Sun StorEdge™ T3 Dual Storage Array - Part 2

Configuration

By Mark Garner - Enterprise Engineering

Sun BluePrints™ OnLine - March 2001

http://www.sun.com/blueprints

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303 USA
650 960-1300 fax 650 969-9131

Part No.: 816-0089-10
Revision 01, 02/08/01
Edition: March 2001

Please
Recycle

Sun StorEdge™ T3 Dual Storage Array - Part 2

Configuration

The intended audience for this article is customers that already have several single T3 arrays or T3 array partner groups and wish to reconfigure them, and any other parties responsible for the design and configuration of Sun StorEdge™ T3 arrays.

This is the second of three-article series that will provide a roadmap for the configuration of the Sun StorEdge T3 storage array partner group by discussing the following points:

- Installation planning and design—Part 1
- Configuration—Part 2
- Basic Management—Part 3

Information presented in this article is discussed in greater detail in the Sun StorEdge T3 array, VERITAS Volume Manager, and Sun™ Management Center documentation.

The following best practices are featured in this series of articles:

- Layout of storage (manageability)
- Using VERITAS Volume Manager (VxVM) (maintainability)
- Wide Thin striping (performance)
- Basic systems management and monitoring (availability)

A Sun StorEdge T3 partner group has been split into two single T3 storage arrays. This article looks at the reconfiguration of these two single T3 storage arrays to reform the partner group. Because installation, planning, and design was covered in Part 1, this article starts at the configuration stage.

Note – The third part of this article, Basic Management, will look at automated configuration of a T3 array partner group using Expect scripts.

Configuration

Hardware reconfiguration and filesystem creation are detailed in this section. Please pay particular attention to the cautions and notes.

Caution – The following procedure destroys data. Make sure that any data residing on the T3 arrays to be connected has been backed up prior to reconfiguration.

Hardware Reconfiguration

This specific section discusses the procedure for configuring two Sun StorEdge T3 arrays to form a partner group of redundant controller units. This information is excerpted from Chapter 2,“Installation,” of the *Sun StorEdge T3 Installation, Operation, and Service* manual, with additional tips and annotations.

Connecting Single Controller Arrays To Form a Partner Group

Caution – The following procedure destroys data. Ensure that any data residing on the T3 arrays to be connected has been backed up prior to reconfiguration.

The T3 array controller units are connected with two interconnect cables. To accomplish this:

- Prepare the T3 arrays
- Cable a partner group
- Define and mount volumes on the alternate master

▼ Preparing the T3 Arrays

1. **Identify which T3 array will be the master controller and which will be the alternate master.**
2. **Back up data on both T3 arrays.**

Caution – Make sure you back up data on both arrays *before* proceeding! The volume(s) will need to be re-created on the alternate master after cabling the arrays together. Hence the volume(s) should survive on the array nominated to be the master.

---

3. **Ensure the data path between the host and both T3 arrays is quiescent.**

Note – There should be no I/O activity during configuration

---

4. **Start a telnet session with each T3 array.**
   a. **On a host, use the `telnet` command with the T3 array network name (or IP address) to connect to the T3 array.**
   b. **Log into the T3 arrays as the `root` user—enter the password at the prompt.**

      The T3 array prompt is displayed.
5. **Verify that the levels of all T3 array firmware are the same on both the master array and the alternate master array.**

On each T3 array:

a. Enter the `ver` command to display the controller firmware level.

For example:

```
# telnet T3_array_name
Trying 129.150.47.101...
Connected to 129.150.47.101.
Escape character is '^]'.
pSOSystem (129.150.47.101)
blueprint:/:<1>ver
T300 Release 1.00 1999/12/01 10:43:22 (129.150.47.135)
Copyright (C) 1997-1999 Sun Microsystems, Inc.
All Rights Reserved.
```

In the previous example, the controller firmware level is listed as `Release 1.00`.

b. Enter the `fru list` command to display the EEPROM, disk drive, and interconnect card firmware levels.

