NetApp Post Installation Checklist

From UptimeWiki

Jump to: navigation, search

Contents

[hide]

Administrativa

[edit]

Contact Information

[edit]

This is only required for ASP customers!

Description	OK

Get correct contact information from Customer:
• IT Manager
  ◦ Email Address
  ◦ Phone Number
• System Administrator(s)
  ◦ Email address
  ◦ Phone Number
• Site Address (for every filer)	

Hand over a ServiceDesk contact card (contact info: servicedesk@uptime.be, tel. +32 (0)3 451 23 74, fax +32 (0)3 451 23 79)	
Inform *nagiosadmins@uptime.be* and ServiceDesk of the new machine(s). Send a test autosupport mail (see later)	

[edit]

NOW Site

Description	OK
Create customer account on NOW site: http://now.netapp.com/	
Link customer account to system IDs of customer filers (linking the customer ID to only one system should be sufficient; check if all systems are visible from the same account afterwards. In case of problems, contact Tania Dermul or Axel Breens from NetApp	
Are we ASP for this customer ? (Check with Roger Sels or Tania Dermul/Axel Breens from NetApp) • If so, make sure machine(s) show(s) up under our NOW account ("uptime-now")	

[edit]

Licenses

Description	OK

Check if licenses are present on filer(s) and nearstore(s)
`filer> license`
`cifs .......`
`cluster .......`
`cluster_remote not licensed`
`...`	

Check with customer's account on NOW site. The same licenses should be present from this link: https://now.netapp.com/eservice/agree.do Also make sure the necessary software licenses (SnapDrive, SnapManager, DFM, ...) are available from the NOW site for the customer	

[edit]

Upgrade to Latest Versions

Description	OK

Upgrade to correct motherboard firmware and diagnostics for this system. Check firmware version with the commands:
`filer> sysconfig -a`
`NetApp Release 7.1.1: Sun Jun 25 03:57:55 PDT 2006`
`System ID: 0084181441 (na-demo01); partner ID: 0084186271 (na-demo02)`
`System Serial Number: 2018924 (na-demo01)`
`System Rev: E1`
`slot 0: System Board 650 MHz (TSANTSA D0)`
`Model Name: FAS270`	

```
                Part Number:        110-00046
                Revision:           D0
                Serial Number:      295361
                Firmware release:   CFE 1.2.0
...
filer> version -b
1:/mips/kernel/primary.krn: OS 7.1.1
1:/backup/mips/kernel/primary.krn: OS 7.1
1:/mips/diag/diag.krn:  Diagnostic_4.6.6
1:/mips/firmware/tsantsa/firmware.img: Firmware 1.2.0
```

Check against the latest firmwares on this page: http://now.netapp.com/NOW/download/tools/serviceimage/. The upgrade procedure is briefly explained here (check official NetApp documentation): NetApp How to Upgrade Firmwares#System/Motherboard Firmware and Diagnostics

Upgrade to correct disk firmwares for this system. The firmware version can be checked with the command:

```
filer> sysconfig -a
...
        slot 0: FC Host Adapter 0b (Dual-channel, QLogic 2312 rev. 2, 32-bit, L-port, <UP>)
                Firmware rev:  3.3.220
                Host Loop Id:  7       FC Node Name:   5:00a:098000:001e80
                Cacheline size:        8       FC Packet size: 2048
                SRAM parity:   Yes     External GBIC:  No
                Link Data Rate:        1 Gbit
                 21: NETAPP   X270_HJURE036F10 NA14  34.0GB 520B/sect (433K8115)
                 22: NETAPP   X270_HJURE036F10 NA14  34.0GB 520B/sect (433K8087)
                 23: NETAPP   X270_HJURE036F10 NA14  34.0GB 520B/sect (433F3696)
                 24: NETAPP   X270_HJURE036F10 NA14  34.0GB 520B/sect (433K7917)
                 25: NETAPP   X270_SCHT6036F10 NA08  34.0GB 520B/sect (3JA23XLL000073529CE5)
                 26: NETAPP   X270_HJURE036F10 NA14  34.0GB 520B/sect (433K8057)
                 27: NETAPP   X270_HJURE036F10 NA14  34.0GB 520B/sect (43358015)
                 28: NETAPP   X270_SCHT6036F10 NA08  34.0GB 520B/sect (3JA24336000073528MLT)
                 29: NETAPP   X270_HJURE036F10 NA14  34.0GB 520B/sect (433K7941)
                 16: NETAPP   X270_HJURE036F10 NA14  34.0GB 520B/sect (433K3477)
                 17: NETAPP   X270_HJURE036F10 NA14  34.0GB 520B/sect (433K0278)
                 18: NETAPP   X270_HJURE036F10 NA14  34.0GB 520B/sect (433J6350)
                 19: NETAPP   X270_HJURE036F10 NA14  34.0GB 520B/sect (433K0174)
                 20: NETAPP   X270_HJURE036F10 NA14  34.0GB 520B/sect (433K4383)
                Shelf 1: EFH  (DS14-Mk2-FC Based Filer)
                       I/O base 0xfe00, size 0x100
                       memory mapped I/O base 0x41000000, size 0x1000
...
```

Look up the latest firmwares for the different disktypes via this link: http://now.netapp.com/NOW/download/tools/diskfw/. The upgrade procedure is briefly explained here (check official NetApp documentation): NetApp How to Upgrade Firmwares#Disk Firmwares

Upgrade to correct shelf firmwares for this system (can be done non-disruptively for FCAL shelves, but disruptively for (S)ATA shelves). Check firmware version with the command:

```
filer> sysconfig -a
...
       slot 0: FC Host Adapter 0a (Dual-channel, QLogic 2322 rev. 3, 64-bit, L-port, <UP>)
               Firmware rev:   3.3.10
               Host Loop Id:   7       FC Node Name:   5:00a:098200:00abce
               Cacheline size: 16      FC Packet size: 2048
               SRAM parity:    Yes     External GBIC:  No
               Link Data Rate: 2 Gbit
                19: NETAPP   X274_HPYTA146F10 NA02 136.0GB 520B/sect (V5WGK0HA)
...
               Shelf 1: ESH2  Firmware rev. ESH A: 16  ESH B: 16
               Shelf 2: ESH2  Firmware rev. ESH A: 16  ESH B: 16
                       I/O base 0xde00, size 0x100
                       memory mapped I/O base 0xa1740000, size 0x1000
```

or with the command:

```
filer> environment status
...
       Channel: 0c
       Shelf: 2
       SES device path: local access: 0c.32
       Module type: ESH2; monitoring is active
       Shelf status: normal condition
       SES Configuration, via loop id 32 in shelf 2:
        logical identifier=0x50050cc00211ef4e
        vendor identification=XYRATEX
        product identification=DS14-Mk2-FC
        product revision level=1616
```

The latest firmwares for the different shelves can be found here: http://now.netapp.com/NOW/download/tools/diskshelf/. The upgrade procedure is briefly explained here (check official NetApp documentation): NetApp How to

Upgrade Firmwares#Shelf Firmwares	

Upgrade to correct Data ONTAP version for this system. The current release can be checked with the commands:
 	
`filer> sysconfig`
`NetApp Release 7.1.1: Sun Jun 25 03:57:55 PDT 2006`
`...`
`filer> version -b`
`1:/mips/kernel/primary.krn: OS 7.1.1`
`1:/backup/mips/kernel/primary.krn: OS 7.1`
`1:/mips/diag/diag.krn: Diagnostic_4.6.6`
`1:/mips/firmware/tsantsa/firmware.img: Firmware 1.2.0`
 	
Select your NetApp filer type and desired DATA ONTAP version from this link: http://now.netapp.com/NOW/cgi-bin/software	

Install the documentation on the system. Download the documentation from the link: http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml Select you DATA ONTAP version, download the .zip or .tar file, then go to the filer via http (http://filer/na_admin/) and upload & install the documentation	

See NetApp How to Upgrade Firmwares for more information.

Data ONTAP Upgrade Cautions [edit]

http://now.netapp.com/NOW/knowledge/docs/ontap/rel7121/html/ontap/upgrade/2upgra13.htm#1327835

Disks and Aggregates [edit]

General comments [edit]

Description	OK
For a discussion on how to add different-sized disks to an existing aggregate, see this KB article: https://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb21090	

Run the **sysconfig -c** command and check for possible errors:
 	
`filer> sysconfig -c`
`sysconfig: There are no configuration errors.`	

Don't mix SATA disks with FC disks on the same loop
 	
`filer> sysconfig -a`
`...`	

Don't mix disks of different sizes in the same aggregate or traditional volume:
 	
`filer> sysconfig -r`
`...`	

[edit]

Software based disk ownership systems: FAS250/270(c) and all New Filers

Description	OK
The following Filer models require software ownership to work properly: • FAS250 • FAS270 • FAS270c • FAS20x0 (all models) • FAS3040 • FAS3070 • FAS60x0 (all models) Make sure all disks are visible:	

```
filer> sysconfig -r
...
```

If not, use the following command to assign disk ownership to a filer:

```
filer> disk assign -s unowned
...
```

[edit]

RAID group sizes

Description	OK

Use 14 disks in a RAID group for (S)ATA disks, use 16 for FCP & SAS disks (14 disks is acceptable):
`filer> aggr status -v`
`Aggr State Status Options`
`aggr0 online raid_dp, aggr root, raidsize=14`
(Use **vol status** for traditional volumes)
Use **aggr options** *aggregate_name* **raidsize 16** to change this	

