

Split Mode for Patch D

This appendix provides an alternative example split procedure to that documented in “Split Example” on page 64, which is designed to be used when running a Netra ft 1800 system with upgrade patch 108145-10 (patch D) installed.

Split Example

The following example demonstrates how to split a Netra ft 1800, upgrade a component such as a software patch, and then remerge the system. Each stage of the split process is documented, with a description of the state of the system at the end of each stage.

The procedure documented here is an instructor-led demonstration of split mode on the Netra ft 18000 system. As such, it is not intended as a tutorial, but provides one example of the use of split mode. The procedure should be used in conjunction with the *Netra ft 1800 User's Guide* (part number 805-4529-10).

This procedure assumes that you have some user-level knowledge of SEVM 2.5.

Assumptions

- If NIS is running, edit the `/etc/nsswitch.conf` file so that the line:

```
hosts: xfn nis [NOTFOUND=return] files
```

reads

```
hosts: files xfn nis [NOTFOUND=return]
```

`files` is now the first option on the line, which will force the system to read the local `/etc/hosts` file first.

- Before the split, the combined system has the host name `seagoon` and IP address 129.156.203.152. After the split, side A will be the split winner with host name `seagoon` and side B will be the loser with the host name `seagoon-2` and IP address 129.156.203.18.
- OBP (PROM) variables:
 - `use-nvramrc?=true`
 - `auto-boot?=true`
 - `boot-device=a-dsk0 b-dsk0`
 - `diag-device=a-dsk0 b-dsk0`
 - `diag-switch?=false`

Check using the `eeeprom` command and, if necessary, set using:

```
eeeprom use-nvramrc?=true
eeeprom auto-boot?=true
eeeprom boot-device=a-dsk0 b-dsk0
eeeprom diag-device=a-dsk0 b-dsk0
eeeprom diag-switch?=false
```

The `boot-device` PROM variable is concatenated with an NVRAM variable, `boot-list`, to produce a sequence of disks from which the system tries to boot. The `boot-list` variable is constructed from a motherboard EEPROM variable, which comprises the HDDs in the SEVM root disk group. Setting `boot-device` to the value above ensures that the system will try to boot from A-DSK0 and, if this fails, from B-DSK0. Since these HDDs are mirrored, either should produce a current copy of the boot disk.

- All devices on the losing side should be quiesced, that is. no devices using losing side FRUs should be used by any applications.
- Disk configuration:
 - `rootdisk` is the SEVM representation of `a-dsk0`, which is the disk device `c2t0d0`;
 - `disk01` is the SEVM representation of `b-dsk0`, which is the disk device `c3t0d0`.

Initial Configuration

Machine: 2P 512Mb Netra ft 1800

OpenBoot PROM version: release 24 (part no. 258-7354-10)

FPGA version: release 19 (part no. 258-7134-07)

Operating System: Netra ft 1800 Solaris 6.7 CD, supplemental CD , patches CD, SEVM 2.5 (including patch 105463)

Patch Level: Patch D-10

FRUs present: HDD in A-DSK0 (/dev/rdisk/c2t0d0) called rootdisk in SEVM, mirrored in B-DSK0 (/dev/rdisk/c3t0d0) called disk01 in SEVM. No other HDDs are present in the machine. No PCI devices are attached.

Node name: seagoon (IP address 129.156.203.152)

Initial configuration (cmsconfig): As shown in FIGURE E-1

Item	Name	Fault Loc	State	Page 1 of 1
0	A-MBD 0	A-MBD	enabled	
1	B-MBD 0	B-MBD	enabled	
2	CAF 0	A-CAF	enabled	
3	CAF 1	B-CAF	enabled	
4	CPU 0	A-CPU	enabled	
5	CPU 1	B-CPU	enabled	
6	DSK 0	A-DSK	enabled	
7	DSK 1	B-DSK	enabled	
8	HDD 0	A-DSK0	enabled	
9	HDD 7	B-DSK0	enabled	
10	ft_alarm 0		usable	
11	ft_core 0		enabled	Fault tolerant
12	ft_network 0		online_up	A (online) & B (online)
13	ft_serial 0		online	
14	icn_system 0		disabled	System is not split

(H)elp, (I)nclude, (E)xclude, (S)elect, (V)iew, (Q)uit or <Number> ?

FIGURE E-1 Initial CMS Configuration

Notice the presence of:

- two motherboards (MBDs)
- two CAFs
- two CPUs
- two disk chassis (DSKs)

- two hard disks (HDDs)
- an `ft_alarm` (the alarms subsystem)
- an `ft_core` (the fault-tolerant virtual machine)
- `ft_network 0` (a fault-tolerant network)
- `ft_serial` (the serial subsystem)
- `icn_system`

Check that the system (`seagoon`) is in combined mode using the command:

```
seagoon# splitinfo -a
attributes=combined
```

Ensure that the Netra ft 1800 combined system is not acting as a router. This is accomplished by the existence of the file `/etc/notrouter`, which can be created by issuing the command:

```
seagoon# touch /etc/notrouter
```

Configuring for Split Mode

This example shows how to split a Netra ft 1800 with nodename `seagoon` into `seagoon` and `seagoon-2`. While `seagoon` is left running on side A, `seagoon-2` will be updated on side B, and then the two sides merged to reform a combined, or fault-tolerant, Netra ft 1800 system.

