



**EMC® VNX® Series**

**Release 8.1**

**Parameters Guide for VNX® for File**

**P/N 300-151-171 Rev 01**

**EMC Corporation**

*Corporate Headquarters:*  
Hopkinton, MA 01748-9103  
1-508-435-1000  
[www.EMC.com](http://www.EMC.com)

---

Copyright © 1998 - 2013 EMC Corporation. All rights reserved.

Published August 2013

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date regulatory document for your product line, go to the Technical Documentation and Advisories section on EMC Powerlink.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All other trademarks used herein are the property of their respective owners.

Corporate Headquarters: Hopkinton, MA 01748-9103

|   |           |
|---|-----------|
| <b>Preface</b> .....  | <b>5</b>  |
| <b>Chapter 1: Introduction</b> .....                                      | <b>7</b>  |
| System requirements.....  | 8         |
| Cautions.....   | 8         |
| User interface choices.....   | 9         |
| Related information.....  | 10        |
| <b>Chapter 2: Concepts</b> .....  | <b>11</b> |
| Server parameters.....  | 12        |
| System parameters.....  | 12        |
| <b>Chapter 3: View and Modify by Using the CLI</b> .....                  | <b>13</b> |
| View a list of facilities.....  | 15        |
| List server parameters for a facility.....                                | 17        |
| View details for a specific server parameter.....                         | 18        |
| View detailed server parameter information for one or all facilities..... | 19        |
| Modify a server parameter in the CLI.....                                 | 20        |
| Reboot a Data Mover in the CLI.....                                       | 21        |
| Restart services.....   | 22        |
| <b>Chapter 4: View and Modify by Using Unisphere</b> .....                | <b>23</b> |
| View server parameters.....   | 24        |
| Modify a server parameter value.....                                      | 24        |
| Reboot a Data Mover.....  | 24        |
| <b>Chapter 5: Edit Server or System Parameter Files</b> .....             | <b>25</b> |

|   |            |
|---|------------|
| Add or modify a server parameter for a single Data Mover..... | 27         |
| Add or modify a global server parameter.....                  | 28         |
| Add or modify a system parameter.....                         | 29         |
| <b>Chapter 6: Parameter Facilities.....</b>                   | <b>31</b>  |
| <b>Chapter 7: VNX Server Parameters.....</b>                  | <b>33</b>  |
| Parameter facility A to C.....                                | 34         |
| Parameter facility D to H.....                                | 72         |
| Parameter facility I to L.....                                | 78         |
| Parameter facility M to P.....                                | 85         |
| Parameter facility Q to S.....                                | 98         |
| Parameter facility T to V.....                                | 107        |
| <b>Chapter 8: VNX System Parameters.....</b>                  | <b>123</b> |

## Preface

*As part of an effort to improve and enhance the performance and capabilities of its product lines, EMC periodically releases revisions of its hardware and software. Therefore, some functions described in this document may not be supported by all versions of the software or hardware currently in use. For the most up-to-date information on product features, refer to your product release notes.*

*If a product does not function properly or does not function as described in this document, please contact your EMC representative.*

---

## Special notice conventions

EMC uses the following conventions for special notices:

---

Note: Emphasizes content that is of exceptional importance or interest but does not relate to personal injury or business/data loss.

---

 Identifies content that warns of potential business or data loss.

 Indicates a hazardous situation which, if not avoided, could result in minor or moderate injury.

 Indicates a hazardous situation which, if not avoided, could result in death or serious injury.

 Indicates a hazardous situation which, if not avoided, will result in death or serious injury.

---

## Where to get help

EMC support, product, and licensing information can be obtained as follows:

Product information—For documentation, release notes, software updates, or for information about EMC products, licensing, and service, go to EMC Online Support (registration required) at <http://Support.EMC.com>.

Troubleshooting—Go to EMC Online Support at <http://Support.EMC.com>. After logging in, locate the applicable Support by Product page.

Technical support—For technical support and service requests, go to EMC Customer Service on EMC Online Support at <http://Support.EMC.com>. After logging in, locate the applicable Support by Product page, and choose either **Live Chat** or **Create a service request**. To open a service request through EMC Online Support, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

---

Note: Do not request a specific support representative unless one has already been assigned to your particular system problem.

---

---

## Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications.

Please send your opinion of this document to:

[techpubcomments@EMC.com](mailto:techpubcomments@EMC.com)

The *Parameters Guide for VNX for File* enables you to view and modify parameter values, and to control the system settings according to the requirements of your site. You can control the attributes for the various server and system parameters.

This document is a part of the VNX documentation set and is intended for system administrators implementing changes to the parameter values.

Topics included are:

- ◆ [System requirements on page 8](#)
- ◆ [Cautions on page 8](#)
- ◆ [User interface choices on page 9](#)
- ◆ [Related information on page 10](#)

## System requirements

Table 1 on page 8 describes the EMC® VNX® software and hardware configurations.

**Table 1. System requirements**

|          |                          |
|----------|--------------------------|
| Software | VNX Software version 8.1 |
| Hardware | VNX                      |

## Cautions

Modifying server or system parameters can produce adverse effects. Before modifying a parameter, take the following precautions:

- ◆ Change parameter values in small increments.
- ◆ Unless the parameters are interdependent, change server parameters one at a time and thoroughly test each change before making additional changes.
- ◆ Facility and parameter names are case-sensitive. Type them exactly as shown in this guide. A typographical error, the omission of a dot, or an incorrect capitalization renders the entry invalid.
- ◆ An invalid value for a parameter renders the entry in the server\_log invalid. Both command line interface (CLI) and EMC Unisphere® ignore invalid server parameter values, and such parameters do not appear in either interface. If you make an error while editing a server parameter file, you will not be able to correct it through CLI or Unisphere.
- ◆ When you use CLI or Unisphere to set a server parameter, a process detects any anomalies in the parameter file and removes or adjusts them. This process cannot detect an out-of-range value or wrong case usage for a facility or parameter name.
- ◆ Some server parameter changes do not take effect until the Data Mover is rebooted or services are restarted. When editing a parameter file, you do not receive notification that such an action is required.
- ◆ Rebooting does not alter the parameter file; that is, the incorrect parameters persist.
- ◆ Keep a detailed log of each system or server parameter change and the reasons for making that change. This information can be useful if a problem arises.
- ◆ If you are uncertain about the effects of the change, consult EMC Customer Service before changing the parameter. EMC cannot be responsible for any adverse effects that result from changes made on your own.

## User interface choices

VNX offers flexibility in managing parameters based on your support environment and interface preferences.

You can view and modify server parameters (but not system parameters) by:

- ◆ Using the CLI `server_param` command
- ◆ Using EMC Unisphere

The *EMC VNX Command Line Interface Reference for File* provides detailed information about using CLI commands.

Unisphere online help provides information about using the EMC Unisphere software.

Alternatively, you can set or modify server or system parameters by editing the following files on the Control Station:

- ◆ `/nas/server/slot_x/param` (server parameters for each Data Mover; for example, `slot_2/param` contains the parameters for `server_2`)
- ◆ `/nas/site/slot_param` (global server parameters)
- ◆ `/nas/site/nas_param` (system parameters)

To manage server parameters, use either the `server_param` command or Unisphere.

The only method for adjusting system parameters is to edit the `file/nas/site/nas_param`, as described in [Add or modify a system parameter on page 29](#).

[Table 2 on page 9](#) summarizes the benefits of using the CLI or Unisphere rather than editing the server parameter files directly. Read the [Cautions on page 8](#) before modifying a server or system parameter.

**Table 2. Interface comparison**

| Benefit   | Recommended methods                          |           | Non-recommended methods        |
|---|--|-----------|--------------------------------|
|   | CLI command<br>( <code>server_param</code> ) | Unisphere | Editing server parameter files |
| Displays range of values  | Yes  | Yes       | No                             |
| Confirms action   | Yes  | Yes       | No                             |
| Validates syntax and value  | Yes  | Yes       | No                             |
| Overrides existing settings for individual Data Movers through use of global settings | Yes  | Yes       | No                             |

**Table 2. Interface comparison** *(continued)*

| Benefit               | Recommended methods           |           | Non-recommended methods        |
|-----------------------|-------------------------------|-----------|--------------------------------|
|                       | CLI command<br>(server_param) | Unisphere | Editing server parameter files |
| Reports command error | Displayed                     | Displayed | Entry in server_log            |

## Related information

For specific information related to the features and functionality described in this document:

- ◆ *EMC VNX Operating Environment for File Release Notes*
- ◆ *EMC VNX Command Line Interface Reference for File*
- ◆ *VNX for File man pages*

### EMC VNX documentation on EMC Online Support

The complete set of EMC VNX series customer publications is available on EMC Online Support. To search for technical documentation, go to <http://Support.EMC.com>. After logging in to the website, click **Support by Product** and type **VNX series** in the Find a Product text box. Then search for the specific feature required.

VNX software sets specific system attributes by default. You can modify the parameter settings that control these attributes according to the requirements of your site. Server parameters control attributes for a specific Data Mover, while system parameters control attributes for the entire server.

Topics included are:

- ◆ [Server parameters on page 12](#)
- ◆ [System parameters on page 12](#)

---

## Server parameters

Server parameters apply to Data Movers and can be set independently on each Data Mover. A global server parameter is a setting that applies to all Data Movers on the server.

Modifying server parameters can produce adverse effects. Before changing a server parameter, read [Cautions on page 8](#).

[Chapter 7](#) lists the server parameters that you can modify.

---

## System parameters

System parameters affect the entire system.

Modifying system parameters can significantly affect the function of the server. Before changing a system parameter, read [Cautions on page 8](#).

[Chapter 8](#) lists the system parameters that you can modify.

## View and Modify by Using the CLI

To view or modify server parameters in the CLI, use the `server_param` command.

The command response displays requested information, confirms a parameter value change, reports errors, and provides related information, such as additional steps necessary for the parameter change to take effect.

To view or modify system parameters, you need to edit the system parameter file as described in [Add or modify a system parameter on page 29](#).

The following examples show incorrect command line entries and their resulting error messages:

- ◆ The facility name (Portable Archive Interchange (PAX)) is typed incorrectly. Facility names are case-sensitive:

```
$ server_param server_2 -facility pax -modify paxStatBuff -value 64
```

```
server_2 :  
Error 2213: server_2 : pax : invalid name specified
```

- ◆ The parameter name (paxStatBuff) is misspelled:

```
$ server_param server_2 -facility PAX -modify paxStatBuf -value 64
```

```
server_2 :  
Error 2213: server_2 : paxStatBuf : invalid name specified
```

- ◆ An invalid value is given for the parameter:

```
$ server_param server_2 -facility PAX -modify paxStatBuff -value 100000
```

```
server_2 :  
Error 4418: server_2 : 100000 is not in range (1, 512)
```

Use the `server_param` command in CLI to:

- ◆ [View a list of facilities on page 15](#)
- ◆ [List server parameters for a facility on page 17](#)
- ◆ [View details for a specific server parameter on page 18](#)
- ◆ [View detailed server parameter information for one or all facilities on page 19](#)
- ◆ [Modify a server parameter in the CLI on page 20](#)
- ◆ [Reboot a Data Mover in the CLI on page 21](#)
- ◆ [Restart services on page 22](#)

## View a list of facilities

| Action   |
|--|
| <p>To view a list of facilities that have adjustable server parameters, use this command syntax:</p> <pre>\$ server_param &lt;movername&gt; -info</pre> <p>where:</p> <p>&lt;movername&gt; = name of the Data Mover</p> <p>Example:</p> <p>To view the list of facilities on server_2, type:</p> <pre>\$ server_param server_2 -info</pre>   |
| Output   |
| <pre>server_2 : facility      description FLRCompliance File Level Retention NDMP         Network Data Management Protocol PAX          Data Backup/Restore facility for NDMP/server_archive RCP          Replication Control Protocol arp          Address Resolution Protocol cfs          Common File System cifs         Common internet FS (Microsoft Network File Protocol) config       Misc. configuration parameters for CIFS cvfs         Celerra Virtual File System dedupe       DEDUPE: File level Data Deduplication dns          Domain Name Service ds           Domain Controller service for CIFS fcTach       Agilent Fibre Channel Controller file         Overall file system parameters filesystem   File system ftpd         File Transfer Protocol Daemon http         Hypertext Transfer Protocol ip           Internet protocol iscsi        Internet Scsi Protocol kernel       THREADS deadlock detection ldap         Lightweight Directory Access Protocol lockd        Network Lock Manager lockmgr      CFS lock manager mount        NFS Mount Protocol</pre> |

|          |                                       |
|----------|---------------------------------------|
| nbs      | Network Block Service Protocol        |
| nfs      | Network File System                   |
| nfsv4    | NFS version 4 protocol                |
| quota    | File system quota management          |
| security | Security/Credential parameters        |
| shadow   | Cross Protocol naming support         |
| ssl      | SSL security network protocol         |
| statd    | Host status demon                     |
| statmon  | Statistics Framework                  |
| streamio | Streaming tape I/O support            |
| tcp      | Transmission Control Protocol         |
| tftp     | Trivial File Transfer Protocol Daemon |
| trunk    | Network trunking support              |
| ufs      | Dart native file system               |
| usrmap   | User name mapping support             |
| vbb      | Volume Based Backup                   |
| vdevice  | Virtual IP Device Parameters          |
| viruschk | Virus checking service                |

## List server parameters for a facility

| Action  |
|---|
| <p>To list all server parameters associated with a facility, use this command syntax:</p> <pre>\$ <b>server_param</b> {&lt;movername&gt; ALL} <b>-facility</b> &lt;facility_name&gt; <b>-list</b></pre> <p>where:</p> <p>&lt;movername&gt; = name of the Data Mover (or keyword ALL for all Data Movers)</p> <p>&lt;facility_name&gt; = name of the facility (case-sensitive) to which the parameter belongs</p> <p>Example:</p> <p>To view a list of PAX parameters on server_2, type:</p> <pre>\$ <b>server_param</b> server_2 <b>-facility</b> PAX <b>-list</b></pre> <p>Note: The command output includes the default and current value for each parameter, and if applicable, the configured value that is not yet in effect.</p>        |
| Output  |
| <pre>server_2 : param_name          facility    default  current  configured filter.numDirFilter PAX        5        5 scanOnRestore       PAX        1        1 allowVLCRestoreToUFS PAX        0        0 filter.numFileFilter PAX        5        5 writeToArch         PAX        1        1 paxStatBuff         PAX        128     128 filter.dialect      PAX        ''       '' writeToTape         PAX        1        1 checkUtf8FileNames PAX        1        1 readWriteBlockSizeInKB PAX        64     64 axWriteBuff         PAX        64     64 noFileStreams       PAX        0        0 dump                PAX        0        0 nRestore            PAX        8        8 filter.caseSensitive PAX        1        1</pre> |
| Note  |
| <p>where:</p> <p><i>param_name</i> = name of the parameter</p> <p><i>facility</i> = facility to which the parameter belongs</p> <p><i>default</i> = default value for the parameter</p> <p><i>current</i> = value currently in effect on the Data Mover</p> <p><i>configured_value</i> = new value that is not yet in effect (requiring either a Data Mover reboot or services restart) or an invalid value</p>   |

## View details for a specific server parameter

| Action  |
|---|
| <p>To view information about a server parameter, including the default and current values, use this command syntax:</p> <pre>\$ <b>server_param</b> {&lt;movername&gt; ALL} <b>-facility</b> &lt;facility_name&gt; <b>-info</b> &lt;param_name&gt; [<b>-verbose</b>]</pre> <p>where:</p> <p>&lt;movername&gt; = name of the Data Mover (or keyword ALL for all Data Movers)</p> <p>&lt;facility_name&gt; = name of the facility (case-sensitive) to which the parameter belongs</p> <p>&lt;param_name&gt; = name of the parameter (case-sensitive)</p> <p>Example:</p> <p>For detailed information about the paxStatBuff parameter on server_2, type:</p> <pre>\$ <b>server_param server_2 -facility PAX -info paxStatBuff -verbose</b></pre> |
| Output  |
| <pre>server_2 :   name                = paxStatBuff   facility_name        = PAX   default_value        = 128   current_value        = 64   configured_value     = 64   user_action          = none   change_effective     = immediate   range                = (1,512)   description          = Number of buffers in stat queue per backup   session.  detailed_description Specifies the number of buffers between the threads that provide metadata (NASS) and read data from disk (NASA). More buffers can increase the speed with which metadata is provided.</pre>  |

**Note**

where:

*facility\_name* = facility to which the parameter belongs

*default\_value* = default value for the parameter

*current\_value* = value currently in effect on the Data Mover

*configured\_value* = new value that is not yet in effect (requiring either a Data Mover reboot or services restart) or an invalid value

*user\_action* = user action necessary for the configured value to take effect (either none, reboot Data Mover, or restart Service)

*change\_effective* = when the parameter change takes effect (either immediate, reboot Data Mover, or restart Service)

*range* = range of possible values; either (<minimum>, <maximum>) for numeric values or '\*' for a string

*description* = brief description of parameter

*detailed\_description* = full description (only if -verbose is included)

## View detailed server parameter information for one or all facilities

**Action**

To view information about all server parameters for a facility or for all facilities on a Data Mover, use this command syntax:

```
$ server_param <movername> -facility {<facility_name>|-all} -info -all
```

where:

<movername> = name of the Data Mover

<facility\_name> = name of the facility (or -all for all facilities)

Example:

To view information about all server parameters for all facilities on a Data Mover, type:

```
$ server_param server_2 -facility PAX -info -all
```

## Output

```

server_2 :
name = filter.numDirFilter
facility_name = PAX
default_value = 5
current_value = 5
configured_value =
user_action = none
change_effective = immediate
range = (0,50)
description = Specifies the number of directory filters that can be speci-
fied for a single backup.

name = scanOnRestore
facility_name = PAX
default_value = 1
current_value = 1
configured_value =
user_action = reboot DataMover
change_effective = reboot DataMover
range = (0,1)
description = Whether to scan virus on restore
.
.
(additional output omitted)
.
.
name = filter.caseSensitive
facility_name = PAX
default_value = 1
current_value = 1
configured_value =
user_action = none
change_effective = immediate
range = (0,1)
description = 1=Filter is case-sensitive, 0=Filter is not case-sensitive

```

## Modify a server parameter in the CLI

In the CLI, when you set a parameter value for a specific Data Mover, that value overrides any existing global value for the same parameter. Similarly, if you set a global value for a Data Mover parameter with an existing value, it applies to all Data Movers. The global value replaces existing parameter values, if any.

## Action

To modify a server parameter value, use this command syntax:

```

$ server_param {<movername>|ALL} -facility <facility_name> -modify
<param_name> -value <new_value>

```

where:

<movername> = name of the Data Mover (or keyword ALL for all Data Movers)

<facility\_name> = name of the facility (case-sensitive) to which the parameter belongs

| Action   |
|--|
| <p><code>&lt;param_name&gt;</code> = name of the parameter (case-sensitive)</p> <p><code>&lt;new_value&gt;</code> = new value for the parameter. Assign numeric values as either a decimal value or in hexadecimal format (0x followed by one to eight hexadecimal characters). To specify a null string, type double quotation marks ("").</p> <p>Example:</p> <p>To set the PAX facility <code>paxStatBuff</code> parameter to 128, type:</p> <pre>\$ server_param server_2 -facility PAX -modify paxStatBuff -value 128</pre> |
| Output   |
| <pre>server_2 : done</pre> <p>If you set a global parameter (by specifying ALL as the movername), command output appears for each Data Mover and indicates if the new value is replacing an existing setting for that parameter on the Data Mover.</p> <p>For example:</p> <pre>Warning 17716815732: server_2 : replacing server_2: paxStatBuff=64 with 128 server_3 : done server_4 : done server_5 : done</pre>  |

**Note:** If the output includes a message about rebooting the Data Mover or restarting a service, perform the task [Reboot a Data Mover in the CLI on page 21](#).

## Reboot a Data Mover in the CLI

If a server parameter change requires the Data Mover to be rebooted before the change takes effect, the command response includes the message:

```
Warning 17716815750: <movername> : You must reboot <movername> for
<param_name> changes to take effect.
```

For example:

```
Warning 17716815750: server_2 : You must reboot server_2 for bufksz changes
to take effect.
```

| Action   |
|--|
| <p>To reboot a Data Mover, use this command syntax:</p> <pre>\$ server_cpu {&lt;movername&gt; ALL} -reboot now</pre> <p>where:</p> <p><code>&lt;movername&gt;</code> = name of the Data Mover (or keyword ALL for all Data Movers)</p> <p>Example:</p> <p>To reboot <code>server_2</code>, type:</p> |

|   |
|---|
| Action  |
| <code>\$ server_cpu server_2 -reboot now</code> |
| Output  |
| <code>server_2 : done</code>                    |

## Restart services

If a server parameter change requires the associated service to be restarted before the change takes effect, the command response includes the message:

```
Warning 17716815753: <movername> : You must stop and start <facility_name>
for <param_name> changes to take effect.
```

For example:

```
Warning 17716815753: server_2 : You must stop and start cifs for acl.checkacl
changes to take effect.
```

|  |
|--|
| Action   |
| To restart services, use this command syntax:  |
| <code>\$ server_setup {&lt;movername&gt; ALL} -Protocol {cifs viruschk} -option stop</code>  |
| <code>\$ server_setup {&lt;movername&gt; ALL} -Protocol {cifs viruschk} -option start</code> |

## View and Modify by Using Unisphere

Unisphere provides procedures to manage the various server parameters.