This is an example of the output of the `fru list` command:

```
blueprint:/:<2>fru list
ID TYPE VENDOR MODEL REVISION SERIAL
------ ----------------- ----------- ----------- -------------
u1ctr controller card SCI-SJ 375-0084-01- 0209 000887
u1d1 disk drive SEAGATE ST318203FSUN D44A LR170377
u1d2 disk drive SEAGATE ST318203FSUN D44A LR223933
u1d3 disk drive SEAGATE ST318203FSUN D44A LR194267
u1d4 disk drive SEAGATE ST318203FSUN D44A LR225714
u1d5 disk drive SEAGATE ST318203FSUN D44A LR104406
u1d6 disk drive SEAGATE ST318203FSUN D44A LR171635
u1d7 disk drive SEAGATE ST318203FSUN D44A LR171507
u1d8 disk drive SEAGATE ST318203FSUN D44A LR230196
u1d9 disk drive SEAGATE ST318203FSUN D44A LR235722
u1l1 loop card SCI-SJ 375-0085-01- 4.11 Flash 001733
u1l2 loop card SCI-SJ 375-0085-01- 4.11 Flash 002128
u1pcu1 power/cooling unit TECTROL-CAN 300-1454-01( 0000 003137
u1pcu2 power/cooling unit TECTROL-CAN 300-1454-01( 0000 001762
u1mpn mid plane SCI-SJ 370-3990-01- 0000 000563
```

In the previous example:

- EEPROM firmware version is listed as `Controller card, Revision 0209`
- Disk drive firmware version is listed as `Revision D44A`
- Interconnect card (`loop card`) firmware version is listed as `Revision 4.11 Flash`

Upgrade firmware (if necessary).

- If the firmware levels between the arrays are identical, proceed to Step 6.
- If the firmware levels for any of the four types are different between the master and the alternate master, upgrade the firmware to be the identical on both arrays. See the instructions in the "Upgrading Firmware" section of the *Sun StorEdge T3 Installation, Operation, and Service* manual.

6. On each array, use the `set -z` command to reset all system settings to factory defaults.

Note – This will erase the T3 array's password, IP address, netmask, and gateway. You must reconfigure these values on the array that you wish to be the master before logging out.

When prompted to respond, enter `y` (yes), to reset the system settings to the factory defaults. For example:

```
blueprint:/:<3>set -z
WARNING - Resetting system NVRAM to default, are you sure? [N]: y
blueprint:/:<4>
```

The **`set -z`** command will erase the T3 array's password, IP address, netmask, and gateway—these values need to be reconfigured on the **master array**. Therefore, prior to shutting down and powering off the arrays reset the following parameters:

- Root Password
- IP address, netmask, and gateway
- Hostname
- Date and Time
- Time Zone off-set

Refer to section, "Cabling a Partner Group" on page 7, to determine which array will be the master (the lower array if they are stacked and cabled as described). Then set the parameters as follows:

Setting Root Password

Set the root password of the T3 array by using the following command:

```
blueprint:/:<5>passwd
```

Note – Initially, there is no root password set on the T3. To log on, use the username root and press return when prompted for the password.

IP Configuration

Configure the IP address, gateway, and netmask on the T3 array with the following commands:

```
:/:<1>set ip 192.168.49.183
:/:<2>set gateway 192.168.49.254
:/:<3>set netmask 255.255.255.0
:/:<4>set hostname blueprint
```

▼ Time Settings

Set the date on the T3 array by entering the following command (with the appropriate time substituted). The syntax is date yyyymmddHHMM.SS:

```
blueprint:/:<6>date 200007071320.15
```

Set the time zone as an off-set from Greenwich Mean Time (GMT), also known as Universal Time Coordinated (UTC), by entering the following command (with the appropriate off-set substituted):

```
blueprint:/:<7>tzset -8000
```

Note – The previous example command would be used if the T3 array was located on the west coast of the United States of America.

The T3 arrays have had their system settings reset to the factory defaults. The master array's parameters, root password, IP address, netmask, gateway, hostname, date, time, and time zone have been set. The arrays can now be shutdown and powered off as described below prior to cabling them as a partner group.

1. **Power off both arrays.**

   Enter the following commands:

```
blueprint:/:<5>shutdown
Shutdown the system, are you sure? [N]: y
```

2. **Press the power button once on each power and cooling unit to turn the switch off.**

▼ Cabling a Partner Group

After the T3 array settings on the alternate master have been reset to the factory default, and the firmware levels on both arrays have been verified as being identical, the T3 arrays are ready to be connected.

1. **Place the alternate master next to or on top of the master array.**

   If the arrays are installed in a cabinet, make sure the alternate master is installed in the slot directly above the master array.

