



NetApp HCI Disaster Recovery with Cleondris

NetApp Solutions

NetApp
March 06, 2021

Table of Contents

- TR-4830: NetApp HCI Disaster Recovery with Cleondris 1
 - Overview of Business Continuity and Disaster Recovery 1
 - Business Impact Assessment 1
 - Application Catalog 1
 - What Not to Protect 1
 - Product Overview 1
 - Installing Cleondris: NetApp HCI DR with Cleondris 2
 - Configuring Cleondris: NetApp HCI DR with Cleondris 4
 - Disaster Recovery Pairing: NetApp HCI DR with Cleondris 5
 - Recovery Planning: NetApp HCI DR with Cleondris 6
 - Failover: NetApp HCI DR with Cleondris 9
 - Best Practices: NetApp HCI DR with Cleondris 10
 - Where to Find Additional Information: NetApp HCI DR with Cleondris 14

TR-4830: NetApp HCI Disaster Recovery with Cleondris

Michael White, NetApp

Overview of Business Continuity and Disaster Recovery

The business continuity and disaster recovery (BCDR) model is about getting people back to work. Disaster recovery focuses on bringing technology, such as an email server, back to life. Business continuity makes it possible for people to access that email server. Disaster recovery alone would mean that the technology is working, but nobody might be using it; BCDR means that people have started using the recovered technology.

Business Impact Assessment

It is hard to know what is required to make a tier 1 application work. It is usually obvious that authentication servers and DNS are important. But is there a database server somewhere too?

This information is critical because you need to package tier 1 applications so that they work in both a test failover and a real failover. An accounting firm can perform a business impact assessment (BIA) to provide you with all the necessary information to successfully protect your applications: for example, determining the required components, the application owner, and the best support person for the application.

Application Catalog

If you do not have a BIA, you can do a version of it yourself: an application catalog. It is often done in a spreadsheet with the following fields: application name, components, requirements, owner, support, support phone number, and sponsor or business application owner. Such a catalog is important and useful in protecting your applications. The help desk can sometimes help with an application catalog; they often have already started one.

What Not to Protect

There are applications that should not be protected. For example, you can easily and cheaply have a domain controller running as a virtual machine (VM) at your disaster recovery site, so there is no need to protect one. In fact, recovering a domain controller can cause issues during recovery. Monitoring software that is used in the production site does not necessarily work in the disaster recovery site if it is recovered there.

It is usually unnecessary to protect applications that can be protected with high availability. High availability is the best possible protection; its failover times are often less than a second. Therefore, disaster recovery orchestration tools should not protect these applications, but high availability can. An example is the software in banks that support ATMs.

You can tell that you need to look at high-availability solutions for an application when an application owner has a 20-second recovery time objective (RTO). That RTO is beyond replication solutions.

Product Overview

The Cleondris HCI Control Center (HCC) adds disaster recovery capabilities to new and existing NetApp HCI deployments. It is fully integrated with the NetApp SolidFire storage engine and can protect any kind of data and applications. When a customer site fails, HCC can be used to recover all data at a secondary NetApp HCI

site, including policy-based VM startup orchestration.

Setting up replication for multiple volumes can be time consuming and error prone when performed manually. HCC can help with its Replication Wizard. The wizard helps set up the replication correctly so that the servers can access the volumes if a disaster occurs. With HCC, the VMware environment can be started on the secondary system in a sandbox without affecting production. The VMs are started in an isolated network and a functional test is possible.

Installing Cleondris: NetApp HCI DR with Cleondris

This section will detail the prerequisites and deployment steps for installing Cleondris.

Prerequisites

There are several things to have ready before you start with the installation.

This technical report assumes that you have your NetApp HCI infrastructure working at both your production site and your disaster recovery site.