---

Note: Unisphere does not provide access to system parameters.

---

Use Unisphere to:

- ◆ [View server parameters on page 24](#)
- ◆ [Modify a server parameter value on page 24](#)
- ◆ [Reboot a Data Mover on page 24](#)

## View server parameters

To view the server parameters by using Unisphere, select **Settings > Data Mover Parameters**.

Use the **Show Server Parameters for** drop-down list to view:

- ◆ A list of server parameters that apply to all Data Movers (global parameters) or the server parameters for a selected Data Mover.
- ◆ A list of parameters for a selected facility.

## Modify a server parameter value

1. From Unisphere, select **Settings > Data Mover Parameters**.
2. Select the parameter to display and click **Properties**.
3. In the **Value** field, type the new value.
4. Click **OK**.



Changing parameters can adversely affect other processes running on the system. Ensure that you understand the ramifications of the change before modifying a parameter.

## Reboot a Data Mover

1. From Unisphere, select **System > Hardware > Data Movers**.
2. Select the Data Mover, and click **Reboot**.
3. Click **OK** to send a reboot message to the selected Data Movers.

While the reboot is in progress, refreshing Data Movers shows the rebooting Data Mover in various states as they appear in the Status column.

If the operation takes more than 30 seconds, a message that includes a task ID appears at the top of the page.

To view the status of the operation, select **System > Monitoring and Alerts > BackgroundTasks for File**, and then double-click the task ID.

The *VNX System Operations* provides more information about rebooting Data Movers.

## Edit Server or System Parameter Files

Parameter settings are stored in the following files on the Control Station:

- ◆ `/nas/server/slot_x/param` (server parameters for each Data Mover; for example, `server_2` provides information about `slot_2`)
- ◆ `/nas/site/slot_param` (global system parameters that apply to all Data Movers)
- ◆ `/nas/site/nas_param` (system parameters)

These are text files that you can modify by using a text editor, such as vi.

The parameter files include only the parameters that have been modified from their default values. If the file does not contain an entry for a given parameter, the server or system uses the default value for that parameter.

---

Note: The `nas_param` file in `/nas/sys` contains information about system parameters, but this file is overwritten whenever the software is upgraded. To set system parameters, edit the `nas_param` file in `/nas/site`.

---

Each server and system parameter entry appears on a separate line in the following format:

- ◆ Server parameter entry:

```
param <facility> <parameter>=<paramvalue>
```

where:

<facility> = name of the facility (case-sensitive) to which the parameter belongs

<parameter> = name of the parameter (case-sensitive)

<paramvalue> = value to which the parameter is set

For example:

```
param cifs gpo=1
```

- ◆ System parameter entry:

`<parameter>:<paramvalue>:[<paramvalue>:]`

where:

`<parameter>` = name of the parameter (case-sensitive)

`<paramvalue>` = value to which the parameter is set

For example:

`ckpt:10:100:20:`

To edit server or system parameter files, use:

- ◆ [Add or modify a server parameter for a single Data Mover on page 27](#)
- ◆ [Add or modify a global server parameter on page 28](#)
- ◆ [Add or modify a system parameter on page 29](#)

## Add or modify a server parameter for a single Data Mover

Steps to add or modify a server parameter for a single Data Mover are:

1. Log in to the Control Station.
2. Go to the directory that contains the server parameter file for the Data Mover by using this command syntax:

```
$ cd /nas/server/slot_<x>
```

where:

<x> = slot number for the Data Mover

Example:

To go to the directory that contains the parameter file for server\_2, type:

```
$ cd /nas/server/slot_2
```

3. Open the param file by using a text editor, such as vi.
4. Add a parameter by appending the following line to the file:

```
param <facility> <parameter>=<value>
```

where:

<facility> = name of the facility (case-sensitive) to which the parameter applies

<parameter> = name of the server parameter (case-sensitive) to set

<value> = value for the parameter

Or to modify a parameter, locate the entry for the parameter and replace the current value with the new value.

Example:

To set the cifs facility gpo parameter to 1, add or modify the following line in the param file:

```
param cifs gpo=1
```

5. Save your changes and close the param file.
6. Confirm the parameter entry in the file, by typing:

```
$ more param
```

The contents of the param file appear.

7. If the parameter change does not take effect until the Data Mover is rebooted, use the procedure [Reboot a Data Mover on page 24](#).

## Add or modify a global server parameter

The following procedure applies to all Data Movers, unless the server parameter file (slot\_x) for a specific Data Mover contains another value for the same parameter. In other words, manually typing a global parameter does not override existing, nondefault settings for that parameter on individual Data Movers.

Steps to add or modify a global server parameter are:

1. Log in to the Control Station.
2. Go to the directory that contains the global server parameters file by typing:

```
$ cd /nas/site
```

3. Open the slot\_param file by using a text editor, such as vi.
4. Append the following line to the file:

```
param <facility> <parameter>=<value>
```

where:

<facility> = name of the facility (case-sensitive) to which the parameter applies

<parameter> = name of the server parameter (case-sensitive) to set globally

<value> = value for the parameter

Or to modify a parameter, locate the entry for the parameter and replace the current value with the new value.

Example:

To set the cifs facility gpo parameter to 1, add or modify the following line in the slot\_param file:

```
param cifs gpo=1
```

5. Save your changes and close the slot\_param file.
6. Confirm the parameter entry by typing:

```
$ more slot_param
```

The contents of the slot\_param file appear.

7. If the parameter change does not take effect until the Data Movers are rebooted, use the procedure [Reboot a Data Mover on page 24](#), and specify the keyword ALL to reboot all of the Data Movers.

## Add or modify a system parameter

**CAUTION** Do not modify any system parameter that is not listed in [Chapter 8](#). Do not enter a blank line in the `nas_param` file.

Steps to add or modify a system parameter are:

1. Log in to the Control Station.
2. Go to the directory that contains the system parameters file by typing:

```
$ cd /nas/site
```

3. Open the `nas_param` file by using a text editor, such as `vi`.
4. Append the following line to the file:

```
<parameter>:<value>:
```

where:

`<parameter>` = name of the system parameter (case-sensitive)

`<value>` = value of the parameter. For some parameters, the value can be a comma-separated list. If no value is specified, the default value of the parameter is used.

Or to modify a system parameter, locate the entry for the parameter and replace the current value with the new value.

Example:

To set the replication system parameter to 15, add or modify the following line in the `nas_param` file:

```
replication:15:
```

5. Save your changes and close the `nas_param` file.
6. Confirm the parameter entry by displaying the contents of the `nas_param` file by typing:

```
$ more nas_param
```

---

Note: You do not need to reboot the Control Station after modifying a system parameter.

---



[Table 3 on page 31](#) lists the facilities associated with the parameters in this guide.

**Table 3. Parameter facilities**

| Facility   | Definition  |
|------------|---|
| ana        | Adaptec network adapter                                       |
| arp        | Address Resolution Protocol                                   |
| cfs        | Common file system  |
| cifs       | Common Internet File System (Microsoft Network File Protocol) |
| config     | Miscellaneous configuration parameters for CIFS               |
| vufs       | VNX Virtual File System                                       |
| dns        | Domain Name Service   |
| ds         | Domain controller service for CIFS                            |
| fcTach     | Agilent Fibre Channel controller                              |
| file       | Overall file system parameters                                |
| filesystem | File system   |
| ftpd       | File Transfer Protocol daemon                                 |
| http       | Hypertext Transfer Protocol                                   |
| ip         | Internet Protocol   |
| iscsi      | Internet SCSI protocol  |
| ldap       | Lightweight Directory Access Protocol                         |
| lockd      | Network lock manager  |
| mount      | NFS mount protocol  |
| nbs        | Network Block Service protocol                                |

**Table 3. Parameter facilities** *(continued)*

| Facility    | Definition  |
|-------------|---|
| NDMP        | Network Data Management Protocol                        |
| nfs         | Network File System                                     |
| nfsv4       | NFS version 4 protocol                                  |
| PAX         | Data backup/restore facility for NDMP/server_archive    |
| quota       | File system quota management                            |
| RCP         | Replication Control Protocol                            |
| replication | Data replication facility                               |
| shadow      | Cross-protocol naming support                           |
| ssl         | SSL security network protocol                           |
| statd       | Host status daemon                                      |
| streamio    | Streaming tape I/O support                              |
| tcp         | Transmission Control Protocol                           |
| tftp        | Trivial File Transfer Protocol daemon                   |
| trunk       | Network trunking support                                |
| ufs         | VNX operating environment (for file) native file system |
| usrmap      | Username mapping support                                |
| vbb         | Volume Based Backup                                     |
| vdevice     | Virtual IP Device Parameters                            |
| viruschk    | Virus checking service                                  |
| VRPL        | Volume replication control                              |

## VNX Server Parameters

This section provides the following list of adjustable server parameters:

- ◆ [Parameter facility A to C on page 34](#)
- ◆ [Parameter facility D to H on page 72](#)
- ◆ [Parameter facility I to L on page 78](#)
- ◆ [Parameter facility M to P on page 85](#)
- ◆ [Parameter facility Q to S on page 98](#)
- ◆ [Parameter facility T to V on page 107](#)

## Parameter facility A to C

Table 4 on page 34 describes the parameters for facilities from A to C.

**Table 4. Parameter facility A to C**

| Facility | Parameter              | Values  | Comments/description  |
|----------|------------------------|---|---|
| ana      | rxburst                | 1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024<br>Default: 256 | <p>Sets the maximum number of received packets to be processed.</p> <p>If this limit is reached and more received packets require processing, the system creates another driver invocation to handle the packets.</p> <p>Increasing this value raises the number of packets processed per interruption, which affects overall system resource utilization.</p> <p>Lowering this value increases the number of interruptions required to process a given number of packets.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/> |
| arp      | holddown               | 600–6000<br>Default: 600                                    | <p>Specifies the hold-down time (in seconds) of an ARP entry since it was last referenced.</p> <p>600 = 10 minutes.</p>   |
| cfs      | commonthrdPoolMinCount | Default: 512  | Specifies the default minimum number of threads.  |
| cfs      | commonthrdPoolMaxCount | Default: 1024   | Specifies the default maximum number of threads.  |
| cfs      | deleteDelay            | 0 or 1<br>Default: 1  | <p>Controls whether CIFS file deletion requests are handled immediately or asynchronously in the background.</p> <p>0 = Handle CIFS deletion requests immediately.</p> <p>1 = CIFS deletion requests receive an immediate response, but the actual deletion occurs asynchronously.</p>  |
| cfs      | readwritesharing       | Default: 0  | Specifies the ability to use several VNX Operating Environments for File to access the same File System. This feature has been deprecated. This parameter by default, disables the feature.   |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter       | Values                      | Comments/description  |
|----------|-----------------|-----------------------------|---|
| cfs      | showChildFsRoot | 0 or 1<br>Default: 0        | <p>Controls the VNX Virtual File System (VVFS) version 1 NFS client access to checkpoints in the root directory of the production file system.</p> <p>0 = Checkpoint subdirectories do not appear in the root directory.</p> <p>1 = Each mounted checkpoint of a production file system is visible to NFS clients as a subdirectory of the file system root directory.</p> <p>The .ckpt directory is hidden by design and cannot be made visible.</p> |
| cfs      | showHiddenCkpt  | 0 or 1<br>Default: 1        | <p>Controls the VVFS version 2 client access to checkpoints.</p> <p>0 = Disable VVFS v2 and prevent accessing virtual checkpoint subdirectories.</p> <p>1 = Enable VVFS v2 so each mounted checkpoint of a production file system is accessible to clients as a special virtual subdirectory in the production file system's user directory.</p> <p>The .ckpt directory is hidden by design and cannot be made visible.</p>                           |
| cifs     | acl.archive     | 0, 1, 2, or 3<br>Default: 3 | <p>Controls whether a change to the security descriptor (SD) of a file or directory sets the DOS archive bit.</p> <p>0 = Set the DOS archive bit when the SD of a file or directory changes.</p> <p>1 = Set the DOS archive bit when the SD of a file changes.</p> <p>2 = Set the DOS archive bit when the SD of a directory changes.</p> <p>3 = Do not set the DOS archive bit when the SD of a file or directory changes.</p>                       |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter                 | Values               | Comments/description   |
|----------|---------------------------|----------------------|--|
| cifs     | acl.checkacl              | 0 or 1<br>Default: 1 | <p>Controls whether Microsoft Windows access control lists (ACLs) are checked during file system object access checking.</p> <p>0 = Do not verify ACLs.</p> <p>1 = Verify ACLs during access checking or else access checking is done only with UNIX rights.</p> <p>The parameter is always 0 if CIFS is not started (except if it is manually set). When CIFS is started in NT security, this parameter is set to 1; when CIFS is stopped, it is set to 0.</p> <p>The default value depends on the state of the CIFS service and security mode on the Data Mover.</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/> |
| cifs     | acl.checkAclConsistency   | 0 or 1<br>Default: 1 | <p>Controls whether ACLs are checked for consistency. This parameter is intended to troubleshoot the inconsistent ACL.</p> <p>0 = Do not verify ACLs for consistency.</p> <p>1 = Verify the consistency of ACLs during ACL creation and on ACL access. A panic is generated when this detects an inconsistent ACL.</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/>   |
| cifs     | acl.checkUnixXForCifsOpen | 0 or 1<br>Default: 0 | <p>Controls whether the X (execute) bit (on UNIX mode bits) is verified for CIFS access in UNIX or SECURE access policy.</p> <p>0 = Do not verify the X bit for CIFS access.</p> <p>1 = Verify the X bit for CIFS access.</p> <p>A file created by CIFS protocol uses the umask bits to create the mode bit on UNIX without knowing whether the file would be created for exec (usually the exec bit is set by CIFS client later in the process).</p>  |

Table 4. Parameter facility A to C (continued)

| Facility | Parameter  | Values              | Comments/description  |
|----------|------------|---------------------|---|
| cifs     | acl.extacl | 0–127<br>Default: 0 | <p>This parameter is a bit list that enables special capabilities for ACL management.</p> <p>Two kinds of capabilities are:</p> <ul style="list-style-type: none"> <li>◆ Backup or restore specific UNIX attributes like access rights, UNIX mode, and UNIX name and symbolic link by using a regular CIFS-based backup tool (like NTbackup).</li> <li>◆ View or change UNIX access rights from an ACL management tool, such as Windows Explorer. The bit list consists of seven binary bits (bits 0 through 6, right to left). Any combination of bits is allowed. Each bit is 1 when set, otherwise 0.</li> </ul> |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter | Values | Comments/description  |
|----------|-----------|--------|---|
|          |           |        | <ul style="list-style-type: none"> <li>◆ Bit 0 set (0000001 or +1) — VNX presents the UNIX metadata associated with files and directories to CIFS backup clients by using a special ACE (access control entry) in the file or directory's ACL. This ACE can take either of two forms. If bit 0 is not set, VNX uses an ACE type (CIFS allows vendors to define their own ACE types). If bit 0 is set, VNX uses a standard ACE and encodes the information in the SID associated with that ACE.</li> <li>◆ Bit 1 set (0000010 or +2) — If bit 1 is set, then a Windows client can view and modify UNIX permissions on files and directories on VNX. The UNIX permissions are presented as three additional ACEs in the ACL of each file and directory. You can view and modify these ACEs by using any CIFS ACL management application, such as Windows Explorer.</li> <li>◆ Bit 2 set (0000100 or +4) — If bit 2 is set, VNX presents the UNIX permissions associated with files and directories in the ACL of the files so that CIFS network backup applications can back up and restore them to and from a VNX file system.</li> <li>◆ Bit 3 set (0001000 or +8) — If bit 3 is set, VNX presents UNIX symbolic links as zero byte files with a special ACL that captures the information associated with the symbolic link (for example, its target). If bit 3 is not set, VNX may follow symbolic links on behalf of CIFS clients and, hence, a CIFS backup application does not back up the symbolic links, but instead, the files they point to. If this bit is set, VNX follows symbolic links on behalf of CIFS network backup clients. This means that the CIFS backup application backs up the symbolic links, and not the files and directories they point to.</li> </ul> |

Table 4. Parameter facility A to C (continued)

| Facility | Parameter | Values | Comments/description  |
|----------|-----------|--------|---|
|          |           |        | <ul style="list-style-type: none"> <li>◆ Bit 4 set (0010000 or +16) — Any file or directory on a VNX can have as many as three names in the file system: a UNIX name, a long Windows or M256 name, and a DOS 8.3 name. If bit 4 is set, VNX encodes the UNIX name of files and directories in a special ACE in the ACL of the files so that CIFS network backup applications can back up and restore all three names of files and directories.</li> <li>◆ Bit 5 set (01000000 or +32) — By default, there is no way for NFS v2 and v3 clients to view or modify the ACLs associated with files and directories on VNX. The tool emcsetsd allows NFS v2 and v3 clients to view and, if the user has permission to do so, modify the ACLs associated with files and directories on VNX. Bit 5 must be set for the emcsetsd client tool to work.</li> <li>◆ Bit 6 set (1000000 or +64) — If set, bit 6 modifies the functions enabled with bit 1. If bit 6 is not set, UNIX rights applied to the file are the granted rights plus the rights not denied by the discretionary ACL (DACL).</li> </ul> <p>If bit 6 is set, UNIX rights applied are the granted rights less the denied rights by the DACL. In addition, the request is rejected if one of the three special ACEs is inheritable. This is because when changing rights on a directory, the client propagates rights down the tree to all nodes (files and directories), which is typically not a desired behavior. Setting this bit prevents that. In practice, this means that ACLs for directories must be set by using the Advanced panel in the security properties within Windows Explorer.</p> |

**Table 4. Parameter facility A to C (continued)**

| Facility | Parameter          | Values               | Comments/description   |
|----------|--------------------|----------------------|--|
|          |                    |                      | <p>Examples of bit string settings:</p> <p>0000010 — (bit 1 only = 2) Allows CIFS clients to view and modify the UNIX permissions on files and directories by using Windows Explorer. All other settings use the default values.</p> <p>1000010 — (bit 1 + bit 6 = 66) Changes the way that the ACL is translated into UNIX permissions on files and directories. As a result, UNIX permissions applied to files and directories are the rights granted by any grant ACE for the UNIX user/group/other less any rights explicitly denied in any deny ACE for the UNIX user/group/other.</p> <p>1100010 — (bit 1 + bit 6 + bit 5 = 98) Enables NFS v2 and v3 clients to view and modify the ACLs on files and directories by using the emcsetsd tool.</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/> |
| cifs     | acl.extendExtraGid | 0 or 1<br>Default: 0 | <p>Controls whether CIFS user credentials and NFS NT credentials include information about the UNIX groups to which the user belongs. This applies in a multiprotocol environment only.</p> <p>0 = Base the user credential only on the Windows Security IDs (SIDs) for the user and the Windows groups to which the user belongs.</p> <p>1 = Include in the user credential, the SIDs that correspond to the UNIX groups to which the user belongs.</p> <p><i>Managing a Multiprotocol Environment on VNX</i> provides more information.</p>  |