2. **Ensure the 10BASE-T cables are connected to a network with the same Management host.**

3. **Connect the interconnect cables to the cards as shown in** FIGURE 13.**.**

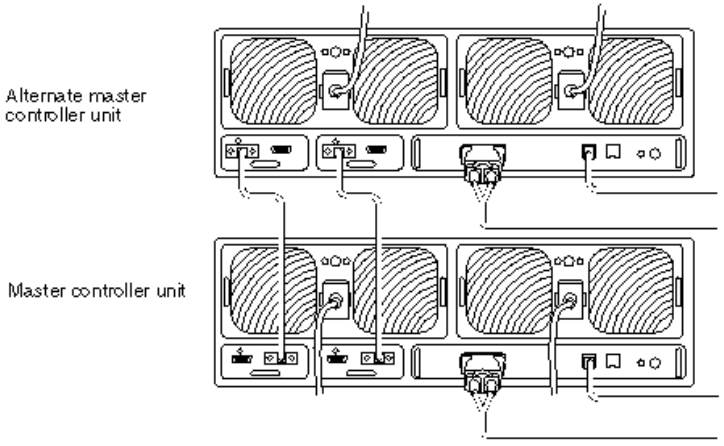

FIGURE 1 Connecting the Interconnect Cables

Note – Ensure the interconnect cables are connected to the correct interconnect cards *exactly* as shown in 3.FIGURE 1 because this cable connection determines which array is the master and which is the alternate master.

4. **Power on the master array by pressing the power on/off button on both the power and cooling units.**

   The power and cooling unit LEDs on both units will turn green when powered up.

Note – If the power and cooling unit LEDs on the alternate master do not light, press the power on/off switches on the alternate master.

5. **Connect and log onto the T3 array partner group using telnet with the IP address of the master array and enter the following command:**

```
blueprint:/:<1>sys mp_support rw
```

   This command ensures that multipathing support is enabled on the T3 to permit the use of the dynamic multipathing feature within VxVM.

6. **Enable cache mirroring by entering the following command:**

```
blueprint:/:<2>sys mirror auto
```

Note – If the T3 configuration needs to be reconfigured from its factory default settings for any reason (for example, the IP address was not set in the previous section), follow the instructions in the Sun BluePrints article “Sun StorEdge T3 Single Storage Array Design and Installation,” September, 2000.

▼ Defining and Mounting Volumes on the Alternate Master Array

After the arrays are cabled and power has been restored to both arrays, define and mount the volume(s) residing on the alternate master array. The volumes residing on the master array should survive during the reconfiguration process.

Note – Ensure that both arrays are available—that is, all LEDs are green. It may take several minutes after powering on for the arrays to be ready for use.

---

1. **Connect to the T3 array partner group by entering the following command on the host:**

```
# telnet <IP address or hostname of T3 master array>
```

2. **Log into the T3 array as the `root` user.**

3. **Confirm that the two arrays have successfully created a partner group by entering the following command to confirm all revision levels for each FRU type are identical:**

```
blueprint:/:<1> fru list
ID  TYPE   VENDOR MODEL REVISION SERIAL
------ ---------------- ----------- ----------- --------------------
u1ctr controller card SCI-SJ 375-0084-01- 0209 000887
u2ctr controller card SCI-SJ 375-0084-01- 0209 000626
u1d1 disk drive SEAGATE ST318203FSUN D44A LR170377
u1d2 disk drive SEAGATE ST318203FSUN D44A LR223933
u1d3 disk drive SEAGATE ST318203FSUN D44A LR194267
u1d4 disk driveSEAGATE ST318203FSUN D44A LR225714
u1d5 disk drive SEAGATE ST318203FSUN D44A LR104406
u1d6 disk drive SEAGATE ST318203FSUN D44A LR171635
u1d7 disk drive SEAGATE ST318203FSUN D44A LR171507
u1d8 disk drive SEAGATE ST318203FSUN D44A LR230196
u1d9 disk drive SEAGATE ST318203FSUN D44A LR235722
u2d1 disk drive SEAGATE ST318203FSUN D44A LR238482
u2d2 disk drive SEAGATE ST318203FSUN D44A LR162445
u2d3 disk drive SEAGATE ST318203FSUN D44A LR178429
u2d4 disk drive SEAGATE ST318203FSUN D44A LR163240
u2d5 disk drive SEAGATE ST318203FSUN D44A LR222181
u2d6 disk drive SEAGATE ST318203FSUN D44A LR239365
u2d7 disk drive SEAGATE ST318203FSUN D44A LR222571
u2d8 disk drive SEAGATE ST318203FSUN D44A LR170798
u2d9 disk drive SEAGATE ST318203FSUN D44A LR223756
u1l1 loop card SCI-SJ 375-0085-01- 4.11 Flash 001733
u1l2 loop card SCI-SJ 375-0085-01- 4.11 Flash 002128
u2l1 loop card SCI-SJ 375-0085-01- 4.11 Flash 002024
u2l2 loop card SCI-SJ 375-0085-01- 4.11 Flash 002029
u1pcu1 power/cooling unit TECTROL-CAN 300-1454-01( 0000 003137
u1pcu2 power/cooling unit TECTROL-CAN 300-1454-01( 0000 001762
u2pcu1 power/cooling unit TECTROL-CAN 300-1454-01( 0000 001430
u2pcu2 power/cooling unit TECTROL-CAN 300-1454-01( 0000 001429
u1mpn mid plane SCI-SJ 370-3990-01- 0000 000563
u2mpn mid plane SCI-SJ 370-3990-01- 0000 000512
```