Make sure the aggregates use **double parity** for all RAID groups:
`filer> aggr status`
`Aggr State Status Options`
`aggr0 online raid_dp, aggr root, raidsize=14`
(Use **vol status** for traditional volumes)
Use **aggr options** *aggregate_name* **raidtype raid_dp** to change this	

[edit]

Spare Disks

Description	OK
Make sure the system has enough spare disks. A typical recommendation: • One shelf = one spare disk • Two shelves or more = two spare disks • More than four shelves = provide extra spare disks • Note that (S)ATA disks are generally more prone to errors • Note that bigger disks take a longer time to reconstruct • Take into account the environmental conditions of the server room	
If disks of different sizes are used, or if both FC and (S)ATA disks are used, keep enough spare disks of every kind	

When adding disks to an aggregate:
• Don't wait until the aggregate is 90+% full, this will lead to fragmenation and overuse of the newly added disk. Performance will suffer
• Add more than one disk at a time
• Perform a volume defragment of all flexible volumes on an aggregate after having expanded an aggregate:
`filer> reallocate start -f /vol/vol0`
`...`
`filer> reallocate start -f /vol/data_vol`
`...`
You may even consider running the reallocation jobs more than once. Check the fragmentation levels afterwards:
`filer> priv set advanced`
`filer*> wafl scan measure_layout vol0`
`...`
`filer*> priv set`
`filer>`	

Discuss the **raid.timeout** option with the customer. If needed, change it. Default = 24 hours. When the system is degraded (=no more spares and enough failed disks so that one additional disk failure could lead to data loss), the filer will shutdown after 24 hours. Some customers may want to raise this to 72 hours ...	

To change this, use the **options** command:

```
filer> options raid.timeout 72
```

[edit]

Logfiles/Web Interface Autoindexing

Description	OK
Set this option: `filer> options httpd.autoindex.enable on` When this option is enabled, one can browse to this URL: http://filer/na_admin/logs/ to get a list of the contents of the /etc/log/ directory of the root volume.	

[edit]

Volumes

Root Volume (vol0) Size

[edit]

Description	OK

The root volume is typically too big on newly shipped filers - set its size to 20 GB for smaller models and 50 GB for bigger models:
 	
`filer> vol size vol0 20g`
 	
Check with:
 	
`filer> vol size vol0`
`vol size: Flexible volume 'vol0' has size 20g.`
`filer> df vol0`
`Filesystem kbytes used avail capacity Mounted on`
`/vol/vol0/ 16777216 726932 16050284 4% /vol/vol0/`
`/vol/vol0/.snapshot 4194304 13272 4181032 0% /vol/vol0/.snapshot`	

[edit]

Volume Unicode Support

Description	OK

Set create_ucode and convert_ucode to on for all volumes (especially vol0, as newly created volumes will inherit the same options as vol0):
 	
`filer> vol options vol0 create_ucode on`
`filer> vol options vol0 convert_ucode on`
`...`
 	
(Repeat for all volumes)
 	
Check with:
 	
`filer> vol options vol0`
`root, diskroot, nosnap=off, nosnapdir=off, minra=off,`
`no_atime_update=off, nvfail=off, snapmirrored=off, create_ucode=on,`
`convert_ucode=on, maxdirsize=10470, fs_size_fixed=off,`
`guarantee=volume, svo_enable=off, svo_checksum=off, svo_allow_rman=off,`
`svo_reject_errors=off, no_i2p=off, fractional_reserve=100, extent=off,`
`try_first=volume_grow`
`...`
 	
(Repeat for all volumes)	

[edit]

Security style

Description	OK

Set the correct security style for new volume creation (*ntfs* or *unix*), especially on vol0 To set it correctly for newly created volumes, use this command:
`filer> options wafl.default_security_style ntfs`
or
`filer> options wafl_default_security_style unix`
To change a volume's or qtree's security style, use this command:
`filer> qtree security /vol/vol_name ntfs`
(or *unix*)	

[edit]

Snapshots

Description	OK
Delete any manually created snapshots (that you may have created on the filer during the training) before leaving	
Set a proper snapshot schedule for all volumes: • Disable snapshotting if not needed for particular volumes • Set correct schedules • ~snapshot/.snapshot directories visible for all volumes ? `filer> vol options vol_name nosnapdir off` (See also CIFS section)	

For snapshots scheduled from the filer, The snapshot names hourly.0, hourly.1, hourly.2, ... may confuse users who are browsing the ~snapshot directory. You may want to prefer a naming scheme that includes the snapshot creation time in the snapshot name. Do this by setting this volume option:
`filer> vol options volumename schedsnapname create_time`
Set it back to "old-school" naming via:
`filer> vol options volumename schedsnapname ordinal`	

[edit]

LUNs

Space Reservation

[edit]

Description	OK

Explain the importance of space reservation with LUNs to the customer:
• LUNs must have space reservation enabled:
`filer> lun set reservation /vol/vol_name/lun_name enabled`
• Volumes must have guarantee set to *volume* or *file*
`filer> vol options vol_name guarantee volume`
• Volumes must have the *fractional_reserve* option to 100%
`filer> vol options vol_name fractional_reserve 100`	

Notes:

- There is also a "file reservation" command: eg. you want to protect a large database file on an NFS mounted volume:
  - Set volume reservation to file
  - Use the "file reservation" command to enable reservation for this particular file

Avoiding Volume Full Warnings on Non-snapshot Volumes

[edit]

Description	OK
When you don't intent to take volume snapshots and create a maximum-sized LUN inside a volume, the web interface will complain about the volume being full. This is annoying, as the green status light is now gone. See this KB article on how to solve this: https://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb10170	

[edit]

LUN Configuration Check

Description	OK

Run a LUN configuration check cia the **lun config_check -v** command:
 	
`filer> lun config_check -v`
`Checking igroup ostype & fcp cfmode compatibility`
`==========`
`No Problems Found`
 	
`Checking local and partner cfmode`
`==========`
`No Problems Found`
 	
`Checking for down fcp interfaces`
`==========`
`No Problems Found`
 	
`Checking initiators with mixed/incompatible settings`
`==========`
`No Problems Found`	

[edit]

SnapDrive

Description	OK
Install the necessary MS patches before installing SnapDrive	
Disable SnapDrive Monitoring if needed: • Even when no snapshots are taken, SnapDrive will, by default every 60 minutes, monitor reserved space and trigger autosupport mails • Option 1: MMC, Computer Management, right-click on "SnapDrive", choose "Notification Settings ..." Image:Snapdrive notification settings.jpg SnapDrive Monitor Notification Settings (1/2) Image:Snapdrive notification settings dialog.jpg SnapDrive Monitor Notification Settings (2/2) • Option 2: MMC, Computer Management, SnapDrive, Disks, right-click, choose "Properties", "Virtual Disk Monitor" tab, set the interval to 0 Image:Snapdrive disk properties.jpg SnapDrive Monitor Notification Settings (1/2) Image:Snapdrive disk properties vdm.jpg SnapDrive Monitor	

Notification Settings (2/2)	
Set preferred addresses for target IP and initiator IP address: • MMC, Computer Management, SnapDrive, Disks, right-click, choose "Properties", "Preferred Filer IP Addresses" tab Image:Snapdrive disk properties.jpg SnapDrive Monitor Notification Settings (1/2) Image:Snapdrive disk properties preferred ip.jpg SnapDrive Monitor Notification Settings (2/2)	

[edit]

Scheduled LUN Defragmentation

Description	OK

Set up scheduled LUN defragmentation jobs. Use the default values:
 	
• Perform a fragmentation scan for a LUN every 24 hours
• Start a defragmentation if the fragmentation level exceeds 4
 	
For every LUN, execute:
 	
`filer> reallocate start /vol/vol_name/lun_name`
`Mon Jul 24 00:08:41 CEST [filer: wafl.scan.start:info]: Starting WAFL layout measurement on volume vol_name.`
`Reallocation scan will be started on /vol/vol_name/lun_name.`
`Monitor the system log for results.`
 	
(Repeat this for every LUN on the system)
 	
Check with this command:
 	
`filer> reallocate status`
`Reallocation scans are on`
`/vol/vol_name/lun_name:`
`State: Idle`
`Schedule: n/a`
`Interval: 1 day`
`Optimization: 2`
`...`
 	
Note: Make sure you have enough volume space to perform the defragmentations	

[edit]

Host Utilities (Support Kits)

Description	OK
Install the host support kit to set proper I/O timeout values (see NetApp Best Practices for SAN Management)	

[edit]

Set Correct Timeout Registry Settings

Description	OK
The set_tunables.exe executable from the Host Support Kit (Windows) currently contains at least one value that is set incorrectly (HKLM\SYSTEM\CurrentControlSet\Services\Disk\TimeOutValue - the docs mention different values, MUST BE SET TO 190 (seconds)). Make sure the registry values for Windows hosts are set to these values: • For iSCSI: ◦ HKLM\SYSTEM\CurrentControlSet\Control\Class*{identifier}**instance*\Parameters	

\MaxRequestHoldTime - should be 120 seconds
    - *identifier* is the key which has the default name value "SCSI and RAID controllers".
    - *instance* is the controller instance which has the DriverDesc value "Microsoft iSCSI Initiator".
  - HKLM\SYSTEM\CurrentControlSet\Control\Class*{identifier}**instance*\Parameters\LinkDownTime - should be 5 seconds
  - HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk\TimeOutValue - MUST BE 190 seconds
  - HKLM\SYSTEM\CurrentControlSet\Services\mpio\Parameters\PathVerifyEnabled - should be 0
  - HKLM\SYSTEM\CurrentControlSet\Services\vnetapp\Parameters\PathVerifyEnabled - should be 0
  - HKLM\SYSTEM\CurrentControlSet\Services\msiscsidsm\Parameters\PathVerifyEnabled [2K3] - should be 0
  - HKLM\SYSTEM\CurrentControlSet\Services\msdsm\Parameters\PathVerifyEnabled [2K8] - should be 0

Other parameters, read http://now.netapp.com/NOW/knowledge/docs/hba/iscsi/win/iscsiwinhu41/pdfs/setup.pdf

- For FCP:

(still need to be added here: read FCP Host Utilities docs if you need them)

[edit]

Other Operating Systems

Description	OK
ESX, Solaris, Linux, ... Use the appropriate Support Kits for your OS: http://now.netapp.com/NOW/cgi-bin/software Read the docs for these Operating Systems	
Run the Unix Configuration Checker (part of SnapDrive for Unix 4.0 and higher, but also available as a separate download	

ADJUST TIMEOUT SETTINGS FOR VMWARE GUESTS !!!