Table 4. Parameter facility A to C (continued)

| Facility | Parameter                | Values                   | Comments/description   |
|----------|--------------------------|--------------------------|--|
| cifs     | acl.FailOnSDRestoreError | 0, 1, or 2<br>Default: 1 | <p>Determines what to do when a mapping error occurs during a security descriptor restore or a cifs update &lt;path&gt; setacl command. During a restore process from a backup tool or regular migration tool such as emcopy, the file must be opened with the backup bit option that requires at least backup operator rights.</p> <p>0 = No error is returned, wrong ACEs are discarded, and a wrong owner or primary group of the node is replaced by the current user ID (user performing the restore).</p> <p>1 = Error is returned, wrong ACEs are discarded, and a wrong owner or primary group of the node is replaced by one of the connected users.</p> <p>2 = Error is returned and the security descriptor is rejected, but the work has been done in the way described.</p>                           |
| cifs     | acl.mappingErrorAction   | 0 - 15 Default: 0        | <p>Defines the rules for unknown mapping between security, user, and group identifiers (SID/UID/GID) on ACL settings.</p> <p>Two kinds of errors might occur:</p> <ul style="list-style-type: none"> <li>◆ The SID set in the ACL is unknown to the domain controllers being used.</li> <li>◆ The username is not yet mapped to a UID/GID.</li> </ul> <p>The bit list consists of four binary bits (bits 0 through 3, right to left). Each bit is 1 when set; otherwise 0.</p> <p>Bit 0 (0001 or +1): Store unknown SID.</p> <p>Bit 1 (0010 or +2): Store SID with no UNIX mapping.</p> <p>Bit 2 (0100 or +4): Enable debug traces.</p> <p>Bit 3 (1000 or +8): Do lookup only in cache (secmap or global SID cache or per connection SID cache).</p> <p>Refer to the cifs facility acl.retryAuthSid parameter.</p> |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter                   | Values                      | Comments/description   |
|----------|-----------------------------|-----------------------------|--|
| cifs     | acl.mngt                    | Default: 1                  | Enables ACL management from a CIFS client. By default, Celerra supports NTFS behavior, which includes ACL management from a CIFS client. It is possible by disabling this parameter, to disable ACL management on Celerra. This could be useful if you want to manage access right modification only from UNIX.  |
| cifs     | acl.ntCheckModebits         | Default: 1                  | Valid mode bits checking in NT access right policy. Sets to valid the checking of mode bits in the NT access right policy. If the parameter value is 0, only ACLs are taken into account in access right checking for UNIX access.   |
| cifs     | acl.otherFromEveryoneOnly   | Default: 0                  | In MIXED access policy, generate other mode bits from EVERYONE only. In MIXED access policy, when ACL is modified, modebits are generated from the ACL. If the parameter value is 0, all ACEs corresponding to a user or group different from the file owner and its group, are used to generate other parts of the mode bits. If the parameter value is 1, only the EVERYONE ACE is used to generate other bits.  |
| cifs     | acl.restrictedTakeOwnership | 0 or 1<br>Default: 0        | Controls ownership rights in SECURE or UNIX access policy.<br><br>1 = Privilege. Take ownership. Ownership rights are taken into account only for the account mapped on UID 0. Only UID 0 can take ownership (except for backup or restore purposes).  |
| cifs     | acl.retryAuthSid            | 1-0xffffffff<br>Default: 10 | Sets the delay (in seconds) between attempts to resolve a given SID. This avoids overloading the Data Mover and networking by sending too many requests.<br><br>This occurs if a SID is stored in an ACL that has not been mapped to a UID/GID (see cifs facility acl.mappingErrorAction parameter).<br><br>On ACL access, if an unmapped SID is detected, the VNX server tries to resolve the mapping. If it encounters an error, it tries again only after this delay. |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter              | Values               | Comments/description   |
|----------|------------------------|----------------------|--|
| cifs     | acl.sdcompactthreshold | Default: 16300       | When trying to set an SD which is bigger than the specified size in bytes, an attempt is made to compact the ACEs in the SACL and DACL without altering the semantic. The format of an SD limits its size to 64KB. However, in case of memory fragmentation, the VNX Operating Environment for File may not be able to allocate a big chunk of memory for big SDs. That is the reason why, when a new SD is set with a size bigger than the one specified by this parameter, the VNX Operating Environment for File tries to find and remove any duplicate ACEs in the DACL and SACL, but without altering the semantic of the DACL or SACL. |
| cifs     | acl.sortAces           | 0 or 1<br>Default: 0 | Sorts the access control entries (ACEs) in an ACL, in the order expected by a Windows client.<br><br>0 = Store ACEs in the order received.<br><br>1 = Store the ACEs of an ACL set by CIFS or NFSv4 in the following order:<br><br>1. First, DENY an inherited ACE.<br>2. Then, ALLOW an inherited ACE.<br>3. Then, DENY a non-inherited ACE.<br>4. Then, ALLOW a non-inherited ACE.   |
| cifs     | acl.takegroupship      | 0 or 1<br>Default: 0 | Sets the primary group based on the cifs facility acl.useUnixGid parameter during a Windows take ownership request, from the CIFS protocol, without primary group information in the request. The new primary group is set by using CIFS or UNIX user primary group, depending on the value of the acl.useUnixGid parameter.<br><br>0 = Disable this setting.<br><br>1 = Enable this setting.  |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter           | Values               | Comments/description   |
|----------|---------------------|----------------------|--|
| cifs     | acl.umask           | 0 or 1<br>Default: 0 | <p>Synchronizes the UNIX execute bit with the corresponding execute right from the CIFS DACL, at the time of file or directory creation only.</p> <p>0 = If this value is set, the UNIX execute bit will not be computed from the DACL at the time of creation.</p> <p>1 = If this value is set, the UNIX execute bit will be computed from the DACL at the time of creation.</p> <p>_____</p> <p>Note: This parameter is not supported in the 6.0.x NAS code family.</p> <p>_____</p> |
| cifs     | acl.unixCheckAcl    | 0 or 1<br>Default: 1 | <p>Specifies the type of check for the UNIX access rights policy.</p> <p>0 = Verify only UNIX mode bits.</p> <p>1 = Verify ACL and UNIX mode bits.</p>   |
| cifs     | acl.useUnixGid      | 0 or 1<br>Default: 0 | <p>Sets the group identifier (GID) mapping for file system objects created on a Windows client.</p> <p>0 = Assign the GID of the Windows primary domain group to which the user belongs.</p> <p>1 = Assign the Windows user's GID (as found in the GID field of the .etc/passwd file, NIS database entry, or Active Directory).</p>  |
| cifs     | admin.adminsAreRoot | 0 or 1<br>Default: 1 | <p>Controls whether Windows NT administrators are granted the same privileges as a UNIX root account on the Data Mover.</p> <p>0 = Do not grant Windows NT administrators the same privileges as UNIX root.</p> <p>1 = Grant Windows NT administrators the same privileges as UNIX root.</p>   |

Table 4. Parameter facility A to C (continued)

| Facility | Parameter              | Values                       | Comments/description   |
|----------|------------------------|------------------------------|--|
| cifs     | admin.shareC_RO        | 0 or 1<br>Default: 0         | From C\$ share access, you are able to write on file systems mounted on /, but not on the root file system.<br><br>0 = Set directory / as read-only, but allow users to write on subdirectories.<br><br>1 = Set the C\$ share as read-only. A security descriptor is created on the share to prevent writes, and the umask is set to 0222.<br><br>_____<br>Note: Restart the CIFS service for the changes to take effect.<br>_____ |
| cifs     | allowSnapSureVss       | 0 or 1<br>Default: 1         | Controls the Microsoft Shadow Copies of Share Folders (SCSF) feature, which can be used to access previous versions of a file or directory stored by using a checkpoint directly from a CIFS client.<br><br>0 = Disable SCSF checkpoint access.<br>1 = Enable SCSF checkpoint access.  |
| cifs     | audit.auditThreadCount | Default: 8                   | Defines the number of threads that write Windows event logs. Protocol threads generate Windows event logs and pass them to dedicated audit threads which write the event logs to disk. This parameter defines the number of audit threads.   |
| cifs     | audit.defaultCacheSize | 8 - 8192<br>Default: 64      | Specifies maximum number of entries in the event log cache used for seek requests. A seek request for a specific event can be very slow because the event log file is sequential. The cache stores (eventId, offset) tuples so that seek requests can start from the closest event stored in it.   |
| cifs     | audit.eventsCountFlush | 1-0xffffffff<br>Default: 100 | Specifies the number of Windows event logs cached before being flushed on a disk.<br><br>This parameter affects memory. Losing data is a risk.   |
| cifs     | audit.eventsFormat     | 0 or 1<br>Default: 0         | Sets the format for entries in the security event log to provide compatibility with third-party applications.<br><br>0 = Set all bits to 0. The default format is used.<br><br>1 = Enclose the value of the Client Logon ID field in event ID 560(0x00000001) in parentheses, and replace the dash with a comma.   |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter             | Values                          | Comments/description  |
|----------|-----------------------|---------------------------------|---|
| cifs     | audit.maxQueuedEvents | Default: 512                    | Defines the maximum number of pending event records per Windows event log file. To enhance performance, Windows event records are written asynchronously on disk.   |
| cifs     | capabilities          | 0-0xffffffff<br>Default: 0xa3f9 | <p>Allows maximum buffer size and maximum I/O size on the client to be negotiated with the server.</p> <p>Capabilities returned when negotiating with the client are:</p> <p>0x1 = READ/WRITE raw<br/>                     0x2 = READ/WRITE Mpx<br/>                     0x4 = UNICODE<br/>                     0x8 = Large files<br/>                     0x10 = NT LM0.12 special<br/>                     0x20 = Remote API<br/>                     0x40 = NT Status<br/>                     0x80 = Level II OpLock<br/>                     0x100 = Lock and Read<br/>                     0x200 = NT Find Cmd<br/>                     0x1000 = DFS support (not supported)<br/>                     0x2000 = Windows 2000/XP (mainly Quota support)<br/>                     0x4000 = Large_ReadX<br/>                     0x8000 = Large_WriteX<br/>                     0x80000000 = SPNEGO (Kerberos authentication)</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/> |
| cifs     | cifsclient.log        | Default: 0                      | Flags to enable or disable some of the cifs client logs. Those flags allow you to enable or disable some of the CIFS client Logs. Current supported flags are: Bit 0: log connect failures.   |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter             | Values                          | Comments/description   |
|----------|-----------------------|---------------------------------|--|
| cifs     | cifsclient.timeout    | 0–0xffffffff<br>Default: 20000  | Specifies the CIFS client timeout (in milliseconds) used in the VNX Data Migration Service (VDMS). This is the time allowed for a response to a CIFS client request from one Data Mover to another CIFS server (other Data Mover, regular Microsoft server, or other CIFS server).<br><br>Example:<br>20000 = 20 seconds   |
| cifs     | comment_filtered      | 0 or 1<br>Default: 0            | On Windows XP and Windows Server 2003 clients, the comment associated with the server is stored in the Registry the first time the client connects to a share.<br><br>0 = VNX stores the release number in the comment. Some clients can become confused if the comment does not change after an upgrade to the VNX code.<br><br>1 = The default comment removes the release number from the default server release.<br><br>_____<br>Note: Restart the CIFS service for the changes to take effect.<br>_____ |
| cifs     | DC.autoRemovalSeconds | 0–0xffffffff<br>Default: 604800 | Sets the time (in seconds) to remove from the DC list any obsolete domain controllers that are not announced by WINS/DNS.<br><br>0 = Do not remove obsolete domain controllers.  |
| cifs     | DC.useFastest         | 0 or 1<br>Default: 1            | Controls which domain controller to use.<br><br>0 = Use the first DC in the same subnet.<br><br>This is true only if DNS is not used to retrieve the DC list.<br><br>1 = Use the DC that answers the fastest during the discover phase for authentication.<br><br>_____<br>Note: Restart the CIFS service for the changes to take effect.<br>_____   |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter      | Values               | Comments/description  |
|----------|----------------|----------------------|---|
| cifs     | djAddAdminToLg | 0 or 1<br>Default: 0 | <p>When joining a CIFS server to a domain, the user performing the join can be added to the Administrators Local Group of the CIFS server.</p> <p>0 = Do not automatically add the user to the Administrators Local Group.</p> <p>1 = Automatically add the user to the Administrators Local Group.</p>   |
| cifs     | djEnforceDhn   | 0 or 1<br>Default: 1 | <p>When joining a CIFS server to a domain, the join process sets the account attribute dNSHostName to the server's DNS hostname.</p> <p>0 = Allow the domain join process to continue without setting the account attribute dNSHostName.</p> <p>1 = Set the account attribute dNSHostName to the server's DNS hostname. If the domain join process is unable to set the attribute, the operation fails and a notification is sent.</p> <p>In certain domain configurations, the domain join process is unable to set the account attribute because of access rights.</p> <p>Use this parameter only in a controlled way to fix a temporary access right problem. If the dNSHostName account attribute is not set, Windows clients authenticate to the Data Mover in NTLMSSP mode instead of Kerberos.</p> |

Table 4. Parameter facility A to C (continued)

| Facility | Parameter           | Values                                  | Comments/description  |
|----------|---------------------|---|---|
| cifs     | djUseKpasswd        | 0 or 1<br>Default: 0                    | <p>When joining a CIFS server to a domain, set the CIFS server password by using the Kerberos set/change password protocol.</p> <p>When performing the domain join, if you are a member of only a domain local group (and not of a domain global group), your Kerberos ticket does not carry the SID for that domain local group, which precludes you from setting the CIFS server password. In this case, set the CIFS server password through a special procedure that uses the MS RPC protocol.</p> <p>0 = The domain join process sets the CIFS server password by using the SAMR protocol. djUseKpasswd is disabled.</p> <p>1 = The parameter djUseKpasswd is enabled and the former kPassword method is used.</p> |
| cifs     | enableFileFiltering | 0x0, 0x1, 0x3, 0x5, 0x7<br>Default: 0x7 | <p>Controls file filtering, pop-up messages, and audits when an extension is denied.</p> <p>The bit list consists of three binary bits (bits 0 through 2, right to left). Each bit is 1 when set, otherwise 0.</p> <p>0x0 (000 — No bits set) = Disable file filtering.</p> <p>0x1 (001 — Bit 0) = Enable file filtering.</p> <p>0x3 (011 — Bits 0 and 1) = Enable file filtering and generate pop-up messages.</p> <p>0x5 (101 — Bits 0 and 2) = Enable file filtering and auditing.</p> <p>0x7 (111 — Bits 0, 1, and 2) = Enable file filtering, pop-up messages, and auditing.</p>   |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter           | Values               | Comments/description   |
|----------|---------------------|----------------------|--|
| cifs     | extSec.capabilities | Default: 2147542013  | <p>Sets CIFS server capabilities for Kerberos authentication (on port 445)</p> <p>This parameter contains capabilities returned by the Data Mover during the negotiation protocol issued by the client when it connects on port 445(Kerberos authentication). These capabilities should not be changed, but in a very particular case, some Celerra capabilities could be hidden from the client for a specific application requirement. See also the cifs.capability parameter for authentication through port 139.</p> <p>The values are defined by a bit list:</p> <p>0x1 = READ/WRITE Raw<br/>                     0x2 = READ/WRITE Mpx<br/>                     0x4 = UNICODE<br/>                     0x8 = Large Files<br/>                     0x10 = NT LM0.12 special<br/>                     0x20 = Remote API<br/>                     0x40 = NT Status<br/>                     0x80 = Level II Oplock<br/>                     0x100 = Lock and Read<br/>                     0x200 = NT Find CMD<br/>                     0x1000 = DFS support<br/>                     0x2000 = Windows 2000/XP (mainly Quota support)<br/>                     0x4000 = Large_ReadX<br/>                     0x8000 = Larage_WriteX<br/>                     0x80000000 = SPNEGO (Kerberos authentication)</p> |
| cifs     | gpo                 | 0 or 1<br>Default: 1 | <p>Controls group policy object (GPO) support.</p> <p>0 = Disable GPO support.<br/>                     1 = Enable GPO support.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/>   |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter                         | Values   | Comments/description  |
|----------|-----------------------------------|--|---|
| cifs     | gpo.fileSecurityThreadsNb         | Default: 4                                       | Number of threads for the GPO file security update. This parameter is the number of threads executing the file system security update.  |
| cifs     | gpo.fileSecurityUpdateInterval    | Default: 17 hours                                | Defines the time interval (in hours) between two file system GPO security applications on files and folders in VNX operating environment (for file).  |
| cifs     | gpocache                          | 0 or 1<br>Default: 1                             | Controls GPO caching support.<br>0 = Disable GPO caching support.<br>1 = Enable GPO caching support. On CIFS start, users can set the GPO even if the DC is not yet available.  |
| cifs     | LanmanServer.disable-NameChecking | 0 or 1<br>Default: 0                             | Controls whether to disable checking the Kerberos ticket of the client for the principal name of the server. Microsoft Knowledge Base Article ID 281308 provides detailed information.<br>0 = Do not disable checking. The client must use the primary computer name in order to connect.<br>1 = Disable checking and allow the client to connect with a DNS alias.<br><br>Note: Reboot the Data Mover for changes to take effect.                          |
| cifs     | LanmanServer.IdleUserAutoLogoff   | 0–0xffffffff<br>Default: 4294967295 (0xffffffff) | Sets the number of minutes after which an idle user with no open files is automatically logged out of the server. Setting this parameter causes VNX to free resources associated with user sessions that are orphaned by the client.<br><br>The default behavior is for the user to be logged out when explicitly requested by the client or when the TCP connection is reset.<br><br>15 = After 15 minutes of idle time, log out the user from the server. |
| cifs     | LanmanServer.MaxMpxCount          | 5–2048<br>Default: 127                           | Sets the maximum number of CIFS client commands allowed without acknowledgment in the Data Mover (for example, Notify request). This value is used by the client machine to limit the number of commands. The value is returned in the negotiate command.   |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter               | Values                   | Comments/description   |
|----------|-------------------------|--------------------------|--|
| cifs     | cLanmanServer.SessUsers | 1–16383<br>Default: 2048 | <p>Sets the maximum number of sessions that a single client (that is, an open TCP connection) is able to open on the Data Mover. This parameter has been set because a Windows 2000 client can issue many sessions if the share where they were connected is unavailable (for example, the file system is unmounted).</p> <p>These values are not checked by the software.</p>   |
| cifs     | listBlocks              | 1–255<br>Default: 255    | <p>Specifies the number of 256-entry blocks used for File ID/Tree ID/User ID (FID/TID/UID) number creation.</p> <p>64 = You can use 16,000 IDs simultaneously for each type of ID. If a heavily loaded site requires more, an error message in the log file recommends increasing this number.</p> <p>Context management uses more memory than data caching. Therefore, changing these values can degrade performance.</p> <p>The actual number of IDs is limited by the available memory in the system.</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/> |
| cifs     | lookup.secmappingOnly   | 0 or 1<br>Default: 1     | <p>Specifies the resolvers used to translate UID to SID when the SID mapping file is generated. This parameter is used for the quota report command.</p> <p>0 = Use NIS, usermapper, and domain controller to map users who are not in secmap. The time to execute the command is much greater, especially with many users, but all the SIDs are returned.</p> <p>1 = Use secmap only. Only the SIDs of Windows users who have logged in to the server at least once are returned.</p>   |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter         | Values                 | Comments/description  |
|----------|-------------------|------------------------|---|
| cifs     | lookup.traceLevel | 1–11<br>Default: 3     | <p>Specifies the level of Data Mover logs needed. This can be used for debugging purposes.</p> <p>1 = Take action immediately.<br/>                     2 = Log critical conditions.<br/>                     3 = Log error messages only.<br/>                     4 = Log warning conditions.<br/>                     5 = Log normal but significant conditions.<br/>                     6 = Log informational data.<br/>                     7 = Log debugging information.<br/>                     8 = Only used by CS.<br/>                     9 = Debug Level 1<br/>                     10 = Debug Level 2<br/>                     11 = Enable all available Data Mover logs.</p> |
| cifs     | lsarpc.maxDomains | 32–2048<br>Default: 32 | <p>Sets the maximum number of domains accessible by the Data Mover. If the number of different domains (contained within the same ACL) exceeds 32, you can increase the value of this parameter.</p> <p>Memory is the only theoretical limitation. This parameter refers to the number of domains visible to the Data Mover, not the number of domains to which the Data Mover can belong.</p> <p>If the value is too low, you cannot set the ACL with as many domains.</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/>   |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter             | Values                     | Comments/description  |
|----------|-----------------------|----------------------------|---|
| cifs     | majorVersion          | 0–0xffffffff<br>Default: 0 | <p>Explicitly sets the major server version returned to a client when it queries for the version of the server. To operate properly, some client applications require a specific version number to be returned.</p> <p>0 = The server returns the version number based on other server settings.</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/>  |
| cifs     | maxForceDCInfoRefresh | 0–20<br>Default: 10        | <p>When no domain controller is available, the list of available domain controllers is refreshed. This parameter defines the number of times information is requested from WINS, or is broadcast for domain controllers.</p>  |
| cifs     | maxLockXPending       | 0–0xffffffff<br>Default: 0 | <p>Allows changing the maximum number of pending lock requests. When an incoming lockX request conflicts with a specified timeout, the Data Mover retries the pending lock for the timeout. To avoid performance impact against other requests, this parameter limits the number of pending requests.</p> <p>0 = The system allows a maximum of requests equal to half the number of CIFS threads.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/>   |
| cifs     | maxMpxCount           | 5–2048<br>Default: 127     | <p>Sets the maximum number of commands allowed without acknowledgment in the Data Mover (for example, notify requests). The client uses this value to limit the number of commands. The value is returned in the negotiate command. Each user must disconnect in order to be taken into account.</p> <p>You must stop and then restart the CIFS service so that all users take the new value.</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/> |

Table 4. Parameter facility A to C (continued)

| Facility | Parameter    | Values                     | Comments/description  |
|----------|--------------|----------------------------|---|
| cifs     | maxVCThreads | 1-0xffffffff<br>Default: 3 | <p>Sets the maximum number of threads held in reserve to break deadlocks between file access requests and virus checking requests for the files. The maximum is one less than the number of CIFS threads set for the Data Mover.</p> <p>Setting reserved thread for VC does not mean the VC is unable to use another thread. This is only the number of threads used exclusively by the VC.</p> <p>The reserved threads are taken from the total available threads. Be careful not to use too high a number of threads for VC compared to the total CIFS thread. This could adversely affect performance.</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/> |
| cifs     | minFreeFS    | 0-100<br>Default: 10       | <p>Specifies the minimum amount of available free space (as a percentage). If you are connecting to the share and the space limit is reached, you receive a pop-up message.</p> <p>Prior to this, set the bit 1 value to allow pop-up messages to be sent. The mask is 0x2. Windows 9x requires special software to enable pop-up messaging.</p>  |
| cifs     | minorVersion | 0-0xffffffff<br>Default: 0 | <p>Explicitly sets the minor server version, returned to a client, when it queries for the version of the server.</p> <p>Some client applications require a specific version number to be returned in order to operate properly.</p> <p>0 = The server returns the version number based on other server settings.</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/>   |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter              | Values                      | Comments/description   |
|----------|------------------------|-----------------------------|--|
| cifs     | nanoroundoff           | 0 or 1<br>Default: 0        | <p>CIFS protocol can set time by using a 64-bit value (time in 100th of nanosecond since January 1, 1601 in the Gregorian calendar). The system, however, does not store timestamps to this granularity.</p> <p>0 = VNX rounds timestamps down to the nearest second. Thus, the timestamp of a file can be older (by up to 1 second) than the client expects.</p> <p>1 = VNX rounds timestamps up to the nearest second. This helps in applications that compare timestamps to keep files synchronized between CIFS clients and VNX. A typical application is when Windows user roaming profiles are stored on VNX. With this parameter set, the user logout time can be significantly reduced because clients do not attempt to recopy files that have not changed.</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/> |
| cifs     | NTBufSz                | 512–65535<br>Default: 65535 | <p>Sets the maximum buffer size negotiated with Windows NT 4.0 clients. The optimum value for large data transfer is 32832 (0x8040).</p> <p>Windows NT is negotiated 4356 (0x1104).</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/>  |
| cifs     | NTsec.DCConnectTimeout | 100–60000<br>Default: 5000  | <p>Sets the maximum time (in milliseconds) that the Data Mover waits for a domain controller response while trying to connect to a domain controller after connection.</p> <p>If the domain controller does not respond within this period, the Data Mover considers it unavailable and searches for another domain controller.</p>  |