4. **Confirm all FRU are available by entering the following command:**

```
blueprint:/:<2> fru stat
CTLR STATUSSTATE ROLE PARTNER TEMP
---------------------- ---------- ------- ----
u1ctr ready enabled master u2ctr 32.5
u2ctr ready enabled alt master u1ctr 32.5
DISK STATUS STATE ROLE PORT1 PORT2 TEMP VOLUME
------------ ---------- ---------- ----------------- ------
u1d1 ready enabled data disk ready ready 37 v0
u1d2 ready enabled data disk ready ready 39 v0
u1d3 ready enabled data disk ready ready 36 v0
u1d4 ready enabled data disk ready ready 38 v0
u1d5 ready enabled data disk ready ready 37 v0
u1d6 ready enabled data disk ready ready 35 v0
u1d7 ready enabled data disk ready ready 40 v0
u1d8 ready enabled data disk ready ready 36 v0
u1d9 ready enabled standby ready ready 37 v0
u2d1 ready enabled unassigned ready ready 42 -
u2d2 ready enabled unassignedready ready 36 -
u2d3 ready enabled unassignedready ready 35 -
u2d4 ready enabled unassignedready ready 35 -
u2d5 ready enabled unassignedready ready 40 -
u2d6 ready enabled unassignedready ready 36 -
u2d7 ready enabled unassignedready ready 35 -
u2d8 ready enabled unassignedready ready 37 -
u2d9 ready enabled unassignedready ready 33 -
LOOP STATUS STATE MODE CABLE1 CABLE2 TEMP
------ ------- ---------- ------- --------- --------- ----
u2l1 ready enabled master installed - 29.5
u2l2 ready enabled slave installed - 33.0
u1l1 ready enabled master - installed 28.0
u1l2 ready enabled slave - installed 32.0
POWER STATUS STATE SOURCEOUTPUT BATTERY TEMP  FAN1   FAN2
------------ --------- ------------------ ----- ------ -----
u1pcu1ready enabled line normalnormal normal normal normal
u1pcu2ready enabled line normalnormal normal normal normal
u2pcu1ready enabled line normalnormal normal normal normal
u2pcu2ready enabled line normalnormal normal normal normal
```

The batteries in the power and cooling units recharge after powering on the array. During recharge, a *fault* message is displayed in the `fru stat` output against the batteries, and the write-behind cache is disabled. The charging time will depend on the state of the batteries. A refresh after a power cycle can take up to 2 hours. A complete recharge can take up to 12 hours.

The batteries are drained and recharged every 14 days by default. The write-behind cache is disabled while the batteries are drained, which can take up to 20 minutes for each battery. Because there is a performance degradation to write activity during this time, it may be desirable to specify when the refresh occurs in the `/etc/schd.conf` file on the T3 array. See the "Sun StorEdge T3 Disk Tray Administrator's Guide" for more information. A refresh can be invoked manually by using the `refresh` command.