You may find it useful to have two xterms open, one for each of the following:

- A-side console, to monitor the progress of the split on the A-side
- B-side console, to monitor the progress of the split on the B-side

The following steps are required to prepare to split a system.

1. Edit the `splitd.conf` file
2. Set up ICN
3. Set the split master
4. Add other IP addresses to `/etc/hosts`
5. Configure `ft_network 0`
6. Check SEVM status
7. Check CPUs are in sync

8. Check split daemon status
9. Turn off boot disk checksumming

Editing the `split.conf` File

FIGURE E-2 shows the contents of `/etc/splitd.conf`, naming `seagoon` and `seagoon-2` as the two host names to be used following the split.

```
# The domain-address section.
# Add the primary and alternative hostnames
# (used for inter-daemon comms) of the two sides here.
# Example:
# host_prim host-i0 host-2-i0
# host_alt thishost thishost-2
host_prim seagoon-i0 seagoon-2-i0
host_alt seagoon seagoon-2

# The port-number section
port_address a b 3501
port_address b a 3500
port_address a b 3502
port_address b a 3503

# The misc-section
timeout 60
dolog yes

# Hostnames section
# add the hostnames of the two sides here. Example:
# hostnames thishost thishost-2
hostnames seagoon seagoon-2
```

FIGURE E-2 The `/etc/splitd.conf` File

Setting Up ICN

Use `cmsconfig` to set the `hostname` entry in `icn0` to the IP address of the split winner's ICN network. In this case, `hostname` is set to `seagoon-i0`

```
Select: icn 0
```

Item	Name	Value	Page 1 of 1
0	state	not_present	
1	icn_cmd	disable	
2	description	ICN controller	
3	hostname	seagoon-i0	
4	interface_cmd	up	
5	user_label		
6	memory_size	524288	
7	descriptor_limit	255	
8	info		
9	software_fault	no	
10	interface_state	unknown	
11	busylock	no	
12	action	disable	

(H)elp, (Number)to modify, (S)elect, (T)op or (Q)uit ?

Setting the Split Master

The split master negotiates the transfer of ownership of FRUs from one side to another. Only one side of a Netra ft 1800 can be the master, and this can be changed in combined or split mode. In this case, the following command is used to change the master to side A, which also happens to be the split winner:

```
seagoon# splitconf -m A  
master = A
```

Note – The master does not have to be the same as the split winner or merge winner. The split or merge winner is the side that continues running during the split or merge operation, respectively, whereas the master negotiates FRU ownership transfers between sides.

Adding Other IP Addresses to /etc/hosts

Here, entries for both `seagoon-2`, and host names for both side's ICN 0 network instances have been added to the `/etc/hosts` file:

```
#
# Internet host table
#
127.0.0.1      localhost      loghost
129.156.203.152 seagoon
129.156.203.18 seagoon-2
192.168.1.1    seagoon-i0
192.168.1.2    seagoon-2-i0
```

In addition, `localhost` has been set as the `loghost` to ensure that logging is performed on the local machine after a split. Otherwise, logging would be directed across the network to the other side of the Netra ft 1800.

Configuring `ft_network 0`

Check that `ft_network 0` is operating in fault-tolerant mode, that is, both A and B controllers are online.

```
Select: ft_network 0

Item  Name                               Value                               Page 1 of 1
-----
0     state                               online_up
1     description                          network multiplexor
2     user_label
3     hostname                             seagoon
4     preferred_controller                 none
5     controllerA_FRU                     A-CAF
6     controllerA_Funct                   Net_0
7     controllerB_FRU                     B-CAF
8     controllerB_Funct                   Net_0
9     info
10    busylock                             no
11    devpath                              pnet0
12    link                                100 Mbps half-duplex link up
13    transceiver                         internal transceiver selected
14    usable_controllers                  A (online) & B (online)
15    controller_in_use                   A
```

The split mode software ensures that when the losing (B) side is rebooted after the split, it will (in this example) take on the host name `seagoon-2`.

Any other configured `ft_networks` should set to NULL any controllers using interfaces on the losing side. This will prevent the split loser from bringing online networks with host names that clash with host names on the split winner.

Checking the SEVM State

Check the state of the SEVM by issuing the following commands:

```
seagoon# ps -ef -oargs | grep vx
vxconfigd -m boot

seagoon# vxctl mode
mode: enabled

seagoon# vxdisk list
DEVICE      TYPE      DISK      GROUP      STATUS
c2t0d0s2    sliced    rootdisk  rootdg      online
c3t0d0s2    sliced    disk01    rootdg      online
```

The first command ensures that the SEVM configuration daemon is running.