Table 4. Parameter facility A to C (continued)

| Facility | Parameter                 | Values  | Comments/description  |
|----------|---------------------------|---|---|
| cifs     | NTsec.DCTimeout           | 5000–60000<br>Default: 20000  | <p>Sets the maximum time (in milliseconds) the Data Mover waits while trying to contact a domain controller after connection. This value is greater than the NTsec.DCConnectTimeout value because the domain controller could send requests to other domain controllers, for example, in a Trust relation case.</p> <p>If the domain controller does not connect within this period, the Data Mover considers it unavailable and searches for another domain controller.</p>  |
| cifs     | NTsec.getDCfromADServices | 0 or 1<br>Default: 1  | <p>This parameter applies only for interfaces connected to a Windows 2000 forest (Kerberos).</p> <p>0 = Use WINS to build the list sorted by using the response time (SamLogon) and subnet info.</p> <p>1 = Use DNS exclusively to get the list of domain controllers. This list is sorted to get the best DC candidate, by using services available on this DC, site information, response time (using LDAP ping), and subnet info. The DCs that are not in DNS, but detected by broadcast, are not included in this list.</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/> |
| cifs     | nullSession               | 0 or 1<br>Default: 1<br>Prior to Celerra Network Server v4.2, nullSession was not supported, so, the default value was 0. | <p>Allows client to use nullSession (no username, no password) to connect to a share. By default, you will be able to login, but not read/write on the file system.</p> <p>The mapped UID for such a user is –2.</p> <p>0 = Prevent clients from connecting to a share when using nullSession.</p> <p>1 = Allow clients to connect to a share without username and password.</p> <p>Refer to cifs.nullSessionNotOnFS parameter.</p>   |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter            | Values                   | Comments/description   |
|----------|----------------------|--------------------------|--|
| cifs     | nullSessionNotOnFS   | 0 or 1<br>Default: 1     | <p>0 = Enables user read/write access on the file system parameter.</p> <p>1 = Disables user read/write access on the file system, if their UID is -2.</p> <p>When you use cifs.nullSession and turn off this parameter, the quota will not work correctly.</p> <p>This parameter works along with cifs.nullSession = 1 and SysAccessNetworkLogon privilege, which must be granted to everyone on the server.</p> <p>Use authenticated user ACE's in the share's ACL.</p>                              |
| cifs     | popupToConnectedUser | 0, 1, or 3<br>Default: 0 | <p>Controls where Windows pop-ups are sent.</p> <p>0 = Send Windows pop-ups to the workstation.</p> <p>1 = Send Windows pop-ups to the connected user.</p> <p>3 = Send Windows pop-ups to the connected user. If the workstation refuses the pop-up because the user is invalid, an error appears in the server_log.</p> <p>Limitations: The username connected to VNX must be the same as the user on the workstation.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/> |

Table 4. Parameter facility A to C (continued)

| Facility | Parameter         | Values                     | Comments/description  |
|----------|-------------------|----------------------------|---|
| cifs     | prealloc          | 0–0xffffffff<br>Default: 4 | <p>When a CIFS client wants to write data, it asks the file server if there is enough space to write 1 byte of data (regardless of how much it wants to write). The file server then responds with an amount of data the client can write before it must ask again (the prealloc size).</p> <p>The larger the prealloc size the file server returns to the client, the less often the client will ask the question. Fewer questions mean better performance.</p> <p>Allocation size is computed from a virtual block size allocation and real file size.</p> <p>0 = 8 KB block size<br/>1 = 16 KB block size<br/>2 = 32 KB block size<br/>3 = 64 KB block size<br/>4 = 128 KB block size (default)<br/>5 = 256 KB block size<br/>6 = 512 KB block size</p> <p>Preallocation size is reserved temporarily from total disk space. This could temporarily reduce the available disk space if many files are opened simultaneously, and induce unexpected file system full or quota exceeded errors.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/> |
| cifs     | quotas.queryNames | 0 or 1<br>Default: 0       | <p>Specifies how to find a valid SID for Windows 2000 quotas enumerations.</p> <p>0 = Use security mapping and usermappers.<br/>1 = Use NIS or local password and group files to query user/group names.</p>  |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter       | Values                   | Comments/description  |
|----------|-----------------|--------------------------|---|
| cifs     | ReadOnly.Comp   | 0, 1, or 2<br>Default: 0 | <p>Controls support of the CIFS read-only bit behavior with other file server vendors.</p> <p>0 = No correlation is made between UNIX permissions and the DOS read-only attribute on a file or directory.</p> <p>1 = Enable compatibility with Samba:</p> <ul style="list-style-type: none"> <li>◆ If the UNIX owner write bit is off, the DOS read-only bit is considered on.</li> <li>◆ Only the owner can manipulate the DOS read-only bit.</li> <li>◆ If the DOS read-only bit is set from CIFS, all UNIX write bits are reset.</li> <li>◆ If the DOS read-only bit is reset from CIFS, the write bit is set by an exclusive OR with the umask on all UNIX write bits.</li> </ul> <p>2 = Enable compatibility with Network Appliance:</p> <ul style="list-style-type: none"> <li>◆ If any UNIX write bit is on, the DOS read-only bit is considered off.</li> <li>◆ If no UNIX write bit is on, the DOS read-only bit is considered on.</li> <li>◆ If the DOS read-only bit is set from CIFS, all UNIX write bits are reset.</li> <li>◆ If the DOS read-only bit is reset from CIFS, the UNIX owner write bit is set.</li> </ul> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/> |
| cifs     | ReadOnly.Delete | 0 or 1<br>Default: 0     | <p>This parameter applies only if the cifs facility ReadOnly.Comp parameter has a nonzero value.</p> <p>0 = Prevent NFS clients from deleting files and directories on which the DOS read-only attribute is set (except for directories in Network Appliance mode).</p> <p>1 = Ignore the DOS read-only attribute and allow the deletion of files and directories.</p>  |

Table 4. Parameter facility A to C (continued)

| Facility | Parameter     | Values                      | Comments/description  |
|----------|---------------|-----------------------------|---|
| cifs     | resolver      | 0 or 1<br>Default: 0        | <p>Controls whether a username or group name is checked for an appended domain name (username.domain or groupname.domain) when retrieving entries from local password or group files, NIS, or Archive Directory with SFU/IdMu.</p> <p>0 = Use the domain extension when retrieving a UID or GID.</p> <p>1 = Retrieve UIDs or GIDs without the appended domain extension. If retrieval fails, look for the name with the appended domain extension.</p>  |
| cifs     | secmap.enable | 0 or 1<br>Default: 1        | <p>Enables or disables CIFS secure mapping (Secmap) cache. Secmap stores all the UID/GID mappings that correspond to SIDs to ensure that they are consistent.</p> <p>Each time a new mapping is used in the system, Secmap adds it to the database. All mappings are then retrieved directly from the Secmap cache without having to query external name resolution mechanisms (such as DC, NIS, and usermapper).</p> <p>Secmap is enabled only at CIFS start and disabled only at CIFS stop.</p> <p>Disable Secmap only if absolutely necessary.</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/> |
| cifs     | sendMessage   | 0, 1, 2, or 3<br>Default: 1 | <p>Controls the type of messages sent when quota limits are reached.</p> <p>0 = Disable all pop-up messages.</p> <p>1 = Enable only error pop-up messages.</p> <p>2 = Enable only warning pop-up messages.</p> <p>3 = Enable both warning and error pop-up messages.</p>  |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter           | Values                                   | Comments/description   |
|----------|---------------------|--|--|
| cifs     | set_eas_ok          | 0 or 1<br>Default: 0                     | <p>CIFS allows applications to apply extended attributes to files. Some applications use this mechanism to store proprietary information. By default, VNX rejects any request to set the extended attributes of a file. Some applications, however, rely on this mechanism.</p> <p>0 = Prevent VNX from accepting requests to set the extended attributes on files.</p> <p>1 = Enable VNX to accept requests to set extended attributes on files. Note that this does not mean that VNX stores the extended attributes; it simply means that it does not reject a request to set them. Backup and migration applications can then restore or migrate files with extended attributes to VNX. The file data and standard attributes of the files are restored or migrated to VNX, but not the extended attributes.</p> |
| cifs     | share.default.umask | 0-511 (0777, 0x1FF)<br>Default: 18 / 022 | <p>Determines the default umask value for CIFS share when a umask is not specified. This value is the UNIX umask value that is used to determine UNIX mode bits that follow standard UNIX rules when creating a file or directory from the CIFS side.</p> <p>The default value is 18 (octal 22).</p> <p>Umask with 022 means rw-r--r-- for file creation and rwx-r-xr-x for directories.</p> <p>Alternatively, you can set a specific umask value for a specific share by using the <code>server_export</code> command. If umask is defined in the share command declaration, it overrides the default value.</p>  |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter                   | Values                  | Comments/description  |
|----------|-----------------------------|-------------------------|---|
| cifs     | sidcache.globalSidCacheSize | 0–32768<br>Default: 811 | <p>Specifies the maximum number of entries stored in the global SID cache for groups. You can use this parameter to increase connection performance when there are many users.</p> <p>0 = No cache</p> <p>Maximum value is limited by the memory. Reasonable value is no more than 16K.</p> <p>For cache efficiency, use a prime number.</p> <hr/> <p>Hint: To obtain the procedure for generating a list of prime numbers in Microsoft Excel, go to the Microsoft Support website.</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/> |
| cifs     | sidcache.size               | 0–2500<br>Default: 53   | <p>Specifies the maximum number of entries in the SID mapping cache local to each connection used for ACL mapping.</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/>  |
| cifs     | simulateNTFS                | 0 or 1<br>Default: 1    | <p>Sets the file system type identifier returned to a CIFS client.</p> <p>0 = UxFS (native VMX file system)</p> <p>1 = NTFS (Windows NT file system)</p>  |

**Table 4. Parameter facility A to C** *(continued)*

| Facility | Parameter   | Values               | Comments/description   |
|----------|-------------|----------------------|--|
| cifs     | smb signing | 0 or 1<br>Default: 1 | <p>Controls Server Message Block (SMB) signing globally on the Data Mover.</p> <p>0 = Disable SMB signing. The Data Mover overrides any SMB signing GPO that is set for the domain. SMB signing must also be disabled on Windows Server 2003 clients.</p> <p>1 = Enable SMB signing. The SMB signing relies on the GPO, if defined. When the GPO is not defined, it relies on the CIFS server Registry that is present on the Data Mover. GPOs override the CIFS server Registry settings.</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/> |

Table 4. Parameter facility A to C (continued)

| Facility | Parameter         | Values              | Comments/description  |
|----------|-------------------|---------------------|---|
| cifs     | smb2.capabilities | 0-7F<br>Default: 7F | <p>This parameter is a mask of bits that specify the SMB2 capabilities, which are enabled for the CIFS server. It allows you to tune the SMB2 capabilities that VNX supports. The format of this parameter complies with the SMB2 negotiation message:</p> <hr/> <p>Note: The Windows Registry manages the Distributed File System (DFS) option. The value of the DFS bit that is specified in this parameter does not affect the behavior of the data mover.</p> <hr/> <ul style="list-style-type: none"> <li>◆ 0x00000001 = Indicates that the server supports the DFS.</li> <li>◆ 0x00000002 = Indicates that the server supports leasing. This flag is only valid for the SMB 2.1 dialect.</li> <li>◆ 0x00000003 = Indicates that the server supports both DFS and LEASING capabilities.</li> <li>◆ 0x00000004 = Indicates that the server supports large MTU. This flag is only valid for the SMB 2.1 dialect.</li> <li>◆ 0x00000008 = Indicates that the server supports Multi Path IO (MPIO). This flag is only valid for the SMB 3.</li> <li>◆ 0x00000010 = Indicates that the server supports persistent handles. This flag is only valid for the SMB 3.</li> <li>◆ 0x00000020 = Indicates that the server supports directory leasing. This flag is only valid for the SMB 3.</li> <li>◆ 0x00000040 = Indicates that the server supports encryption. This flag is only valid for the SMB 3.</li> </ul> |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter                 | Values                                | Comments/description  |
|----------|---------------------------|---------------------------------------|---|
| cifs     | smb2.durableHdlAliveDelay | Default: 960                          | <p>Determines the delay in seconds, for which the CIFS server keeps the durable open, after the SMB2 client is disconnected.</p> <p>0 = Disable the durable handle feature.</p> <hr/> <p>Note: Change is effective immediately.</p> <hr/>   |
| cifs     | smb2.listBlocks           | Default: 511                          | <p>Number of blocks used for SMB2 FID number creation. The number of 256 entry blocks used for File ID (FID) number creation. This allows you to be able to use more than 16K IDs at the same time for each type of ID. If on a very loaded site, you are required to use more, an error message in the log file will invite you to increase this number. Minimum value = 1 Maximum value = 1023 (256K). Context management uses more memory than data caching. Therefore, changing these values could negatively impact performance. The number of FIDs is limited by the memory space used in the system.</p> |
| cifs     | smb2.maxCaTimeout         | 0-180 seconds<br>Default: 120 seconds | <p>Specifies maximum amount of time, which the CIFS server keeps opens (directories or files with associated locks) in server valid after a failover with clients. This parameter specifies the timeout in seconds. (0 disables the durability feature)</p>   |
| cifs     | smb2.maxResiliencyTimeout | Default: 300                          | <p>Maximum resiliency timeout in seconds that the CIFS server allows when an SMB2 client requests a file to be resilient after disconnection. Allows changing the maximum resiliency timeout that the CIFS server allows when a SMB2 client requests a file to be resilient. This parameter specifies the timeout in seconds. (0 disables the resiliency feature).</p>  |
| cifs     | spncheck                  | 0 or 1<br>Default: 0                  | <p>Disables Service Principal Name (SPN) verification and the server's Fully Qualified Domain Name (FQDN) cache for incoming CIFS connections from the client to the Data Mover.</p> <p>You can disable this functionality if you do not use Kerberos, or have several clients that are not joined to a domain, or for security reasons.</p> <p>By default, the connection verification and FQDN cache are enabled.</p>   |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter                 | Values               | Comments/description  |
|----------|---------------------------|----------------------|---|
| cifs     | srvmgr.diskdrive          | 0 or 1<br>Default: 0 | Provides disk drive information to show the available disk space on a CIFS server to third-party applications.<br><br>0 = VNX returns one drive (drive C:) mapped on the root file system.<br><br>1 = VNX returns a disk drive list based on the share name definition. If a share name <A through Z>\$ exists, then a list of drive equivalents to the share name (without the \$) exist. Number of drives is limited to 26.<br><br>All drives with share names ending in a \$ assume the "Hidden Directory" property. |
| cifs     | srvmgr.globalShares       | 0 or 1<br>Default: 0 | Determines share creation policy for Windows tools.<br><br>0 = Create NetBIOS shares. Attach all shares, created by using Windows tools, to the CIFS server used to create them. You can access these shares from this server only.<br><br>1 = Create global shares. Set all the shares created from server manager or MMC as global shares. These shares are accessible from all the virtual servers defined in the Data Mover (NetBIOS name or computer names).   |
| cifs     | srvpwd.encryptAccountFile | 0 or 1<br>Default: 1 | Controls whether the account file is encrypted.<br><br>0 = Do not encrypt the account file.<br><br>1 = Encrypt the account file.  |
| cifs     | srvpwd.maxHistory         | 1-10<br>Default: 2   | Specifies the number of passwords to retain for Kerberos authentication after a CIFS server password changes.<br><br>1 = Retain only the current password.<br><br>2-10 = Retain the most recent passwords including the current password.<br><br>Example:<br><br>3 = Retain the current password and the two previous passwords.  |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter          | Values                                       | Comments/description  |
|----------|--------------------|--|---|
| cifs     | srvpwd.updtMinutes | 0–0xffffffff<br>Default: 0                   | <p>Defines the minimum time interval (in minutes) between CIFS server password changes.</p> <p>0 = Disable server password changes.</p> <p>Examples:</p> <p>720 = Allow a server password change after 720 minutes (12 hours).</p> <p>1440 = Allow a server password change after 1440 minutes (24 hours).</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/>   |
| cifs     | tcpkeepalive       | 0xff000000–0xffffffff<br>Default: 0xff01030a | <p>Sets the TCP timeout. Use the hexadecimal format: 0xff &lt;mm&gt;&lt;pp&gt;&lt;ss&gt;</p> <p>where:</p> <p>&lt;mm&gt; = first timeout in minutes.</p> <p>&lt;pp&gt; = number of probes after mm minutes.</p> <p>&lt;ss&gt; = number of seconds between probes after mm minutes.</p> <p>Example:</p> <p>0xff01030a = 1 minute timeout, 3 probes, and 10 seconds between probes.</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/> |

Table 4. Parameter facility A to C (continued)

| Facility | Parameter     | Values                      | Comments/description  |
|----------|---------------|-----------------------------|---|
| cifs     | useADMap      | 0 or 1<br>Default: 0        | <p>Enables or disables CIFS management on the Data Mover.</p> <p>0 = Disable use of Active Directory usermapper.</p> <p>1 = Enable use of Active Directory usermapper. The mapping resolution is done in this order:</p> <ul style="list-style-type: none"> <li>◆ Local files (/etc/passwd or /etc/group)</li> <li>◆ Network Information Service (NIS)</li> <li>◆ Active Directory</li> <li>◆ Usermapper</li> </ul> <p>Do not set this parameter to 1 unless you have extended the Active Directory schema.</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/> |
| cifs     | vnoderpercent | 0–0xffffffff<br>Default: 80 | <p>Limits the number of files and directories that CIFS can open to a percentage of the maximum number of vnodes in the system.</p> <p>Example:</p> <p>80 = CIFS can open up to 80% of the total number of vnodes.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/>   |
| cifs     | W2KBufSz      | 512–65535<br>Default: 65535 | <p>Sets the maximum buffer size (in bytes) negotiated with Windows 2000 clients.</p> <p>The optimum value for large data transfer is 61504. Windows 2000 clients are negotiated with 16644.</p> <p>Because of limitations on the requirement for space for header information, the minimum is 512 bytes. This is similar to the Microsoft recommendation.</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/>   |

**Table 4. Parameter facility A to C** (continued)

| Facility      | Parameter               | Values                              | Comments/description   |
|---------------|-------------------------|-------------------------------------|--|
| cifs          | W95BufSz                | 512–65535<br>Default: 65535         | Sets the maximum buffer size (in bytes) negotiated with Windows 95 clients.<br><br>The optimum value for large data transfer is 57404 (0xE03C).<br><br>Note: Restart the CIFS service for the changes to take effect.  |
| cifs          | wins.UseDCs             | 0 or 1<br>Default: 1                | Controls whether the Data Mover queries the domain controller list from the WINS server.<br><br>0 = Do not query the domain controller list. Use broadcast to access the domain controller or NET LOGON.<br><br>1 = Query the DC list.<br><br>Note: Restart the CIFS service for the changes to take effect.   |
| ckpt_schedule | auto_cvfsname_prefix    | Yes or No<br>Default: No            | Controls whether relative VVFS name is enabled for schedules created through the Unisphere software.   |
| ckpt_schedule | cvfsname_delimiter      | "." or "_"<br>Default: "."          | Specifies the delimiter for the VVFS name.   |
| ckpt_schedule | cvfsname_starting_index | 0 or 1<br>Default: 0                | Specifies the starting index for the VVFS name when mounting a checkpoint on a checkpoint schedule.<br><br>Once the checkpoint is mounted, the latest snapshot cvfsname index assumes the default index value.<br><br>When the checkpoint is temporarily unmounted, the information exists in the Control Station database. The inactive checkpoint index value is preserved till it becomes active again. |
| config        | cs_external_ip          | Text string<br>Default: null string | Specifies the external IP address for the Control Station.   |
| config        | cs_external_ipv6        | Default: *                          | This is the CS External IPV6 address. It is used for XML API, and is not meant for users to set.   |

**Table 4. Parameter facility A to C** (continued)

| Facility | Parameter          | Values                       | Comments/description  |
|----------|--------------------|------------------------------|---|
| config   | userMapper.version | 1 or 3<br>Default: 3         | <p>Refers to the RPC protocol version to be used between the usermapper and the Data Mover.</p> <p>1 = Use name.domain for mapping purposes.</p> <p>3 = Use SID (of a group or user) to map. The user.domain (or group.domain) associated with the SID also has the same mapping UID or GID as that assigned to its SID.</p> <p>2 = Invalid value.</p> <hr/> <p>Note: Restart the CIFS service for the changes to take effect.</p> <hr/>                      |
| cvfs     | virtualDirName     | Text string<br>Default: ckpt | <p>VVFS version 2 allows users to traverse mounted checkpoint from a hidden virtual directory. This parameter defines a user-specified virtual directory name.</p> <p>The actual directory name is the specified string preceded by a dot.</p> <p>ckpt = Use .ckpt for the virtual directory name.</p> <p>Example:<br/>snapshot = Use .snapshot as the virtual directory name.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/> |

## Parameter facility D to H

Table 5 on page 72 describes parameters for facilities from D to H.

