

Environnement

Les modes

```
|# set -privilege advanced | Enter into privilege mode |  
|# set -privilege johnnastic | Enter into johnnastic mode |  
|# set -privilege admin | Enter into admin mode |
```

Vue d'ensemble

```
|# storage aggregate show -state !online | show all aggregates that are not online |  
|# system node image show | Show the running Data Ontap versions and which is the default boot |  
|# dashboard performance show | Shows a summary of cluster performance including interconnect traffic |  
|# node run * environment shelf | Shows information about the Shelves Connected including Model Number |
```

Configuration

```
|# system node run - node local sysconfig -a | Run sysconfig on the local node |  
|# node run -node -command sysstat -c 10 -x 3 | Running the sysstat performance tool with cluster mode |
```

Utilisateurs

Information

```
|# suseradmin whoami | Displays the user name of the account you are currently using |  
|# suseradmin user list | Lists all administrative users configured for this storage system|  
|# suseradmin user list user_name | Lists the extended information for a specific administrator |  
|# suseradmin user list -x | Lists the extended information for all administrators |  
|# suseradmin user list -g grp_name | Lists information for all users assigned to a specified group |  
|# suseradmin domainuser list -g group_name | Lists the Security Ids of all domain administrative users assigned to a specified group |  
|# suseradmin group list | Lists all the administrative user groups configured for this storage system |  
|# suseradmin group list group_name | Lists the extended details for a specified single group |  
|# suseradmin role list | Lists all the roles configured for this storage system. Each role entry lists the role name, comment information, and allowed capabilities  
|# suseradmin role list role_name | Lists the information for a single specified role name |
```

Configuration

```
|# security login unlock -username john | Unlock the john user |  
|# security login password -username john | Set a password for the john user |  
|# security login show -username john | Show the john user |
```

Session

```
|# system timeout modify 30 | Sets system timeout to 30 minutes |  
|# system timeout show | displays the timeout value for CLI sessions |
```

Backup

```
|# system configuration backup create -backup-name node1-backup -node node1 | Create a cluster backup from node1 |  
|# system configuration backup create -backup-name node1-backup -node node1 -backup-type node | Create a node backup of node1 |  
|# system configuration backup upload -node node1 -backup node1.7z -destination ftp://username:password@ftp.server.com | Uploads a backup file to ftp |
```

LOGS

```
|# debug vreport show | must be run in priv -set john, shows WAFL and VLDB consistency |
|# event log show -messagename scsiblade.* | show that cluster is in quorum-service |
To look at the logs within clustered ontap you must log in as the john user to a specific node
...

```

```
# set -privilege advanced
# systemshell -node
  username: john - password:
# cd /mroot/etc/mlog
# cat command-history.log | grep volume // searches the command-history.log file for the keyword volume
exit
...

```

Service Processor

```
|# system node image get -package http://webserver/306-02765_A0_SP_3.0.1P1_SP_FW.zip -replace-package true | Copies the firmware file from the webserver into the
mroot directory on the node |
|# system node service-processor image update -node node1 -package 306-02765_A0_SP_3.0.1P1_SP_FW.zip -update-type differential | Installs the firmware package to
node1 |
|# system node service-processor show | Show the service processor firmware levels of each node in the cluster |
|# system node service-processor image update-progress show | Shows the progress of a firmware update on the Service Processor |

```

Cluster

```
|# cluster statistics show | shows statistics of the cluster - CPU, NFS, CIFS, FCP, Cluster Interconnect Traffic |
|# cluster ring show -unitname vldb | check if volume location database is in quorum-service |
|# cluster ring show -unitname mgmt | check if management application is in quorum-service |
|# cluster ring show -unitname vifmgr | check if virtual interface manager is in quorum-service |
|# cluster ring show -unitname bcomd | check if san management daemon is in quorum-service |
|# cluster unjoin | must be run in priv -set admin, disjoins a cluster node. Must also remove its cluster HA partner |
|# debug vreport show | must be run in priv -set john, shows WAFL and VLDB consistency |
|# event log show -messagename scsiblade.* | show that cluster is in quorum-service |

