How do I configure a RAID 5 array using the Virtual Disk Manager?

Keywords

unix openserver openserver5 v5 enterprise 5.0 virtual disk manager vdm raid5 raid 5 configure array

Release

SCO OpenServer Enterprise System Release 5.0.0
SCO OpenServer Desktop System Release 5.0.0
SCO OpenServer Host System Release 5.0.0
Virtual Disk Manager Release 1.1.0

Problem

How do I configure a RAID 5 array using the Virtual Disk Manager?

RAID 5 is based on data striping (as in RAID 0) and parity (as in RAID 4). Data striping distributes data blocks across pieces stored on multiple disks. The difference between RAID 5 and RAID 4 is that parity is striped across all disks. Because parity information is distributed, no one disk bears an excessive I/O load. Thus, RAID 5 is preferred to RAID 4. Parity is defined as redundancy (or error-correction) information generated during writes. If a disk in the array fails, the parity information and the partial data on the remaining disks are used to reconstruct the data. It requires a minimum of three pieces on separate disks. You should not create three pieces on one disk, as disk failure would disallow any method to rebuild it with RAID 5.

The following example assumes three 525 megabyte SCSI disks are attached to one SCSI host adapter. None are part of the root or bootable hard drive, and the entire disk space for each of the drives will be used in the RAID 5 array.

Solution

Virtual Disk Manager Release 1.1.0 includes a Release Supplement disk. Before proceeding, make sure it is installed.

Before proceeding, obtain the following information about the SCSI hard disks:

* host adapter type and its prefix (spad, eiad, arad, alad, and so on)

* host adapter number

* bus id

Note: You will be asked for the bus id. For a dual channel host adapter, Channel A will be bus id =0 and Channel B will be bus id=1. The host adapter number will be the same for either channel. For example, when trying to add a SCSI disk that is attached to a host adapter, and the host adapter is the first of its type in the system, the host adapter number is 0, regardless of the channel the device is attached to.

* target id of the SCSI disk.

After obtaining the above information you are ready to proceed. You must be in System Maintenance mode to continue. Type:

```
# mkdev hd
```

You will see:

```
1. Add a hard disk to IDE controller
2. Add a hard disk to SCSI controller
3. Add a hard disk to IDA controller
```

Choose 1, 2 or 3.

```
Enter the prefix of the SCSI host adapter that supports
this device or press  <Return> for the default of ad.
```

Note: This may not always read 'ad' as the default.

Enter the prefix. You will see:

```
Which xxxx SCSI host adapter supports this device?
```

where 'xxxx' is the prefix you entered previously.

Enter the host adapter number.

```
The following parameters will be used to configure xxxx SCSI
host adapter z.
Change these parameters y/n
```

where 'xxxx' is the prefix you entered previously and 'z' is the host adapter number.

You will see the setup parameters. If they are correct, choose 'n'. If you need to change them, choose 'y' and enter the hardware details about the host adapter card; confirm that you want to save these values.

```
What SCSI bus is this device attached to? Press <Return>
to use the default:0
```

Enter bus id. You will see:

```
What is the target id for this device?
```

Note: This is the target id of the hard drive, not the host adapter.

Enter target id. You will see:

```
What is the LUN?
```

Enter 0. You will see:

```
You are about to add the following SCSI device.

Host Adapter Type     Device   Adapter    Id   Lun    Bus
-------------------------------------------------------
xxxx                    Sdsk     0         1     0      0
```

where 'xxxx' is the prefix added previously.

The assumed host adapter is 0, target id of scsi disk is 1, lun is always 0, and, in this case, the bus id was 0. You will see:

```
Update SCSI configuration? (y/n)
```

Choose 'y' if all is as it should be. You will see:

```
The SCSI configuration file has been updated.

A new kernel must be built and rebooted before disk
configuration can continue.
Would you like to relink at this time (y/n)
```

Choose 'n' and add the other disks following the same directions.

After adding each drive, you will select 'y' to have the kernel relinked. You can continue on at this point. You will see:

```
Do you want the kernel to boot by default?
```

Choose 'y'. You will see:

```
Do you want the kernel environment rebuilt?
```

Choose 'y'.

