

VCS 4.0 Command Quick Reference

Installation

```
installsf      VCS 4.0/Storage Foundation
installer      4.0
installvcs     Pre-4.0
```

Status

```
lltconfig
lltstatus -nvv|-c|-l
gabconfig -a
hastatus [-sum]
```

Communications Config

```
eeprom "local-mac-address=true"
```

```
/etc/llthosts:
  0 sys_name1
  1 sys_name2
```

```
/etc/llttab:
  set-node      1
  set-cluster   10
  link qfe0 /dev/qfe:0 - ether - -
  link qfe4 /dev/qfe:4 - ether - -
  link-lowpri hme0 /dev/hme:0 - ether - -
```

```
/etc/gabtab:
  gabconfig -c -n 2
  [-j : panic on split brain]
```

```
/etc/VRTSvcs/conf/config/main.cf
Optional: /etc/VRTSvcs/conf/sysname
```

Config Manipulation

```
haconf -makerw
haconf -dump
haconf -dump -makero
hacf -verify /etc/VRTSvcs/conf/config
hasnap -backup -f /dir/file.tar -n -m myclus
hacf -cftocmd /etc/VRTSvcs/conf/config
```

Start/Stop Commands

```
lltconfig -c      Start LLT
gabconfig -c -x   Ignore -n option
/etc/init.d/vxfen start Start IO fencing
hastart
hastop -all       Offlines all SGs
hastop -all -force Leaves SG apps running
hasys -force sys_name Clear STALE_ADMIN_WAIT
hastart -stale    Force remote build
hastop -local -evacuate Move SGs off this node
/etc/init.d/vxfen stop Stop fencing driver
gabconfig -U     Stop GAB
lltconfig -U     Stop LLT
```

User Accounts

```
hauser -add username
hauser -display
hauser -delete username
```

Set User's Password:

```
hauser -update username
```

Map Unix users to VCS users:

```
(VCS & Unix usernames must exist)
haclus -modify -change AllowNativeCliUsers=1
```

Cluster-Wide Users:

```
haclus -modify Administrators|Operators \
        -add username
```

Service Group Users:

```
hagr -modify sg_name \
        Administrators|Operators -add username
```

Logging

```
In /var/VRTSvcs/log,
engine_A.log, hashadow_A.log, agent_name_A.log
```

Service Groups

```
hagr -add sg_name
hagr -modify sg_name SystemList sys_name1 0 \
        sys_name2 1
hagr -modify sg_name AutoStartList \
        sys_name1 sys_name2
hagr -display sg_name
hagr -online|-offline sg_name -sys sys_name
hagr -online|-offline sg_name -any
```

Temporary Freeze:

```
hagr -switch sg_name -to sys_name
hagr -freeze|-unfreeze sg_name
```

Persistent Freeze:

```
haconf -makerw
hagr -freeze|-unfreeze sg_name -persistent
haconf -dump -makero
hagr -display sg_name -attribute Frozen
```

Dependencies:

```
hagr -link sg_name1 sg_name2 online local \
        soft|firm|hard
hagr -link sg_name1 sg_name2 online \
        global|remote soft|firm
hagr -link sg_name1 sg_name2 offline local
```

See tables in course documentation for explanation of categories.
Maximum 3 levels, only one child per parent.
Use FileOnOff/ElifNone for additional links.

Flushing Stuck Service Groups:

```
hagr -flush sg_name -sys sys_name
```

Fault Handling:

Set SG parameters, refer to flowcharts:
ManageFaults All|None
FaultPropagation 0|1
AutoFailover 0|1

Clearing ADMIN_WAIT:

```
hagr -clearadminwait sg_name -sys sys_name
hagr -clearadminwait -fault sg_name \
        -sys sys_name
```

The latter form causes resource to be faulted.

Clear AUTODISABLE:

```
hagr -autodisable sg_name -sys sys_name
```

Resources

```
hares -modify res_name Enabled 0|1
hares -modify res_name Critical 0|1
hares -online|-offline res_name -sys sys_name
hares -clear res_name [-sys sys_name]
hares -probe res_name -sys sys_name
hares -delete res_name
```

Dependencies:

```
hares -link parent_res child_res
hares -dep
```

NIC:

```
hares -add nic_name NIC sg_name
hares -modify nic_name Device qfel
```

IP:

```
hares -add ip_name IP sg_name
hares -modify ip_name Device qfel
hares -modify ip_name Address 192.168.1.1
```

DiskGroup:

```
hares -add dg_res_name DiskGroup sg_name
hares -modify dg_res_name Diskgroup dg_name
```

Volume:

```
hares -add vol_res_name Volume sg_name
hares -modify vol_res_name Volume vol_name
hares -modify vol_res_name DiskGroup dg_name
```

Mount:

```
hares -add res_name Mount sg_name
hares -modify res_name MountPoint /somewhere
hares -modify res_name BlockDevice /dev/vx...
hares -modify res_name FSType ufs
hares -modify res_name FsckOpt %y
hares -modify res_name MountOpt logging
```

Process:

```
hares -add process_name Process sg_name
hares -modify process_name PathName /bin/sh
hares -modify process_name Arguments \
        "/somewhere/filename args"
```

Proxy:

```
hares -add proxy_name Proxy sg_name
hares -modify proxy_name TargetResName \
        real_res_name
```

Phantom (Parallel SGs Only):

```
hares -add phantom_name Phantom sg_name
```

Notifier (ClusterService SG only):

```
hares -add notifier_name NotifierMgr \
ClusterService
hares -modify notifier_name SmtServer \
localhost|ip_address
hares -modify notifier_name SmtpRecipients \
-add username level
Level can be:
Information, Warning, Error, SevereError.
To send emails to a resource or SG "owner":
hares -modify res_name ResourceOwner username
hagrp -modify sg_name GroupOwner username
```

MultiNICB (Use with IPMultiNICB):

Configure IPMP as normal, but with test addresses only (no failovers).

```
hares -add res_name1 MultiNICB sg_name
hares -modify res_name1 Device qfe0 0 qfe1 1
hares -add res_name2 IPMultiNICB sg_name
hares -modify res_name2 Address 192.168.2.1
hares -modify res_name2 BaseResName res_name1
hares -modify res_name2 Address 192.168.2.1
To switch address between interfaces:
haipswitch res_name1 192.168.2.1 qfe0 qfe1
```

Disk Heartbeats:

Add ServiceGroupHB resource to the bottom of the resource dependency tree. Use a dedicated raw disk.