The second command ensures that the daemon is enabled.

The third command shows the status of the HDDs. In this case there are two HDDs under SEVM control, `rootdisk` and `disk01`, which are both members of the `rootdg` disk group. This SEVM HDD configuration is necessary for split mode to work correctly.

In addition, issue the `vxprint` command to determine the redundant state of SEVM as shown in CODE EXAMPLE E-1:

CODE EXAMPLE E-1 StorEdge Volume Manager State

```
seagoon# vxprint
Disk group: rootdg

TY NAME      ASSOC      KSTATE  LENGTH  PLOFFS  STATE  TUTILO  PUTILO
dg rootdg    rootdg     -        -        -        -        -        -
dm disk01    c3t0d0s2  -        17538444 -        -        -        -
dm rootdisk  c2t0d0s2  -        17538444 -        -        -        -
v opt        fsgen     ENABLED  4097331 -        ACTIVE  -        -
```

CODE EXAMPLE E-1 StorEdge Volume Manager State (Continued)

pl	opt-01	opt	ENABLED	4097331	-	ACTIVE	-	-
sd	rootdisk-04	opt-01	ENABLED	4097331	0	-	-	-
pl	opt-02	opt	ENABLED	4097331	-	ACTIVE	-	-
sd	disk01-01	opt-02	ENABLED	4097331	0	-	-	-
v	rootvol	root	ENABLED	1106028	-	ACTIVE	-	-
pl	rootvol-01	rootvol	ENABLED	1106028	-	ACTIVE	-	-
sd	rootdisk-B0	rootvol-01	ENABLED	1	0	-	-	Block0
sd	rootdisk-02	rootvol-01	ENABLED	1106027	1	-	-	-
pl	rootvol-02	rootvol	ENABLED	1106028	-	ACTIVE	-	-
sd	disk01-02	rootvol-02	ENABLED	1106028	0	-	-	-
v	swapvol	swap	ENABLED	1052163	-	ACTIVE	-	-
pl	swapvol-01	swapvol	ENABLED	1052163	-	ACTIVE	-	-
sd	rootdisk-01	swapvol-01	ENABLED	1052163	0	-	-	-
pl	swapvol-02	swapvol	ENABLED	1052163	-	ACTIVE	-	-
sd	disk01-03	swapvol-02	ENABLED	1052163	0	-	-	-
v	usr	fsgen	ENABLED	7185591	-	ACTIVE	-	-
pl	usr-01	usr	ENABLED	7185591	-	ACTIVE	-	-
sd	rootdisk-05	usr-01	ENABLED	7185591	0	-	-	-
pl	usr-02	usr	ENABLED	7185591	-	ACTIVE	-	-
sd	disk01-04	usr-02	ENABLED	7185591	0	-	-	-
v	var	fsgen	ENABLED	4097331	-	ACTIVE	-	-
pl	var-01	var	ENABLED	4097331	-	ACTIVE	-	-
sd	rootdisk-03	var-01	ENABLED	4097331	0	-	-	-
pl	var-02	var	ENABLED	4097331	-	ACTIVE	-	-
sd	disk01-05	var-02	ENABLED	4097331	0	-	-	-

Notice that there are five volumes in the `rootdg` disk group, each containing two plexes. All plex `KSTATE` entries should be in the `enabled` state, and all plex `STATE` entries should be in the `active` state. These states are required in order for split mode to succeed. Recover any plexes that are not in these states so that they are in the `active` state.

Confirming the CPUs Are In Sync

Enter the command `cmsconfig`:

```
ps -elf -oargs | grep u4ftsplitt cmsconfig
```

Item	Name	Fault Loc	State	Page 1 of 1
0	A-MBD 0	A-MBD	enabled	
1	B-MBD 0	B-MBD	enabled	
2	CAF 0	A-CAF	enabled	
3	CAF 1	B-CAF	enabled	
4	CPU 0	A-CPU	enabled	
5	CPU 1	B-CPU	enabled	
6	DSK 0	A-DSK	enabled	
7	DSK 1	B-DSK	enabled	
8	HDD 0	A-DSK0	enabled	
9	HDD 7	B-DSK0	enabled	
10	ft_alarm 0		usable	
11	ft_core 0		enabled	Fault tolerant
12	ft_network 0		online_up	A (online) & B (online)
13	ft_serial 0		online	
14	icn_system 0		disabled	System is not split

(H)elp, (I)nclude, (E)xclude, (S)elect, (V)iew, (Q)uit or <Number> ?

Notice that the `ft_core` object reads `Fault tolerant`, indicating that the CPUs are in sync.

