Case Study: Configuration and Administration of HP-UX for HA

Daniel Hartline
AT&T Broadband and Internet Services
4100 E. Dry Creek Rd.
Littleton, CO 80122
hartline.dan@tci.com

Table of Contents

Forward

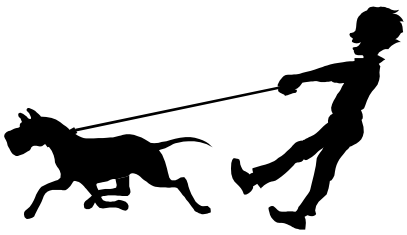

- **Installation, administration and security procedures**
- **Administration philosophy**
- **Assumes a working knowledge of HP-UX 10.20**
- **It's not enough to simply perform the provided procedures. You are required to LEARN!!!!!!!!!!!!**
- **Develop an increasing level of understanding of the procedures and philosophies documented here.**

This document provides detailed HP-UX 10.20 (9000/800) installation, administration and security procedures; it also documents the managerial philosophies and standards applied to the finance environment as a whole.

This document does not offer highly detailed explanations of the purpose or impact of each command or process. A working knowledge of a wide variety of HP-UX administration and management concepts is assumed.

It is not enough to simply perform the procedures provided, you are required to develop an increasing level of understanding of the information contained in this document. If you are not comfortable with a command or procedure read the appropriate man pages, ask a team member or open a call with the HP response center. If you are still not comfortable, request assistance, a team member will perform the procedure with you or provide the appropriate guidance.

Forward Con't

- **Provide reasonable documentation**
- **Provide overview of environment**
- **Procedures performed in similar sequence**
- **Achieve and maintain standardization**
- **Function as a learning tool**

The purpose of this document is to:

- Provide reasonable documentation of the environment. By itself this document does not provide specific documentation for each system, it does offer documented standards that are applied to each system. System specific documentation can be viewed on-line in the /usr/local/bin/system_configs directories.
- Provide a general overview of the environment. This overview will equip each administrator with a high level understanding of the configuration standards we utilize and what is expected on a daily basis. This understanding will facilitate the success of each administrator.
- Ensure that each procedure is performed in similar sequence. Many of the procedures contained in this document are very complex. To ensure success and consistency, each procedure will be performed as presented.
- Ensure that system standardization is achieved and maintained. The configuration standards we have implemented are essential to the consistent success of patch management, security, troubleshooting and daily support.
- Function as a learning tool. Regardless of the experience level of each administrator this document offers a learning road map for the environment. In other words if you learn/understand all the topics/commands included in this document you will be a better-equipped administrator and therefore your contribution to this environment will be enhanced.

Rules of the Road

- **Expectations**
- **Environment is Unique**
- **Mistakes**
- **Standardization**
- **Fighting Fires**

The following guidelines are good rules for administrators to keep in mind and they are a partial list of what is expected of each administrator in this environment.

- All configuration changes must be communicated to and approved by the team lead.
- All configuration changes must be communicated to the team.
- Do not implement new configuration standards just prior to or during a close period.
- Mistakes are a fact of life: they will however, be the exception and not the rule.
- No development or testing will occur on production machines. This rule applies to developers, DBAs and administrators.
- When you issue a command its success or failure is your responsibility. Whether the command was given to you by a co-worker, a HP engineer or it was derived from this document the result is your responsibility.
- It is highly recommended that you use a logbook to manage your task lists and document commands/procedures when troubleshooting. This is a recommendation, we prefer to allow each administrator to perform the job in the ways that they have found successful. However if tasks are frequently incomplete, incorrect or if you are unable to communicate the steps/commands you executed on the systems while performing a procedure or troubleshooting a more structured work arrangement may be required.
- Configuration standards are paramount to the success this environment enjoys. Every effort will be made to ensure that the standards are maintained on every system. If for instance, you are troubleshooting a problem with swinstall and identify a configuration change that will resolve it, the change you identified will be implemented on all systems (with approval) in the environment even if the problem is not displayed on every system.
- All configuration changes will be deployed on development and test machines; thoroughly tested and then after approval is obtained they will be moved into production.

- We do not spend a great deal of time fighting fires. If your approach to administration is to fight fires rather than plan for the future because you're too busy …change your approach. We move forward, we identify problems and we resolve them, we establish goals and we meet them. A while ago a new employee suggested we create a script to find core files and delete them. I responded with a NO and explained that while we do come across the occasional core file they're fairly rare. If core files become a problem we will figure out why and fix the problem, we will not spend time simply deleting them. You should have seen the look on his face; he never did understand that when we identify a problem we resolve it, no buts about it.
- Each administrator will participate in the on-call pager rotation. All on-call pages will be responded to within 15 minutes of receipt. While carrying the on-call pager it is the administrator's responsibility to monitor the service level of the pager. If the pager is not in a "full service" mode the administrator will power cycle and or replace the batteries. If the pager is still not in full service mode the administrator will notify the team lead and operations manager. Until the problem is resolved the administrator will call Skytel every 2 hours to check for any pages that may have been sent.

Installing and Configuring HP-UX 10.20

- **These installation instructions are specific to cold installs of HP-UX 10.20 on HP 9000/800 T and K class systems.**
- **Boot procedures described here are specific to K class systems.**

Selecting Appropriate Devices for the Installation

- **Assumes Jamaica Storage Enclosure Utilizing 4 GB Devices**
  Minimum # of Devices
  9 GB Devices
- **K, T5XX and T600**

These installation instructions assume that HP-UX 10.20 is being installed on a Jamaica Storage Enclosure (Model A3312A) that contains 4GB drives. A minimum of 6 4GB drives is required to support this configuration.

Set the SCSI ids.

