



Best Practices: NetApp HCI DR with Cleondris

NetApp Solutions

Kevin Hoke
May 19, 2021

Table of Contents

- Best Practices: NetApp HCI DR with Cleondris 1
 - Recommendations for Success 1
 - Additional Uses for Disaster Recovery Orchestration Tools 2
 - Active-Active Site 2
 - Allowing Extra Resources in Test Failover 2
 - Syslog 3
 - VM State 3
 - Add an Execute-Only Account 4
 - Idle Time Out 4
 - Inventory Rescan 5
 - General Support 5

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.

Edit Event Receiver

Type: SYSLOG

Hostname: 10.193.136.33

Message Format: DEFAULT

Filter: *

Buttons: Send Test Event, Save, Cancel

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.

CDM-09670	Default	User creates BCDR plan	User %(u) creates BCDR plan %(s)
CDM-09671	Default	User updates BCDR plan	User %(u) updates BCDR plan %(s)
CDM-09672	Default	User deletes BCDR plan	User %(u) deletes BCDR plan %(s)
CDM-09680	Default	User executes BCDR plan	User %(u) executes BCDR plan %(s)
CDM-09681	Default	User tests BCDR plan	User %(u) tests BCDR plan %(s)

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.

vCenter		Hosts	VMs	
✓	sfps-megatron-vcsa.rtp.openenglab.netapp.com	5	133	
✓	sfps-primus-vcsa.rtp.openenglab.netapp.com	2	6	

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.

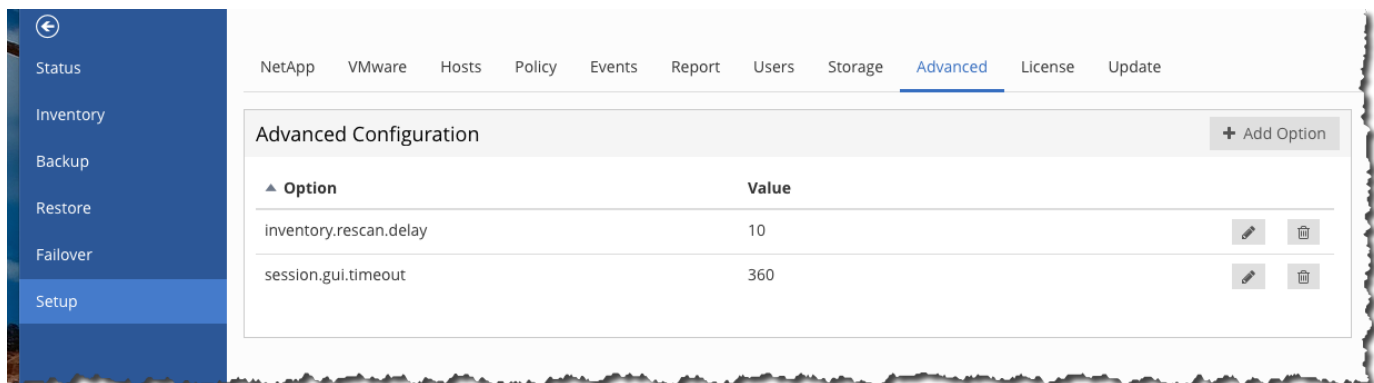
Select the Setup option and then select the Advanced tab to see the Advanced Configuration window.

Advanced Configuration		+ Add Option	
Option	Value		
session.gui.timeout	360		

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.



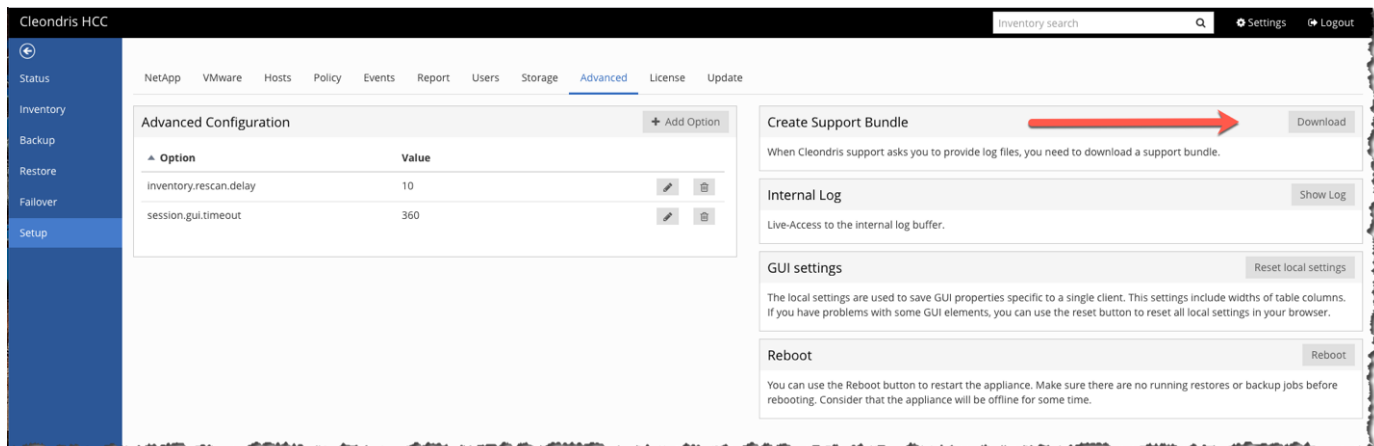
In the previous screenshot, the interval is set for 10 minutes.

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.



- 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.

Advanced Configuration + Add Option

▲ Option	Value		
inventory.rescan.delay	10		
log.level	debug		
session.gui.timeout	360		

- 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.

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.