Checking the Split Daemon Status

Ensure that the split daemon, `u4ftsplitt`, is running by issuing the command:

```
seagoon# ps -elf -oargs | grep u4ftsplitt
/usr/platform/SUNW,Ultra-4FT/SUNWcms/lib/u4ftsplitt
```

and noting that the output contains `u4ftsplitt`. If it is not running, restart it with the command:

```
seagoon# /etc/init.d/u4ftsplitt start
```

Turning off Boot Disk Checksumming

With patch D installed, it is necessary to turn off disk checksumming on the loser's boot disk to avoid SEVM errors swamping the split loser's console after a split. To do this, use `cmsconfig` to turn off checksumming in the `B-DSK0` FRU.

Quick Checklist

At this stage, confirm the following before proceeding:

- The CPUs are running in sync.
- There are HDDs present in `A-DSK 0` and `B-DSK 0`, which are mirrored using SEVM. In addition, these two HDDs are both members of the SEVM disk group `rootdg`. The SEVM configuration contains exactly *one* disk group (`rootdg`) containing *one* mirrored pair.
- `/etc/splitd.conf` names both `seagoon` and `seagoon-2`, and at least one of the ICN networks.
- `/etc/hosts` contains all IP addresses (for `seagoon`, `seagoon-2` and the ICN networks configured in `/etc/splitd.conf`).
- The `ft_network 0` is configured in fault-tolerant mode.
- For each ICN network instance configured in `/etc/splitd.conf`, the `hostname` field of the appropriate object in `cmsconfig` should be set. For example, if ICN network instance 0 is configured, the host name for ICN 0 in `cmsconfig` should be set.
- The boot disks are mirrored and the mirroring process has completed. A fully-mirrored volume has an `ACTIVE` state for the `volume` and `plex` fields in the output from the `vxprint` command.

Splitting the System

You are now ready to split the system, but, before proceeding, note the following caution:



Caution – Do not use motherboard resets (`AAres` and `BBres`) or any other mechanism in which the motherboard is powered off, or perform a power cycle or an operation that results in a power cycle, on either side when in split mode. Doing so will result in undesirable failures occurring in certain FRUs on the other side. This is because some resources on a motherboard may belong to the other side and these resources will be removed from service if the motherboard is reset.

If the system is split naturally, that is, if all the resources residing on the side-A motherboard belong to side A and all resources residing on side-B motherboard belong to side B, or if the resources on the motherboard that needs to be reset are not deemed to be critical for providing service to the clients of the other side, it is possible to use a motherboard reset.

However, to avoid failure of the bridges, which will cause subsequent merge operations to fail, in the CMS first disable the motherboard that is to be reset on the side the side that is *not* going to be reset. this will prevent bridge failures and will enable the system to merge successfully when it is brought back into fault-tolerant operation.

Issue the `split` command from either side:

```
splitadm -w a split
```

Note – This will cause many warning messages from SEVM as it loses sight of its mirror.

FIGURE E-3 shows the output on the A-side console after the `split` command has been issued from the A-side.

```

seagoon# splitadm -w a split
NOTICE: vxdmp: Path failure on 32/124
WARNING: vxvm:vxio: error on Plex var-02 while writing volume var offset 119388
length 2
WARNING: vxvm:vxio: Plex var-02 detached from volume var
WARNING: vxvm:vxio: disk01-05 Subdisk failed in plex var-02 in vol var
vxvm:vxconfigd: NOTICE: Offlining config copy 1 on disk c3t0d0s2:
    Reason: Disk write failure
vxvm:vxconfigd: NOTICE: Detached disk disk01
vxvm:vxconfigd: NOTICE: Detached plex opt-02 in volume opt
vxvm:vxconfigd: NOTICE: Detached plex rootvol-02 in volume rootvol
vxvm:vxconfigd: NOTICE: Detached plex swapvol-02 in volume swapvol
vxvm:vxconfigd: NOTICE: Detached plex usr-02 in volume usr
vxvm:vxconfigd: NOTICE: Detached plex var-02 in volume var
vxvm:vxassist: ERROR: Cannot allocate space to replace subdisks
vxvm:vxassist: ERROR: Cannot allocate space to mirror 1106028 block volume
vxvm:vxassist: ERROR: Cannot allocate space to mirror 1052163 block volume
vxvm:vxassist: ERROR: Cannot allocate space to replace subdisks
vxvm:vxassist: ERROR: Cannot allocate space to replace subdisks
icn network setupWARNING: Out-of-sync on CPUset B

(via CMS): icn.keepalive done.
domain = A
attributes = split,master
master = A

```

FIGURE E-3 A-Side Console After split Command

The B-side console display is shown in FIGURE E-4, which is a copy of the A-side console output before the split causes a reboot of the B-side.

```
NOTICE: vxdmp: Path failure on 32/124
WARNING: vxvm:vxio: error on Plex var-02 while writing volume var offset 119388
length 2
WARNING: vxvm:vxio: Plex var-02 detached from volume var
WARNING: vxvm:vxio: disk01-05 Subdisk failed in plex var-02 in vol var
vxvm:vxconfigd: NOTICE: Offlining config copy 1 on disk c3t0d0s2:
    Reason: Disk write failure
vxvm:vxconfigd: NOTICE: Detached disk disk01
vxvm:vxconfigd: NOTICE: Detached plex opt-02 in volume opt
vxvm:vxconfigd: NOTICE: Detached plex rootvol-02 in volume rootvol
vxvm:vxconfigd: NOTICE: Detached plex swapvol-02 in volume swapvol
vxvm:vxconfigd: NOTICE: Detached plex usr-02 in volume usr
vxvm:vxconfigd: NOTICE: Detached plex var-02 in volume var
```