**Table 5. Parameter facility D to H**

| Facility | Parameter        | Value                        | Comments/description  |
|----------|------------------|------------------------------|---|
| dns      | bindUpdInterface | Default: 0                   | Enforces the source IP address of the dynamic update message. Allows you to enforce the dynamic update message to be sent with the source IP address of the interface being updated in the DNS server. By default, the dynamic update message is sent with the outgoing interface IP address based on the routing network table of the data mover. When set to 1, the source address of the dynamic update message is the IP address of the interface to update. Setting this parameter to 1 may cause dynamic update failures when the DNS server to update cannot ping the data mover interface. This might be needed if the IP used for the default route does not have the permission to perform the dynamic updates. |
| dns      | updateMode       | 0, 1, or 2<br>Default: 2     | By default, the Data Mover issues secure Dynamic DNS updates to the DNS server for the DNS domain it joins.<br><br>0 = Do not issue updates.<br>1 = Issue nonsecure updates.<br>2 = Issue secure updates.   |
| dns      | updatePTRrecord  | 0 or 1<br>Default: 0         | Controls whether the DNS client of the Data Mover updates the PTR record for all CIFS servers.<br><br>0 = The DNS client does not update the PTR record.<br>1 = Allow the DNS client to update the PTR record.  |
| ds       | RefreshDC        | 0–0xffffffff<br>Default: 900 | Sets the number of seconds between DC list updates from the DNS.  |
| ds       | RetryRefreshDC   | 0–0xffffffff<br>Default: 40  | Specifies the timeout (in seconds) between retries when a problem occurs in getting DC lists from DNS.  |

Table 5. Parameter facility D to H (continued)

| Facility | Parameter              | Value                | Comments/description  |
|----------|------------------------|----------------------|---|
| ds       | useADSite              | 0 or 1<br>Default: 1 | Controls what DC information is gathered.<br><br>0 = All the DCs of the domain including those outside the Data Mover's site.<br><br>1 = Only the DC of the Data Mover's Active Directory site.<br><br>_____<br>Note: Restart the CIFS service for the changes to take effect.<br>_____                   |
| ds       | useDCLdapPing          | 0 or 1<br>Default: 1 | Specifies how the service list is ordered.<br><br>0 = Order by services only.<br><br>1 = LDAP requests use site and response time information to order the list.<br><br>_____<br>Note: Restart the CIFS service for the changes to take effect.<br>_____  |
| ds       | useDSFile              | 0 or 1<br>Default: 0 | Specifies how LDAP, Kerberos, and Kerberos password servers are discovered.<br><br>0 = Use the standard discovery mechanisms.<br><br>1 = Use the servers listed in the directory services file on the Data Mover.<br><br>_____<br>Note: Restart the CIFS service for the changes to take effect.<br>_____ |
| et       | hostThreads.perSession | Default: 20          | ETrace IP to Hostname Resolution Threads. This parameter represents the number of parallel threads that can be used for IP to Hostname resolution per session.  |
| et       | journalSize            | Default: 1024        | ETrace Aggregator journal entries per pass. This parameter represents the number of journal entries to process per pass.  |

**Table 5. Parameter facility D to H** (continued)

| Facility | Parameter                              | Value  | Comments/description   |
|----------|--|--|--|
| et       | maxContainerPages                      | (4, 393216)<br><br>Default: Depending on the memory of the blade (Total main memory) | <p>Sets an upper limit on the number of pages an ETRACE container may hold.</p> <p>This is a memory size dependent parameter.</p> <p>The parameter's default and maximum values are different depending on the size in the memory in the blade (Total main memory)</p>   |
| fcTach   | linx_speed_aux0<br><br>linx_speed_aux1 | 0x1000,0x2000,<br>0x4000, 0x8000<br><br>Default: 0x8000                              | <p>Sets the tape backup link speed for the AUX0 or AUX1 Fibre Channel port—this applies to the VNX NS series only.</p> <p>0x1000 = Tape backup link speed set to 4 Gb/s.</p> <p>0x2000 = Tape backup link speed accommodates a 2 Gb Fibre Channel-Arbitrated Loop interface.</p> <p>0x4000 = Tape backup link speed accommodates an LC optical 1 Gb interface.</p> <p>0x8000 = Controller detects the link speed.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/> |
| fcTach   | linx_speed_be0                         | Default: 0x00008000  | <p>Sets the link speed for the Fibre Channel backend port, BE0. This parameter sets the backend port BE0 to a link speed that corresponds to one of the following hexadecimal values:</p> <p>&lt;link_speed_mask&gt;s: 0x8000 = auto-negotiate the link speed</p> <p>0x4000 = 1Gbps</p> <p>0x2000 = 2Gbps</p> <p>0x1000 = 4Gbps</p> <p>0x0800 = 8Gbps</p> <p>Reboot the Data Mover for these changes to take effect.</p>   |

**Table 5. Parameter facility D to H** (continued)

| Facility | Parameter                | Value                | Comments/description  |
|----------|--------------------------|----------------------|---|
| fcTach   | linx_speed_be1           | Default: 0x00008000  | <p>Sets the link speed for the Fibre Channel backend port BE1. This parameter sets the backend port BE1 to a link speed that corresponds to one of the following hexadecimal values:</p> <p>&lt;link_speed_mask&gt;s: 0x8000 = auto-negotiate the link speed</p> <p>0x4000 = 1Gbps</p> <p>0x2000 = 2Gbps</p> <p>0x1000 = 4Gbps</p> <p>0x0800 = 8Gbps</p> <p>Reboot the Data Mover for these changes to take effect.</p> |
| file     | asyncthreshold           | 1–32<br>Default: 128 | <p>Sets the maximum number of blocks that are cached by NFSv3 asynchronous writes.</p> <p>Examples:</p> <p>1 = Improves the accuracy of file system quota enforcement.</p> <p>32 = Provides optimal performance for system throughput.</p> <p>Setting this parameter to any value below the default adversely affects async performance.</p>  |
| file     | asyncThresholdPercentage | 20–90<br>Default: 50 | Refers to the total number of cached dirty blocks for Network File System (NFS) asynchronous writes.  |
| file     | cachedNodes              | Default: 1536000     | The in-memory vnode cache enables inode and block cache for files and directories on the filesystem. It is a global cache, not a per file system cache. A large vnode cache will provide better performance, but will consume more memory. Any changes made to this parameter will be applicable after reboot of the Data Mover.  |

**Table 5. Parameter facility D to H** (continued)

| Facility   | Parameter        | Value                | Comments/description   |
|------------|------------------|----------------------|--|
| file       | dnlcNents        | Default: 2304000     | The in-memory DNLC cache provides the mapping between filenames and files on the file system. It is a global cache, not a per file system cache. A large DNLC cache size will provide better performance, but will consume more memory. Any changes made to this parameter will be applicable after reboot of the Data Mover.  |
| file       | fsInodeThreshold | 10–99<br>Default: 90 | <p>Sets the percentage of used inodes in a file system, so that an event can be logged. In a root file system, a call home is generated for the event. This parameter applies to all file systems of the Data Mover or cabinet depending on how the parameter is set.</p> <p>The number of inodes used is basically equivalent to the number of files and directories that exist in a file system. This parameter is useful if you want to be notified when you are running out of inodes, but not necessarily out of disk space. This parameter triggers a call home when inode usage in the root file system reaches the percentage threshold.</p> |
| file       | fsSizeThreshold  | 10–99<br>Default: 90 | <p>Specifies the file system usage threshold as a percentage of total space. When the threshold is exceeded, an event is issued:</p> <pre>file system size threshold (nn%) crossed</pre> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/>  |
| filesystem | rstchown         | 0 or 1<br>Default: 1 | <p>Sets restricted file ownership.</p> <p>0 = Allow the owner of a file to change the file ownership or group ID to any other owner or group because chown and chgrp follow the less restrictive Portable Operating System Interface for Unix (POSIX) semantics.</p> <p>1 = Allow only the superuser to change the owner of a file. The current owner can change the group ID only to a group to which the owner belongs.</p> <hr/> <p>Note: This parameter applies to NFS, but not to CIFS.</p> <hr/>   |

Table 5. Parameter facility D to H (continued)

| Facility      | Parameter       | Value                                      | Comments/description  |
|---------------|-----------------|--|---|
| ftpd          | bounceAttackChk | 0 or 1<br>Default: 1                       | Controls verification of the PORT command.<br><br>FTP servers can allow an attacker to connect to arbitrary ports on machines other than the FTP client (CVE-1999-0017).<br><br>0 = Disable PORT command verification.<br><br>1 = FTPD verifies that the PORT command does not contain an IP address that differs from the IP address of the client.<br><br>If you need to transfer files between FTP servers, this parameter must be set to 0. |
| ftpd          | defaultdir      | Text string (cannot be null)<br>Default: / | Sets the default working directory for FTP users. This default working directory is used when the home directory of the user that is connected is not reachable, or is not specified.   |
| ftpd          | maxCnx          | 0–0xffffffff<br>Default: 65535             | Specifies the default maximum FTPDs, which are hard-coded.<br><br>_____<br>Note: Reboot the Data Mover for changes to take effect.<br>_____   |
| ftpd          | shortpathdir    | 0x00045cc3b0<br>Default: 0x00000000        | Enables (1) or disables (0) the return file name instead of the full pathname in the commands dir or ls. If a wildcard character is used, this parameter is inefficient.  |
| ftpd          | wildcharsInDir  | 0 or 1<br>Default: 0                       | Controls the use of wildcard characters in directory names for the commands dir and ls.<br><br>0 = Do not allow use of wildcard characters.<br><br>1 = Allow use of wildcard characters.<br><br>The wildcard characters are:<br><br>* = matches any sequence of zero or more characters<br><br>? = matches any single character   |
| FLRCompliance | ActivityLogRP   | 7–2190<br>default: 7                       | Specifies the retention period of the activity log. The minimum retention period is 7 days.   |

**Table 5. Parameter facility D to H** (continued)

| Facility      | Parameter      | Value                       | Comments/description  |
|---------------|----------------|-----------------------------|---|
| FLRCompliance | writeverify    | 0 or 1<br>Default: 0        | Controls the writeverify flag for the FLR-C file systems.<br><br>0 = The writeverify flag is disabled.<br>1 = The writeverify flag is enabled.  |
| http          | quarantineTime | 0-0xffffffff<br>Default: 60 | Specifies the time (in seconds) that an unavailable IP address is quarantined—that is, not retried. This delay prevents reuse of the IP address when the HTTP client detects multiple connection problems on an HTTP server during a session. The quarantine period enables the HTTP client to fail over the connection to another IP address (another physical server).<br><br>_____<br>Note: Reboot the Data Mover for changes to take effect.<br>_____ |

## Parameter facility I to L

Table 6 on page 78 describes the parameters for facilities from I to L.

**Table 6. Parameter facility I to L**

| Facility | Parameter     | Value                | Comments/description   |
|----------|---------------|----------------------|--|
| ip       | reflect       | 0 or 1<br>Default: 1 | Controls Packet Reflect for the system.<br><br>0 = Disable Packet Reflect.<br>1 = Enable Packet Reflect.   |
| iscsi    | AsyncEvent    | 0 or 1<br>Default: 1 | Specifies the type of asynchronous message sent to initiators when the inventory of iSCSI LUNs changes.<br><br>0 = Send asynchronous logout message.<br>1 = Send asynchronous event message. |
| iscsi    | CollPerfStats | 0 or 1<br>Default: 0 | Controls the collection of internal iSCSI performance statistics.<br><br>0 = Do not collect internal iSCSI performance statistics.<br>1 = Collect internal iSCSI performance statistics.     |

Table 6. Parameter facility I to L (continued)

| Facility | Parameter   | Value                | Comments/description  |
|----------|-------------|----------------------|---|
| iscsi    | EnableAptpl | 0 or 1<br>Default: 0 | <p>Specifies whether targets on the Data Mover support SCSI persistent reservation (PR) commands that have the Activate Persist Through Power Loss (APTPL) bit set. When APTPL support is enabled, the server supports cluster requests to maintain a copy of the PR information on nonvolatile storage in order to restore information after a power loss.</p> <p>0 = Disable APTPL support. A PR command will fail if the APTPL bit is set.</p> <p>1 = Enable APTPL support for all targets on the Data Mover.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/> |
| iscsi    | ImmedData   | 0 or 1<br>Default: 1 | <p>Specifies whether the iSCSI initiator is allowed to send unsolicited immediate data.</p> <p>0 = Initiator must not send unsolicited immediate data.</p> <p>1 = Initiator may send unsolicited immediate data. (This is also the default setting for the Microsoft iSCSI initiator.)</p>  |
| iscsi    | InitialR2t  | 0 or 1<br>Default: 0 | <p>Specifies whether the iSCSI initiator is allowed to send unsolicited nonimmediate data.</p> <p>0 = Target turns off the default use of R2T and allows the initiator to send unsolicited nonimmediate data. (This is also the default setting for the Microsoft iSCSI initiator.)</p> <p>1 = Enable use of R2T and disallow unsolicited nonimmediate data.</p>  |

**Table 6. Parameter facility I to L** (continued)

| Facility | Parameter          | Value                      | Comments/description   |
|----------|--------------------|----------------------------|--|
| iscsi    | MaxConnections     | 0–0xffffffff<br>Default: 4 | Specifies maximum number of connections supported per iSCSI session. (The Microsoft iSCSI initiator uses the same default value.)<br><br>Example:<br><br>10 = Initiator can create as many as 10 connections in a given session.<br><br>_____<br><br>Note: Reboot the Data Mover for changes to take effect.<br><br>_____  |
| iscsi    | PreferDataDigest   | 0 or 1<br>Default: 0       | Disables or enables data digests.<br><br>0 = Target does not use data digests even if they are offered by the iSCSI initiator.<br><br>1 = Target uses data digests if the iSCSI initiator offers them.<br><br>The server disables data digests by default because they generate significant CPU overhead. Unless your environment requires them, do not enable data digests. |
| iscsi    | PreferHeaderDigest | 0 or 1<br>Default: 1       | Controls header digests.<br><br>0 = Target does not use header digests even if they are offered by the iSCSI initiator.<br><br>1 = Target uses header digests if offered by the iSCSI initiator.   |
| iscsi    | RequireChap        | 0 or 1<br>Default: 0       | Specifies whether the target enforces Challenge Handshake Authentication Protocol (CHAP) authentication during login for every iSCSI normal session.<br><br>0 = Target does not require CHAP authentication.<br><br>1 = Target requires CHAP authentication during login for every iSCSI normal session.   |

Table 6. Parameter facility I to L (continued)

| Facility | Parameter            | Value                | Comments/description  |
|----------|----------------------|----------------------|---|
| iscsi    | RequireDiscoveryChap | 0 or 1<br>Default: 0 | Specifies whether the target enforces CHAP authentication for every iSCSI discovery session.<br><br>0 = Target does not require CHAP authentication for the discovery session.<br><br>1 = Target requires CHAP authentication for the discovery session.  |
| iscsi    | SendTargetsMode      | 0 or 1<br>Default: 0 | Controls the filtering of targets during iSCSI discovery.<br><br>0 = iSCSI returns information about all targets even if LUN masking prevents the initiator from connecting to the LUNs.<br><br>1 = iSCSI returns information only about targets where the initiator can access LUNs. Information about targets where the initiator is blocked from all LUN access is not returned. |
| iscsi    | WinCompat            | 0 or 1<br>Default: 1 | Specifies whether iSCSI LUN numbering is Windows-compatible.<br><br>0 = Allow 256 LUNs per target (numbered 0 through 255).<br><br>1 = Enforce the Windows limit of 255 LUNs per target (numbered 0 through 254).   |
| ldap     | BufferSize           | Default: 256         | Determines how large each LDAP Buffer will be, when the created Default size is 256 (kilobytes), with the maximum being 10240 (10 MB). Changes to this parameter take effect after rebooting the Data Mover.  |

**Table 6. Parameter facility I to L** (continued)

| Facility | Parameter               | Value                        | Comments/description   |
|----------|-------------------------|------------------------------|--|
| ldap     | cacheEnable             | 0,1,or 2<br>Default: 2       | <p>Specifies whether the LDAP cache is enabled, and if so, the mode.</p> <p>It can be useful to disable the cache for testing.</p> <p>When enabled in offline mode, if an expired cache entry is found and the entry cannot be updated because none of the LDAP servers are available, the LDAP service reports the expired entry.</p> <p>Enabling in cache-only mode can be used to enhance security.</p> <p>0 = Disable the LDAP cache.<br/>1 = Enable the LDAP cache in cache-only mode.<br/>2 = Enable the LDAP cache in offline mode.</p> |
| ldap     | cacheMaxGroups          | 10-1000000<br>Default: 10000 | <p>Specifies the maximum number of cached groups.</p> <p>After reaching this limit, each new group entry removes the oldest entry.</p> <p>To reduce the maximum count below the current count of entries, use the <code>server_ldap -clear</code> command or reboot the Data Mover.</p>  |
| ldap     | cacheMaxHosts           | 10-1000000<br>Default: 10000 | <p>Specifies the maximum number of cached hosts.</p> <p>After reaching this limit, each new host entry removes the oldest entry.</p> <p>To reduce the maximum count below the current count of entries, use the <code>server_ldap -clear</code> command or reboot the Data Mover.</p>  |
| ldap     | cacheMaxHostToNetgroups | 10-1000000<br>Default: 10000 | <p>Specifies the maximum number of cached netgroups.</p> <p>After reaching this limit, each new netgroup entry removes the oldest entry.</p> <p>To reduce the maximum count below the current count of entries, use the <code>server_ldap -clear</code> command or reboot the Data Mover.</p>  |

Table 6. Parameter facility I to L (continued)

| Facility | Parameter     | Value                        | Comments/description  |
|----------|---------------|------------------------------|---|
| ldap     | cacheMaxUsers | 10–1000000<br>Default: 10000 | Specifies the maximum number of cached users.<br><br>After reaching this limit, each new user entry removes the oldest entry.<br><br>To reduce the maximum count below the current count of entries, use the server_ldap -clear command or re-boot the Data Mover.  |
| ldap     | cacheTTL      | 10–86400<br>Default: 600     | Sets the duration in seconds for an LDAP cache entry to remain in cache. After this time has elapsed, the Data Mover will query the directory server again and refresh the entry. When the offline mode is enabled for LDAP, the expired entries can be revived until LDAP is in the offline mode.  |
| ldap     | SecurityLayer | 0, 1, 2, or 4<br>Default: 2  | Sets the level of security the Data Mover uses during negotiation with the domain controller during an LDAP BIND session.<br><br>0 = No security layer — Respond with no security layer regardless of what the domain controller proposes. In Windows environments configured to enforce LDAP signing, the BIND procedure fails because of this value.<br><br>1 = Same as LDAP server — Agree with any security layer proposed by the domain controller.<br><br>2 = Integrity protection — Always propose LDAP signing, which checks the contents of the LDAP messages.<br><br>4 = Privacy protection — Always propose LDAP message encryption, which prevents the data in the packets from being sent in clear text.<br><br>_____<br>Note: Changes to this parameter take effect at the next LDAP BIND.<br>_____ |
| lockd    | asyncTO       | 0–100<br>Default: 10         | Sets the delay (in milliseconds) before sending NLM_GRANTED or NLM_GRANTED_MSG. This parameter is used to bypass client bugs. Some clients do not want to receive a blocked status that is quickly followed by a granted status.  |

Table 6. Parameter facility I to L (continued)

| Facility | Parameter        | Value                         | Comments/description  |
|----------|------------------|-------------------------------|---|
| lockd    | gpDuration       | 30–180<br>Default: 45         | Sets the grace period interval (in seconds) after the Data Mover reboots during which clients can reclaim the locks established before the reboot. During this interval, no new lock can be granted.  |
| lockd    | grantCS          | 0–0xffffffff<br>Default: 1024 | Specifies size (in 8 KB units) of the cache that contains NLM_GRANTED_MSG messages that could need to be repeated. (Each message requires an M_DATA part, which has a granularity of 8 KB.)<br><br>This parameter applies only if the lockd facility grantRetry parameter is nonzero.<br><br>Example:<br><br>1024 = Represents more than 8 MB of memory (1024 x 8 KB).<br><br>This resource is actually used if NLM locks are used. |
| lockd    | grantRetry       | 0 or 10<br>Default: 0         | Defines the number of retries to perform on lockd grant. The maximum of 10 is theoretical, because the maximum actually depends on the size of the cache (defined by the lockd facility grantCS parameter) and the rate of grant messages emitted at that time.   |
| lockd    | grantTO          | 0–0xffffffff<br>Default: 4000 | Sets the timeout (in milliseconds) after which the Data Mover reissues the NLM_GRANTED_MSG call to the client.<br><br>0 = Use the default time-out value.<br><br>4000 = Use a 4-second timeout.   |
| lockd    | maxLockPerClient | Default: 30000                | In order to limit the processing time to remove all client locks next to a statd reboot notification (NLM) or client expiration (NFSv4), VNX Operating Environment for File limits the count of file range locks that it may grant to an NFS client. CIFS clients are not impacted by this parameter.   |

**Table 6. Parameter facility I to L (continued)**

| Facility | Parameter                           | Value   | Comments/description  |
|----------|-------------------------------------|---|---|
| lockd    | OpLockTO1<br>OpLockTO2<br>OpLockTO3 | 1-0xffffffff<br>Defaults:<br>OpLockTO1: 5000<br>OpLockTO2: 5000<br>OpLockTO3: 9 | Sets the CIFS OpLock break request time-out values in milliseconds. OpLockTO1 is the length of time the Data Mover waits after sending an OpLock break request to a client before sending a second one if no response is received to the first request.<br><br>OpLockTO2 is the length of time the Data Mover waits after sending a second OpLock break request to a client before sending a third one if no response is received to the second request.<br><br>OpLockTO3 is the length of time the Data Mover waits after a file is opened before sending any OpLock break requests.                                     |
| lockd    | vnodePercent                        | 10 - 99<br>Default: 90  | Specifies the number of nodes that can be used by any protocol for opens, delegations, and range locks as a percentage of the total count of nodes. After this limit is reached, an event is periodically sent to alert the administrator. The event includes a list of clients that are suspected of locking entropy.<br><br>You can define the total count of nodes that the Data Mover can manage in the setup file /nas/server/slot_N/file, by setting file initialize nodes = xxx.<br><br>_____<br>Note: Reboot the Data Mover for changes to take effect.<br><br>_____<br>No GUI or CLI command reports this value. |

## Parameter facility M to P

[Table 7 on page 85](#) describes the parameters for facilities from M to P.