```

Node

```
|# system node rename -node -newname | |
|# system node reboot -node NODENAME -reason ENTER REASON | Reboot node with a given reason. NOTE: check ha policy |

```

FLASH CACHE

```
|# system node run -node * options flexscale.enable on | Enabling Flash Cache on each node |
|# system node run -node * options flexscale.lopri_blocks on | Enabling Flash Cache on each node |
|# system node run -node * options flexscale.normal_data_blocks on | Enabling Flash Cache on each node |
|# node run NODENAME stats show -p flexscale | fashcache configuration |
|# node run NODENAME stats show -p flexscale-access | display flash cache statistics |

```

FLASH POOL

```
|# storage aggregate modify -hybrid-enabled true | Change the AGGR to hybrid |
|# storage aggregate add-disks -disktype SSD | Add SSD disks to AGGR to begin creating a flash pool |
|# priority hybrid-cache set volume1 read-cache=none write-cache=none | Within node shell and john mode disable read and write cache on volume1 |
```

FAIL-OVER

```
|# storage failover takeover -bynode | Initiate a failover |
|# storage failover giveback -bynode | Initiate a giveback |
|# storage failover modify -node -enabled true | Enabling failover on one of the nodes enables it on the other |
|# storage failover show | Shows failover status |
|# storage failover modify -node -auto-giveback false | Disables auto giveback on this ha node |
|# storage failover modify -node -auto-giveback enable | Enables auto giveback on this ha node |
|# aggregate show -node NODENAME -fields ha-policy | show SFO HA Policy for aggregate |
```

AGGREGATES

```
|# aggr create -aggregate -diskcount -raidtype raid_dp -maxraidsize 18 | Create an AGGR with X amount of disks, raid_dp and raidgroup size 18 | |
|# aggr offline | online | Make the aggr offline or online |
|# aggr rename -aggregate -newname | |
|# aggr relocation start -node node01 -destination node02 -aggregate-list aggr1 | Relocate aggr1 from node01 to node02 |
|# aggr relocation show | Shows the status of an aggregate relocation job |
|# aggr show -space | Show used and used% for volume foot prints and aggregate metadata |
|# aggregate show | show all aggregates size, used% and state |
|# aggregate add-disks -aggregate -diskcount | Adds a number of disks to the aggregate |
|# reallocate measure -vserver vmware -path /vol/datastore1 -once true | Test to see if the volume datastore1 needs to be reallocated or not |
|# reallocate start -vserver vmware -path /vol/datastore1 -force true -once true | Run reallocate on the volume datastore1 within the vmware vserver |
```

DISKS

```
|# storage disk assign -disk 0a.00.1 -owner | Assign a specific disk to a node | | | | | | | | | |
|# storage disk assign -count -owner | Assign unallocated disks to a node |
|# storage disk show -ownership | Show disk ownership to nodes |
|# storage disk show -state broken | copy | maintenance | partner | percent | reconstructing | removed | spare | unfail |zeroing | Show the state of a disk |
|# storage disk modify -disk NODE1:4c.10.0 -owner NODE1 -force-owner true | Force the change of ownership of a disk |
|# storage disk removeowner -disk NODE1:4c.10.0 -force true | Remove ownership of a drive |
|# storage disk set-led -disk Node1:4c.10.0 -action blink -time 5 | Blink the led of disk 4c.10.0 for 5 minutes. Use the blinkoff action to turn it off |
```

LUNS

```
|# lun show -vserver | Shows all luns belonging to this specific vserver |
|# lun modify -vserver -space-allocation enabled -path | Turns on space allocation so you can run lun reclaims via VAAI |
|# lun geometry -vserver path /vol/vol1/lun1 | Displays the lun geometry |
```

VSERVER Commands

```
|# vserver setup | Runs the clustered ontap vserver setup wizard |
|# vserver create -vserver -rootvolume | Creates a new vserver |
|# vserver show | Shows all vservers in the system |
|# vserver show -vserver | Show information on a specific vserver |
```