Reboot the system and enter System Maintenance mode.

You must run 'mkdev hd' again and enter the same parameters as you did previously in order to finish configuring the hard drives. You must perform this process separately for each drive. After adding the same parameters for the first drive you will see:

```
During installation you may choose to overwrite all or part
of the present contents of your hard disk.
Continue y/n
```

Choose 'y'. You will see:

```
The hard disk installation program will now invoke /etc/fdisk.
Entering 'q' at the following menu will exit /etc/fdisk,
and the hard disk installation will continue.

1. Display Partition Table
2. Use Entire Disk for Unix
3. Use Rest of Disk for Unix
4. Create Unix Partition
5. Activate Partition
6. Delete Partition
7. Create Partition

Enter choice or q to quit
```

Note: For this example, the entire disk for Unix will be used for each of the three disks.

Choose "2. Use Entire Disk for Unix".

Note: If you previously had data on the disk, you will see this message:

```
Warning! All data on your disk will be lost!
Do you wish to continue (y/n)
```

Choose 'y'. You will see:

```
Total disk size .....
Press <return> to continue

1. Scan entire Unix partition
2. Scan a specified range of blocks
3. Scan a specified filesystem
4. List current bad block table
5. Add entries to bad block table
6. Delete entries from bad block table
7. Clear bad block table
```

```
8. Re-allocate bad blocks

Enter choice or q to quit
```

Choose "1. Scan entire Unix partition". You will see:

```
The virtual disk driver is installed.
Do not create a division table on partitions
which will be used for virtual disk pieces.

Choose one of the following options below or quit.

`y' to create a division table on the active partition

`n' to avoid creating a division table

 Enter `y', `n', or enter `q' to quit
```

Choose 'q' to quit.

IMPORTANT NOTE: YOU MUST CHOOSE Q TO QUIT HERE AS PREVIOUS NOTICE WARNS YOU! CREATING DIVISION TABLES ON PARTITIONS BEING USED FOR RAID 5 WILL PREVENT THEM FROM EASILY BEING CONFIGURED IN THE ARRAY.

When you have completed this for each of the drives, enter multiuser mode and bring up the Virtual Disk Manager. To do this, click on the System Administration ==>Filesystems ==> Virtual Disk Manager.

You will see:

```
+-------------------------------------------------------------+
|       Disk   Piece   Database   Boot   View   Options       |
|                                                             |
+-------------------------------------------------------------+
```

Click on 'Disk'. A list will appear. Click on 'New'. This will bring up the following window:

```
+-------------------------------------------------------------+
|                                                             |
|   Select the virtual disk type and device name              |
|                                                             |
|   Virtual type:   +--------------+                          |
|                   | Raid  5      |                          |
|                   +--------------+                          |
|         Device:   +--------------+                          |
|                   | vdisk 4      |                          |
|                   +--------------+                          |
|   +---Vdisk parameters-----------------------------------+   |
|   |                                                     |   |
|   |   Number of data pieces:   +-----------+            |   |
|   |                            | 3         |            |   |
|   |                            +-----------+            |   |
```

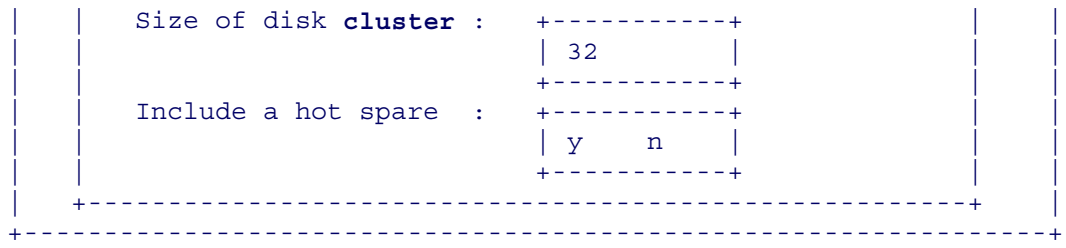

"Number of data pieces" - In this example, there are 3. Three 540 megabyte drives have been added and the entire disk will be used for each in this Raid 5 array example.