**Table 7. Parameter facility M to P**

| Facility | Parameter     | Value                | Comments/description   |
|----------|---------------|----------------------|--|
| mount    | allowNullCred | 0 or 1<br>Default: 0 | Controls whether a client that sends AUTH_NONE credential can access NFS exports.<br><br>0 = Do not allow access by AUTH_NONE credential.<br><br>1 = Allow access by AUTH_NONE credential. |

**Table 7. Parameter facility M to P** (continued)

| Facility | Parameter          | Value                | Comments/description   |
|----------|--------------------|----------------------|--|
| mount    | checkPort          | 0 or 1<br>Default: 0 | Controls mount requests from non-privileged ports (>= 1024).<br><br>0 = Allow mount requests from non-privileged ports.<br>1 = Do not allow mount requests from non-privileged ports.  |
| mount    | forceFullShowmount | 0 or 1<br>Default: 1 | Controls the visibility of the NFS export information. It allows the filtering of entries if the client does not have mount permission for the file system corresponding to that entry. By default, it is disabled.<br><br>0 = Enable NFS export hiding.<br>1 = Disable NFS export hiding. |
| mount    | tcpResponseLimit   | Default: 262144      | Changes the maximum size of the mount tcp response. Sets the parameter according to how many mount-exports are configured in the system. This parameter may limit how many exports can be shown by 'showmount'.  |
| nbs      | extendRplcInfo     | 0 or 1<br>Default: 1 | Controls extended information about iSCSI replication.<br><br>0 = Do not return extended information.<br>1 = Return extended replication information (such as the replication label and alias).  |
| nbs      | sparseTws          | 0 or 1<br>Default: 0 | Controls support of thin provisioned temporary writable snapshots (TWS).<br><br>0 = Create a regular (dense) TWS if the production LUN is not virtually provisioned.<br>1 = Always create a virtually provisioned (sparse) TWS.  |

Table 7. Parameter facility M to P (continued)

| Facility | Parameter             | Value                          | Comments/description  |
|----------|-----------------------|--------------------------------|---|
| NDMP     | bufsz                 | 64–1168<br>Default: 128        | <p>Sets the size (in kilobytes) of the buffer allocated for reading and writing the tape requested by TAPE_READ/TAPE_WRITE from the NDMP client.</p> <p>If you use EMC NetWorker®, the NetWorker documentation provides appropriate buffer size for your model of tape drive.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/>  |
| NDMP     | CDBFsinfoBufSizeInKB  | Default: 1024                  | <p>Can be used for a Celerra system to adjust buffer size for SCSI/TAPE_EXECUTE_CDB and CONFIG_GET_FS_INFO. It can be increased if the system has a big tape library with thousands of slots and the backup software does not divide the CDB read_element_status into multiple requests. It can also be increased if the system has thousands of file systems mounted. It may be decreased if the system cannot boot up due to fragmented memory.</p> |
| NDMP     | convDialect           | Text string<br>Default: 8859-1 | <p>Sets the conversion dialect, which is required when restoring an ASCII-mode Data Mover image to the Data Mover with Internationalization turned on. This applies only if the conversion dialect is a non-Latin1 dialect.</p> <p>To enter a null string, use double quotation marks ("").</p> <p>Examples:</p> <p>big5</p> <p>latin1</p>  |
| NDMP     | concurrentDataStreams | 1–8<br>Default: 4              | <p>Displays the maximum number of concurrent backup or restore streams that are set. This parameter also enables the user to change the concurrent backup sessions from the default value of four up to a maximum of eight, provided the system has at least 8 GB memory. The default value is 4.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/>  |

**Table 7. Parameter facility M to P** (continued)

| Facility | Parameter               | Value                               | Comments/description  |
|----------|-------------------------|-------------------------------------|---|
| NDMP     | portRange               | Default: 1024–65535                 | <p>Specifies the port range that NDMP uses to listen for data connections.</p> <p>If an invalid port range is set, the server_param command fails with an error message.</p> <p>If the port range is set and no port is available during NDMP backup, the backup fails with an error message. This error is logged in server_log and DMA log.</p> <p>It is recommended to set the NDMP port range using dynamic or private ports (49152–65535) of IANA.</p> |
| NDMP     | dialect                 | Text string<br>Default: null string | <p>Sets the dialect on the Data Mover when operating in Internationalization mode. This value should match the dialect used by the Data Mover clients.</p> <p>Null string ("") = Data Mover uses UTF-8.</p>   |
| NDMP     | excludeSvtlFs           | 0 or 1<br>Default: 1                | <p>Specifies whether config_get_fs_info excludes SVTL file systems.</p> <p>0 = Include SVTL file systems.<br/>1 = Exclude SVTL file systems.</p>  |
| NDMP     | forceRecursiveForNonDAR | 0 or 1<br>Default: 0                | <p>Controls whether to override the environment variable settings (DIRECT and RECURSIVE) and force a recursive restore for directories in non-DAR restoration operations.</p> <p>0 = Do not override environment variable settings for recursive restore.<br/>1 = Force a recursive restore even if the DIRECT and RECURSIVE environment variables are set (or default) to N.</p>   |
| NDMP     | includeCkptFs           | 0 or 1<br>Default: 1                | <p>Specifies whether config_get_fs_info includes ckpt file systems.</p> <p>0 = Exclude ckpt file systems.<br/>1 = Include ckpt file systems.</p>  |
| NDMP     | maxProtocolVersion      | 2–4<br>Default: 4                   | <p>Sets the maximum NDMP version used by the NDMP server and the Data Management Application. This parameter overrides autonegotiation between the server and the NDMP client to determine the highest NDMP version used.</p>   |

Table 7. Parameter facility M to P (continued)

| Facility | Parameter   | Value                | Comments/description   |
|----------|-------------|----------------------|--|
| NDMP     | md5         | 0 or 1<br>Default: 1 | <p>Specifies whether MD5 encryption is used when sending the password during NDMP authentication.</p> <p>0 = Do not use MD5 authentication (send password in plain text).</p> <p>1 = Enable MD5 encryption for sending the password.</p> <p>To use the text method, type the command (as root):</p> <pre># /nas/sbin/server_user &lt;mover name&gt; -add -passwd &lt;username&gt;</pre> <p>To use the MD5 method, type the command (as root):</p> <pre># /nas/sbin/server_user &lt;mover name&gt; -add -md5 -passwd &lt;username&gt;</pre> |
| NDMP     | scsiReserve | 0 or 1<br>Default: 0 | <p>Determines whether the SCSI reserve command is used for tape open.</p> <p>0 = Do not use SCSI reserve.</p> <p>1 = Enable SCSI reserve.</p> <p>_____</p> <p>Note: Reboot the Data Mover for changes to take effect.</p> <p>_____</p>   |
| NDMP     | snapsure    | 0 or 1<br>Default: 0 | <p>Determines whether SnapSure™ is used for backup if the NDMP environment variable SNAPSURE is not set or not supported by the DMA software. The NDMP environment variable, if set, overrides this parameter.</p> <p>0 = Do not use SnapSure for backup.</p> <p>1 = Use SnapSure for backup.</p>  |
| NDMP     | snapTimeout | 2–60<br>Default: 5   | <p>Specifies the timeout (in minutes) for creation of the backup SnapSure file system.</p>   |

**Table 7. Parameter facility M to P** (continued)

| Facility | Parameter                 | Value                      | Comments/description  |
|----------|---------------------------|----------------------------|---|
| NDMP     | tapeSilveringStr          | Text string<br>Default: ts | Specifies the end of the string used for indicating tape silvering backup. The complete string is made by applying the prefix "/.". Example:<br><br>ts = Strings of the form /ts/ will trigger tape silvering backups.<br><br>If the Data Management Application checks for the file system, strings of the form //ts can be specified.   |
| NDMP     | v4OldTapeCompatible       | 0 or 1<br>Default: 1       | Sets the NDMP v2/v3 tape operation behaviors as the NDMP v4 spec.<br><br>0 = Do not use v4 tape behaviors for NDMP v2/v3.<br>1 = Enable v4 tape behaviors for NDMP v2/v3.<br><br>Note: Reboot the Data Mover for changes to take effect.  |
| nfs      | checkForExistingFileFirst | 0 or 1<br>Default: 0       | Specifies the type of error message to return to a non-root user who lacks write permission when trying to create a directory that already exists.<br><br>0 = Indicates only that permission is denied.<br>1 = Indicates that the file already exists.  |
| nfs      | ntAllSpecialUsersD        | Default: 0                 | Allows returning most of the commonly used Windows special users including those that are not defined in RFC 3530.<br><br>0 = All the Windows special users not listed in the RFC 3530 are returned as 'nobody', and an ACL containing such an ACE cannot be modified using NFSv4.<br><br>1 = These special users are translated to a unique identifier, which allows modifying the ACL by adding or removing some ACE. |
| nfs      | NTcred.LDAP               | Default: 0                 | This will give access to nested and universal groups if the domain controller has access to the global catalog.   |

Table 7. Parameter facility M to P (continued)

| Facility | Parameter    | Value                         | Comments/description  |
|----------|--------------|-------------------------------|---|
| nfs      | NTcred.size  | 0-0xffffffff<br>Default: 1009 | Specifies the maximum number of entries allowed in the NT credential cache for the Data Mover. The oldest unused entry is removed first. This parameter can be adjusted to avoid rebuilding credential on each NFS request.<br><br>Note: Reboot the Data Mover for changes to take effect.  |
| nfs      | NTcred.trace | 0-7<br>Default: 3             | Specifies the level of detail that is logged in to the Data Mover server_log associated with the use of the NT credential cache.<br><br>Examples:<br><br>3 = Only a few informational messages appear in the log. Increasing the value increases the level of detail and hence the number of messages that appear in the server_log.<br><br>7 = The system logs the maximum amount of information.<br><br>Changing this parameter should be necessary only when troubleshooting the system. |
| nfs      | NTcred.TTL   | 0-0xffffffff<br>Default: 20   | Sets the time-to-live expiration stamp (in minutes) for a Windows NT entry in the NT credential cache. When failed mapping entries expire, the system retries to map the UID to the SID.  |

Table 7. Parameter facility M to P (continued)

| Facility | Parameter            | Value  | Comments/description   |
|----------|----------------------|--|--|
| nfs      | NTcred.winDomain     | Text string: any valid NetBIOS domain name<br>Default: null string | <p>Specifies the default Windows NetBIOS domain name of the Data Mover to be used for NFS users, accessing a file system where the ntcredential option has been used. This value is used only if several different SIDs match the UNIX UID of the user, or the UID-to-name reverse mapping returns an ambiguous username (no domain).</p> <p>In order to support the nfs NTcred feature for the NFS multiple domains feature, the following Windows registry entry is used to specify the equivalent parameter per VDM value:</p> <pre>[KEY_LOCAL_MACHINE\Software\EMC\UnixNTCredential] "NetbiosDomainName"="&lt;domain_name&gt;"</pre> <p>When the VDM name resolution is confined to the VDM, the nfs NTcred mechanism uses the domain specified in the registry and a CIFS server configured in the VDM.</p> |
| nfs      | nthreads             | Default: 512   | <p>This parameter represents the number of threads dedicated to serve nfs requests. This parameter is memory dependent. Ensure that the system memory can support your configuration.</p> <p>The update of the nfs thread count needs several minutes to take effect, and the nfs.nthreads value cannot be changed until the update is done.</p>   |
| nfs      | ofCachesize          | Default: 1080000   | <p>This parameter represents the size of the NFS open file cache. Reboot the Data Mover for the changes made to this parameter to take effect.</p>   |
| nfs      | rpcgss.discardReplay | 0 or 1<br>Default: 1   | <p>Controls how lost NFS/RPCGSS requests are handled. If the NFS/RPCGSS request is lost on the network, some secure NFS clients resend the NFS/RPCGSS request with the same sequence number.</p> <p>0 = Resend the NFS/RPCGSS request with the same sequence number.</p> <p>1 = Discard the replay of RPCGSS requests.</p> <p>This problem has been seen on Solaris 8 and 9.</p>   |

Table 7. Parameter facility M to P (continued)

| Facility | Parameter        | Value                                   | Comments/description   |
|----------|------------------|---|--|
| nfs      | secureExportMode | 0 or 1<br>Default: 0                    | Specifies how to resolve access conflicts between read-only and read/write access when using the server_export command with multiple netgroups. If a netgroup is specified twice, once on a read-write list and again on a read-only list, an access conflict occurs.<br><br>0 = Allow read/write access if there is a conflict.<br>1 = Allow read-only access if there is a conflict.   |
| nfs      | spExportLookup   | 0 or 1<br>Default: 0                    | Determines the access permissions that are granted by matching the hostname, IP address, subnet, and netgroup membership of a NFS client.<br><br>0 = Access permissions are determined by matching the hostname and subnet, which are given higher priority, followed by permissions determined by matching the netgroup membership of the NFS client.<br><br>1 = Access permissions are determined by matching the hostname, IP address, or netgroup membership, followed by the subnet of the NFS client. Here the hostname, IP address, or netgroup membership is given priority over access permissions determined by matching the subnet of the NFS client. |
| nfs      | transChecksum    | 0 or 1<br>Default: 0                    | Specifies whether the Data Mover supports Oracle Direct NFS (DNFS) for clients that use Oracle Database 11g with NFSv3. When support is enabled, the Data Mover ensures that each transaction carries a unique ID and avoids the possibility of conflicting IDs that result from the reuse of relinquished ports.<br><br>0 = Do not support DNFS.<br>1 = Support Oracle 11g DNFS clients that use NFSv3.   |
| nfs      | v3xfersize       | 8192, 16384, or 32768<br>Default: 32768 | Specifies the default transfer size for NFSv3 reads and writes.  |

**Table 7. Parameter facility M to P** (continued)

| Facility | Parameter       | Value                               | Comments/description   |
|----------|-----------------|-------------------------------------|--|
| nfsv4    | 32bitClient     | 0 or 1<br>Default: 1                | Sets the bit size of attribute values.<br><br>1 = Require the VNX NFSv4 server to return attribute values in 32 bits only.<br><br>0 = Allow the VNX NFSv4 server to return attribute values in 64 bits.  |
| nfsv4    | advisoryLocking | Default: 0                          | Sets file locking mode to advisory if the value is 1. Advisory locking is compatible with NFSv3. Applications using file locks must be cooperative. A file lock does not prevent conflicting I/Os. The default is mandatory locking (like CIFS). NFSv4 provides mandatory locking, that is, the ability to block operations by other applications on a locked file. Mandatory locking is set by default for NFSv4, but using this parameter set to 1, you can switch to an advisory model for NFSv4. |
| nfsv4    | domain          | Text string<br>Default: null string | Specifies the NFS version 4 domain name. Use the format: xxx.xxx<br><br>The NFSv4 domain name is necessary for the service to map user UIDs and group GIDs to UTF-8 encoded user and group names.<br><br>To replace a value with a null string, type double quotation marks ("").  |
| nfsv4    | leaseDuration   | 20–180<br>Default: 18               | Defines the duration during which the server maintains client states in the absence of client activity.<br><br>This value must be less than the grace period duration specified by the lockd facility gpDuration parameter.<br><br>_____<br>Note: Reboot the Data Mover for changes to take effect.<br>_____   |
| nfsv4    | maxOwners       | Default: 52428                      | This parameter defines the maximum number of nfsv4 open owners and lock owners allowed on the nfsv4 server.  |

**Table 7. Parameter facility M to P** (continued)

| Facility | Parameter            | Value                | Comments/description   |
|----------|----------------------|----------------------|--|
| nfsv4    | recallTimeout        | 5–60<br>Default: 10  | <p>Sets the timeout to the seconds the server should wait before recalling delegations. Delegations are recalled when access conflict occurs and when the file system is unmounted. The parameter also defines the quota returned on write delegations. The client must be able to commit all its unwanted data within this time.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/>   |
| nfsv4    | removeOpened         | Default: 1           | <p>1 = When a file is opened by the same client or another client, it can be removed. The open state is preserved, so that the application that has opened it can still access it. The file will be removed when it is closed.</p> <p>0 = The error NFS4ERR_FILE_OPEN is returned on an attempt to remove an opened file.</p>  |
| nfsv4    | vnodePercent         | 10–80<br>Default: 80 | <p>Specifies the maximum percentage of nodes usable by NFSv4 for opens and delegations. After the limit is reached, new opens are denied.</p> <p>You can define the total count of nodes that the Data Mover can manage in the setup file <code>/nas/server/slot_N/file</code>, by setting <code>file initialize nodes = xxx</code>.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/> <p>No GUI or CLI command reports this value.</p> |
| PAX      | allowVLCRestoreToUFS | 0 or 1<br>Default: 0 | <p>Specifies whether to allow Volume-Level Copy (VLC) backup to be restored to UFS file system.</p> <p>0 = Do not allow VLC backup to be restored to UFS file system.</p> <p>1 = Allow VLC backup to be restored to UFS file system.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/>  |

**Table 7. Parameter facility M to P** (continued)

| Facility | Parameter            | Value                               | Comments/description   |
|----------|----------------------|-------------------------------------|--|
| PAX      | checkUtf8FileNames   | 0 or 1<br>Default: 1                | Controls whether filenames are checked for valid UTF-8 format.<br><br>0 = Disable the verification.<br>1 = Enable the verification.<br><br>_____<br>Note: Reboot the Data Mover for changes to take effect.<br>_____ |
| PAX      | dump                 | 0 or 1<br>Default: 0                | Specifies whether to force dump format for backup.<br><br>0 = Disable the dump format.<br>1 = Force the dump format.<br><br>_____<br>Note: Reboot the Data Mover for changes to take effect.<br>_____                |
| PAX      | filter.caseSensitive | 0 or 1<br>Default: 1                | Specifies whether the directory and file filters used to exclude directories and files from a backup are case-sensitive.<br><br>0 = Filters are not case-sensitive.<br>1 = Filters are case-sensitive.               |
| PAX      | filter.dialect       | Text string<br>Default: null string | Specifies the dialect to use for filter parameters when VNX operates in Internationalization mode. This value must match the dialect used by the backup client.<br><br>Example:<br>UTF-8                             |
| PAX      | filter.numDirFilter  | 0-50<br>Default: 5                  | Specifies the maximum number of directory filters that can be used for a backup job. A higher number means more directory types can be excluded from a backup.   |
| PAX      | filter.numFileFilter | 0-50<br>Default: 5                  | Specifies the maximum number of file filters that can be used for a backup job. A higher number means more file types can be excluded from a backup.   |

Table 7. Parameter facility M to P (continued)

| Facility | Parameter              | Value   | Comments/description   |
|----------|------------------------|---|--|
| PAX      | noFileStreams          | 0 or 1<br>Default: 0  | Specifies whether to exclude CIFS file streams from backup/restore operations.<br><br>0 = Do not exclude CIFS file streams from restore/backup.<br>1 = Exclude CIFS file streams from restore/backup.<br><br>_____<br>Note: Reboot the Data Mover for changes to take effect.<br>_____ |
| PAX      | nRestore               | 510 Data Movers and higher: 1–16<br>507 Data Movers and lower: 1–16<br>Default: 8 | Specifies the number of restore threads for each restore job. More threads can increase the speed with which data is restored.   |
| PAX      | nThreads               | Default: 64   | Specifies the number of backup threads. Increasing the number of backup threads could increase the speed with which the data is retrieved from disk.   |
| PAX      | paxStatBuff            | 1–512<br>Default: 128   | Specifies the number of buffers between the threads that send and receive metadata (NASA), and read data from disk (NASS).<br><br>More buffers can increase the speed with which metadata is provided.   |
| PAX      | readWriteBlockSizeInKB | 64–256<br>Default: 64   | Specifies the maximum allowed PAX buffer size for NDMP read/write.<br><br>While using EMC Data Manager (EDM) for backup and restore operations, the value must be set to 128.<br><br>_____<br>Note: Reboot the Data Mover for changes to take effect.<br>_____                         |

**Table 7. Parameter facility M to P** (continued)

| Facility | Parameter     | Value                | Comments/description  |
|----------|---------------|----------------------|---|
| PAX      | scanOnRestore | 0 or 1<br>Default: 1 | Specifies whether to scan for viruses after a file is restored.<br><br>0 = Do not scan restored files for viruses.<br>1 = Scan restored files for viruses.<br><br>_____<br>Note: Reboot the Data Mover for changes to take effect.<br>_____ |
| PAX      | writeToArch   | 0 or 1<br>Default: 1 | Determines whether the backup writes to a file when using the server_archive command.<br><br>0 = server_archive discards the data without writing to a file.<br>1 = server_archive writes to a file.  |
| PAX      | writeToTape   | 0 or 1<br>Default: 1 | Determines whether the backup writes to tape.<br><br>0 = The backup data is discarded to test the backup speed while eliminating tape-drive speed as a performance factor.<br><br>1 = The backup writes to tape.                            |

## Parameter facility Q to S

Table 8 on page 98 describes the parameters for facilities from Q to S.