FIGURE E-4 B-Side Console After `split` Command

Assuming all has gone well, check that the winning host is running on side A:

```
seagoon# u4ftvmctl -c
This is CPUset_A (0)
```

Issuing the `cmsconfig` command on `seagoon` should display information similar that in FIGURE E-5:

```

seagoon# cmsconfig
-----
Item  Name          Fault  Loc      State
-----
0     A-MBD 0          A-MBD   enabled
1     B-MBD 0          B-MBD   enabled
2     CAF 0           A-CAF   enabled
3     CAF 1           B-CAF   disabled
4     CPU 0           A-CPU   enabled
5     CPU 1           X B-CPU   enable_failed  Could not disable: busy
6     DSK 0           A-DSK   enabled
7     DSK 1           B-DSK   disabled
8     HDD 0           A-DSK0  enabled
9     HDD 7           B-DSK0  disabled
10    ft_alarm 0
11    ft_core 0
12    ft_network 0
13    ft_serial 0
14    icn 0
15    icn 1
16    icn 2
17    icn 3
18    icn_system 0

(H)elp, (I)nclude, (E)xclude, (S)elect, (V)iew, (Q)uit or <Number> ?

```

FIGURE E-5 CMS Configuration After Split

Note that all the B-side FRUs, apart from the B-side motherboard, are in a disabled state.

Note – With patch D installed, the state of the B-CPU is reported incorrectly. Use `cmsconfig` to bring it to a disabled state.

Bringing Up the Losing Side

The losing (B) side will attempt a reboot from B-DSK0, finally resulting in a prompt to enter SEVM license keys:

CODE EXAMPLE 5-1 B-Side Reboot

```
Sun Ultra 4FT UPA/PCI(2 X UltraSPARC-II 296MHz), No keyboard
OpenBoot 3.7 [PROTO-Plb-sd_st: Fusion-B2], 256 MB memory installed, Serial #9539444.
Ethernet address 8:0:20:91:8f:74, Host ID: 80918f74.

Running preboot tests:                SUCCESS
Boot device: b-dsk0 File and args:
Loading: ufs-file-system package 1.4 04 Aug 1995 13:02:54
Fcode UFS Reader 1.10 96/10/15 00:57:29
Loading: /platform/SUNW,Ultra-4FT/ufsboot
Loading: /platform/sun4u/ufsboot
SunOS Release 5.6 Version 108145-10 [UNIX(R) System V Release 4.0]
Copyright (c) 1983-1997, Sun Microsystems, Inc.
WARNING: forceload of drv/scsi failed
NOTICE: Ultra-4FT DDI extensions installed
WARNING: forceload of drv/scsi failed
WARNING: forceload of drv/ssd failed
WARNING: forceload of drv/sf failed
WARNING: forceload of drv/pln failed
WARNING: forceload of drv/soc failed
WARNING: forceload of drv/socal failed
}rconsconfig: device at /pci@6,2000/u4ftbus@0/u4ioslot@4/ebus@4/se@14,400000:b is
inaccessible.
rconsconfig: device at /pci@6,2000/u4ftbus@0/u4ioslot@4/ebus@4/se@14,400000:a is
inaccessible.
VxVM starting in boot mode...
vxvm:vxconfigd: WARNING: Detaching plex rootvol-01 from volume rootvol
vxvm:vxconfigd: WARNING: Detaching plex usr-01 from volume usr
vxvm:vxconfigd: WARNING: Disk rootdisk in group rootdg: Disk device not found
vxvm: NOTE: Setting partition /dev/dsk/c3t0d0s1 as the dump device.
VxVM starting special volumes ( swapvol var )...
dumpvp_setup: Setting partition /dev/dsk/c3t0d0s1 as the dump device.
The / file system (/dev/vx/rdisk/rootvol) is being checked.
/dev/vx/rdisk/rootvol: 3023 files, 77365 used, 441552 free
/dev/vx/rdisk/rootvol: (336 frags, 55152 blocks, 0.0% fragmentation)
The /usr file system (/dev/vx/rdisk/usr) is being checked.
/dev/vx/rdisk/usr: 25591 files, 436045 used, 3099937 free
/dev/vx/rdisk/usr: (2785 frags, 387144 blocks, 0.0% fragmentation)
FIRST LOSER BOOT
YOU MUST ENTER THE VERITAS LICENCE KEYS HERE
Please enter your key: 7141 4164 7746 8669 8047 077

vrts:vxserial: INFO: Feature name: CURRSET [95]
```

