



***HIGH-AVAILABILITY***

***Sun 3310 Storage Array (SCSI) for Clustering***

## Introduction

This guide explains the setup of the 3310 (SCSI) for use with clustering on Solaris and is intended to be used by customers and partners when setting up the system hardware, prior to High-Availability.Com or trained engineers beginning to attempt to install the clustering software.

This guide may also be used to guide the setup of other RAID devices, although the specifics are targeted only at 3310 hardware, the theory applies to all configurations.

This guide is not intended to replace training and documentation provided by Sun and others but to be supplementary information to accelerate the successful deployment of clustering.




## Collecting Information

The 3310 is usually setup to accommodate an application(s) and to set this up correctly the consultant installing and configuring the application should be consulted prior to setting up the 3310 disks.

In addition, **RSF-1 requires a dedicated logical disk (logical partition)** of absolute minimal size to be configured, on which no other partitions etc will be configured. It does not require mirroring but can be mirrored. This allows RSF-1 to have heartbeats setup, which will not have any conflicts with disk management products, like Solstice Disk Suite (SDS), now known as Solstice Volume Manager.

## Hardware Configuration


The clustering software relies on the fact that all disks can be used and managed by all attached systems. This means that the 3310 array must be configured with a single bus. This requires both hardware and software configuration.



**Figure 1.**  
Single Bus, Dual Attached

## Physical Connections

The hardware setup, to support single bus, dual attach setup of the 3310 SCSI array requires the connection of a 'SCSI Jumper Cable', which is supplied with the 3310. The cable must be connected between Channel 0 and 'Single Bus Configuration' connectors, as shown in Figure 2. below.



**Figure 2.**  
SCSI Jumper Cable

## Logical Setup

The RAID must be setup in the normal way to provide logical drives to the hosts. The drives must then be added to the SCSI channels 1 & 3. **This MUST be done using Primary Channel ONLY.** This can be verified by checking the setup on the 3310, 'host luns' menu option, which is used to assign the disks to the hosts. If channel 3 is shown as a secondary/redundant channel then delete the SCSI id for the channel and re-add it so that is on the primary channel and **MATCHES** the SCSI id given to channel 1.

To setup the RSF-1 Heart Beat (HB) logical disk;

A logical drive can not be created using less than an entire disk, which means that we must use logical volumes to provide a partition for disk heartbeats that will not be affected by disk reservations required by some disk management software. Using one of the 'logical drives' per 3310, create a 'logical volume' for use as the heartbeat



disk. This should be of minimal size. Use the rest of this logical drive to create another logical volume for normal data.

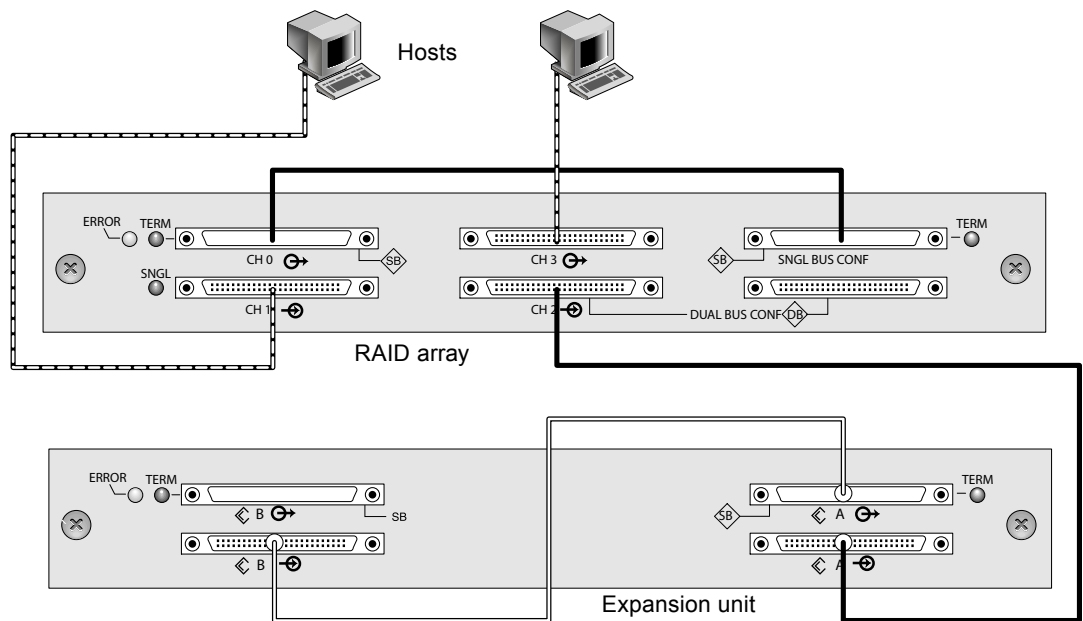
Now assign the logical disks to BOTH channels in the same order. This ensures that the visible path from the Sun machines will be the same. This avoids confusion and greatly simplifies setup later, especially with SDS.