- **DNS.** You should have DNS prepared for your HCC disaster recovery tool when you install it.
- **FQDN.** A fully qualified domain name for the disaster recovery tool should be prepared before installation.
- **IP address.** The IP will be part of the FQDN before it is put into DNS.
- **NTP.** You need a Network Time Protocol (NTP) server address. It can be either your own internal or external address, but it needs to be accessible.
- **Storage location.** When you install HCC, you must know which datastore it should be installed to.
- **vCenter Server service account.** You will need to have a service account created in vCenter Server on both the disaster recovery and production side for HCC to use. It does not require administrator-level permissions at the root level. If you like, you can find exactly what is required in the HCC user guide.
- **NetApp HCI service account.** You need a service account in your NetApp HCI storage for both the disaster recovery and production side for HCC to use. Full access is required.
- **Test network.** This network should be connected to all your hosts in the disaster recovery site, and it should be isolated and nonrouting. This network is used to make sure applications work during a test failover. The built-in test network that is temporary only is a one-host network. Therefore, if your test failover has VMs scattered on multiple hosts, they will not be able to communicate. I recommend that you create a distributed port group in the disaster recovery site that spans all hosts but is isolated and nonrouting. Testing is important to success.
- **RTOs.** You should have RTOs approved by management for your application groups. Often it is 1 or 2 hours for tier 1 applications; for tier 4 applications, it can be as long as 12 hours. These decisions must be approved by management because they will determine how quickly things work after a critical outage. These times will determine replication schedules.
- **Application information.** You should know which application you need to protect first, and what it needs to work. For example, Microsoft Exchange needs a domain controller that has a role of Global Catalog to start. In my own experience, a customer said that they had one email server to protect. It did not test well, and when I investigated, I discovered the customer had 24 VMs that were part of the email application.

Download Information

You can download HCC from the [Cleondris site](#). When you buy it, you receive an email with a download link as well.

License

Your license will arrive in an email when you purchase or if you get a not-for-resale (NFR) version. You can get a trial license through the [Cleondris Support Portal](#).

Deployment

You download an OVF file, so it is deployed like many other things.

1. Start by using the Actions menu available at the cluster level.

[Error: Missing Graphic Image]

2. Select the file.

[Error: Missing Graphic Image]

3. Name the appliance and select the location for it in the vCenter infrastructure.

[Error: Missing Graphic Image]

4. Select the Compute location.

5. Confirm the details.

6. Accept the license details.

7. Select the appropriate storage location.

8. Select the network that our appliance will work on.

9. Review the details again and click Finish.

10. Now wait for the appliance to be deployed, and then power it up. As it powers up, you might see a message saying that VMware tools are not installed. You can ignore this message; it will go away automatically.

Initial Configuration

To start the initial configuration, complete the following steps:

1. This phase involves doing the configuration in the Appliance Configurator, which is the VM console. So, after the appliance powers up, change to work in the console by using the VMware Remote Console (VMRC) or the HTML5 VMRC version. Look for a blue Cleondris screen.

[Error: Missing Graphic Image]

2. Press any key to proceed, and configure the following:

- The web administrator password
- The network configuration: IP, DNS, and so on
- The time zone
- NTP

3. Select the Reboot and Activate Network/NTP Settings. You will see the appliance reboot. Afterward, do a ping test to confirm the FQDN and IP.

Patching Cleondris

To update your Cleondris product, complete the following steps:

1. When you first log in to the appliance, you see a screen like the following:

[Error: Missing Graphic Image]

2. Click Choose File to select the update you downloaded from the Cleondris website.

[Error: Missing Graphic Image]

3. Upload the patch. After the appliance reboots, the following login screen is displayed:

[Error: Missing Graphic Image]

4. You can now see the new version and build information; confirming that the update was successful. Now you can continue with the configuration.

Software Used

This technical report uses the following software versions:

- vSphere 6.5 on production
- vSphere 6.7 U3 on DR
- NetApp Element 11.5 on production
- NetApp Element 12.0 on DR
- Cleondris HCC 8.0.2007 Build 20200707-1555 and 8.0.2007X2 build 20200709-1936.

Configuring Cleondris: NetApp HCI DR with Cleondris

You now configure Cleondris to communicate with your vCenter Servers and storage. If you have logged out, returned, and log in again to start here, you are prompted for the following information:

1. Accept the EULA.
2. Copy and paste the license.
3. You are prompted to perform configuration, but skip this step for now. It is better to perform this configuration as detailed later in this paper.
4. When you log back in and see the green boxes, you must change to the Setup area.