"Size of disk **cluster**" - A **cluster** consists of a set of contiguous blocks of data written to a physical hard disk within a disk array. In other words, it is the size of the data chunks written to the pieces in an array. The **cluster** size is the parameter that has the most impact on performance of array systems. Improperly matched I/O and **cluster** sizes can adversely affect performance. By default, you will see 32. This figure is an approximation based on test results, but may not be ideal for all applications. Modifying **cluster** size will require reconfiguration. That will mean backing up the data, reconfiguring the array with the new **cluster** size and restoring from backups.

"vdisk4" - Assuming the release supplement has been installed, vdisk4 will show in device if this is the first array you are configuring.

Click on 'Continue' after making any necessary modifications. You will see this table:

Click on 'Allocate piece' with 'Piece 1' highlighted.

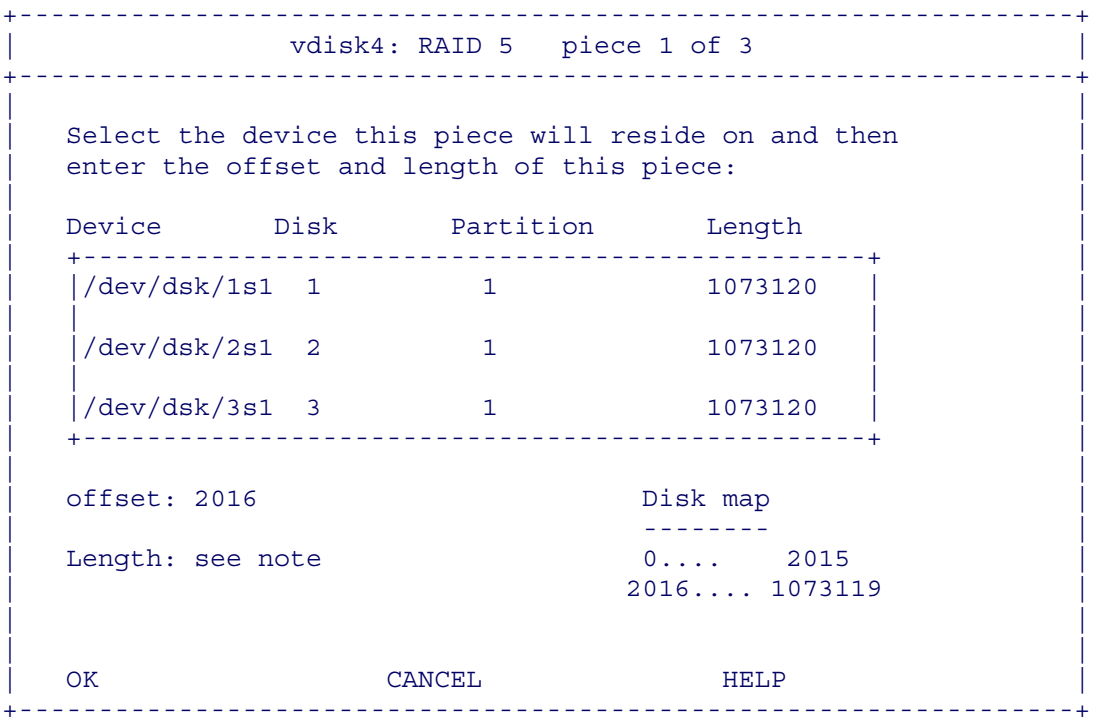

"Offset" - This space is used by the Virtual Disk Manager for protecting the fdisk table and for bookkeeping.

"Length" - Creating disks with parity data (mirror, RAID 5, RAID 4) where whole disks are being used as pieces in the array, the default piece size should not be accepted. Some arrays give disk errors during restoration and the virtual disk goes off-line. As a workaround, select a piece size of about 5% smaller than the total disk size.

```
Take 1073119-2016=1071103X.05=53555.15
Now 1071103-53555.15=1017548
Change the default length to 1017548
```

You may see a message that the length is not a multiple of the **cluster** size and it will be changed automatically.