**Table 8. Parameter facility Q to S**

| Facility | Parameter                 | Value                | Comments/description  |
|----------|---------------------------|----------------------|---|
| quota    | countRootUsagelnQuotaTree | 0 or 1<br>Default: 1 | Controls whether root users are included in the count of files and directories for the tree quotas.<br><br>0 = Exclude root users from the tree quota count.<br>1 = Include root users (and nonroot users) in the tree quota count.<br><br>_____<br>Note: Restart the CIFS service for the changes to take effect.<br>_____ |

Table 8. Parameter facility Q to S (continued)

| Facility | Parameter         | Value                                 | Comments/description   |
|----------|-------------------|---------------------------------------|--|
| quota    | maxuid            | 0–0xffffffff<br>Default: 0            | <p>Sets the highest user ID to which quotas apply. You can use this parameter to prevent problems caused by accidentally imposing quotas on very large UIDs.</p> <p>Set this parameter to the highest UID the site expects to support. You can increase the value later, if requirements change.</p> <p>Example:<br/>0 = Enforce quota limits for all users (UIDs).</p>  |
| quota    | policy            | filesize or blocks<br>Default: blocks | <p>Specifies the quota checking policy used to track disk usage.</p> <p>blocks = Quota is based on the number of file system blocks (8 KB) allocated.</p> <p>filesize = Quota is based on file usage in 1 KB increments.</p> <p>Before changing this parameter, turn off quotas. After you change this parameter, reboot the Data Mover and turn on quotas. <i>Using Quotas on VNX</i> provides specific procedures.</p> <p>Switching from filesize to blocks might put users over their quota if they were close to their quota limit before the change.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/>   |
| quota    | useQuotasInFsStat | 0 or 1<br>Default: 0                  | <p>Controls whether quotas are included when displaying file system free-space statistics to NFS clients that use the UNIX <code>df -k</code> command to view statistics.</p> <p>0 = Exclude quotas when a disk quota verification is done by using <code>df</code>. The actual available space can be less than the "space available" shown in the command output.</p> <p>1 = Include quotas. The <code>df</code> command run by a nonroot user reports only the space available to the user. This means the "space available" displays factors in the space that are preallocated to quota trees, users, and groups. <code>rquota</code> is not supported for tree quotas.</p> |

**Table 8. Parameter facility Q to S** (continued)

| Facility | Parameter | Value                  | Comments/description  |
|----------|-----------|------------------------|---|
| RCP      | tcpwindow | 0–2097152<br>Default:0 | <p>Specifies the window size used by nas_copy operations on the Data Mover. This value indicates the number of bytes that can be sent before the data is acknowledged by the receiving site.</p> <p>If you change the parameter after these processes start, you must reset the connection for the change to take effect.</p> <p>If a request exceeds 2 MB, the system sets the window size to 2 MB.</p> <p>0 = Set the window size to 128 KB.<br/>2097152 = Set the window size to 2 MB.</p> <p>Increasing the window size is most effective while working with latent, low-loss networks.</p> <p>Use this equation to calculate window size:<br/>Window Size = Round Trip Delay x Desired Rate</p> <p>Example:<br/>If the round-trip delay is 100 ms and you plan to send 10 MB/s across the IP network, you need a window size of 1 MB (0.1 s x 10 MB/s = 1 MB). So you set this parameter to 1048576 (1 MB = 1024 x 1024 bytes).</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/> |

Table 8. Parameter facility Q to S (continued)

| Facility    | Parameter      | Value                   | Comments/description   |
|-------------|----------------|-------------------------|--|
| RCP         | tcpwindowlowat | 0–2097152<br>Default: 0 | <p>Sets the LO_WAT (low water mark) of TCP used by replication. This value indicates the amount of data that can be unacknowledged with TCP before the data is sent again after blocking because of hitting the HI_WAT (high water mark).</p> <p>TCP xmits are halted from last window size byte segment until the last LO_WAT window size segment TCP ACK. With extremely large window size settings, this xmit-stop time approaches RTT.</p> <p>Increasing the LO_WAT must correspond to the increase of tcpwindow size. For the change to take effect, the connection must be reestablished.</p> <p>Use this equation to calculate LO_WAT window size:<br/>LO_WAT window size = HI_WAT – number of bytes ACK'd when TCP xmit stops.</p> <p>For example, if data ACK at 100 KB/s and the rate of transfer for application is 400 KB/s and window size is set to 1 MB, then this HI_WAT is reached in about 4 seconds (300 KB/s unack). So, the LO_WAT must be increased to around 600 KB (1 MB – 400 KB &lt;&lt;&lt; data ACK in 3 seconds), with a value of 614400.</p> <p>0 = Sets the window size to the TCP default of 64 KB. This parameter is applied to nas_copy for Replication V2.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/> |
| replication | highcapacity   | 0 or 1<br>Default: 1    | <p>Controls whether high-capacity mode is used for easier testing of upgrades.</p> <p>0 = Do not use high-capacity mode.</p> <p>1 = Use high-capacity mode.</p> <p>This parameter is applied to VNX Replicator V2.</p>   |
| security    | aesSupport     | Default: Disabled       | <p>This parameter is provided in the event a user wishes to enable AES support. The Data Mover must be rebooted when this parameter is changed, to reload the keys (with the updated encryption types) for each CIFS server.</p>   |

**Table 8. Parameter facility Q to S** (continued)

| Facility | Parameter                 | Value                 | Comments/description   |
|----------|---------------------------|-----------------------|--|
| security | maxgroups                 | 16–128<br>Default: 16 | In UNIX, a user belongs to many groups. This information is stored in the user credential when the user is authenticated. VNX supports 16 groups by default. Using this parameter you can extend the number of groups supported by VNX from 16 to 128. It also impacts the credential passed to VNX for NFS requests in RPC. Extending the parameter value allows a client (for example, AIX or RedHat) to pass more than 16 groups in the RPC credential. For Kerberos authentication, to support more than 16 groups, set the value of this parameter to more than 16. |
| security | maxNISCacheGroupsCount    | Default: 10000        | If the number of the groups exceeds this value, the group will not be inserted to the cache.   |
| security | maxNISCacheNetgListsCount | Default: 10000        | If the number of the netgroup list exceeds this value, the netgroup list will not be inserted to the cache.  |
| security | maxNISCacheUsersCount     | Default: 10000        | If the number of users exceeds this value, the user will not be inserted to the cache.   |
| security | maxNISGroupCount          | Default: 65536        | On Celerra, the maximum number of primary groups applicable per file system is 64K. To avoid too many cycles to read the NIS group map and mitigate the impact on the memory cache you can limit the number of entries reads on the NIS map. If the number of entries exceeds the limit, you can still use the groups already read.<br><br>The NIS group database is read using yp_first and yp_next. If the database contains the same group with different GIDs, the Celerra may loop indefinitely on yp_next. This parameter avoids this issue.                       |
| security | NISCacheTTL               | Default: 600          | Time to Live of a NIS Cache Entry, in seconds.   |

**Table 8. Parameter facility Q to S** (continued)

| Facility | Parameter   | Value                         | Comments/description  |
|----------|-------------|-------------------------------|---|
| shadow   | asciifilter | 0 or 1<br>Default: 0          | <p>Controls ASCII filtering.</p> <p>0 = Do not use ASCII filtering.</p> <p>1 = Use ASCII filtering.</p> <p>While ASCII filtering is enabled, you cannot use non-ASCII characters in filenames.</p> <p>When using Kerberos authentication in Windows 2000 networks, either ASCII filtering or I18N (internationalization) must be enabled. If you do not enable I18N, ASCII filtering is automatically enabled when you join a CIFS server to a Windows 2000 domain.</p> <p>If ASCII filtering is enabled and you have created at least one compname (the name for a CIFS server in a Windows 2000 network), you cannot disable ASCII filtering until you delete all the compnames.</p> <p>The software automatically sets this parameter to 1 when there is a compname and the file system is non-I18N.</p> <p>Set I18N mode for Unicode on all Data Movers.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/> |
| shadow   | flushCount  | 1-0xffffffff<br>Default: 1024 | <p>Specifies the flush count. When building (or rebuilding) a shadow directory from scratch, the system loops as it reads in blocks and puts others in a queue to be written to disk. The flush count is the number of files the system gets from the UNIX directory before flushing the shadow to disk. By flushing these changes to disk, the thread gives up the CPU and lets other threads run.</p> <p>This parameter is relevant for upgraded systems only.</p>  |

**Table 8. Parameter facility Q to S** (continued)

| Facility | Parameter         | Value                | Comments/description   |
|----------|-------------------|----------------------|--|
| shadow   | followabsolutpath | 0–3<br>Default: 0    | <p>Controls whether Windows clients can follow a symbolic link that contains absolute paths (full pathnames). The bit list consists of two binary bits (bits 0 and 1, right to left). Each bit is 1 when set; otherwise 0.</p> <p>Bit 0:</p> <p>0 = Do not allow following symbolic links that contain an absolute path.</p> <p>1 = Allow following symbolic links that contain an absolute path. There are no restrictions on access, allowing users to easily go outside the share and also into another file system. When users go outside of the share, the security set on the original share to which the user is connected applies.</p> <p>Bit 1:</p> <p>0 = Allow only absolute symbolic links owned by root (UID 0) to be followed.</p> <p>1 = Allow any absolute symbolic links to be followed. Bit 1, if set, creates a potential security issue for NFS access because the NFS client can create an absolute symbolic link to any location in the Data Mover.</p> <p>Example:</p> <p>1 (01) = Allow following symbolic links, but only those owned by root.</p> <p>0 through 3 assume their binary values.</p> <p><i>Managing a Multiprotocol Environment on VNX</i> provides more information about this parameter.</p> |
| shadow   | followdotdot      | 0 or 1<br>Default: 0 | <p>Controls symbolic link following within the current share if the target path includes the dot-dot component (.).</p> <p>0 = Do not allow following symbolic links.</p> <p>1 = Allow following symbolic links.</p>   |
| shadow   | stream            | 0 or 1<br>Default: 1 | <p>Controls support of alternate data stream—this parameter is relevant in a CIFS environment only.</p> <p>0 = Do not support alternate data streams.</p> <p>1 = Support alternate data streams.</p>   |

**Table 8. Parameter facility Q to S** (continued)

| Facility | Parameter | Value  | Comments/description  |
|----------|-----------|--|---|
| ssl      | cipher    | Text string<br>Default:<br>ALL:!ADH:!SSLv2<br>:@STRENGTH | Specifies a keyword defined by the open SSL project. The supported SSL cipher suites define the default list of cryptographic algorithms. SSL applications on the Data Mover may use these algorithms during an SSL connection if not specified by the application.<br><br>Examples:<br><br>ALL:@STRENGTH = All ciphers, sorted by security level. ALL:!LOW:@STRENGTH = All strong ciphers, sorted by security level.<br><br>For the complete list of valid keywords, go to the OpenSSL website.<br><br>_____<br>Note: Reboot the Data Mover for changes to take effect.<br>_____ |
| ssl      | protocol  | 0, 1, or 2<br>Default: 0                                 | Sets the preferred protocol for applications during an SSL connection, if a protocol is not specified by the application.<br><br>0 = Support all SSL/TLS protocols.<br><br>1 = Support SSLv3 only.<br><br>2 = Support TLSv1 only.<br><br>Use the default value (both SSLv3 and TLSv1) for compatibility. SSLv2 is not supported because it does not provide adequate security.<br><br>_____<br>Note: Reboot the Data Mover for changes to take effect.<br>_____   |
| ssl      | threads   | 4–256<br>Default: 256                                    | Specifies the number of threads dedicated to the SSL tasks.<br><br>_____<br>Note: Reboot the Data Mover for changes to take effect.<br>_____  |

**Table 8. Parameter facility Q to S** (continued)

| Facility | Parameter | Value                      | Comments/description  |
|----------|-----------|----------------------------|---|
| ssl      | timeout   | 0–30000<br>Default: 10     | Specifies the maximum time (in milliseconds) that is required for a response to an SSL handshake message.<br><br>Note: Reboot the Data Mover for changes to take effect.  |
| ssl      | trace     | 0–0xffffffff<br>Default: 0 | Specifies which traces to display in the server_log for SSL.<br><br>10 = Display low level dumps.<br><br>100 = Display information level.<br><br>1000 = Display warning level.  |
| statd    | hostname  | Text string<br>Default: /  | Lists remote client names separated by comma. Typically, the local hostname is server_n, but remote clients refer to the Data Mover by other hostnames as defined in its network name service.<br><br>When notifying clients after a reboot, the Data Mover must present the hostname that the client used for its locking requests. Because the locked request does not contain this hostname, statd cannot use the same name in its notification. Thus, statd successively notifies clients under all possible Data Mover names to be sure that each client recognizes its locking context and reclaims its locks.<br><br>/ = Local hostname<br><br>Missing names in this parameter can prevent some clients from recovering all their locks after the Data Mover reboots.<br><br>This parameter is global to the Data Mover and cannot be used for each VDM, because each VDM should notify a different hostname for their attached interface(s).<br><br>By default, the VDM notifies the VDM name (internal name managed with the nas_server command). The new file named .etc/hostname, stored in the VDM root file system, allows the administrator to specify a hostname(s) per VDM, which must be used for server reboot notifications. |

Table 8. Parameter facility Q to S (continued)

| Facility | Parameter        | Value                     | Comments/description  |
|----------|------------------|---------------------------|---|
| statd    | maxAddrPerClient | 4–512<br>Default: 100     | <p>Specifies the maximum number of different IP addresses that a remote client can use to monitor file locks through lockd/statd. The statd service on the Data Mover registers one entry for each client name and client source address. Each address is notified when the Data Mover reboots, as required by some client operating systems.</p> <p>The default value provides reasonable protection from network attacks or from faulty network switch equipment creating false statd registrations.</p> <p>Do not change this parameter unless a client machine requires more than the default number of IP addresses.</p> |
| streamio | timeout          | 180–7200<br>Default: 1800 | <p>Sets the read/write tape time-out period (in seconds) when a tape drive takes too long to respond to a single tape command.</p> <p>7200 = 7200 seconds (2 hours)</p>   |

## Parameter facility T to V

Table 9 on page 107 describes the parameters for facilities from T to V.

Table 9. Parameter facility T to V

| Facility | Parameter | Value                | Comments/description  |
|----------|-----------|----------------------|---|
| tcp      | ackpush   | 0 or 1<br>Default: 0 | <p>Controls the response to a pause packet. Some operating systems (for example, Windows prior to Windows 98) expect a response to the pause packet before continuing communication.</p> <p>0 = Do not send a response to the pause packet.</p> <p>1 = Send an empty frame to the pause packet immediately after the packet is received. Communication can continue thereby, enhancing performance.</p> <p>Changes to this parameter affect all newly opened TCP connections.</p> |

**Table 9. Parameter facility T to V** (continued)

| Facility | Parameter               | Value                 | Comments/description  |
|----------|-------------------------|-----------------------|---|
| tcp      | backlog                 | 0–256<br>Default: 100 | <p>Specifies the number of connections that can be in the process of being established.</p> <p>If you receive the message stating that there are too many connections in the listen queue of port xxx, increase the value of this parameter. Then, restart the affected application.</p> <p>TCP maxstreams setting limits the value of this parameter. Use the maximum value because of memory usage.</p>   |
| tcp      | do_newreno              | 0 or 1<br>Default: 1  | <p>Controls whether the TCP NewReno algorithm is used. NewReno enhances TCP performance in lossy network environments, and can operate with the SACK algorithm or as a stand-alone recovery algorithm.</p> <p>0 = Do not use NewReno algorithm.<br/>1 = Use NewReno algorithm.</p> <hr/> <p>Note: Any change to the parameter affects all existing connections.</p> <hr/>   |
| tcp      | do_sack                 | 0 or 1<br>Default: 1  | <p>Controls whether the TCP Selective Acknowledgment (SACK) algorithm is used. SACK enhances TCP performance in lossy network environments. SACK is a TCP negotiated option.</p> <p>Both peers in the TCP connection must agree to use the SACK option during connection establishment to utilize the SACK recovery algorithm.</p> <p>0 = Do not use SACK algorithm.<br/>1 = Use SACK algorithm.</p> <hr/> <p>Note: Any change to the parameter affects all existing connections.</p> <hr/> |
| tcp      | dosattackadminoptionset | Default: 0            | <p>Once DOS is detected, the administrator can intervene and enable this option to trigger the DOS defense mechanism. The defense mechanism when triggered, will release all the half open connection resources at a faster rate than the standard 75 seconds timeout period.</p>   |

Table 9. Parameter facility T to V (continued)

| Facility | Parameter                        | Value       | Comments/description  |
|----------|----------------------------------|-------------|---|
| tcp      | dosattackadminrecoveryoption-set | Default: 0  | After the DOS attack when the system is stable, the administrator notices an increase in the number of established connections. The administrator can enable this option, which triggers the recovery algorithm, which keeps track of the established connection count and also whether this count value is greater than the total number of maximum allowed half open connections. If so, it resets the system back to normal and all the DOS related options and counts are reset (Value = 0) |
| tcp      | dosattackautodefenseoption-set   | Default: 0  | Once DOS is detected, the automatic defense mechanism will trigger the DOS defense mechanism without administrator intervention. The defense mechanism when triggered, will release all the half open connection resources at a faster rate than the standard 75 seconds timeout period.  |
| tcp      | dosattackautorecoveryoption-set  | Default: 0  | After the DOS attack when the system is stable, this parameter if set, will automatically trigger the recovery algorithm without administrator intervention.<br><br>The algorithm keeps track of the established connection count and whether this count value is greater than the total number of maximum allowed half open connections. If so, it resets the system back to normal and all the DOS related options and counts are reset.  |
| tcp      | dosfeatureoptionset              | Default: 0  | The DOS option is disabled by default. Enabling this option will monitor the system for any Denial Of Service Attacks.  |
| tcp      | dosmonitortimeoutvalue           | Default: 5  | This parameter enables initializing the DOS Monitor timer value. Initializing this timer value is effective only if the dosfeatureoption is enabled.<br><br>This timer is used to keep track of half open connection counts and also monitor the system for any potential DOS attacks.  |
| tcp      | dossynreceivedtimeoutvalue       | Default: 35 | This parameter enables initializing the syn received timer value. Initializing this timer value is effective only if the dosfeatureoption is enabled.<br><br>When this timer fires off, the TCB entries pertaining to the the half open connections are cleared and the connection is dropped earlier than the 75 seconds timeout value.  |

**Table 9. Parameter facility T to V** (continued)

| Facility | Parameter              | Value                          | Comments/description  |
|----------|------------------------|--------------------------------|---|
| tcp      | fastRTO                | 0 or 1<br>Default: 0           | <p>Determines which TCP timer to use when calculating the retransmission time-out value.</p> <p>0 = Use the TCP slow timer (500 ms). With this timer, the first time-out retransmission occurs in 1 to 1.5 s.</p> <p>1 = Use Fast RTO. The TCP fast timer (200 ms) results in the first timeout to occur in 400 to 600 ms.</p> <p>Changing the retransmission time also affects the time a connection stays open when there is no response from the remote connections.</p> <p>Fast RTO can increase network traffic. Enable this feature with caution and recognize that it cannot improve performance.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/> |
| tcp      | maxhalfopenconnections | Default: 65535                 | <p>Specifies the maximum number of TCP half open connections that a Data Mover can have. The maximum value is 64K, or 65536.</p>  |
| tcp      | maxStreams             | 0-0xffffffff<br>Default: 65535 | <p>Sets the maximum number of TCP connections allowed for the Data Mover.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/>  |
| tcp      | pcbHashLock            | Default: 0                     | <p>This is a flag to enable TCP Secondary PCB Hash locking for MTNT. For releases not yet supporting multi-thread capable networking, this parameter defaults to Off. For multi-thread capable networking, this parameter will default to On.</p> <p>This flag controls the locking protection for the secondary TCPCB cache and must be enabled for multi-thread capable networking to function correctly.</p>   |