CODE EXAMPLE 5-1 B-Side Reboot (Continued)

```
vrts:vxserial: INFO: Number of licenses: 1 (non-floating)
vrts:vxserial: INFO: Expiration date: Sun 14 Nov 1999 08:00:00 (33.9 days from now)
vrts:vxserial: INFO: Release Level: 20
vrts:vxserial: INFO: Machine Class: All
vrts:vxserial: INFO: Key successfully installed in /etc/vx/elm/95.
Please enter your key: 8919 6476 4493 0199 2931 280

vrts:vxserial: INFO: Feature name: RAID [96]
vrts:vxserial: INFO: Number of licenses: 1 (non-floating)
vrts:vxserial: INFO: Expiration date: Sun 14 Nov 1999 08:00:00 (33.9 days from now)
vrts:vxserial: INFO: Release Level: 20
vrts:vxserial: INFO: Machine Class: All
vrts:vxserial: INFO: Key successfully installed in /etc/vx/elm/96.
Licensed features:
  Mirroring
  Concatenation
  Disk-spanning
  Striping
  RAID-5
*** REBOOT THE SYSTEM ***
syncing file systems... done
rebooting...
Resetting ...
```

This will be followed by a second reboot of side B.

When side B has successfully rebooted, you can log into both sides. The winning side (A) retains the identity of seagoon, and the losing side (B) has the new identity seagoon-2.

Check that the system is actually in split mode by issuing the command `splitinfo` from both sides:

```
seagoon# splitinfo -d -a
domain = A
attributes = split, master
master = A
```

```
seagoon-2# splitinfo -d -a
domain = B
attributes = split
master = A
```

Issuing `cmsconfig` on side B (seagoon-2) should now produce the following:

```

seagoon# cmsconfig
-----
Item  Name          Fault  Loc    State                                     Page 1 of 1
-----
0     A-MBD 0          A-MBD  enabled
1     B-MBD 0          B-MBD  enabled
2     CAF 0           X A-CAF  enable_failed  Unable to power on FRU
3     CAF 1           B-CAF  enabled
4     CPU 0           X A-CPU  enable_failed  Unable to power on FRU
5     CPU 1           B-CPU  enabled
6     DSK 0           X A-DSK  enable_failed  Unable to power on FRU
7     DSK 1           B-DSK  enabled
8     HDD 0           A-DSK0 enable_failed  Cannot be enabled until DSK 0
9     HDD 7           B-DSK0 enabled
10    ft_alarm 0
11    ft_core 0          enabled          Running on B-CPU
12    ft_network 0       online_up        B (online)
13    ft_serial 0       online           1 (unusable) 0 (unusable)
14    icn 0          enabled
15    icn 1          disabled
16    icn 2          disabled
17    icn 3          disabled
18    icn_system 0     enabled
(H)elp, (I)nclude, (E)xclude, (S)elect, (V)iew, (Q)uit or <Number> ?

```

Note – When running patch D, the losing (A) side FRUs are recorded as failed and are in an `enable_failed` state. In fact, the FRUs are not faulty and you should use `cmsconfig` to bring them to a disabled state.

Note that the `icn_system` is enabled. To ensure that the sides are correctly communicating, issue the following commands on both sides:

```

seagoon# ifconfig -a
lo0: flags=849<UP,LOOPBACK,RUNNING,MULTICAST> mtu 8232
    inet 127.0.0.1 netmask ff000000
pnet0: flags=863<UP,BROADCAST,NOTRAILERS,RUNNING,MULTICAST> mtu 1500
    inet 129.156.203.152 netmask ffffffff00 broadcast 129.156.203.255
    ether 8:0:20:91:8f:54
icn0: flags=8843<UP,BROADCAST,RUNNING,MULTICAST,PRIVATE> mtu 40945
    inet 192.168.1.1 netmask ffffffff00 broadcast 192.168.1.255
    ether 8:0:20:91:0:52

```

```
seagoon# u4ftctl -d /dev/icn status
(icn#0): Online + Exporting + Importing
(icn#1): Not initialised
(icn#2): Not initialised
(icn#3): Not initialised
```

```
seagoon-2# ifconfig -a
lo0: flags=849<UP,LOOPBACK,RUNNING,MULTICAST> mtu 8232
    inet 127.0.0.1 netmask ff000000
pnet0: flags=863<UP,BROADCAST,NOTRAILERS,RUNNING,MULTICAST> mtu 1500
    inet 129.156.203.18 netmask fffffff0 broadcast 129.156.203.255
    ether 8:0:20:91:8f:76
icn0: flags=8843<UP,BROADCAST,RUNNING,MULTICAST,PRIVATE> mtu 40945
    inet 192.168.1.2 netmask fffffff0 broadcast 192.168.1.255
    ether 8:0:20:91:0:74
```

```
seagoon-2# u4ftctl -d /dev/icn status
(icn#0): Online + Exporting + Importing
(icn#1): Not initialised
(icn#2): Not initialised
(icn#3): Not initialised
```

Note that icn0 should have inet set to 192.168.1.1 on seagoon and 192.168.1.2 on seagoon-2. The netmasks should both be fffffff0, that is, 255.255.255.0.