5. **To create a 7 + 1 disk (RAID-5) volume with one hot spare on the alternate master array, enter the following commands:**

```
blueprint:/:<3> vol add v1 data u2d1-8 raid 5 standby u2d9
blueprint:/:<4> vol init v1 sysarea rate 16
blueprint:/:<5> vol init v1 data rate 16
blueprint:/:<6> vol mount v1
```

The `sysarea` qualifier creates a system area on the volume used for configuration and status information pertaining to the T3 array. The `data` qualifier creates the area on the volume that will hold the user data.

Note – Drive # 9 is the only drive that can be configured as a hot spare. Only one hot spare is allowed per T3 array and can only be used by that array, not by the partner array.

---

6. **Confirm the volumes are laid out correctly by entering the following command:**

```
blueprint:/:<7> vol list
```

You should get the following output, assuming the master array originally contained a 7+1 disk (RAID-5) volume with one hot spare disk:

```
volume capacity raid data standby
v0  125.3 GB 5 u1d1-8 u1d9
v1  125.3 GB 5 u2d1-8 u2d9
```

Note – The volume v0 is a volume that existed on the master array prior to reconfiguration as a partner pair—it may survive during the reconfiguration process, while the volume on the alternate master will not. Be prepared to restore of all previous data.

---

7. **After the volumes have been successfully created, enter the following command to reboot the T3 partner group:**

```
blueprint:/:<8> reset
```

8. **After the T3 partner group has rebooted, log onto the T3 partner group and issue the `vol list`, `fru list`, `fru stat`, `vol mode`, `sys list`, `port list`, and `port listmap` commands.**

   Verify that the configuration is correct by checking the output of these commands.

   The following is an example of the output after using the `port listmap` command—note how the volumes are accessed through a primary and failover fibre channel (FC-AL) host connection:

```
blueprint:/:<1>port listmap
port targetedaddr_typelunvolumeowneraccess
u1p11   hard0   v0  u1  primary
u1p11   hard1   v1  u2  failover
u2p12   hard0   v0  u1  failover
u2p12   hard1   v1  u2  primary
```

▼ Device Configuration

After the T3 array partner group has been connected to the host, the new devices need to be configured. Operating system versions prior to the Solaris™ 8 Operating Environment require a reconfiguration boot, or use of the `luxadm` command to create the new devices, see below:

```
{0} ok boot -r
```

If the `luxadm` command is present, the new devices can be created by the root user issuing the following command, negating the need for a system reboot on Operating system versions prior to the Solaris™ 8 Operating Environment:

```
$ luxadm insert
```

After creating the new devices on the host, it can be seen by entering the `format` command, that each volume on the T3 array partner group generates two targets, one on each controller. See below, c1t2d0 and c2t1d0 are one volume, c1t2d1 and c2t1d1 are the other:

```
$ format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
       0. c0t0d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
          /sbus@3,0/SUNW,fas@3,8800000/sd@0,0
       1. c0t1d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
          /sbus@3,0/SUNW,fas@3,8800000/sd@1,0
       2. c0t2d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
          /sbus@3,0/SUNW,fas@3,8800000/sd@2,0
       3. c0t3d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
          /sbus@3,0/SUNW,fas@3,8800000/sd@3,0
       4. c1t2d0 <SUN-T300-0116 cyl 34145 alt 2 hd 56 sec 128>
          /sbus@2,0/SUNW,socal@d,10000/sf@0,0/
           ssd@w50020f2300000907,0
       5. c1t2d1 <SUN-T300-0116 cyl 34145 alt 2 hd 56 sec 128>
          /sbus@2,0/SUNW,socal@d,10000/sf@0,0/
           ssd@w50020f2300000907,1
       6. c2t1d0 <SUN-T300-0116 cyl 34145 alt 2 hd 56 sec 128>
          /sbus@6,0/SUNW,socal@d,10000/sf@0,0/
           ssd@w50020f2300000b99,0
       7. c2t1d1 <SUN-T300-0116 cyl 34145 alt 2 hd 56 sec 128>
          /sbus@6,0/SUNW,socal@d,10000/sf@0,0/
           ssd@w50020f2300000b99,1
Specify disk (enter its number):
```

Label the newly created volume on the alternate master array by using the `format` command. The existing volume on the master array already has a label.

Note – If both volumes are newly created, they *both* need to be labeled, otherwise they may not be recognized by VxVM.