- http://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb41511
- http://now.netapp.com/NOW/download/software/sanhost_esx/3.1/download.shtml

[edit]

FCP

Description	OK
Start FCP as soon as possible since it requires a reboot to activate: `filer> fcp start`	

On a 30x0/60x0, make sure the onboard FC adapters are configured correctly, i.e. are set to either "target" or "initiator" depending on their function:
 	
- target = the adapter will be used to connect to servers/HBAs (typically via a FC switch), and the adapter will be used to access LUNs on the filer
- initiator = the adapter will be used to connect to disk shelves
 	
Command:
 	
`filer> fcadmin config`
`Local`
`Adapter Type State Status`
`---`
`0a target CONFIGURED offline`
`0b target CONFIGURED offline`
`0c target CONFIGURED offline`
`0d target CONFIGURED offline`
`filer> fcadmin config -t { initiator | target } adapter_name ...`
 	
Note that any change requires a REBOOT in order to become active	

Avoid "Partner path misconfigured" errors & autosupport mails	

lun stats -o show LUN statistics, look at *Partner KB*. **lun stats -z** zeroes statistics **lun config_check –v** shows which initiator is generating partner traffic (very handy eg. for troubleshooting ESX problems)	
On Brocade 200E switches, set port speeds manually: http://now.netapp.com/NOW/knowledge/docs/ontap/rel73rc/html/ontap/rnote_rc2/rel_notes/concept/c_oc_rn_lim-net-200e.html#c_oc_rn_lim-net-200e	
Watch out for partition alignment problems on Linux systems: http://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb8190	

[edit]

Autosupport

Description	OK
Explain the importance of triggering an ASUP to the customer, eg. let's say filer is unresponsive to CIFS or NFS and -in a panic- you want to reboot the system: please trigger an ASUP first so that possibly vital statistics are collected and sent out to NetApp & Uptime	
Use a unique "From:" address for every filer: `filer> options autosupport.from filername@domainname.com`	
Add *netapp@uptime.be* to the recipients for email autosupport notifications, example: `filer> options autosupport.to email@domain.com,netapp@uptime.be`	
You can use more than one mail relay host, Data ONTAP will try the next one if the first one is down: `filer> options autosupport.mailhost ws03-ex1,ws03,ex2,172.16.7.44`	
Change the autosupport protocol to SMTP instead of HTTP(S) `filer> options autosupport.support.transport smtp` Note: this cannot be done via the web interface but must be done via the **options autosupport.transport** command	

Enable verbose autosupport content for CIFS and DAFS:
 	
`filer> options autosupport.cifs.verbose on`
`filer> options autosupport.dafs.verbose on`
 	
Note: Recent Data ONTAP version do not have the *autosupport.dafs.verbose* option anymore	

Test autosupport while logged in to the filer's console so the possible error messages will inform you why the autosupport mail is not being sent: `filer> options autosupport.doit test123` If the subject contains the string "test", NetApp will send a delivery receipt to everyone in the "To:" list	
In ONTAP 7.0.5 and later, it is possible to eliminate mail sending to customer email addresses, see **options autosupport.notify_threshold**. We will keep the traditional notification settings (all mails sent to everyone): `filer> options autosupport.notify_threshold debug` Note: This option is no longer present in Data ONTAP 7.2.x	

You may want to change the retry counter and retry interval for autosupport retries in case of failed deliveries; these are the default values:
 	
`filer> options autosupport.retry.count 15`
`filer> options autosupport.retry.interval 4m`
 	
If you don't change the values, the system will only try to reach a mailserver for 1 hour.	

You may want to specify multiple mail servers. Data ONTAP will try subsequent mailservers if the previous ones cannot be reached.	

[edit]

RLM

- This is available on FAS30x0 models (on lower-end models it is not offered by default) and FAS60x0 models
- On FAS20x0 models, the **RLM** (**R**emote **L**AN **M**odule) is called **BMC** (**B**aseboard **M**anagement **C**ontroller)
- Only ssh access is allowed
- Log in as *naroot*, not *root*, on RLM modules; additional users can be created
- Log in as *naroot*, *root*, or *Administrator* on BMC modules; no additional users can be created

Description	OK

Configure the RLM, either during the normal setup, or via:
 	
`filer> rlm setup`
 	
The BMC can be configured via:
 	
`filer> bmc config ...`
 	
(example **bmc config ipaddr**)	

Test the RLM by ssh'ing to it and by testing the autosupport mail: • Can you connect to the mail relay server from the RLM board's IP address range ?	

[edit]

Clustered Machines

Description	OK

On a clustered machine, run the cluster configuration checker from the NOW site (note: or use our NetApp tool). The latest versions can also be run from Windows machines.
 	
Example output from the cluster configuration checker should look like this (run from a linux host):
 	
`$ perl cf-config-check.cgi -l 10.0.10.130 10.0.10.131`
`10.0.10.130 rsh login: root`
`Password:`
`10.0.10.131 rsh login: root`
`Password:`
`== NetApp Cluster Configuration Checker v1.4.4 ==`
 	
`Checking rsh logins ...`
 	
`OK`
 	
`Checking OS versions...`
 	
`OK`
 	
`Checking licenses...`
 	
`OK`
 	
`Checking cluster identity...`
 	
`OK`
 	
`Checking cf status...`
 	
`OK`
 	
`Checking fcp cfmode settings...`	

```
N/A

Checking options...

OK

Checking Network Configuration...


Checking network config in /etc/rc

OK
No Cluster Issues Found
Done.
```

When issues are encountered, the output looks somewhat like this:

```
$ perl cf-config-check.cgi -l 10.7.0.7 10.7.0.8
10.7.0.7 rsh login: 'root
Password:
10.7.0.8 rsh login: root
Password:
== NetApp Cluster Configuration Checker v1.4.4 ==


Checking rsh logins ...

OK

Checking OS versions...

OK

Checking licenses...

a_sis exists on 10.7.0.7, not on 10.7.0.8
snapmanagerexchange exists on 10.7.0.8, not on 10.7.0.7

Checking cluster identity...

OK

Checking cf status...

OK

Checking fcp cfmode settings...

OK

Checking options...

Option timed.servers                10.7.20.1
 in 10.7.0.8 has no match in 10.7.0.7.
Option timed.servers                10.0.10.6,10.0.10.15
 in 10.7.0.7 has no match in 10.7.0.8.

Checking Network Configuration...

vif_frontend (10.7.0.7) on 10.7.0.7 does not have a partner on 10.7.0.8
vif_backend (10.8.0.8) on 10.7.0.8 does not have a partner on 10.7.0.7
vif_frontend (10.7.0.8) on 10.7.0.8 does not have a partner on 10.7.0.7

Checking network config in /etc/rc

NO PARTNER FOR vif_frontend (10.7.0.7) ON na-demo03 IN /etc/rc
NO PARTNER FOR vif_backend (10.8.0.8) ON na-demo04 in /etc/rc
NO PARTNER FOR vif_frontend (10.7.0.8) ON na-demo04 in /etc/rc
Cluster Issues Found above.  Please fix them.
Done.
```

Unless the issues are resolved and the cluster configuration checker is run again, cluster failover may not work properly!

Test failover and giveback (while connected via a console cable on both machines):

```
filer> cf takeover
```

and

```
filer> cf giveback
```

[edit]

Network

Name Resolution [edit]

General [edit]

Description	OK

On systems with more than one IP address, in different subnets, make sure only one IP address uses WINS. The other interfaces must have the setting **NOWINS** in the output of the *ifconfig* command. This can be done with the -wins option to the ifconfig command in /etc/rc. Note that the WINS configuration also controls dynamic DNS registrations ! Typically, you don't want eg. iSCSI or SnapMirror/SnapVault interfaces to register with WINS or DNS:
 	
`filer> rdfile /etc/rc`
`...`
`ifconfig vif_lan `hostname`-vif_lan mediatype auto netmask 255.255.0.0 partner vif_lan`
`ifconfig vif_iscsi `hostname`-vif_iscsi mediatype auto netmask 255.255.0.0 -wins partner vif_iscsi nfo`
`ifconfig vif_nfs `hostname`-vif_nfs mediatype auto netmask 255.255.0.0 -wins partner vif_nfs nfo`
`...`
`filer> ifconfig -a`
`...`
`lo: flags=1948049<UP,LOOPBACK,RUNNING,MULTICAST,TCPCKSUM> mtu 8160`
`inet 127.0.0.1 netmask 0xff000000 broadcast 127.0.0.1`
`ether 00:00:00:00:00:00 (VIA Provider)`
`vif_lan: flags=948043<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM> mtu 1500`
`inet 10.10.1.66 netmask 0xffff0000 broadcast 10.10.255.255`
`partner vif_lan (not in use)`
`ether 02:a0:98:0a:eb:a7 (Enabled virtual interface)`
`vif_iscsi: flags=4948043<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM,NOWINS> mtu 1500`
`inet 10.250.2.66 netmask 0xffff0000 broadcast 10.250.255.255`
`partner vif_iscsi (not in use)`
`ether 02:a0:98:0a:eb:a6 (Enabled virtual interface)`
`nfo enabled`
`vif_nfs: flags=4948043<UP,BROADCAST,RUNNING,MULTICAST,TCPCKSUM,NOWINS> mtu 1500`
`inet 10.250.1.66 netmask 0xffff0000 broadcast 10.250.255.255`
`partner vif_nfs (not in use)`
`ether 02:a0:98:0a:eb:a5 (Enabled virtual interface)`
`nfo enabled`	