Note also that the u4ftctl command should show any ICN networks in an Online+Exporting+Importing state. If they are not, disable the icn_system FRUs on both sides and then re-enable them. Check the output of the above commands again. If the ICN network instances are still not in the Online+Exporting+Importing state, stop and restart the split daemon by issuing the command:

```
seagoon# /etc/init.d/e4ftspltd stop
```

This will stop and restart the split daemon.

Check that the split daemon is running on each side:

```
seagoon# ps -efo args | grep u4ftspltd
/usr/platform/SUNW,Ultra-4FT/SUNWcms/lib/u4ftspltd
```

```
seagoon-2# ps -efo args | grep u4ftspltd
/usr/platform/SUNW,Ultra-4FT/SUNWcms/lib/u4ftspltd
```

If the daemon is not running, start it using the `u4ftspltd` command:

```
seagoon# /etc/init.d/u4ftspltd start
```

```
seagoon-2# /etc/init.d/u4ftspltd start
```

Updating seagoon-2

You can now make changes to `seagoon-2` as required. In this example, a few files are created on `seagoon-2` that should survive the merge.

```
seagoon-2# touch /toto /usr/toto /var/toto /opt/toto /etc/toto
```

You can also install a new operating system, new packages and patches, or upgrade an application on the split loser.

Merging the Systems

The steps required for merging the two systems are as follows:

1. Check the split daemons on both sides are communicating with each other.
2. Issue the `merge` command.
3. Bring the losing side FRUs back online.
4. Wait for the CPUsets to recombine.
5. Remirror the new boot disk.

Confirming the Split Daemons are Communicating

Before a merge can occur, the split daemons on the two sides of the machine must be communicating with each other over the chosen network (set up in `/etc/splitd.conf`), which is typically ICN. Confirm this by attempting to change split mastership on both sides:

```
seagoon# splitconf -m <side>  
master = A
```

`<side>` is dependent on the existing master: if the existing master is A, `<side>` should be B; if the existing master is B, `<side>` should be A.

- If the change of mastership proceeds without producing an error and the `splitinfo` command shows `<side>` as the master on both sides, the split daemons are communicating.
- If this is not the case, stop the split daemon on each side using the command `/etc/init.d/u4ftsplit stop`, which should result in the split daemons stopping and then restarting. If necessary, recheck split daemon communication using the above procedure.

Issuing the merge Command

You can issue the `merge` command from either side. In this case, it is issued from the merge winner:

```
seagoon-2# splitadm -w B merge  
icn network shutdown (via CMS): icn.keepalive done.  
domain = C  
attributes = combined  
master = A
```

This will force the losing side (A) to reboot. and the side A console displays the following:

```
icn network shutdown (via CMS): icn.keepalive done.

INIT: New run level: 6
The system is coming down. Please wait.
System services are now being stopped.
Print services stopped.
Stopping the syslog service.
syslogd: going down on signal 15
Killed download daemon.
Oct 11 11:26:33 snmpdx: received signal 15
The system is down.

INIT: failed write of utmpx entry:"s6"

INIT: failed write of utmpx entry:"rb"
syncing file systems... done
rebooting...
Resetting ...
```

At this point the side A console may go dead because the newly-combined system has the A-CAF disabled.

Waiting for the CPUsets to Become Fault Tolerant

Check by running `cmsconfig` and noting that `ft_core` has the value `Fault tolerant`.

Bringing the Losing Side FRUs Back Online

Issue the `cmsconfig` command on the merge winner (seagoon-2, side B):

```
seagoon# cmsconfig

Item  Name          Fault  Loc      State          Page 1 of 1
-----
0     A-MBD 0         A-MBD  enabled
1     B-MBD 0         B-MBD  enabled
2     CAF 0           A-CAF  disabled      WARNING: Power off command
3     CAF 1           B-CAF  enabled
4     CPU 0           A-CPU  busy
5     CPU 1           B-CPU  enabled
6     DSK 0           A-DSK  disabled      WARNING: Power off command
7     DSK 1           B-DSK  enabled
8     HDD 0           A-DSK0 disabled      WARNING: Power off command
9     HDD 7           B-DSK0 enabled
10    ft_alarm 0         usable
11    ft_core 0         enabled      Waiting for A-CPU to be ready
12    ft_network 0       online_up    B (online)
13    ft_serial 0       online      0 (unusable)
14    icn_system 0      disabled

(H)elp, (I)nclude, (E)xclude, (S)elect, (V)iew, (Q)uit or <Number> ?
```

Re-enable merge-losing side FRUs (in this case, A-CAF, A-DSK and A-DSK0). After re-enabling the A-CAF, the A-side console should return.