NFS

```
|# vserver modify -4.1 -pnfs enabled | Enable pNFS. NOTE: Cannot coexist with NFSv4 |
```

FCP

```
|# storage show adapter | Show Physical FCP adapters |
|# fcp adapter modify -node NODENAME -adapter 0e -state down | Take port 0e offline |
|# node run fcadmin config | Shows the config of the adapters - Initiator or Target |
|# node run -t target 0a | Changes port 0a from initiator or target - You must reboot the node |
```

CIFS Shares Configuration

```
|# vserver cifs create -vserver -cifs-server -domain | Enable Cifs |
|# vserver cifs share create -share-name root -path / | Create a CIFS share called root |
|# vserver cifs share show | |
|# vserver cifs show | |
```

SNAPSHOTS

```
|# volume snapshot create -vserver vserver1 -volume vol1 -snapshot snapshot1 | Create a snapshot on vserver1, vol1 called snapshot1 |
|# volume snapshot restore -vserver vserver1 -volume vol1 -snapshot snapshot1 | Restore a snapshot on vserver1, vol1 called snapshot1 |
|# volume snapshot show -vserver vserver1 -volume vol1 | Show snapshots on vserver1 vol1 |
```

DP MIRRORS AND SNAPMIRRORS

```
|# volume create -vserver -volume vol10_mirror -aggregate -type DP | Create a destinaion Snapmirror Volume |
|# snapmirror create -vserver -source-path sysadmincluster://vserver1/vol10 -destination -path sysadmincluster://vserver1/vol10_mirror -type DP | Create a
snapmirror relationship for |# sysadmincluster |
|# snapmirror initialize -source-path sysadmincluster://vserver1/vol10 -destination-path sysadmincluster://vserver1/vol10_mirror -type DP -foreground true |
Initialize the snapmirror example |
|# snapmirror update -source-path vserver1:vol10 -destination-path vserver2:vol10_mirror -throttle 1000 | Snapmirror update and throttle to 1000KB/sec |
|# snapmirror modify -source-path vserver1:vol10 -destination-path vserver2:vol10_mirror -throttle 2000 | Change the snapmirror throttle to 2000 |
|# snapmirror restore -source-path vserver1:vol10 -destination-path vserver2:vol10_mirror | Restore a snapmirror from destination to source |
|# snapmirror show | show snapmirror relationships and status |
```

NOTE: You can create snapmirror relationships between 2 different clusters by creating a peer relationship

VOLUMES

```
|# volume create -vserver -volume -aggregate -size 100GB -junction-path /eng/p7/source | Creates a Volume within a vserver |
|# volume move -vserver -volume -destination-aggregate -foreground true | Moves a Volume to a different aggregate with high priority |
|# volume move -vserver -volume -destination-aggregate -cutover-action wait | Moves a Volume to a different aggregate with low priority but does not cutover |
|# volume move trigger-cutover -vserver -volume | Trigger a cutover of a volume move in waiting state |
|# volume move show | shows all volume moves currently active or waiting. NOTE: You can only do 8 volume moves at one time, more than 8 and they get queued |
|# system node run - node vol size 400g | resize volume_name to 400GB | OR
|# volume size -volume -new-size 400g | resize volume_name to 400GB |
|# volume modify -vserver -filesystem-size-fixed false -volume | Turn off fixed file sizing on volumes |
```

SNAPVAULT

```
|# snapmirror create -source-path vserver1:vol5 -destination-path vserver2:vol5_archive -type XDP -schedule 5min -policy backup-vspolicy | Create snapvault relationship with 5 min schedule using backup-vspolicy |
```

NOTE: Type DP | asynchronous | , LS | load-sharing mirror | , XDP | backup vault, snapvault | , TDP | transition | , RST | transient restore |

NETWORK INTERFACE

```
|# network interface show | show network interfaces |
|# network port show | Shows the status and information on current network ports |
|# network port modify -node * -port -mtu 9000 | Enable Jumbo Frames on interface vif_name>
|# network port modify -node * -port -flowcontrol-admin none | Disables Flow Control on port data_port_name |
|# network interface revert * | revert all network interfaces to their home port |
```