Add vCenter Servers

To add the vCenter Servers, complete the following steps:

1. Change to the VMware tab and add your two vCenter Servers. When you are defining them, add a good description and use the Test button.

[Error: Missing Graphic Image]

This example uses an IP address instead of an FQDN. (This FQDN didn't work at first; I later found out that I had not entered the proper DNS information. After correcting the DNS information, the FQDN worked fine.) Also notice the description, which is useful.

2. After both vCenter Servers are done, the screen displays them.

[Error: Missing Graphic Image]

Add NetApp HCI Clusters

To add the NetApp HCI clusters, complete the following steps:

1. Change to the NetApp tab and add your production and disaster recovery storage. Again, add a good description and use the Test button.

[Error: Missing Graphic Image]

2. When you have added your storage and vCenter Servers, change to the Inventory view so that you can see the results of your configuration.

[Error: Missing Graphic Image]

Here you can see the number of objects, which is a good way to confirm that things are working.

Replication

You can use HCC to enable replication between your two sites. This allows us to stay in the HCC UI and decide what volumes to replicate.

Important: If a replicated volume contains VMs that are in two plans, only the first plan that fails over works because it will disable replication on that volume.

I recommend that each tier 1 application have its own volume. Tier 4 applications can all be on one volume, but there should be only one failover plan.

Disaster Recovery Pairing: NetApp HCI DR with Cleondris

1. Display the Failover page.
2. On the diagram of your vCenter Servers and storage, select the Protection tab.

[Error: Missing Graphic Image]

The far side of the screen displays some useful information, such as how many protected VMs you have. (In this example, none right now.) You can also access the Replication Wizard [here](#).

[Error: Missing Graphic Image]

This wizard makes the replication setup easy.

[Error: Missing Graphic Image]

3. You can select the volumes that are important to you, but also make sure that you have the proper vCenter Server selected at the top in the cluster field.

At the far right, you see the pairing type, and only Sync is allowed or supported.

After you click Next, the destination area is displayed.

[Error: Missing Graphic Image]

4. The default information is normally right, but it's still worth checking. Then click Next.

[Error: Missing Graphic Image]

It is important to make sure that the disaster recovery site vCenter Server is displayed and that all hosts are selected. After that is complete, use the Preview button.

5. Next you see a summary. You can click Create DR to set the volume pairing and start replication.

Depending on your settings, replication might take a while. I suggest that you wait overnight.

Recovery Planning: NetApp HCI DR with Cleondris

This section discusses successful failover of applications in a crisis or in a planned migration. It first looks at protecting complex multitier applications, and then simpler applications. You can build disaster recovery plans that are slow or fast, so this section provides examples of the highest-performing plans.

Multitier Applications

1. From the Failover page, select the Plans tab.

[Error: Missing Graphic Image]

2. On the far right is an +Add Failover Group button.

[Error: Missing Graphic Image]

In this example, we called this plan Multi-Tier. We will use the network mapping in the bottom left to change the virtual switch that is in use on production to the one in use on DR.

[Error: Missing Graphic Image]

The previous screenshot shows how you can choose the network switch in production and then in DR, use the Map button to select them, and then use Save. You can have more than one mapping if necessary.

3. To select the VMs to protect, click Add Failover Group.

Because this plan will protect multitier applications, the first group will be for databases.

[Error: Missing Graphic Image]

Notice how this example enables Wait for VMware Tools. This setting is important, because it helps make sure that the applications are running. We used the Add VM button to add VMs that are databases. We didn't enable Unregister Source VMs, because it will slow down the failover. We now use the Add Failover button to protect the applications.

4. Do the same thing for web servers. When that is done, the screen resembles the following example.

[Error: Missing Graphic Image]

The important part of this plan is to get all the databases working; then the applications start, find the databases, and start working. Then the web servers start, and the applications are complete and working. This approach is the fastest way to set up this sort of recovery.

5. Click Save before you continue.

Simple or Mass Applications to Fail Over

The order in which the VMs start is important, so that they work; that is what the previous section accomplished. Now we will fail over a set of VMs for which order is unimportant.

Let's create a new failover plan, with one failover group that has several VMs. We still need to do the network mapping.

[Error: Missing Graphic Image]

Notice that there are several VMs in this plan. They will also start at different times, but that is OK because they are not related to each other.