[edit]

NetBIOS Aliases

Description	OK
You can use NetBIOS aliases via the options cifs.netbios_aliases option	

[edit]

WINS

DNS [edit]

Description	OK
If running DATA ONTAP 7.0 or earlier, make sure the filer names are added to DNS (7.0 and earlier do not perform automatic DNS updates), both the A records and PTR records	

Check DNS configuration on filer: Is the filer able to resolve names:
 	
`filer> ping vmdemodc1`
`vmdemodc1.demoroom.local is alive`	

Run the command:
 	
`filer> dns info`
`DNS is enabled`
 	
`DNS caching is enabled`
 	
`1 cache hit`
`77 cache misses`
`3 cache entries`
`71 expired entries`
`71 cache replacements`
 	
`IP Address State Last Polled Avg RTT Calls Errs`
`10.7.70.21 UP Sun Jul 23 23:49:54 CEST 200 0 146 8`
`10.0.10.6 DOWN 0 8 8`
 	
`Default domain: demoroom.local`
`Search domains: demoroom.local`	

Is reverse DNS resolution working as expected (eg. no long login timeouts when using ssh to the filer ?)	

[edit]

VIF Configuration

Description	OK

Are all network cables connected ?
 	
`filer> ifconfig -a`
`e0a: flags=48043<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500`
`ether 02:a0:98:01:92:49 (auto-unknown-down) flowcontrol full`
`trunked eth0`
`e0b: flags=848043<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500`
`ether 02:a0:98:01:92:49 (auto-1000t-fd-up) flowcontrol full`
`trunked eth0`
`lo: flags=1948049<UP,LOOPBACK,RUNNING,MULTICAST,TCPCKSUM> mtu 8160`
`inet 127.0.0.1 netmask 0xff000000 broadcast 127.0.0.1`
`ether 60:9a:cf:37:04:00 (VIA Provider)`
`eth0: flags=848043<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500`
`inet 10.7.0.3 netmask 0xffff0000 broadcast 10.7.255.255`
`partner eth0 (not in use)`
`ether 02:a0:98:01:92:49 (Enabled virtual interface)`
 	
Interface e0a is not connected!	

Are all network speeds set correctly ?
 	
`filer> ifconfig -a`
`...`
`e0b: flags=848043<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500`
`ether 02:a0:98:01:92:49 (auto-1000t-fd-up) flowcontrol full`
`trunked eth0`
`...`
 	
This interface is on *1 Gbps (Gigabit)*	

Are all network duplex settings correct ?
 	
`filer> ifconfig -a`
`...`
`e0b: flags=848043<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500`
`ether 02:a0:98:01:92:49 (auto-1000t-fd-up) flowcontrol full`
`trunked eth0`
`...`
 	
This interface is on *Full Duplex*	

netdiag -v command ok ? See manpage
 	
`filer> netdiag -v`
`Performing physical layer diagnostics.....`
`Checking interface e0a....`
`Checking interface e0a for hardware status....`
`Checking interface e0a for link status....`
 	
`The interface e0a does not seem to have a good link.`
`This may be because of one of the following:`
`1. There is no cable connecting the interface e0a to your network.`
`2. The cable connecting the interface e0a to your network is bad.`
`3. The switch/hub connected to the interface e0a has a problem.`
`4. The interface e0a has a hardware problem. If you suspect this last possibility please contact Network Appliance Global Services.`
`Press enter to continue`
 	
`Checking interface e0b....`
`Checking interface e0b for hardware status....`
`Checking interface e0b for link status....`
`Interface e0b OK`
`Checking interface eth0....`
`Checking interface eth0 for hardware status....`
`Checking interface eth0 for link status....`
`Interface eth0 OK`
`Checking interface e0a to see if it is wedged`
`Checking interface e0a for misc errors`
`Checking interface e0a for capacity problems`
`Checking interface e0b to see if it is wedged`
`Checking interface e0b for misc errors`
`Checking interface e0b for capacity problems`
`Performing network layer diagnostics.....`
`IP is OK`
`Performing transport layer diagnostics.....`
`Checking for problems with current TCP connections...`
`Checking for problems with recent TCP connections...`
`TCP is OK`
`UDP is OK`	

LACP VIFs are preferred over static multimode VIFs, since the LACP protocol implements keepalive frames. More network failures are detected on LACP-enabled VIFs	
vif status & multimode VIFs: if received packets for a child interface is zero, the switch is not configured correctly	

vif status & multimode VIFs: if sent packets stays zero for a child interface, check RR/IP/MAC based multi-mode VIF setup. This cannot be set via the webinterface but must be added to /etc/rc file to the **vif create** command (eg. **-b rr** for round robin, default is IP based). Check manpage
 	
`filer> vif status`
`default: transmit 'IP Load balancing', VIF Type 'multi_mode', fail 'log'`
`eth0: 1 link, transmit 'none', VIF Type 'single_mode' fail 'default'`
`VIF Status Up Addr_set`
`up:`
`e0b: state up, since 13Jul2006 13:19:16 (10+11:46:23)`
`mediatype: auto-1000t-fd-up`
`flags: enabled`
`input packets 2156480, input bytes 570063466`
`output packets 2105628, output bytes 2122504403`
`up indications 1, broken indications 0`
`drops (if) 0, drops (link) 0`
`indication: up at boot`
`consecutive 906395, transitions 1`
`broken:`
`e0a: state broken, since 13Jul2006 13:19:30 (10+11:46:09)`
`mediatype: auto-unknown-down`
`flags: disabled`
`input packets 0, input bytes 0`
`output packets 0, output bytes 0`
`up indications 0, broken indications 0`
`drops (if) 0, drops (link) 0`
`indication: broken at boot`
`consecutive 0, transitions 1`
 	
This traffic distribution is entirely disproportionate (It's ok for a single mode VIF though!)	

On clustered systems, is takeover of interface happening correctly ? May cause problems if physical interfaces are in different VLANs	
Compare output of **ifconfig -a** and **vif status** with the contents of **/etc/rc** and the web interface view. We have observed inconsistencies between them when fiddling around with parameters. Note that **/etc/rc** is run upon startup, so it must contain correct information	

[edit]

CIFS

Description

Run **cifs testdc** and make sure it runs smoothly.
 	
`filer> cifs testdc`
`Using Established configuration`
`Current Mode of NBT is B Mode`
 	
`Netbios scope ""`
`Registered names...`
`FILER < 0> Broadcast`
`FILER < 3> Broadcast`
`FILER <20> Broadcast`
`DEMOROOM < 0> Broadcast`
 	
`Testing all Primary Domain Controllers`
`found 1 unique addresses`
 	
`found PDC VMDEMODC1 at 10.7.70.21`
 	
`Testing all Domain Controllers`
`found 1 unique addresses`
 	
`found DC VMDEMODC1 at 10.7.70.21`
 	
If **cifs testdc** does not run smootly (eg. it tries to contact DCs that are not in the domain anymore, or DCs that are unreachable, you may get i problems for CIFS)

If needed, set preferred domain controllers:
 	
`filer> cifs prefdc print`
`No preferred Domain Controllers configured.`
`DCs will be automatically discovered.`
`filer> cifs prefdc add <domain> <dclist>`
`...`
 	
Don't forget to reload the DCs afterwards:

```
filer> cifs resetdc
        Disconnecting from domain DEMOROOM...
        Reconnecting to domain DEMOROOM...
Sun Jul 23 20:11:29 CEST [filer: auth.dc.trace.DCConnection.statusMsg:info]: AUTH: TraceDC- Starting DC address d
Sun Jul 23 20:11:29 CEST [filer: auth.dc.trace.DCConnection.statusMsg:info]: AUTH: TraceDC- Filer is not a member
Sun Jul 23 20:11:29 CEST [filer: auth.dc.trace.DCConnection.statusMsg:info]: AUTH: TraceDC- Found 1 addresses usi
Sun Jul 23 20:11:29 CEST [filer: auth.dc.trace.DCConnection.statusMsg:info]: AUTH: TraceDC- DC address discovery
addresses found.
Sun Jul 23 20:11:29 CEST [filer: auth.dc.trace.DCConnection.statusMsg:info]: AUTH: TraceDC- Connection with \\VMD
Reconnection succeeded
Sun Jul 23 20:11:29 CEST [filer: auth.ldap.trace.LDAPConnection.statusMsg:info]: AUTH: TraceLDAPServer- Starting
discovery for DEMOROOM.LOCAL.
Sun Jul 23 20:11:29 CEST [filer: auth.ldap.trace.LDAPConnection.statusMsg:info]: AUTH: TraceLDAPServer- Found 1 A
using generic DNS query.
Sun Jul 23 20:11:29 CEST [filer: auth.ldap.trace.LDAPConnection.statusMsg:info]: AUTH: TraceLDAPServer- AD LDAP s
DEMOROOM.LOCAL complete. 1 unique addresses found.
```