Write the setup to NVRAM if required and you may need to reset to apply some of the changes.

### Setup with Expansion

The use of the 3310 allows additional units to be added which don't have the controller fitted. These expansion units can be connected and again must be configured in a single bus as shown in Figure 3. below.

**Figure 3.**  
3310 SCSI Array  
With Expansion Units



### Solaris Setup

Solaris needs to be configured to provide support for additional LUNs and /kernel/drv/sd.conf will need to be modified.

Depending on the system/adapters etc you may need to modify another file. Full details are available at Sun's web site [here](#).

Add the following lines to the end of /kernel/drv/sd.conf

```
name="sd" class="scsi"
    target=0 lun=1;
name="sd" class="scsi"
    target=0 lun=2;
name="sd" class="scsi"
    target=0 lun=3;
name="sd" class="scsi"
    target=0 lun=4;
name="sd" class="scsi"
    target=0 lun=5;
name="sd" class="scsi"
    target=0 lun=6;
name="sd" class="scsi"
    target=0 lun=7;
```

If an expansion unit is fitted you will need to add extra lines (six per unit).

This configuration example is based on the assumption that you have used target 0 as the SCSI id when setting up the mapping of the logical drives to SCSI controllers but you may need to change the target= value if this is not the case.

The following lines are normally included but you should check they exist. If adding for a different target SCSI id then these lines may need to be modified and added.

```
name="sd" class="scsi"
```

```
target=0 lun=0;
```

Ensure that Sun StorEDGE Traffic Manager Software (STMS) is enabled and set to auto failback by editing;

```
/kernel/drv/scsi_vhci.conf
```

By default the setting for **mpxio-disable** is "yes", this must be changed to "no";

```
mpxio-disable="no";
```

By default the setting for **auto-failback** is "disable", this must be changed to "enable";

```
auto-failback="enable"
```

*If you change these settings, you **MUST** reboot the machine.*

You can check STMS is installed by verifying there is an entry in /etc/name\_to\_major;

```
# grep scsi_vhci /etc/name_to_major
```

If is not installed you can download it from Sun. You will need the following packages; SUNWsan, SUNWcpl, SUNWcplx and will need the appropriate patches from Sun's SunSolve.

Now force the OS to read this modified configuration;

For Solaris 9+;

```
update_drv -f sd ; devfsadm
```

For Solaris 8 and before;

```
reboot -- -r
```

Don't forget that this needs to be done on both systems. Use 'format' to verify that all disks can be accessed and that they have the same path.