Localizing Attributes:

```
hares -local res_name attribute_name
hares -modify res_name attribute_name \
local_value1 -sys sys_name1
hares -modify res_name attribute_name \
local_value2 -sys sys_name2
```

Restarts:

Set RestartLimit, ConfInterval and ToleranceLimit for the resource

Changing Global Resource Attributes:

```
hatype -modify res_type attribute_name value
Overriding Resource Attributes Per-Resource:
hares -override res_name res_type
attribute_name
hares -modify res_name attribute_name value
hares -display -ovalues res_name
hares -undo_override res_name attribute_name
```

Configuring Storage

```
vxdisksetup -i disk_dev
vxdg init demoDG demoDG01=disk_dev
vxdg adddisk, vxassist make, newfs/mkfs etc
mkdir /demo
No auto mount entry in /etc/vfstab!
vxdg deport demoDG
```

IO Fencing

Use 3 dedicated SCSI3 LUNs for co-ordinator disks. Disks must be h/w mirrored - cannot be

replaced online.

Set up dedicated disk group:

```
vxdisksetup -i disk1 [disk2, disk3]
vxdg init fen_dg disk1 disk2 disk3
vxfsentsthdw -g fen_dg
vxdg deport fen_dg
vxdg import -t -g fen_dg
vxdg deport fen_dg
echo "fen_dg" > /etc/vxfendg
/etc/init.d/vxfendg start
```

Test SCSI3 on data disks - SGs must be offline or they will fault:

```
vxfsentsthdw -rg data_dg
```

Stop HAD & SGs with hastop -all

Set "UseFence = SCSI3" in cluster section of main.cf.

Restart HAD with hastart [-stale]

```
haclus -value UseFence
```

To view keys:

```
vxfsenadm -g all -f /etc/vxfentab
vxfsenadm -g|-r /dev/rdisk/clt0d0s2
```

To clear keys after a double failure:

```
vxfsenclearpre
```

Removing A Node From A Running Cluster

Stop and freeze the node:

```
hasys -freeze persistent \
-evacuate sys_name
hastop -sys sys_name
```

Remove any disk heartbeats.

Stop communications and fencing:

```
/etc/init.d/vxfen stop
gabconfig -U
lltconfig -U
modinfo | grep gab
modunload -i gab_id
modinfo | grep llt
modunload -i llt_id
```

Disconnect Interconnects.

Uninstall VCS software, if necessary.

On a node still in the cluster:

```
hagrp -modify sg_name SystemList \
-delete sys_name
hagrp -modify sg_name AutoStartList \
-delete sys_name
hasys -delete sys_name
haconf -dump -makero
```

On all nodes still in cluster:

```
Remove node from /etc/llthosts
Reduce value in /etc/gabtab
```

Adding A New Node To A Running Cluster

Connect Interconnects.

Install VCS using "installvcs -installonly".

Install licence keys.

Configure /etc/llttab, /etc/llthosts,

/etc/gabtab on new system.

Configure fencing driver diskgroup on new system, if used.

Update /etc/llthosts, /etc/gabtab on existing systems.

Install any Enterprise Agents on new system.

Copy trigger scripts, custom agents, etc from cluster to new system.

Start the new node:

```
lltconfig -c
gabconfig -c -n 3
hastart
gabconfig -a
```

Update SGs to use new system as required:

```
hagrp -modify sg_name SystemList \
-add sys_name priority
hagrp -modify sg_name AutoStartList \
-add sys_name
```

Verify SGs are switchable to new system:

```
hagrp -switch sg_name -to sys_name
```

Oracle Resources

```
hares -add ora_res Oracle sg_name
hares -modify ora_res Res Home /u01
hares -modify ora_res Owner oracle
hares -modify ora_res Sid SID
hares -modify ora_res Pfile \
"/u01/admin/SID/pfile/initSID.ora"
```

Listener

```
hares -add lsnr_res Netlsnr sg_name
hares -modify lsnr_res Home /u01
hares -modify lsnr_res Owner oracle
hares -modify lsnr_res TnsAdmin \
/u01/network/admin
hares -modify lsnr_res Listener SID
Edit listener.ora & tnsnames.ora to use the SG's virtual IP, and use a named listener.
```

Oracle Detail Monitoring

```
sqlplus "/" as sysdba"
create user vcs identified by password
default tablespace users quota 500K on users;
grant create session to vcs;
create table vcs.vctest (tstamp date);
insert into vcs.vctest (tstamp) values
(SYSDATE);
```

```
/opt/VRTSvcs/bin/vcsencrypt -agent password
hares -modify ora_res Pword enc_password
hares -modify ora_res User vcs
hares -modify ora_res DetailMonitor 1
hares -modify ora_res Table vctest
hares -modify ora_res MonScript \
/opt/VRTSvcs/bin/Oracle/SqlTest.pl
```

Check that probe is updating:

```
select TO_CHAR(tstamp, 'MON DD, YYYY HH:MI:SS
AM') tstamp from vcs.vctest;
```