Check DNS and WINS info. **IMPORTANT**: on filers with multiple interfaces, only one interface can register the hostname with any WINS se

Make sure the ~snapshot directory is visible to CIFS clients (if required by customer)

- Global option:

```
filer> options cifs.show_snapshot on
```

- For every volume:

```
filer> vol options vol_name nosnapdir off
```

Make sure "previous versions" tab is present for XP and 2003 clients (if required by customer):

```
filer> options cifs.ms_snapshot_mode xp
```

You may have to do a "regsvr32 twext.dll" on Windows XP to get the previous versions tab to show. See also:

- https://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb34602
- http://support.microsoft.com/kb/888603 (Terminal Services clients)

[edit]

NFS

Description	OK
For VMware over NFS, Netapp recommends this setting: `filer> options nfs.tcp.recvwindowsize 64240`	

Time server configuration

[edit]

Description	OK
Enable time logging: `filer> options timed.log on`	
Set scheduling interval to 1 minute and watch logging messages to see if time synchronization is actually taking place `filer> options timed.sched 1m`	
Set time protocol to ntp (used most widely) `filer> options timed.proto ntp`	
Set time server(s). • You should be able to set the Windows 2000/2003 domain name instead of an IP address or host name, if applicable: filer will find domain controller(s)/time servers itself `filer> options timed.servers demoroom.local`	

When clocks are synchronized, set interval back to 1 hour: `filer> options timed.sched 1h`	
Set the window to 5 minutes: `filer> options timed.window 5m` (if set to 0, there is a problem if the filer's internal clock is too fast: ntp will reset clock backwards on the hour and scheduled snapshots may be taken twice. This windows will randomize time sync +/- 5 mins every hour)	
Set timed.max_skew to 3 hours (allows adjustments of daylight saving time if timezone is accidentally set to GMT) `filer> options timed.max_skew 3h`	
Turn off logging if everything is working fine: `filer> options timed.log off`	

[edit]

Security

See also this NetApp document for more info:

http://media.netapp.com/documents/tr-3649.pdf

Do Your Homework

[edit]

Description	OK
Do you have a security policy in-place	
Do you have a network security policy ? • Are all networks documented ? • Are all ethernet switches documented and configured securely ? • Are all hosts accessing the filer documented and configured securely ? • Are all of the filer's network interfaces documented ?	
Is physical access to the storage system restricted to authorized personnel only ?	
Is management access to the storage system restricted to authorized personnel only ?	
Investigate encryption of data	
Employ strong user-level authentication (eg. Kerberos)	
Ensure you have unique user IDs and allow those user IDs to be traced back to a specific user	
When creating volumes, qtrees or LUNs, organize data by security requirements	

[edit]

Patch Possibly Known Issues

Description	OK

If running DATA ONTAP 7.0.4 or earlier, the following vulnerability (http://now.netapp.com/NOW/products/vulnerability_0601/) needs this fix. Type this command:
`filer> useradmin role modify audit -a api-snmp-get,api-snmp-get-next`
`Role <audit> modified.`
`Sun Jul 23 19:43:28 CEST [na-demo01: useradmin.added.deleted:info]: The role 'audit' has been modified.`	

[edit]

Network and Protocol Considerations

General

[edit]

Description	OK

Are all networks physically or logically separated/segmented ? • Routing • Subnetting • Switching • VLANs	
NetApp recommendations: • Employ strong user-level authentication by using Kerberos with NFS or CIFS • Use LDAP over SSL for centralized authentication and authorization • Enable LDAP signing and sealing with SASL • Enable CIFS signing to ensure the integrity of CIFS data transmission • Set CIFS authentication levels to accept only Kerberos authentication • Use NFSv4 whenever possible and limit NFSv3 usage • Enable NFSv4 ACLs and make sure that those ACLs are designed and assigned correctly	

[edit]

Network and IP Options

Description	OK

• options ip.match_any_ifaddr off
  ◦ Checks incoming packets for correct addressing. If this option is on, the NetApp storage system accepts any packet that is addressed to it, even if that packet came in on the wrong interface
• options ip.fastpath.enable off
  ◦ The NetApp storage system attempts to use MAC address and interface caching (fastpath) to try to send back responses to incoming network traffic by using the same interface as the incoming traffic and (in some cases) the destination MAC address equal to the source MAC address of the incoming data
• options ip.ping_throttle.drop_level 150
  ◦ Specifies the maximum number of ICMP echo or echo reply packets that Data ONTAP accepts per second. Any further packets within 1 second are dropped to prevent ping flood denial of service attacks
• options ip.ping_throttle.alarm_interval 5
  ◦ Specifies how often dropped pings are logged, in minutes. This prevents a ping flood denial of service attack from flooding the audit log with messages
• ip.icmp_ignore_redirect.enable on
  ◦ Disable icmp redirects	

You can turn the routing daemon off: `routed off` Make sure to do this in the /etc/rc file, otherwise a reboot or takeover will reenable the routing daemon again. *routed* enables IDRP router discovery and listening for RIP packets. You can safely disable routed if you do not rely on IRDP or RIP for routing updates. Note: Enabling the MultiStore license in Data ONTAP automatically disables the routed process	
You can enable IPsec for the filer: `filer> options ip.ipsec.enable on` See the networking guide on IPsec	

For Data ONTAP 7.3 and later:
Set up protocol access control, eg.:
`filer> options interface.blocked.cifs e5b`
`filer> options interface.blocked.nfs e1a,e1b`
`filer> options interface.blocked.iscsi e5b`
`filer> options interface.blocked.ftpd e5b,e1a,e1b`
`filer> options interface.blocked.snapmirror e4a,e4b`
`filer> options interface.blocked.cifs ""`
This sets a comma-separated list of interface names for which a specific protocol is blocked	

[edit]

Protocols

Description	OK

Set up protocol access, eg.
• options rsh.access "host = gnesha"
  ○ Allows remote shell access for only one host, named gnesha.
• options telnet.access host=10.42.69.0/24
  ○ Allows telnet access for subnet 10.42.69.0.
• options ssh.access "host=abc,xyz AND if=e0"
  ○ Allows SSH access for hosts abc and xyz when on network interface e0.
• options snmp.access “if=e0,e1,e2”
  ○ Allows SNMP access for network interfaces e0, e1, and e2.
• options httpd.access "if != e3"
  ○ Don't allow access to HTTPD for network interface e3.
• options httpd.admin.access “host=champagne,tequila”
  ○ Allows administrative HTTPD access for hosts champagne and tequila.
• options telnet.access "host=-"
  ○ Disallows all access to telnet.
• options snapmirror.access legacy
  ○ Use /etc/snapmirror.allow to check access to SnapMirror sources.
• options snapvault.access all
  ○ Allows a SnapVault server to accept any client requests.
• options ndmpd.access ”host = backup”
• Allows an NDMP server to accept a connection request from a single backup server.	

CIFS
 	
Things to take into account:
• Do you want to set up virus scanning ? See the *Data ONTAP Data Protection Online Backup and Recovery Guide*.
• NetApp recommends the following best practices to securely implement CIFS:
  ○ Active Directory authentication using Kerberos
  ▪ Select a Microsoft® Active Directory domain during CIFS setup
  ▪ See the NetApp technical report TR-3457 for more information on setting up NetApp storage systems with Active Directory authentication by using Kerberos
  ▪ See the NetApp technical report TR-3458 for more information on setting up NetApp storage systems with Active Directory authentication using Kerberos to support CIFS and NFS clients
  ○ LDAP signing and sealing with SASL and LDAP transport over SSL
  ▪ In conjunction with setting up LDAP for authentication and authorization, LDAP signing provides another level of security, and LDAP sealing provides encryption of all LDAP packets. Recommended Setting:
  ▪ Enable LDAP signing and sealing with SASL.
  ▪ Enable LDAP over SSL.
 	
`filer> options ldap.security.level 2`
 	
• ○ CIFS signing to ensure integrity of CIFS traffic
 	
On the NetApp storage system:
 	
`filer> options cifs.signing.enable on`
 	
On the Windows client:
 	
Enable EnableSecuritySignature and RequreSecuritySignature parameters in the Windows registry:
 	
`HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SeCEdit\`
`  Reg Values\MACHINE/System/CurrentControlSet/Services/LanManServer/Parameters/EnableSecuritySignature`
 	
`HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SeCEdit\`
`  Reg Values\MACHINE/System/CurrentControlSet/Services/LanManServer/Parameters/RequireSecuritySignature`
 	
• Session authentication level
  ○ Determines which challenge/response authentication protocol is used for Windows net logon. The following levels of authentication are supported:
  ▪ Level 1: Accept LM, NTLM, NTLMv2 session security, NTLMv2, Kerberos (default)
  ▪ Level 2: Accept NTLM, NTLMv2 session security, NTLMv2, Kerberos
  ▪ Level 3: Accept NTLMv2 session security, NTLMv2, Kerberos
  ▪ Level 4: Accept NTLMv2, Kerberos
  ▪ Level 5: Accept Kerberos only
  ○ Recommended Setting: For the highest session authentication level, set this option to 5 to accept only Kerberos authentication. Setting the option to this level supports only Windows 2000 and later versions of Windows.
  ○ Procedure
 	
`filer> options cifs.LMCompatibilityLevel 5`	

- Set up share-level permissions, if needed, eg.

```
filer> cifs access <sharename> [-g] <user|group> <rights>
```