File System Creation

At this point in the configuration process, the volumes have been created on the T3 array partner group. These volumes are seen as devices on the host. These new devices must be added to the VxVM configuration—VxVM will refer to these as “Disks.”

From these disks, VxVM can create volumes on which file systems will be created. Finally, the file systems will be mounted on the host and used to store data. These steps are described in the following sections.

Volume Manager

VxVM is used to compliment the high availability features of the T3 array—without a volume manager, the file systems may need to be taken off-line for reorganization. The dynamic multipathing (DMP) feature of VxVM is used to allow failover between the dual fibre channel connections to the host.

Configure the volumes within VxVM by creating a disk group and volumes as described in the following sections.

Creating Disk Groups

If you have just installed VxVM on the host, you must now configure VxVM using the `vxinstall` command, as shown below. Refer to the *VERITAS Volume Manager for Solaris Installation Guide* for more information on `vxinstall`.

After VxVM is installed, enter the following command:

```
# vxinstall
```

This command configures VxVM, and enables disks to be added to the `rootdg` disk group.

Note – At least one disk must belong to the `rootdg` disk group. You must encapsulate a disk if you wish to preserve existing data. Encapsulation is only possible if you have two unassigned partitions, and approximately 15 MB of space that is not assigned to a partition. This is required to create the public and private regions for VxVM. If this requirement is not met, the data has to be restored after the disk has been added to VxVM.

Each volume within the T3 partner group generates two targets, one on each controller.

DMP (enabled by default) is aware of the dual paths. Therefore, using the `vxprint` command will not display both targets—the DMP software hides the dual paths behind a single DMP device.

To see how VxVM is handling the dual paths, execute the following commands—note that the controller SBus addresses given by `vxdmpadm listctlr all` (e.g. `/sbus@2,0/SUNW,socal@d,10000/sf@0,0`) are used as the inputs to `vxdmpadm getsubpaths ctlr=`:

```
# vxdmpadm listctlr all
CTLR-NAME       DA-TYPE      STATE      DA-SNO
==========================================
ctlr0           T300         ENABLED    60020f2000000b990000000000000000
ctlr0=/sbus@2,0/SUNW,socal@d,10000/sf@0,0
ctlr1           T300         ENABLED    60020f2000000b990000000000000000
ctlr1=/sbus@6,0/SUNW,socal@d,10000/sf@0,0
ctlr2           SEAGATE      ENABLED    SEAGATE_DISKS
ctlr2=/sbus@3,0/SUNW,fas@3,8800000

# vxdmpadm getsubpaths ctlr=/sbus@2,0/SUNW,socal@d,10000/sf@0,0
NAME       STATE     TYPE       DMPNODENAME  DA-TYPE   DA-SNO
============================================================
c1t2d1s2  ENABLED  PRIMARY     c1t2d1s2    T300  60020f2000000b9900000000
c1t2d0s2   ENABLED   SECONDARY  c1t2d0s2    T300  60020f2000000b9900000000
# vxdmpadm getsubpaths ctlr=/sbus@6,0/SUNW,socal@d,10000/sf@0,0
NAME       STATE     TYPE       DMPNODENAME  DA-TYPE   DA-SNO
============================================================
c2t1d1s2   ENABLED   SECONDARY  c1t2d1s2     T300  0020f2000000b9900000000
c2t1d0s2   ENABLED   PRIMARY c1t2d0s2 T300  60020f2000000b9900000000
```

If one of the fibre channel paths to the T3 partner group fails, DMP will automatically route the connection via the remaining path. When the failed path is repaired, DMP will automatically revert back to using both paths. The following is an example of the output of the `vxdmpadm getsubpaths` command after the fibre channel cable connected to controller 2 has been unplugged:

```
# vxdmpadm getsubpaths ctlr=/sbus@2,0/SUNW,socal@d,10000/sf@0,0
NAME       STATE     TYPE       DMPNODENAME  DA-TYPE   DA-SNO
============================================================
c1t2d1s2   ENABLED   PRIMARY     c1t2d1s2     T300  60020f2000000b990000000
c1t2d0s2   ENABLED   SECONDARY  c1t2d0s2    T300  60020f2000000b990000000
# vxdmpadm getsubpaths ctlr=/sbus@6,0/SUNW,socal@d,10000/sf@0,0
NAME       STATE     TYPE       DMPNODENAME  DA-TYPE   DA-SNO
============================================================
c2t1d1s2   DISABLED   SECONDARY  c1t2d1s2     T300  0020f2000000b990000000
c2t1d0s2   DISABLED   PRIMARY     c1t2d0s2     T300  60020f2000000b990000000
```

If VxVM is already installed and configured on the host, or after executing the `vxinstall` command, add the disks to a disk group by executing the following command:

```
# vxdiskadd  <target>
```

Where target can be `c2`, `c2t3`, `c2t3d4`, `c2c3t5d1`, or all.