Planned Migration

Planned migration is similar to a disaster recovery failover, but because it is not a disaster recovery situation, it can be handled slightly differently. It is still good to practice the planned migration, but you can add something to your failover group: You can unregister the VM from the source. That takes a little more time, but in a planned migration that is not a bad thing.

A planned migration is usually a move to a new data center. Sometimes it is also used if destructive weather is approaching but has not yet arrived.

Plan of Plans

With a plan of plans, you can trigger one plan and it will take care of all the failover plans.

The Plans tab contains a Plan of Plans section. You can use the +Add Sub-Plan to start a plan and add other plans to it.

[Error: Missing Graphic Image]

In this example, the plan of plans is called Master Plan, and we added the two plans to it. Now when we execute a failover, or test failover, we will have the option for the Master Plan too.

This approach is good because it is best to test your application failovers in their own plan. Each plan is much easier to troubleshoot and fix, and when it is working well, you add it to your master plan.

Script Support

You can use scripts as part of a test failover or for a wide range of other purposes. Uses include the following:

- Turning on anti-spam hardware
- Turning on security hardware

- Populating signage
- Updating IPAM hardware
- Changing the language settings in a database

If you edit your plan and then edit your failover group, you will see entries under Scripts.

[Error: Missing Graphic Image]

In the following screenshot, the word Host refers to the VM that executes scripts. Click the edit button to see the Edit Script window:

[Error: Missing Graphic Image]

You should make sure to test your script before you copy and paste it into this dialog box. You should also select Post in the Group Order field. Make sure to use the right credentials.

If you follow the execution, the following screenshot indicates that the script ran successfully.

[Error: Missing Graphic Image]

If the exit code is anything other than 0, then the script was not successful.

Script Troubleshooting

If a script does not execute properly, then check the following issues:

- VMware Tools only let one external process run at a time. Therefore, if VMware Tools is updating itself, then the script will not execute. This can occur if you set your VMs to automatically upgrade VMware Tools. This is done in VM settings > VM Options > VMware Tools.
- Check for credentials issues.
- Check for script issues, such as a prompt or other functionality that requires human input.

It is a best practice to run simple scripts that only perform essential tasks. You might also want to include a log file for troubleshooting purposes.

Environment Variables

Environmental variables allow a running script to pull information from the environment whether the script is running at the production site or a DR site. Environment variables can be entered in Edit Failover Group dialog box. You can first edit your plan and then edit your failover group.

[Error: Missing Graphic Image]

Note that these environment variables are not in the environment that we normally think of, and you cannot use the set command to see them. To see the full list of variables, run the script from the following screenshot. This script contains `Get-Variable * > c:\utils\var_log.txt` to capture all variables.

[Error: Missing Graphic Image]

This lists the 50+ variables available plus any variable that you have added, which are seen at the end of the list.

Failover: NetApp HCI DR with Cleondris

Test Failover

A test failover is important, because it proves to you, your application owner, your manager, and the BCDR people that your disaster recovery plan works.

To test failover, complete the following steps:

1. From the Failover page, click Start Failover.

[Error: Missing Graphic Image]

2. On the Failover page, you have some choices to make.

[Error: Missing Graphic Image]

Carefully specify the plan, where the VMs came from, and where they are going to be recovered.

[Error: Missing Graphic Image]

The screen displays a list of the VMs that are in the plan. In this example, a warning at the top right says that three VMs are not included. That means there are three VMs we did not make part of the plan in the replicated volume.

If you see a red X in the first column on the left, you can click it and learn what the problem is.

3. At the bottom right of the screen, you must choose whether to test the failover (Failover to Sandbox) or start a real failover. In this example, we select Failover to Sandbox.

[Error: Missing Graphic Image]

4. A summary now lists plans in action. For more information, use the magnifying glass in the far left (described in "Monitoring," later in this document).

Running Failover

At first, the failover is the same as the test failover. But the procedure changes when you arrive at the point shown here:

1. Instead of selecting the Failover to Sandbox option, select Start.

[Error: Missing Graphic Image]

2. Select Yes.

[Error: Missing Graphic Image]

3. The screen shows that this is a failover, and it is running. For more information, use the magnifying glass (discussed in the "Monitoring" section).