Click on 'OK' to allocate the piece after making necessary modifications. Continue allocating pieces. You will see /dev/dsk/2s1 highlighted next. There will be no need to adjust the length for the other pieces; it will set the size for you, as all pieces must have the same size. In fact, it will be shaded out to prevent you from modifying it.

Click on 'OK' after making any necessary modifications for each piece. After all have been allocated, click on 'CREATE' to create the virtual disk.

You will then be prompted as to whether you want to restore parity. Click on 'OK' to restore parity and then choose the filesystem type. Click on 'OK' again. You will then receive an informational message stating:

```
Successfully created XXXX filesystem on /dev/dsk/vdisk4
use the Filesystem Manager to administer this filesystem.
```

where 'XXXX' is the filesystem type you chose earlier.

Click on 'OK'.

You must now bring up the Filesystem Manager so you can mount the filesystem and access the data. You can do this on the Desktop by clicking on System Administration ==>Filesystems ==>Filesystem Manager, which brings you to this screen:

```
+-------------------------------------------------------------------+
|          Mount   Status of filesystems on <machine name>          |
|                                                                   |
|    [  /dev/boot            /stand           read-only             |
|                                                                   |
|    [  /dev/root           /stand            read-write            |
|                                                                   |
+-------------------------------------------------------------------+
```

Note: You may see other filesystems listed, depending on your system.

Click on Mount==>Add Mount Configuration.

```
+-------------------------------------------------------------------+
|                                                                   |
|   Device File: /dev/dsk/vdisk4                                    |
|                                                                   |
|   Mount Point: /u1                                                |
|                                                                   |
|   Description optional:                                           |
|                                                                   |
|   +---------------------------------+  +--Mount Options------+  |
|   |   Filesystem Type   HTFS        |  | Check and Repair    |  |
|   |                                 |  |                     |  |
|   |   Access Mode:  Read-Write      |  | Advanced Options    |  |
|   |                                 |  |                     |  |
|   |   Can user mount?               |  | When to mount       |  |
|   |                                 |  |                     |  |
|   |                                 |  | [] Now              |  |
|   |                                 |  | [] At system startup|  |
|   +---------------------------------+  +---------------------+  |
|                                                                   |
|   OK                    CANCEL                  HELP              |
+-------------------------------------------------------------------+
```

Note: In this example, /u1 was chosen as the mount point. So, to access data, you would cd to /u1. To fsck the filesystem, however, you would umount by typing umount /u1 and type

fsck /dev/dsk/vdisk4.

After making any necessary modifications, click on 'OK'.

You will see a message about adding mount configuration. A window will come up showing this new filesystem added. You can now cd to /u1 to access the filesystem.

IMPORTANT NOTE: IT IS ESSENTIAL THAT GOOD RELIABLE BACKUPS OF THE FILESYSTEM EXIST. IF A DRIVE FAILS, YOU CAN JUMPER THE REPLACEMENT DRIVE TO THE SAME ID AS THE FAILED DRIVE. RUN MKDEV HD AS YOU DID THE PREVIOUS DRIVE. SINCE YOU ALREADY HAVE AN ENTRY IN YOUR CONFIGURATION FILES FOR A DRIVE AT THIS ID, YOU WILL BE BROUGHT INTO YOUR SECOND PASS WHICH WILL BRING YOU INTO FDISK. FOLLOW SAME STEPS IN FDISK, CHOOSE QUIT TO PREVENT ENTERING DIVVY AND THEN GO TO MULTI-USER MODE. PARITY SHOULD AUTOMATICALLY BE RESTORED. IN THE UNLIKELY EVENT THAT TWO DRIVES FAIL YOU WILL HAVE TO RESTORE FROM BACKUPS.

See Also

SCO OpenServer 5.0 System Administration Guide, Chapter 8.

TA105005 created on 22 February 1996 , last updated on 21 June 1996
SSL #: 482803 IT #: os2803