- Data ONTAP 7.2 and later releases provide storage system support for Access-Based Enumeration (ABE), a shared resource security feature introduced in Microsoft Windows Server 2003 Service Pack 1. When ABE is enabled on a CIFS share, users who do not have permission to access a shared folder or file underneath it (whether through individual or group permission restrictions) do not see that shared resource displayed in their environment

- Set up access based enumeration on a share, eg.

```
filer> cifs access <sharename> <-accessbasedenum>
```

- Anonymous Connections (Restrict anonymous)
  - Controls access to users with nonauthenticated connections. Permitted values for this option are 0, 1, and 2. 0 sets no special access restrictions, 1 disallows enumeration of users and shares, and 2 fully restricts access. This option corresponds to the RestrictAnonymous registry entry in Windows.
  - Recommended Setting: Disable access to CIFS shares and sharenames from unauthenticated users.
  - Procedure

```
filer> options cifs.restrict_anonymous 2
```

- Disable CIFS guest access.
  - Procedure

```
filer> options cifs.guest_account ""
```

- Storage-Level Access Guard to provide a third layer of security to CIFS and NFS
  - In Data ONTAP 7.2.2 and later, creates a third level of access control for CIFS and NFS shares.
  - Recommended Setting: Enable Storage-Level Access Guard. See the NetApp technical report TR-3596 for more information on enabling and configuring Storage-Level Access Guard.
  - Procedure

```
filer> fsecurity apply <definition file path> [<options>]
```

- Group policy objects
  - A Group Policy Object (GPO) is a set of rules that are applicable to users and computers in an Active Directory environment and defined centrally for ease of administration and increased security. Data ONTAP is able to recognize and process a certain set of GPOs.
  - Recommended Setting: Enable GPO support. Use GPO for file system security, restricted security groups, event login, and audit policy mapping. See the NetApp technical report TR-3367 for more information on Group Policy Objects.
  - Procedure

```
filer> options cifs.gpo.enable on
```

- Windows domain machine password
  - By default, a NetApp storage system in a Windows 2000 domain does not automatically change its machine password. This option enables you to change the machine password weekly.
  - Recommended Setting: Enable weekly changes of the machine password.
  - Procedure

```
filer> options cifs.weekly_W2K_password_change on
```

- NetBIOS over TCP
  - Enables/disables NetBIOS transport over TCP. If disabled, legacy Windows clients and domains do not communicate with the NetApp storage system. This option takes effect when CIFS is started in Data ONTAP. It should not be changed while CIFS is enabled.
  - Recommended Setting: Disable NetBIOS over TCP.
  - Procedure

```
filer> options cifs.netbios_over_tcp.enable off
```

- CIFS auditing to provide very granular logging information

- Audit CIFS access
  - Audits CIFS access
  - Recommended Setting: Enable the auditing of CIFS access to the NetApp storage system. See the NetApp technical report TR-3595 for information on auditing CIFS and NFS protocols with Data ONTAP
  - Procedure

```
filer> options cifs.audit.enable on
```

- Audit CIFS account management events
  - Audits CIFS file access events when a System Access Control List (SACL) matches a request for access.
  - Recommended Setting: Enable the auditing of CIFS file access events. See the NetApp technical report TR-3595 for information on auditing CIFS and NFS protocols with Data ONTAP
  - Procedure

```
filer> options cifs.audit.account_mgmt_events.enable on
```

Audit CIFS file access events

-
  - Audits CIFS account creation, deletion, and modification.
  - Recommended Setting: Enable the auditing of CIFS account management events. See the NetApp technical report TR-3595 for information on auditing CIFS and NFS protocols with Data ONTAP.
  - Procedure

```
filer> options cifs.audit.file_access_events.enable on
```

- Audit CIFS logon events
  - Audits CIFS logons and logoffs, including CIFS session connects and disconnects.
  - Recommended Setting: Enable the auditing of CIFS logon events. See the NetApp technical report TR-3595 for information on auditing CIFS and NFS protocols with Data ONTAP.
  - Procedure

```
filer> options cifs.audit.logon_events.enable on
```

- Microsoft publishes a best practices guide to auditing security events, available at http://technet2.microsoft.com/windowsserver/en/library/5658fae8-985f-48cc-b1bf-bd47dc2109161033.mspx.

NFS

- NetApp recommends a number of best practices to securely deploy NFS:
  - Kerberos authentication
    - Enables Kerberos authentication for NFS. Requires NFS clients to support Kerberos.
    - Recommended Setting: Enable NFS authentication with Kerberos. Refer to the NetApp technical report TR-3481 for information on setting up Kerberos with NetApp storage systems
    - Procedure

```
filer> nfs setup
```

-
  -
    - After performing the nfs setup command, edit /etc/exports on the NetApp storage system to set “sec=krb5”, “sec=krb5i”, or “sec=krb5p” in the options field of the exported file systems.
  - LDAP signing and sealing with SASL and LDAP transport over SSL
    - Enables LDAP directory lookup service for user authorization. SSL is also supported for secure connection.
    - Recommended Setting: Enable LDAP user lookup for authorization. Enable LDAP over SSL or SASL. Refer to the NetApp technical report TR-3464 for information on setting up LDAP with NetApp storage systems.
    - Procedure

```
filer> options ldap.enable on
filer> options ldap.ssl.enable on
```

-
  - Enable NFSv4
    - Enables NFS version 4, allowing use of NFSv4 Access Control Lists (ACLs).
    - Recommended Setting: Enable NFSv4. Where possible, disable NFSv3 at the same time. See the NetApp technical report TR-3580 for information on NFSv4.
    - Currently, not all applications (Oracle, Vware, ...) are supported on NFSv4. Be very careful when implementing this!
    - Procedure

```
filer> options nfs.v4.enable on
filer> options nfs.v4.acl.enable on
```

-
  - Enable NFS over TCP
    - Enables NFS sessions by using TCP packets instead of UDP. TCP is generally more secure than UDP and may facilitate use of NFS across firewall boundaries. However, enabling NFS traffic through a firewall opens up so many ports in both directions that it is better practice to deploy the NFS clients and servers in the same security zone.
    - Recommended Setting: Enable NFS over TCP.
    - Procedure

```
filer> options nfs.tcp.enable on
filer> options nfs.udp.enable off
```

-
  - Restrict NFS to low-numbered ports
    - Enables/disables NFS mount requests over high-numbered ports. Low-numbered ports are restricted to root

users and are considered more secure.
- Recommended Setting: Restrict NFS mounts to low-numbered ports only.
- Procedure

```
filer> options nfs.mount_rootonly on
```

- Secure the */etc/exports* file
  - Use *man na_exports*
  - Make sure that you are using the appropriate security options in the NFS export to prevent unsolicited clients from mounting or gaining elevated access rights to the desired volumes on the NetApp storage system
  - The following NFS export options are related to security
    - **anon**
      - This option specifies the effective user ID (or name) of all anonymous or root NFS client users that access the file system path
      - An anonymous NFS client user is an NFS client user that does not provide valid NFS credentials; a root NFS client user is an NFS client user with a user ID of 0.
      - Data ONTAP determines a user's file access permissions by checking the user's effective user ID against the NFS server's /etc/passwd file. By default, the effective user ID of all anonymous and root NFS client users is 65534.
      - To disable root access by anonymous and root NFS client users, set the anon option to 65535.
      - To grant root user access to all anonymous and root NFS client users, set the anon option to 0. This is equivalent to the no_root_squash option in some other NFS servers.
      - If a name is provided instead of a user ID, that name is looked up according to the order specified in the /etc/nsswitch.conf file, which determines the corresponding user ID to be assigned by the anon option.
    - **nosuid**
      - This option disables the setuid and setgid executables and mknod commands on the file system path.
      - Unless the file system is a root partition of a diskless NFS client, you should set the nosuid option to prevent NFS client users from creating setuid executables and device nodes that careless or cooperating NFS server users could use to gain root access.
    - **sec**
      - Starting with version 6.5, Data ONTAP supports the ability to specify multiple security (sec) options for each exported resource. The administrator can determine how secure NFS access is to the NetApp storage system. Basically, the following two security service types are supported.
        - *UNIX (AUTH_SYS) authentication (sys)*: Does not use strong cryptography and is the least secure of the security services. This is the default security service used by Data ONTAP.
        - Note: AUTH_SYS credentials are basically a user ID and up to 17 group IDs. Once a person is logged in as a superuser on a UNIX system, that person can use the su command to become a user who is allowed full access to a volume. One way to prevent this scenario from happening is to implement strong authentication mechanisms such as Kerberos.
        - *Kerberos 5* Provides the following three security methods:
          - Authentication (krb5): Uses strong cryptography to prove a user's identity to a storage system and to prove a storage system's identity to a user.
          - Integrity (krb5i): Provides a cryptographic checksum of the data portion of each request and the response message to each request. This defends against "man in the middle" tampering with storage system NFS traffic.
          - Privacy (krb5p): Encrypts the contents of packets bidirectionally, including procedure arguments and user data, by using a shared session key established by the client from the storage system.