Remirroring the New Boot Disk

After the merge, you can check the state of the SEVM using the `vxdisk list` command, as follows:

```
seagoon-2# vxdisk list

DEVICE      TYPE      DISK      GROUP      STATUS
c2t0d0s2    sliced    -         -         online
c3t0d0s2    sliced    disk01    rootdg    online
-           -         rootdisk  rootdg    failed was:c2t0d0s2
```

The boot disk used by `seagoon-2`, containing the updated operating system, is `c3t0d0`, known to SEVM as `disk01`. The original disk used by `seagoon` before the upgrade is `c2t0d0`, known to SEVM as `rootdisk`, which has been failed by SEVM.

The merge process results in SEVM losing sight of `rootdisk` because the CMS tells it that it is disabled (which occurs because the merged system inherits its CMS state from the merge winner in which `rootdisk` was disabled).

The procedure now is to inform SEVM that `rootdisk` is still present in the Netra ft 1800. To do this, use the `vxdiskadm` command, main menu option 5 (Replace a failed or removed disk). The sequence of screens is as follows.

```
seagoon-2# vxdiskadm

Volume Manager Support Operations
Menu: VolumeManager/Disk

 1      Add or initialize one or more disks
 2      Encapsulate one or more disks
 3      Remove a disk
 4      Remove a disk for replacement
 5      Replace a failed or removed disk
 6      Mirror volumes on a disk
 7      Move volumes from a disk
 8      Enable access to (import) a disk group
 9      Remove access to (deport) a disk group
10      Enable (online) a disk device
11      Disable (offline) a disk device
12      Mark a disk as a spare for a disk group
13      Turn off the spare flag on a disk
list   List disk information

?      Display help about menu
??     Display help about the menuing system
q      Exit from menus

Select an operation to perform: 5
```

Replace a failed or removed disk
Menu: VolumeManager/Disk/ReplaceDisk

Use this menu operation to specify a replacement disk for a disk that you removed with the "Remove a disk for replacement" menu operation, or that failed during use. You will be prompted for a disk name to replace and a disk device to use as a replacement. You can choose an uninitialized disk, in which case the disk will be initialized, or you can choose a disk that you have already initialized using the Add or initialize a disk menu operation.

Select a removed or failed disk [<disk>,list,q,?] **list**

Disk group: rootdg

DM NAME	DEVICE	TYPE	PRIVLEN	PUBLEN	STATE
dm rootdisk	-	-	-	-	NODEVICE

Select a removed or failed disk [<disk>,list,q,?] **rootdisk**

Select disk device to initialize [<address>,list,q,?] **list**

DEVICE	DISK	GROUP	STATUS
c2t0d0	-	-	online
c3t0d0	disk01	rootdg	online

Select disk device to initialize [<address>,list,q,?] **c2t0d0**

This disk device is currently listed as in use by another host.
If you are certain that the other host is not using the disk, you
can choose to clear the use status. To use the disk the use status
must be cleared.

Output format: [Device_Name,Disk_Access_Name,Hostid]

[c2t0d0,c2t0d0s2,seagoon]

Clear use status? [y,n,q,?] (default: n) **y**

The following disk you selected for use appears to already have
been initialized for the Volume Manager. If you are certain the
disk has already been initialized for the Volume Manager, then you
do not need to reinitialize the disk device.

Output format: [Device_Name]

c2t0d0

Reinitialize this device? [y,n,q,?] (default: y) **y**

The requested operation is to initialize disk device c2t0d0 and
to then use that device to replace the removed or failed disk rootdisk in disk
group rootdg.

Continue with the operation? {y,n,q,?] (default: y) **y**

Replacement of disk rootdisk in group rootdg with disk device
c2t0d0 completed successfully.

Replace another disk? [y,n,q,?] (default: n) **n**

Check the state of SEVM using the `vxdisk list` command:

```
seagoon-2# vxdisk list
DEVICE      TYPE      DISK      GROUP      STATUS
c2t0d0s2    sliced    rootdisk  rootdg     online
c3t0d0s2    sliced    disk01    rootdg     online
```

Note that there are two devices, both of which are online and mapped to two SEVM
disks. At this point, SEVM knows about c2t0d0 and has assigned rootdisk to it,
as before. The mirroring from disk01 to rootdisk can now proceed.

Check for the existence of a process `vxrecover`, which resynchronizes the boot disk mirror:

```
# rip# ps -ef | grep recover
root  9764  4250  0 17:13:10 console  0:00 grep recover
root  9173      1  0 17:12:00 ?          0:00 vxrecover -sb -g rootdg rootdisk
```

If it doesn't exist, execute it by issuing the command `vxrecover`. The progress of this remirroring can be followed using either `vxprint`, or the X-windows graphical interface, `vxva`.

If SEVM is unable to find the disk device `c2t0d0` for `rootdisk`, add the disk and initialize it before you select option 4.-

This is achieved by selecting option 1 from the main menu of `vxdiskadm`