[Error: Missing Graphic Image]

Monitoring During a Failover

1. When a failover or a test failover is running, you can monitor it by using the magnifying glass at the far right.

[Error: Missing Graphic Image]

2. Click the magnifying glass to see much more detail.

[Error: Missing Graphic Image]

3. As the failover or test failover progresses, a VM Screenshots option appears.

[Error: Missing Graphic Image]

Sometimes it is useful to see the screenshots to confirm that the VM is running. It is not logged in, so you cannot tell if the applications are running, but at least you know that the VM is.

Looking at History When No Failover Is Running

To view past tests or failovers, click the Show Historical button on the Activity tab. Use the magnifying glass for more detail.

[Error: Missing Graphic Image]

[Error: Missing Graphic Image]

You can also download a report with the details.

[Error: Missing Graphic Image]

These reports have various uses: for example, to prove to an application owner that you tested the failover of that application. Also, the report can provide details that might help you troubleshoot a failed failover.

You can add text to a report by adding the text to the plan in the comment field.

[Error: Missing Graphic Image]

Best Practices: NetApp HCI DR with Cleondris

Recommendations for Success

The following tips can help you be more successful with your BCDR work.

Applications

Know your applications and what makes them work. The more time you spend on them, the more successful you will be with your real and test failovers. When there are issues, you will be able to solve them faster.

Protect one application first. Choose a relatively simple one, and demo the test failover to your peers and management. The demonstration will help you with management and peer support, and the test will help you learn more before you protect other applications.

Your tier 1 applications should be on their own volume.

Practice

You need to practice often in as realistic a scenario as possible. For example, practice off-site, sometimes with poor internet in a hotel conference room. Practicing often is key, and try changing the teams around so that application team X is recovering application Y; this approach will help with knowledge sharing.

Executive Sponsor

Make sure to have an executive sponsor. You'll need executive support when teams are not working well together, or when you need application teams to be reasonable about recovery time.

Plan for Partial or Full Outage

Most disaster recovery events are partial ones, so make sure your tier 1 applications can be recovered without having to recover everything.

Trigger Time

Practice the failovers, but also practice managing others who are authorized to trigger a failover. They need to practice, and they need to know how successful or unsuccessful the failovers are. Make sure they practice with you in as realistic a scenario as possible. You can do a sand-table-type exercise in which operations people bring up issues and managers discuss their response.

Why Does Disaster Recovery Fail?

There are several possible reasons for a disaster recovery plan failing:

- BCDR is needed.
- Attitude is missing: People do not care as much as they should.
- The executive sponsor is missing or not assigned.
- There isn't enough practice, or it isn't real enough.
- Data from the test gets into the product. This situation is serious and must be avoided.

Additional Uses for Disaster Recovery Orchestration Tools

Over time, customers have found other uses for disaster recovery orchestration tools. For example, they test application and OS upgrades in a test failover. This testing is better than testing in a lab, because it uses the actual production bits—which means that, when done in production, the process will be as smooth as in the test failover. I have also seen security vulnerability testing done as a test failover first to determine what applications might be negatively affected.

Active-Active Site

Currently, to protect an active-active site, you must install HCC on both sites and protect as normal. There is currently no overview of the protection. Active-active is the best model, because you can split your applications over two sites; when there is an outage, you only need to fail over half.

Allowing Extra Resources in Test Failover

Sometimes it is necessary to have more resources in the test failover so that a proper application test can occur. For example, these resources might consist of things like physical anti-spam appliances or load balancers. You can also include things like databases, which has the potential to cause problems, because you

must make sure test data does not get into production. To perform this process reliably, use the following steps as a guideline.

1. A script executes in the disaster recovery test process (or use a manual process if necessary).
2. A separate logical partition (LPAR) is created.
3. A virtual network is added to the separate LPAR, and it is already connected to the test network.
4. A script exports and copies the appropriate data to the separate and new LPAR. It's likely that you'll need to have the application on the separate partition, too.
5. You might need to tweak DNS names or the configuration of the application in the test network to access this new server.
6. The test completes successfully.
7. After the test is done, and the cleanup occurs, another script runs, and it deletes the separate partition. That step keeps anything from getting into production accidentally.