INTERFACE GROUPS

```
|# ifgrp create -node -ifgrp -distr-func ip -mode multimode | Create an interface group called vif_name on node_name |
|# network port ifgrp add-port -node -ifgrp -port | Add a port to vif_name |
|# net int failover-groups create -failover-group data_fg -node -port | Create a failover group - Complete on both nodes |
|# ifgrp show | Shows the status and information on current interface groups |
|# net int failover-groups show | Show Failover Group Status and information |
```

UNIX

```
|# vserver services unix-user show | | | |
|# vserver services unix-user create -vserver vserver1 -user root -id 0 -primary-gid 0 | Create a unix user called root |
|# vserver name-mapping create -vserver vserver1 -direction win-unix -position 1 -pattern |.+| -replacement root | Create a name mapping from windows to unix |
|# vserver name-mapping create -vserver vserver1 -direction unix-win -position 1 -pattern |.+| -replacement sysadmin011 | Create a name mapping from unix to windows |
|# vserver name-mapping show | Show name-mappings |
```

SMB

```
|# vserver cifs options modify -vserver -smb2-enabled true | Enable SMB2.0 and 2.1 |
```

ROUTING GROUPS

```
|# network interface show-routing-group | show routing groups for all vservers |
|# network routing-groups show -vserver vserver1 | show routing groups for vserver1 |
|# network routing-groups route create -vserver vserver1 -routing-group 10.1.1.0/24 -destination 0.0.0.0/0 -gateway 10.1.1.1 | Creates a default route on vserver1 |
|# ping -lif-owner vserver1 -lif data1 -destination www.google.com | ping www.google.com via vserver1 using the data1 port |
```

DNS

```
|# services dns show | show DNS |
```

NIS

```
|# vserver services nis-domain create -vserver vserver1 -domain vmlab.local -active true -servers 10.10.10.1 | Create nis-domain called vmlab.local pointing to 10.10.10.1 |
|# vserver modify -vserver vserver1 -ns-switch nis-file | Name Service Switch referencing a file |
|# vserver services nis-domain show | |
```

Gestion du temps

NTP

```
|# system services ntp server create -node -server | Adds an NTP server to node_name |
|# system services ntp config modify -enabled true | Enable ntp |
|# system node date modify -timezone | Sets timezone for Area/Location Timezone. i.e. Australia/Sydney |
|# node date show | Show date on all nodes |
```

Date

```
|# timezone -timezone Australia/Sydney | Sets the timezone for Sydney. Type ? after -timezone for a list |
|# date 202211010830 | Sets date for yyymmddhhmm |
|# date -node | Displays the date and time for the node |
```

CONVERGED NETWORK ADAPTERS | FAS 8000 |

```
|# ucadmin show -node NODENAME | Show CNA ports on specific node |
|# ucadmin -node NODENAME -adapter 0e -mode cna | Change adapter 0e from FC to CNA. NOTE: A reboot of the node is required |
```

PERFORMANCE

```
|# show-periodic -object volume -instance volumename -node node1 -vserver vserver1 -counter total_ops|avg_latency|read_ops|read_latency | Show the specific counters for a volume | |
|# statistics show-periodic 0object nfsv3 -instance vserver1 -counter nfsv3_ops|nfsv3_read_ops|nfsv3_write_ops|read_avg_latency|write_avg_latency | Shows the specific nfsv3 counters for a vserver |
|# sysstat -x 1 | Shows counters for CPU, NFS, CIFS, FCP, WAFL |
```

Documentation

<https://arkit.co.in/netapp-cluster-mode-commands-cheat-sheet/>

<https://library.netapp.com/ecmdocs/ECMP1155684/html/GUID-89A9ACCA-501C-42DB-949B-B57B9AFBBB98.html>

<https://community.netapp.com/t5/ONTAP-Discussions/Data-ONTAP-Basic-Command-Please/m-p/59641/thread-id/13975?attachment-id=1682>

<https://community.netapp.com/t5/ONTAP-Discussions/Data-ONTAP-Basic-Command-Please/m-p/59641>