Table 9. Parameter facility T to V (continued)

| Facility | Parameter     | Value                       | Comments/description  |
|----------|---------------|-----------------------------|---|
| tcp      | sndcwnd       | 0–0xffffffff<br>Default: 0  | <p>Specifies the maximum TCP transmit window size (in bytes) that the Data Mover can use when sending data to clients, regardless of the receive window size specified by the TCP client.</p> <p>0 = The Data Mover set its maximum TCP transmit window size to the maximum TCP receive window size specified by the TCP client.</p> <p>If the value is set to less than twice the size of the Maximum Transmission Unit (MTU) of the network interface (in bytes), a transmit window of twice the MTU size is used.</p> <p>Use this parameter if switches, routers, or hosts drop packets because they do not have enough buffers to handle the receive window specified by the remote application/client.</p> <p>Use this parameter with care because it affects all TCP connections to the Data Mover, and might severely affect network performance.</p> <p>_____</p> <p>Note: Reboot the Data Mover for changes to take effect.</p> <p>_____</p> |
| tcp      | std_slowstart | 0–1<br>Default: 0           | <p>Controls the RFC-compliant slow start behavior.</p> <p>0 = Increment the snd_cwnd value based on the number of bytes acknowledged by the peer by using ACK (up to a maximum of 3 segments) during slow start.</p> <p>1 = Increment the snd_cwnd value by t_maxseg for each inbound ACK processed.</p>  |
| tcp      | strictsource  | Default: 0                  | This is a temporary parameter to enable or disable the system wide usage of the TCP option strict source.   |
| tftp     | maxretrans    | Default: 3                  | Specifies the TFTP retransmit count.  |
| tftp     | maxthreads    | 0–0xffffffff<br>Default: 32 | <p>Specifies the maximum number of concurrent TFTP transfers.</p> <p>_____</p> <p>Note: Reboot the Data Mover for changes to take effect.</p> <p>_____</p>  |

**Table 9. Parameter facility T to V** (continued)

| Facility | Parameter     | Value                          | Comments/description  |
|----------|---------------|--------------------------------|---|
| tftp     | pathno        | Default: 1                     | Specifies the maximum number of TFTP paths.   |
| tftp     | thrdtimeout   | 0–0xffffffff<br>Default: 3600  | Specifies the maximum number of seconds a TFTP transfer session persists before timing out.<br><br>Note: Reboot the Data Mover for changes to take effect.  |
| tftp     | tidinitial    | Default: 5000                  | Specifies the initial TFTP port seed.   |
| tftp     | tidrange      | Default: 1023                  | Specifies the maximum range of TFTP ports.  |
| tftp     | timeout       | Default: 5                     | Specifies the TFTP select/read timeout.   |
| trunk    | LoadBalance   | ip, mac, or tcp<br>Default: ip | Specifies the default method used for statistics-based load balancing.<br><br>ip = Use the source and destination IPv4 addresses of the outbound packet to select the link. It is effective when remote hosts are on the local LANs, or they are reached through routers.<br><br>mac = Use the source and destination MAC addresses of the outbound packet to select the link. It is effective when the majority of the remote hosts do not go through the routers.<br><br>tcp = Use the source and destination TCP port numbers of the outbound packet to select the link.<br><br>Potentially, different connections on the same remote hosts can select different links.<br><br>Note: Reboot the Data Mover for changes to take effect. |
| ufs      | aclDebugLevel | 0–0xffffffff<br>Default: 0     | Specifies which debug traces appear in the server log for ACL updates.<br><br>0 = No logging.<br>1 = Trace ACL updates.<br>2 = Trace pending operations.<br>4 = Dump ACLs.  |
| ufs      | cgCacheSize   | Default: 196608                | Defines the total CG cache size.  |

Table 9. Parameter facility T to V (continued)

| Facility | Parameter                | Value   | Comments/description   |
|----------|--------------------------|---|--|
| ufs      | dirOffsetHashSize        | Prime number in the range 115933 through 1387189<br><br>Default: 345089 | Defines the size of UxFS directory filename hashq. The value must be a prime number.<br><br>Hint: To obtain the procedure for generating a list of prime numbers in Microsoft Excel, go to the Microsoft Support website.  |
| ufs      | gid32                    | 0 or 1<br><br>Default: 1  | Specifies the type of GID for a newly created file system on the Data Mover. Changing this setting does not affect existing file systems.<br><br>0 = Newly created file systems use 16-bit GIDs.<br>1 = Newly created file systems use 32-bit GIDs.<br><br><i>Managing Volumes and File Systems for VNX Manually</i> describes significant restrictions with this parameter. |
| ufs      | indBlkHashSize           | 1–2048<br><br>Default: 16   | If you have several large files (larger than 200 MB), increasing this value can improve performance.<br><br>Note: Reboot the Data Mover for changes to take effect.  |
| ufs      | inodelimit               | 257949696–4294967295<br><br>Default: 257949696 (0xf6000000)             | Specifies the maximum number of inodes for a new file system or a file system extension.   |
| ufs      | inoHighWaterMark         | Default: 65536  | When the dirty inode buffer count exceeds the inoHighWaterMark value, a filesystem sync is initiated to flush out dirty buffers.   |
| ufs      | maxAclCached             | 2097152–8388608<br><br>Default: 25165824                                | Specifies the maximum number of cached ACL entries for the Data Mover.<br><br>Note: Reboot the Data Mover for changes to take effect.  |
| ufs      | minFSFreeBlksForAclAlloc | 256–1048516<br><br>Default: 1024  | Specifies the minimum number of blocks a file system must have to create a new ACL record.   |

**Table 9. Parameter facility T to V** (continued)

| Facility | Parameter       | Value                      | Comments/description  |
|----------|-----------------|----------------------------|---|
| ufs      | minRsvdAcSlots  | 1-1048516<br>Default: 2048 | Specifies the minimum number of ACL slots to reserve in record and datafiles. The minimum reserved slots ensure that an ACL allocation can complete.  |
| ufs      | skipAcCache     | 0 or 1<br>Default: 0       | Controls whether the ACL cache is initialized.<br><br>0 = Initialize the ACL cache.<br>1 = Do not initialize the ACL cache.<br><br>Changes to this parameter affect only new file system mounts until the Data Mover is rebooted.   |
| ufs      | skipFscck       | 0 or 1<br>Default: 1       | Controls whether fsck is run on a corrupted file system during bootup.<br><br>0 = Run fsck on a corrupted file system while booting up. Bootup time takes longer and all file systems are unavailable.<br><br>1 = Do not run fsck on a corrupted file system while booting up, and the file system is left unmounted. As a result, bootup is faster and downtime for other file systems is minimized. The corrupted file system remains unavailable until fsck is run on it. You must manually mount and export the file system to make it available.     |
| ufs      | skipMapBlock    | 0 or 1<br>Default: 0       | Controls whether MapBlock is run as part of fsck when a media error occurs.<br><br>0 = Allow MapBlock to run with fsck.<br>1 = Disable MapBlock and let fsck continue.  |
| ufs      | skipRootEtcFsck | 0 or 1<br>Default: 0       | Controls whether fsck is run on a corrupted file system during bootup.<br><br>0 = Run fsck on a corrupted file system while booting up. Bootup time takes longer and all file systems are unavailable.<br><br>1 = Do not run fsck run on a corrupted file system while booting up, and the file system is left unmounted. As a result, bootup is faster and downtime for other file systems is minimized. The corrupted file system remains unavailable until fsck is run on it. You must manually mount and export the file system to make it available. |

Table 9. Parameter facility T to V (continued)

| Facility | Parameter         | Value                       | Comments/description   |
|----------|-------------------|-----------------------------|--|
| ufs      | syncwatchdog      | 6–15<br>Default: 6          | Specifies the number of milliseconds allowed for flushing the UFS log. This period should include the time necessary for the system to correct itself. If the log cannot be flushed within the specified period, the Data Mover panics.  |
| ufs      | verifySummary     | 0 or 1<br>Default: 0        | Checks the summary information of the file system to ensure that it is consistent. If it is not consistent, the summary information is corrected before mounting the file system.<br><br>This is a maintenance parameter. Do not change it without direction from EMC.<br><br>0 = Do not verify summary information.<br><br>1 = Verify summary information.  |
| ufs      | xlateIgnoreShadow | Default: 0                  | If set to nonzero, MPDTranslate will ignore invalid Argument errors from getAllNames. Translation will continue without consulting the shadow for synonyms, thus losing all CIFS names.  |
| ufs      | xlateMaxThreads   | 0–0xffffffff<br>Default: 10 | Controls the number of threads allowed to actively translate file systems to the multiprotocol directory (MPD) format.<br><br>Setting the value below the current number of threads does not affect the number of threads, but it puts all but this number to sleep.<br><br>Setting the number to 0 can be used to prevent translation contention during peak periods of server bandwidth demand.<br><br>_____<br>Note: Reboot the Data Mover for changes to take effect.<br>_____ |

**Table 9. Parameter facility T to V** (continued)

| Facility | Parameter             | Value                         | Comments/description   |
|----------|-----------------------|-------------------------------|--|
| ufs      | xlateMinThreads       | 0-0xffffffff<br>Default: 10   | <p>Controls the number of threads used to translate the file system to the multiprotocol directory (MPD) format that will never exit.</p> <p>If the number of threads is too small, new threads cannot be created to handle new translation requests.</p> <p>The parameter can be changed to increase the number of threads available for translation. You cannot decrease the parameter value.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/> |
| ufs      | xlateSpaceHighCeiling | 0-0xffffffff<br>Default: 100  | <p>Sets the percent (in units of 0.01%) of free space available at which translation pauses.</p> <p>100 = 1%</p> <p>Example:</p> <p>150 = Translation pauses when free space drops to 1.5% free.</p>   |
| ufs      | xlateSpaceLowCeiling  | 0-0xffffffff<br>Default: 1000 | <p>Sets the percentage (in units of 0.01%) of free space, needed for translation, to resume after pausing when space drops to xlateSpaceHighCeiling.</p> <p>1000 = 10% free, the point at which events warn that the file system usage is 90%.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/>  |
| ufs      | xlateToInline         | 0 or 1<br>Default: 1          | <p>Controls creation of inline directories during multiprotocol directory (MPD) translation.</p> <p>0 = MPD translation does not create inline directories.</p> <p>1 = MPD translation attempts to create inline directories.</p>  |

Table 9. Parameter facility T to V (continued)

| Facility | Parameter     | Value  | Comments/description   |
|----------|---------------|--|--|
| usrmap   | autobroadcast | 0–0xffffffff<br>Default: 1                             | Enables or disables the autodiscovery mechanism to discover active usermapper if no usermapper is manually set.<br><br>1 = Enable autodiscovery of active usermapper.<br><br>Note: Reboot the Data Mover for changes to take effect. |
| usrmap   | maxgid        | 0–0xffffffff<br>Default:<br>2147483647<br>(0x7fffffff) | Sets the maximum GID value the usermapper on the Data Mover can assign.<br><br>The value of maxgid must be greater than mingid.<br><br>Note: Reboot the Data Mover for changes to take effect.                                       |
| usrmap   | maxuid        | 0–0xffffffff<br>Default:<br>2147483647<br>(0x7fffffff) | Sets the maximum UID value the usermapper on the Data Mover can assign.<br><br>The value of maxuid must be greater than minuid.<br><br>Note: Reboot the Data Mover for changes to take effect.                                       |
| usrmap   | mingid        | 0–0xffffffff<br>Default: 16                            | Sets the minimum GID value the usermapper on the Data Mover can assign.<br><br>The value of mingid must be less than maxgid.<br><br>Note: Reboot the Data Mover for changes to take effect.  |
| usrmap   | minuid        | 0–0xffffffff<br>Default: 16                            | Sets the minimum UID value the usermapper on the Data Mover can assign.<br><br>The value of minuid must be less than maxuid.   |

**Table 9. Parameter facility T to V** (continued)

| Facility | Parameter      | Value                            | Comments/description   |
|----------|----------------|----------------------------------|--|
| vbb      | skipDedupFiles | 0-1<br>Default: 0                | Determines whether a restore from a VBB or NVB backup of a deduplicated file system fails.<br><br>0 = The restore from a VBB or NVB backup of a deduplicated file system fails at the first deduplicated file that it encounters, therefore failing to restore any non-deduplicated files during the restore process.<br><br>1 = Skips all the deduplicated files, therefore restoring any subsequent non-deduplicated files during the restore process. |
| vbb      | tempDir        | Text string<br>Default: .vbbtemp | Specifies the path to use for storing VBB temporary files.<br><br>Examples:<br><br>tempdir = Store temporary files in the file system /tempdir.<br><br>fs/dir = Store temporary files on the file system "fs" in directory "dir".  |
| vdevice  | key            | Text string<br>Default: none     | Specifies the text string to use as the virtual device key.<br><br>Example:<br><br>KEYSTRING   |
| viruschk | chunkQuota     | 0-0xffffffff<br>Default: 65      | Defines the size the virus checker uses to memorize the path of the files not checked during panic.<br><br>The size is a multiple of 32256 bytes.<br><br>65 = 2 MB<br><br>_____<br>Note: Reboot the Data Mover for changes to take effect.<br>_____  |
| viruschk | fsscanburst    | 0 - 1000<br>Default: 20          | Defines the number of files sent in one burst by the scan process to the virus checker engines.<br><br>File system scanning will be suspended until these files are no longer checked.<br><br>0 = The burst will be set dynamically to:<br><br>Number of AV threads * Number of online Common Antivirus Agent (CAVA) servers   |

Table 9. Parameter facility T to V (continued)

| Facility | Parameter | Value                         | Comments/description  |
|----------|-----------|-------------------------------|---|
| viruschk | noRetry   | 0–0xffffffff<br>Default: 0X38 | <p>Defines the errors returned by the CAVA engine that are not followed by a retry. These errors are considered as normal case. This parameter uses the same bits as the audit parameter.</p> <p>0x02 = ERROR_SETUP<br/>0x04 = AV_NOT_FOUND<br/>0x08 = FILE_NOT_FOUND<br/>0x10 = ACCESS_DENIED<br/>0x20 = FAIL<br/>0x20000 = INSUFFICIENT_RESOURCES</p> <p>Example:<br/>0x38 = 0x08 + 0x10 + 0x20</p> |

**Table 9. Parameter facility T to V** (continued)

| Facility | Parameter | Value                             | Comments/description   |
|----------|-----------|-----------------------------------|--|
| viruschk | Notify    | 0, 1, 2, 3, 6, or 7<br>Default: 7 | <p>Controls notifications to users (through Windows pop-up messages) and the admin (through Control Station events) when the CAVA service deletes, modifies, or renames a file because of a virus infection.</p> <p>0 = Send VC file deletion and rename event notifications to Control Station.</p> <p>1 = Send VC file deletion, rename, and modification event notifications to the Control Station.</p> <p>2 = Send VC file deletion and rename event notifications to the Control Station, and to CIFS client in Windows pop-up.</p> <p>3 = Send VC file deletion, rename, and modification event notifications to the Control Station, and send deletion and rename event notifications to CIFS client in Windows pop-up.</p> <p>6 = Send VC file deletion and rename event notifications to the Control Station, and send deletion, rename, and modification event notifications to CIFS client in Windows pop-up.</p> <p>7 = Send VC file deletion, rename, and modification event notifications to the Control Station and to CIFS client in Windows pop-up.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/> |
| viruschk | RPCtype   | Default: 0x00000003               | Defines the RPC type used by virus checking in VNX Operating Environment for File to communicate with the VC servers.  |

Table 9. Parameter facility T to V (continued)

| Facility | Parameter | Value                      | Comments/description  |
|----------|-----------|----------------------------|---|
| viruschk | Traces    | 0-0xffffffff<br>Default: 0 | <p>Defines the traces that appear in the server_log for virus checker:</p> <p>0x00000001 = In CFS: setCheckStatus, setCheckWriter, checkWait</p> <p>0x00000002 = In CIFS: createEvent, sendEvent, mustBeChecked (scan on read)</p> <p>0x00000004 = In Virus Checker:<br/>connectAnyServer, vc_checkfile, stopThreads, exit, start</p> <p>0x00000008 = In CIFS applibnt: open, writeAsyncMsg, readMsg, close, rename</p> <p>0x00000010 = In Virus Checker: heartbeat of the virus checker servers</p> <p>0x40000000 = Warnings</p> <p>0xC0000000 = Warnings and errors</p> |
| viruschk | vnodeHWM  | 0-100<br>Default: 90       | <p>Sets the percentage of total vnodes available in the system pending on virus checking. An event is sent to the Control Station when the maximum is reached.</p> <p>The CIFS thread is blocked until the low water mark is reached again.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/>  |
| viruschk | vnodeLWM  | 0-100<br>Default: 60       | <p>Sets the percentage of total vnodes available in the system that returns the system in a normal state after the high water mark has been reached. An event is sent to the Control Station when the maximum is reached.</p> <hr/> <p>Note: Reboot the Data Mover for changes to take effect.</p> <hr/>  |

**Table 9. Parameter facility T to V** *(continued)*

| Facility | Parameter | Value                           | Comments/description  |
|----------|-----------|---------------------------------|---|
| viruschk | vnodeMax  | 100-0xffffffff<br>Default: 2000 | <p>Defines the maximum number of total vnodes available in the system pending on Virus Checking.</p> <p>An event is sent to the Control Station when the maximum is reached. The CIFS thread is blocked until the low water mark is reached again.</p> <hr/> <p>Note: Restart the virus checker service for changes to take effect.</p> <hr/> |

## VNX System Parameters

Table 10 on page 123 lists the adjustable system parameters. These parameters are defined in the file `/nas/site/nas_param`. System parameters are also defined in `/nas/sys/nas_param`, but that file is overwritten whenever the software is upgraded.

To modify a system parameter, use the procedure [Add or modify a system parameter on page 29](#).

**Table 10. System parameters**

| Parameter | Value                | Comments/description  |
|-----------|----------------------|---|
| ckpt      | 20–99<br>Default: 20 | <p>Specifies the percentage of system space that is allocated to SavVol.</p> <p>The parameter format is:<br/>ckpt:10:100:20:<br/>_____</p> <p>Note: Change only the last value.<br/>_____</p> <p>20 = Percentage of the entire system volume, which is allotted to the creation and extension of all SavVols used by VNX software features.</p> <p>Do not use a value lower than 20% to ensure proper functionality of features that require SavVols.</p> <p>The first value (10) is the Control Station event polling interval rate in seconds. The second value (100) is the maximum rate in MB/s to which a file system is written. Do not change either of these values or system performance problems will result.</p> <p><i>Using VNX SnapSure</i> provides more information about the use of this parameter.</p> |

Table 10. System parameters (continued)

| Parameter                | Value                          | Comments/description  |
|--------------------------|--------------------------------|---|
| clar_rename_luns         | true or false<br>Default: true | <p>Specifies whether to automatically rename VNX for block LUNs that use the default Unisphere (for block) naming convention (LUN &lt;lun-id&gt; or Virtual Disk &lt;lun-id&gt;). LUNs with a user-defined naming convention are not affected.</p> <p>When enabled (true), those LUNs are renamed using the format:</p> <p>Celerra_ &lt;hostname&gt;_&lt;lun-id&gt;_&lt;disk-id&gt;</p> <p>where:</p> <p>&lt;hostname&gt; = hostname of VNX. If the hostname exceeds 24 characters, the VNX ID is used instead.</p> <p>&lt;lun-id&gt; = the LUN number.</p> <p>&lt;disk-id&gt; = the VNX disk volume alias.</p> <p>If the generated name exceeds 64 characters, this value is replaced with the disk ID in the following format:</p> <p>diskID=&lt;disk-id&gt;</p> <p>The tag (NAS/OS) is appended to a VNX reserved LUN.</p> <p>When disabled (false), the default Unisphere (for block) naming convention is retained.</p> <p>If the parameter is specified without a value, the default value (true) is used.</p> <p>Example:</p> <pre>clar_rename_luns:false:</pre> |
| dm_restart_httpd_timeout | 30–60<br>Default: 40           | <p>Specifies the time-out period (in seconds) for communications between the Control Station and Data Movers. This value applies only to those platforms, where changing the HTTPD configuration on the Control Station causes the daemon to be restarted on the Data Mover.</p> <p>If time-out errors occur after modifying the HTTPD configuration, increase this value to allow enough time for the daemons to restart.</p> <p>Example:</p> <pre>dm_restart_httpd_timeout:50:</pre>  |

Table 10. System parameters (continued)

| Parameter         | Value                | Comments/description  |
|-------------------|----------------------|---|
| hidden_interfaces | Network device names | <p>Specifies one or more network devices to be hidden during initialization so that the source and remote standby Data Movers appear to have the same network device configuration. To mask more than one device, type the names as a comma-separated list. The listed devices are hidden for all Data Movers in the system.</p> <p>The specified devices are hidden from all VNX user interfaces, including the CLI and Unisphere.</p> <p>Example:</p> <p>hidden_interfaces:cge5,cge6:</p> <p>The EMC E-Lab™ Interoperability Navigator, available on EMC Online Support, provides detailed information about Data Mover compatibility based on different types of cabinets.</p> <p>Use this tool before trying to resolve incompatibility issues.</p> |