You can use a similar process to get a domain controller into a test failover:

1. Power off the domain controller in the disaster recovery site. Make sure there is another domain controller still running.
2. After the domain controller is off, clone it.
3. Power on the original domain controller.
4. Put the cloned domain controller on the test network.
5. Power on the clone domain controller.
6. You should be able to use the domain controller in the test now, whether for authentication or DNS.
7. When the test is done, delete the cloned domain controller. Don't skip this step, because you don't want that domain database talking to the production domain.

It's best to script these steps and execute the script from the recovery plan. However, to do that, you need a script or batch file that can tell whether it is executing in test or real failover—and in real failover, it does nothing.

Syslog

It is useful to capture events from Cleondris by using syslog. Groups such as security or operations might benefit.

1. To do this, use the Setup page and the Events tab. Then use the Add Receiver button.

[Error: Missing Graphic Image]

1. Specify which event to send. In this example, the best idea might be to send all of them for now. Select the boxes; some do not apply to Cleondris HCC and BCDR, but they will not be generated if not used.

You can see the BCDR events in the Events section at the bottom of the list.

[Error: Missing Graphic Image]

VM State

The VM state is preserved during a failover. A VM that is powered on or off in production remains in the same

state after a failover or during a test failover. However, be aware that HCC scans vCenter every 20 minutes. Therefore, you need to wait for that scan or use the refresh button in HCC to immediately refresh.

[Error: Missing Graphic Image]

Add an Execute-Only Account

An execute-only account can be useful for a manager to trigger a failover without saving the changes. You create this account yourself.

First, create a role that has the following privileges:

- Login
- Inventory_sf_view
- Inventory_vc_view
- Restore_exec_sf_failover
- Failover_view
- Failover_job_modify
- Failover_config_view

When the role is done, create a user with that role; the resulting account is an execute-only account. This set of privileges lets the user look at and change things but not save the changes.

Idle Time Out

This parameter can be set to perform an automatic log out when there is no activity in the browser. Working on a different tab counts as activity.

Select the Setup option and then select the Advanced tab to see the Advanced Configuration window.

[Error: Missing Graphic Image]

Click the Add Option button to add the option and value. In the screenshot above, 360 seconds must pass before a timeout if there is no activity in the browser.

Inventory Rescan

The inventory rescan setting is used when a VM state is not preserved when it should be. For example, a VM should not be powered on in a failover if it is off in production. The value for the rescan interval can be set between 5 minutes and 1440 minutes; it is set to 20 minutes by default.

[Error: Missing Graphic Image]

In the previous screenshot, the interval is set for 10 minutes.

Be aware that this setting changes the vCenter rescan time and also the Solidfire rescan time.

General Support

The following best practices can improve your experience with Cleondris and assist with support.

- Always include a support bundle when you ask for support.

[Error: Missing Graphic Image]

- With certain edge cases, additional logging is very helpful for support. Enable the additional logging, and then perform the action that you are having trouble with again. You can then delete `log.level` because you do not want to routinely debug this level.

[Error: Missing Graphic Image]

- A busy vCenter Server Appliance (VCSA) can cause issues under some conditions. To minimize this problem, add more memory to the VCSA.
- Issues can also be caused by the fact that one or two VMs might not be cleaned up in a test failover. You can clean these VMs up with the following steps:
 - Power off the VMs. This may take some time.
 - Remove the VMs from inventory.
Often, these two steps allow the datastore to disappear. You can then perform a Rescan Storage operation.

Where to Find Additional Information: NetApp HCI DR with Cleondris

To learn more about the information that is described in this document, review the following websites:

- NetApp HCI Documentation Center
<https://docs.netapp.com/hci/index.jsp>
- NetApp HCI Documentation Resources page
<https://www.netapp.com/us/documentation/hci.aspx>
- NetApp Product Documentation
<https://www.netapp.com/us/documentation/index.aspx>
- Cleondris HCC product page
<https://www.cleondris.com/en/hci-control-center.xhtml>
- Cleondris Support Portal
<https://support.cleondris.com/>

Copyright Information

Copyright © 2021 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means-graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system-without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

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

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

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

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.