Creating a Volume

Once the disk has been added to a disk group, a volume is created from that disk. VxVM will automatically select the disks comprising a volume unless they are specified. To create a volume from a specific disk execute the following command:

```
# vxassist make <volume_name> <length> <diskname>
```

For example:

```
# vxassist make data01 3g disk01
```

This command will create a concatenated volume. To create a Striped (RAID-0), RAID-5, or Mirrored (RAID-1) volumes, see the `vxassist` man pages and *VxVM Command Line Interface Administrator's Guide*. In this instance; however, the T3's hardware will take care of the RAID-5 functions, therefore, a concatenated volume is the correct choice.

Disabling VxVM Hot Relocation

Although not strictly applicable to this hardware RAID configuration, because the T3 hardware will handle hardware and data resilience, disabling hot relocation and enabling hot sparing within VxVM is considered a vital recommendation.

Hot relocation attempts to relocate each subdisk of a failed disk. This can have unpredictable and possibly undesirable results; for example, both plexes of a mirror on the same physical disk, because VxVM will place the subdisks wherever it finds free space within the pool of hot relocation spares.

In contrast, hot sparing substitutes a failed disk with an entire spare disk from the pool of hot spares. This yields predictable results, and more importantly, a less complex recovery procedure.

Hot relocation is enabled by default. It can be disabled, and hot sparing enabled, by editing the file `/etc/rc2.d/S95vxvm-recover` so that the last 10 lines of `/etc/rc2.d/S95vxvm-recover` look as follows:

```
# start the watch daemon. This sends E-mail to the administrator
# any problems are found. To change the address used for sending problem
# reports, change the argument to vxrelocd.

# vxrelocd root &

# to enable hot sparing instead of hot relocation.
# ( comment out vxrelocd before uncommenting vxspare )

vxsparecheck root &

exit 0
```

To enable the changes above without rebooting the machine, stop the `vxrelocd` daemon using the `kill` command and start the `vxsparecheck` daemon manually.

Execute the commands below to specify hot spares:

```
# /etc/vx/bin/vxdisksetup –i <disk device name>
# vxdg –g <disk group name> adddisk hotspare=<disk device name>
# vxedit –g <disk group name> set spare=on hotspare
```

File System Creation

After the volume is created on the T3 array and within VxVM, a filesystem can be created and mounted by the host, example commands are shown here:

```
# newfs –v /dev/vx/rdsk/data01
# mount /dev/vx/dsk/data01 /mnt
```

Create directories for the permanent mount points and edit `/etc/vfstab` to include the new file systems.

Summary

This article has presented the process for reconfiguring two single T3 arrays as a partner group comprising the steps:

- Preparation
  - Backup all data
  - Verify firmware levels
  - Reset settings to factory defaults
  - Configure parameters, hostname, IP address, etc
- Cabling the partner group
  - Configure parameters to allow fail over between arrays
- Define and mount volumes on the T3 arrays
- Create new devices on the host
- Create file systems
  - Add Disks to VxVM
  - Operation of Dynamic MultiPathing
  - Create Volumes in VxVM
  - Disable Hot relocation in VxVM
  - Create and Mount File systems

The next article in the series will look at basic management of the T3 arrays, including logging to a remote or log host, Simple Network Management Protocol (SNMP), Sun StorEdge Component Manager and Automated configuration of the T3 array using Expect scripts.

Author's Bio: Mark Garner

Mark Garner is a Systems Engineer for Enterprise Engineering at Sun Microsystems. Prior to joining Sun, Mark spent three years as a Systems Architect specializing in Email and Office Automation architectures and before this over eight years in systems administration. While at Sun Mark has focused on the design and implementation of mission critical business and Internet application's infrastructure principally to provide for high availability and scalability.