The following two examples show how these security services are used: To specify one security type, enter:

```
/vol/volx -sec=sys,rw=host1
```

To specify multiple security types, enter:

```
/vol/volx -sec=krb5:krb5i:krb5p,rw=host1
```

For more information on setting up NFS using Kerberos authentication, refer to these NetApp technical reports:

- TR-3481 for a key distribution center (KDC) based on UNIX
- TR-3457 for a KDC based on Active Directory

iSCSI

- Read The NetApp Block Management Guide (Chapter 6)
- NetApp recommends several best practices to secure iSCSI storage:
  - Enable iSCSI only on necessary interfaces

```
filer> iscsi interface disable [-f ] {-a | <interface>…}
```

- Disable access for initiators with no security method

```
filer> iscsi default -s deny
filer> iscsi security add -i initiator -s CHAP -p password -n name
```

- Use CHAP authentication with random 128-bit passwords

```
filer> iscsi security generate
```

- Use LUN masking to control access to specific initiators
- - Each iSCSI LUN can be restricted to a specified group of iSCSI initiators. NetApp refers to these initiators groups as igroups. This initiator-to-igroup to LUN combination is known as LUN masking
  - Recommended Setting: Use LUN masking to restrict LUN access to specific igroups. Create an igroup, then create the LUN, and finally create the mask
  - Procedure

```
filer> igroup create -i -t windows igroup-name [node-name]
filer> lun create -s size -t windows lun_name
filer> lun map lun_name igroup_name [lun_ID]
```

- Use iSCSI interface access lists to restrict initiators to specific interfaces
- - Each iSCSI initiator can be restricted to specific network interfaces. This is particularly useful in VLAN environments, where an initiator may not be able to access all interfaces on the NetApp storage system. Creating or modifying an access list may cause sessions to be shut down, so use these commands carefully
  - Recommended Setting: Use iSCSI interface access lists to control initiator access.
  - Procedure

```
filer> iscsi interface accesslist add [initiator] [-a interface]
filer> iscsi interface accesslist remove [initiator] [-a interface]
filer> iscsi interface accesslist show [-a]
```

FCP

- See the Block Access Management Guide, especially chapter 7, "Managing FCP Initiator Groups."
- Implement zoning on the Fibre Channel switches that are deployed as part of the configuration
  - See the switch documentation for details

NDMP

Don't allow clear-text passwords

```
options ndmpd.authtype challenge
```

Multiprotocol options

- Ignore ACLs
  - When on, ACLs do not affect root access from NFS. The option defaults to off.
  - Recommended Setting: Disable the ignoring of any ACLs.
  - Procedure

```
filer> options cifs.nfs_root_ignore_acl off
```

- CIFS bypass traverse checking
  - When on (the default), directories in the path to a file are not required to have the X (traverse) permission. This option does not apply in UNIX qtrees.
  - Recommended Setting: Enable traverse checking by turning this option off.
  - Procedure

```
filer> options cifs.bypass_traverse_checking off
```

- CIFS GID checks
  - This option affects security checking for Windows clients of files with UNIX security, where the requester is not the file owner. In all cases, Windows client requests are checked against the share-level ACL. If the requester is the owner, the "user" permissions are used to determine the access permissions.
  - If the requester is not the owner, and if cifs.perm_check_use_gid is on, files with UNIX security are checked using normal UNIX rules; that is, if the requester is a member of the file's owning group, the "group" permissions are used; otherwise, the "other" permissions are used.
  - If the requester is not the owner and if cifs.perm_check_use_gid is off, files with UNIX security style are checked against the file's "group" permissions, and the "other" permissions are ignored. In effect, the "group" permissions are used as if the Windows client was always a member of the file's owning group, and the "other" perms are never used.
  - Recommended Setting: Enable CIFS GID checks to require UNIX-style security

```
filer> options cifs.perm_check_use_gid on
```

- Default Windows user
  - Specifies the Windows domain user account to use when a UNIX user accesses a file with Windows security (has an ACL) and that UNIX user would not otherwise be mapped.
  - Recommended Setting: Set the option to a null string, denying access.
  - Note: Perform this step only on multiprotocol systems that have NFS/CIFS user mapping configured correctly; disabling this access on an NFS-only NetApp storage system results in access problems for legitimate users.
  - Procedure

```
filer> options wafl.default_nt_user ""
```

- Default Unix user
  - Specifies the UNIX user account to use when a Windows domain user attempts to log in and that Windows user would not otherwise be mapped.
  - Recommended Setting: Set the option to a null string, denying access.
  - Note: Perform this step only on multiprotocol systems that have NFS/CIFS user mapping configured correctly; disabling this access on a CIFS-only NetApp storage system results in access problems for legitimate users.
  - Procedure

```
filer> options wafl.default_unix_user ""
```

- Root to admin mappings
  - When on (the default), a Windows domain administrator is mapped to UNIX root.
  - Recommended Setting: Disable root to administrator mappings by default.
  - Procedure

```
filer> options wafl.nt_admin_priv_map_to_root off
```

- Change permissions
  - When enabled, only the root user can change the owner of a file.
  - Recommended Setting: Allow only root access to change permissions to files.
  - Procedure

```
filer> options wafl.root_only_chown on
```

- Cache credentials
  - Specifies the number of minutes a WAFL® credential cache entry is valid. The value can range from 1 through 20160.
  - Recommended Setting: Set the minutes for cache credentials to 10.
  - Procedure

```
filer> options wafl.wcc_minutes_valid 10
```

- Preserve Unix security
  - Preserves UNIX permissions as files are edited and saved by Windows applications that use temporary files. Enabling this option allows UNIX file permissions to be set by using the Security tab on a Windows client. When enabled, this option causes UNIX qtrees to appear as NTFS volumes. This option affects only NFS files in UNIX or mixed-mode qtrees.
  - Recommended Setting: Enable this option if you are in a mixed UNIX and Windows environment where files are edited by cross-platform client applications
  - Procedure

```
filer> cifs.preserve_unix_security on
```

- File Policies
  - File policies specify file operation permissions according to file type. For example, you can restrict certain file types, such as .jpg and .mpg files, from being stored on the storage system. FPolicy requires CIFS to be licensed and running, even in NFS-exclusive environments.
  - Recommended Setting: Enable file policies if required by corporate security policy.
  - See the Data ONTAP File Access and Protocols Management Guide for more information on enabling file policies.

[edit]

Replication

Description	OK
Help for firewall administrators. This article explains how to set up a firewall for SnapMirror/SnapVault replication: https://now.netapp.com/Knowledgebase/solutionarea.asp?id=ntapcs15894 Common ports used by NetApp: • TCP/22 = ssh (secure shell) • TCP/23 = telnet	

• TCP/80 = http • TCP/443 = https (SSL) • TCP/10000 = NDMP (also needed for OSSV/SnapMirror/SnapVault) • TCP/10555 = Vmotion (VMware ESX) • TCP/10565 + 10566 + 10567 + 10568 + 10569 = SM/SV • TCP/10566 = OSSV	
You may find the following option useful for improved security: • options snapmirror.checkip.enable on ◦ Enables IP address-based verification of SnapMirror destination NetApp storage systems by source NetApp storage systems	

[edit]

Configure Passwords & Password Policy

Generally, you will need to change three passwords

- "root" password
- "administrator" password (of the local administrator account)
- /etc/passwd's root account (used by ftp and sometimes CIFS filers in a workgroup)

Don't forget change the /etc/passwd root password, as this may create a security problem.

Description	OK

Change the password for the built-in *root* account:
`filer> passwd`
`Login: root`
`New password:`
`Retype new password:`
`Mon Jul 24 00:42:50 CEST [filer: passwd.changed:info]: passwd for user 'root' changed.`	

Change the password for the built-in *administrator* (CIFS) account:
`filer> passwd`
`Login: administrator`
`New password:`
`Retype new password:`
`Mon Jul 24 00:41:40 CEST [filer: passwd.changed:info]: passwd for user 'administrator' changed.`	

Change the password for the *root* user in the /etc/passwd file:
`filer> cifs passwd new_password`
`password is _J9..Z8TkmBZd2pufcCg`
• Copy/paste the password hash for the root user in the file
• Note that the password will stay in the command-line history until the filer is rebooted. This is a potential security risk!	

Set up a password policy using these options:
`security.admin.authentication internal`
`security.admin.nsswitchgroup`
`security.passwd.firstlogin.enable off`
`security.passwd.lockout.numtries 4294967295`
`security.passwd.rules.enable on`
`security.passwd.rules.everyone off`
`security.passwd.rules.history 0`
`security.passwd.rules.maximum 256`
`security.passwd.rules.minimum 8`
`security.passwd.rules.minimum.alphabetic 2`
`security.passwd.rules.minimum.digit 1`
`security.passwd.rules.minimum.symbol 0`	

[edit]

Manage Administrative Accounts

Description	OK

Administrative users should be created in Active Directory, NIS, or LDAP environments when these methods of authentication are available. A new feature added in Data ONTAP 7.2 allows administrative users to be defined from NIS or LDAP external authentication. You can combine this flexibility with RBAC to limit all aspects of administration in Data ONTAP. One caveat is that only a single NIS or LDAP group is allowed to participate in administration of the NetApp storage system.
 	
To configure centralized administration, do this:
 	
`filer> options security.admin.authentication nsswitch,internal`
 	
To set the administrative group from the authentication method set in /etc/nsswitch.conf on the NetApp storage system, the following option must be set to a valid NIS or LDAP group:
 	
`filer> options security.admin.nsswitchgroup [groupname]`
 	
Additionally, edit your /etc/nsswitch.conf file appropriately	

If needed, create additional accounts, eg. • create account with non-admin privileges: `filer> useradmin user add username`	
If you want to disable the root account (must be done from another administrative user account (with the security-complete-user-control capability)): `filer> options security.passwd.rootaccess.enable off`	

If needed, set up role based access control.
 	
