.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

---

### Trademark information

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

Network Appliance is a licensee of the CompactFlash and CF Logo trademarks.

Network Appliance NetCache is certified RealSystem compatible.

---

## Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, N.Y. 10504-1785  
U.S.A.

For additional information, visit the web at:  
<http://www.ibm.com/ibm/licensing/contact/>

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

**INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.** Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make

improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.





NA 210-06258\_A0, Printed in USA

GA32-1018-04