### Possible Problems With Disk Suite (metasets)

When adding a disk to a multi-node diskset or when trying to add another node to the set you may get an error, indicating that the disk is not common between the nodes. The common fibre controller used is the Qlogic and a numbers of lines with 'qlc' would be key if this is the controller you have. If you get this error then the OS links on the *new* node will need to be adjusted to match those of the existing node(s). **Always make a backup of the /etc/path\_to\_inst file before making changes.**

Here is an example extract from /etc/path\_to\_inst that needs to be changed from;

```
...
"/pci@1d,700000/SUNW,qlc@1" 0 "qlc"
"/pci@1d,700000/SUNW,qlc@1/fp@0,0" 0 "fp"
"/pci@1d,700000/SUNW,qlc@1/fp@0,0/sdd@w216000c0ff87e8d8,0" 0 "ssd"
"/pci@1d,700000/SUNW,qlc@1/fp@0,0/sdd@w256000c0ffc7e8d8,0" 1 "ssd"
"/pci@1d,700000/SUNW,qlc@1/fp@0,0/sdd@w256000c0ffc7e8d8,1" 2 "ssd"
"/pci@1d,700000/SUNW,qlc@1/fp@0,0/sdd@w256000c0ff87e8d8,0" 32 "ses"
"/pci@1d,700000/SUNW,qlc@1/fp@0,0/sdd@w256000c0ffc7e8d8,0" 33 "ses"
"/pci@1d,700000/SUNW,qlc@1/fp@0,0/sdd@w256000c0ff87e8d8,0" 5 "ssd"
...
"/scsi_vhci" 0 "scsi_vhci"
"/scsi_vhci/ssd@g600c0ff00000000007e8d85fe6ae6201" 3 "ssd"
"/scsi_vhci/ssd@g600c0ff00000000007e8d85fe6ae6200" 4 "ssd"
"/scsi_vhci/ssd@g600c0ff00000000007e8d8000000000" 34 "ses"
"/scsi_vhci/ssd@g600c0ff00000000007e8d85fe6ae6202" 6 "ssd"
...
```

Changing to match the other machine;

```
...
"/pci@1d,700000/SUNW,qlc@1" 0 "qlc"
"/pci@1d,700000/SUNW,qlc@1/fp@0,0" 0 "fp"
"/pci@1d,700000/SUNW,qlc@1/fp@0,0/sdd@w256000c0ffc7e8d8,0" 3 "ssd"
"/pci@1d,700000/SUNW,qlc@1/fp@0,0/sdd@w216000c0ff87e8d8,0" 4 "ssd"
...
"/scsi_vhci" 0 "scsi_vhci"
"/scsi_vhci/ssd@g600c0ff00000000007e8d8000000000" 32 "ses"
"/scsi_vhci/ssd@g600c0ff00000000007e8d878ba51fc01" 0 "ssd"
"/scsi_vhci/ssd@g600c0ff00000000007e8d878ba51fc00" 1 "ssd"
"/scsi_vhci/ssd@g600c0ff00000000007e8d85fe6ae6201" 2 "ssd"
"/scsi_vhci/ssd@g600c0ff00000000007e8d85fe6ae6200" 5 "ssd"
...
```



Edit `/etc/path_to_inst` on the *new* node and adjust the appropriate lines to match lines on the other nodes. The Solaris documentation should be consulted if you are unsure about this. Then delete the files referring to the disk, for example, if the controller was 'c4' and was a 'Low Voltage Differential (LVD) SCSI channels' - which is also the path used for Fibre Channels;

```
rm /dev/dsk/c4* /dev/rdisk/c4 /devices/scsi_vhci/*
```


Then reboot to cause the `path_to_inst` changes to take effect. Rebooting with the following syntax will cause the file references to be rebuilt correctly;

```
reboot -- -r
```

Please use caution when performing these tasks. Imprudent modifications to the `path_to_inst` file or deleting the wrong file references may make prevent the system from booting. Always make a backup of the `path_to_inst` file before making changes.

### What Enterprise Demands

Figure 4., below, shows the basic connections that need to be made. If the second network connection is not in use then a crossover network cable should be used as an additional heartbeat path.



**Figure 4.**  
V440s & 3310 Example  
Configuration