Some theory:
 	
There are four parts to RBAC in Data ONTAP:
 	
• **USERS** - An RBAC user is defined as an account that is authenticated on the NetApp storage system. This can be a local user, a Windows domain user, or a user in a specific NIS or LDAP group. Normal users who access data stored on the NetApp storage system are not part of this definition.
• **GROUPS** - A group is simply a collection of RBAC users. Groups are assigned one or more roles. Groups defined in Data ONTAP are separate from Windows, NIS, or LDAP groups; they are defined specifically for the purposes of assigning roles to their users. When you create new users or Windows domain users, Data ONTAP requires that you specify a group membership. It is a best practice to create appropriate groups before creating local users or Windows domain users.
• **ROLES** - Roles are defined as sets of capabilities. Data ONTAP comes with several predefined roles, which you can modify. You can also create new roles. Again, when you create new groups, Data ONTAP requires that you specify roles for the new groups. It is a best practice to create appropriate roles before creating groups or users.
• **CAPABILITIES** - A capability is defined as the privilege granted to a role to execute commands or take other specified actions. Data ONTAP uses four types of capabilities:
  ◦ Login rights: These capabilities have names that begin with “login-“ and are used to control which access methods an administrator is permitted to use for managing the system.
  ◦ CLI rights: These capabilities have names that begin with “cli-“ and are used to control which commands an administrator can use in the Data ONTAP command-line interface.
  ◦ API rights: These capabilities have names that begin with “api-“ and are used to control which application programming interface (API) commands you can used. API commands are usually executed by programs, rather than directly by administrators.
  ◦ Security rights: These capabilities have names that begin with “security-“ and are used to control the ability to use advanced commands or to change passwords for other users.
 	
You should thoroughly plan a complete RBAC implementation before execution. For additional information on role-based access control in Data ONTAP, refer to the NetApp technical report TR-3358.	

[edit]

Configure Autologout

Description	OK

Configure autologout using these options:
 	
`autologout.console.enable on`
`autologout.console.timeout 60`
`autologout.telnet.enable on`	

```
autologout.telnet.timeout    60
```

[edit]

Set Up Logging

Description	OK

Make sure audit logging is enabled:
 	
`auditlog.enable              on`
`auditlog.max_file_size       10000000`	

[edit]

Disable Unused Services

Description	OK
Disable trusted.hosts access: `filer> options trusted.hosts -`	
Disable telnet if customer doesn't want telnet access: `filer>` **`options telnet.enable off`** It is also possible to set access restriction (which IP address/hosts can connect using telnet), examples:: `filer>` **`options telnet.access host=10.0.0.1`** `filer>` **`options telnet.access all`** See **man na_protocolaccess**	
Disable ndmp if not needed: `filer>` **`options ndmpd.enable off`**	
Disable rsh if not needed: `filer>` **`options rsh.enable off`**	
Disable ftp if not needed: `filer>` **`options ftpd.enable off`**	
Disable tftp if not needed: `filer>` **`options tftpd.enable off`**	
Disable NIS if not needed: `filer>` **`options nis.enable off`**	
Disable PCNFS if not needed: `filer>` **`options pcnfsd.enable off`**	
Disable WebDAV if not needed: `filer>` **`options webdav.enable off`**	

[edit]

Set Up SSH & SSL

Description	OK
Enable SSH: `filer>` **`secureadmin setup ssh`** `SSH server supports both ssh1.x and ssh2.0 protocols.`	

```
SSH server needs two RSA keys to support ssh1.x protocol. The host key is
generated and saved to file /etc/sshd/ssh_host_key during setup. The server
key is re-generated every hour when SSH server is running.

SSH server needs a RSA host key and a DSA host key to support ssh2.0 protocol.
The host keys are generated and saved to /etc/sshd/ssh_host_rsa_key and
/etc/sshd/ssh_host_dsa_key files respectively during setup.

SSH Setup will now ask you for the sizes of the host and server keys.
 For ssh1.0 protocol, key sizes must be between 384 and 2048 bits.
 For ssh2.0 protocol, key sizes must be between 768 and 2048 bits.
 The size of the host and server keys must differ by at least 128 bits.

Please enter the size of host key for ssh1.x protocol [768] : <enter>
Please enter the size of server key for ssh1.x protocol [512] : <enter>
Please enter the size of host keys for ssh2.0 protocol [768] : <enter>

You have specified these parameters:
       host key size = 768 bits
       server key size = 512 bits
       host key size for ssh2.0 protocol = 768 bits
Is this correct? [yes] <enter>

Setup will now generate the host keys in the background. It will take a
few minutes. After Setup is finished you can start SSH server with
command 'secureadmin enable ssh'. A syslog message will be generated
when Setup is complete.

SSH Setup: SSH Setup is done. Host keys are stored in /etc/sshd/ssh_host_key, /etc/sshd/ssh_host_rsa_key and
/etc/sshd/ssh_host_dsa_key.
```

Then, enable SSHv2:

```
filer> options ssh2.enable on
```

Note: Don't enable SSHv1, as this is susceptible to man-in-the-middle attacks

You may also want to turn on/off public key authentication or password authentication:

```
options ssh.passwd_auth.enable
options ssh.pubkey_auth.enable
```

Enable SSL if customer requires this:

```
filer>  secureadmin setup ssl
Country Name (2 letter code) [US]: BE
State or Province Name (full name) [California]: Antwerp
Locality Name (city, town, etc.) [Santa Clara]: Kontich
Organization Name (company) [Your Company]: Uptime NV
Organization Unit Name (division):  NetApp Filers
Common Name (fully qualified domain name) [filer.demoroom.local]: <enter>
Administrator email:  netapp@uptime.be
Days until expires [5475] : <enter>
Key length (bits) [512] : <enter>
Sun Jul 23 19:51:23 CEST [filer: rc:info]: Starting SSL with new certificate.
```

[edit]

Set Up HTTP

Description	OK

Set correct values for the following options that control HTTP access:
 	
`httpd.acecss`
`httpd.admin.access`
`httpd.admin.enable`
`httpd.hostsequiv.enable`
`httpd.admin.ssl.enable`
`httpd.enable`	

[edit]

Encryption

To be added

[edit]

SnapVault/SnapMirror/OSSV

set volume to same language as sources - filenames with accents

do not run SnapMirror/SnapVault on iSCSI SAN network segments !

see security for network ports used info

Description	OK
Delete snapshots when not needed anymore. When you break a relationship snapshots will still hang around...	

TODO

[edit]

TODO: take performance baseline

TODO: qtree oplocks & ms access databases

TODO: OSSV & volume language settings

TODO: cifs access based enumeration

TODO: single_image mode (FCP + cluster)

TODO: options cifs.ms_snapshot_mode pre-xp for W2K clients (slight performance impact)

TODO: FlexShare ?

TODO: http://forums.netapp.com/conversation.asp?tid=1&vid=51504&thd=1&cid=214 fastpath versus same subnet etc. ip.fastpath.enable If the option is on, the filer will attempt to use MAC address and interface caching ("Fastpath") so as to try to send back responses to incoming network traffic using the same interface as the incoming traffic and (in some cases) the destination MAC address equal to the source MAC address of the incoming data. This allows for automatic load-balancing between multiple interfaces of a trunk and between multiple filer interfaces on the same subnet. Valid values for this option are on or off. The default value for this option is on. For TCP connections, the system will also automatically detect if this optimization is not feasible in a specific environment or for a specific connection and turn Fastpath off automatically for those connections for which using Fastpath is inappropriate. The netstat command with the -x option can be used to see if Fastpath is enabled for a specific connection.

netstat -x Applicable only to the first form of this command. Shows extended state information for TCP connections in the ESTABLISHED state. This includes information on whether MAC address and interface caching ("Fastpath") is in use for this connection (On, Off, or Partial). For more information on Fastpath, see the description of the option ip.fastpath.enable in the na_options (1) man page

TODO: http://now.netapp.com/NOW/knowledge/docs/ontap/rel724/html/ontap/nag/3routin5.htm

TODO: https://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb7710

TODO: https://now.netapp.com/Knowledgebase/solutionarea.asp?id=ntapcs11740

TODO: http://now.netapp.com/NOW/cgi-bin/bol?Type=Detail&Display=136637

TODO: http://now.netapp.com/NOW/cgi-bin/bol?Type=Detail&Display=32284

save a config dump !!!

preferred plex to read from bij stretch MC = alternate !!! niet local => more performant

```
options cf.takeover.change_fsid off
```

important link to check:

- https://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb37788
- http://solutions.qlogic.com/KanisaSupportSite/search.do?cmd=displayKC&docType=kc&externalId=7889264&sliceId=SAL_INTERNAL_1_4&dialogID=7650890&stateId=0%200%207640620

Does NetApp support iSCSI MPIO with one hardware initiator and one software initiator?

- no:
- https://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb16350

How to setup iSCSI MPIO on Windows 2003 using Microsoft multipathing

- https://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb23911

vmware disk timeouts in guest OSs

- https://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb37986

Uninstall of MPIO drivers causes mscs cluster to lose disk resources

- http://now.netapp.com/NOW/cgi-bin/bol?Type=Detail&Display=113541

- http://now.netapp.com/NOW/cgi-bin/bol?Type=Detail&Display=139183
- https://now.netapp.com/Knowledgebase/solutionarea.asp?id=ntapcs16249
- https://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb16038
- https://now.netapp.com/Knowledgebase/solutionarea.asp?id=ntapcs17368
- http://now.netapp.com/NOW/cgi-bin/bol?Type=Detail&Display=202007

====> PDORemovePeriod bij MPIO, MaxRequestHoldTime bij non-MPIO systems.

SrbTimeoutDelta ??

- https://now.netapp.com/Knowledgebase/solutionarea.asp?id=ntapcs17115
- https://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb2989

Retrieved from "http://wiki.uptime.be/wiki/NetApp_Post_Installation_Checklist"

Categories: NetApp | Checklists