From top to bottom the drives are set to 6, 5, 4, 3 on each side of the enclosure.

POWER SUPPLY	FAN	FAN
	TARGET 6	TARGET 6
POWER SUPPLY	TARGET 5	TARGET 5
	TARGET 4	TARGET 4
	TARGET 3	TARGET 3

Jamaica Enclosure, 4 GB Devices.

The primary boot device will always be located on the left side of the Jamaica enclosure and the secondary and tertiary devices will always be on the right (assumes you are looking at the front of the enclosure).

POWER SUPPLY	FAN	FAN
	TARGET 6 – PRI. BOOT	TARGET 6 – SEC. BOOT
POWER SUPPLY	TARGET 5	TARGET 5
	TARGET 4	TARGET 4
	TARGET 3	TARGET 3 – TERT. BOOT

Jamaica Enclosure, 4 GB Devices.

The buses selected for connectivity to the Jamaica enclosure will vary depending on the class of system you are installing on. The following describes the standard connectivity we utilize for K, T5xx and T600 class systems.

When installing on a K class box the Jamaica enclosure will be connected to the system in the following fashion (requires HSC daughter card installed on system multi-function IO card). If the K box is configured with internal disks (10/0) care should be taken to ensure SCSI ID conflicts do not exist.

	BUS 10/0	BUS 10/8
POWER SUPPLY	FAN	FAN
	10/0.6 – PRI. BOOT	10/8.6 – SEC. BOOT
POWER SUPPLY	10/0.5	10/8.5
	10/0/4	10/8.4
	10/0.3	10/8.3 – TERT. BOOT

Jamaica Enclosure, 4 GB Devices.

When installing on a T600 the Jamaica enclosure will be connected to the system in the following fashion (requires HP-PB F/W cards installed in 0/28/52 and 0/28/36).

	BUS 0/28/52	BUS 0/28/36
POWER SUPPLY	FAN	FAN
	0/28/52.6 – PRI. BOOT	0/28/36.6 – SEC. BOOT
POWER SUPPLY	0/28/52.5	0/28/36.5
	0/28.52.4	0/28/36.4
	0/28/52.3	0/28/36.3 – TERT. BOOT

Jamaica Enclosure, 4 GB Devices.

When installing on a T520 or T500 the Jamaica enclosure will be connected to the system in the following fashion (requires HP-PB F/W cards installed in 0/52 and 0/36).

	BUS 0/52	BUS 0/36
POWER SUPPLY	FAN	FAN
	0/52.6 – PRI. BOOT	0/36.6 – SEC. BOOT
POWER SUPPLY	0/52.5	0/36.5
	0/52.4	0/36.4
	0/52.3	0/36.3 – TERT. BOOT

Jamaica Enclosure, 4 GB Devices.

HP-UX 10.20 Install Utility

CUSTOMIZATION

- **Primary Boot Disk**
- **LVM with VxFS**
- **/HOME CONFIGURATION = NONE**
- **Modify File System Sizes**

Once you have booted the system from the HP-UX 10.20 Install/Core OS CD you will be presented with the "Install Utility Main Menu"

1-select "Install HP -UX"

2-Would like to enable networking now?

"*n*" for no.

3-From the "HP-UX Install Utility - Select System Root Disk" window, select the entry for the primary boot disk.

K class = 10/0.6
T600 = 0/28/52.6
T5xx = 0/52.6

Select "OK" to proceed.

4-From the "HP-UX Install Utility - Select Whole System Configuration" window, select.

"LVM Configuration with VxFS"

Select "OK" to proceed.

5-From the “HP-UX Install Utility - View/Modify Basic Configuration” window, Set the following values:

Swap	512 MB
Secondary Swap	NONE
Software Selection	CDE Runtime Environment
Language	English
Locale Setting	Default
File Name Length	LONG
/HOME Configuration	NONE
# of Disks in Root VG	1
Make Volatile Dirs. Separate	TRUE
Create /export	FALSE

Select “OK” to proceed.

Note: If you are installing from the 10.20 TFC Install disk the parameter “Load ONC+ Networking Enhancements” will also appear in this window; it should be set to FALSE.

6-From the “HP-UX Install Utility - System Configuration” window Select “Modify FS Parameters”.

7-From the “HP-UX Install Utility – Configure File Systems” window set the following values in MBs for each file system.

/	252 MB
/stand	140 MB
Pri swap	512 MB
/opt	856 MB
/tmp	300 MB
/usr	1016 MB
/var	1016 MB

Select “OK” to proceed.

8-From the “HP-UX Install Utility – System Configuration” window.

Select “OK”.

9-If the device you are installing on has been previously used you will receive a warning.

Select “Continue”.

10-Do you want to interact with SD-UX during the install?

Select “NO”

The installation process will take approximately 20 minutes.

11-The system will reboot when the installation is complete. Upon a successful boot you will be asked for the following information.

Would you like to configure networking? NO

Would you like to set the host name and time zone? YES

Enter the chosen host name and appropriate time zone.

If the system date and time are incorrect set them.

Set root's password.

Would you like to set this system up as a font server? SKIP this step.

12-HP-UX 10.20 is now installed, Selecting "OK" will boot the system to multi-user mode.

Installing HP-UX Applications

- **MirrorDisk/UX**
- **Glance Plus, MeasureWare Agent**
- **On-Line JFS**
- **Contributed Tools**
- **Diagnostics and Support Bundle**
- **Lsof**
- **HP Jetadmin**

1-Add the following two lines to the end of the /var/adm/sw/defaults file.

swinstall.polling_interval = 60
swcopy.polling_interval = 60

Save the defaults file.

2-You must now stop and restart the swagentd daemon to reread the defaults file.

ps –ef | grep swagentd
kill process_id_swagentd_daemon
/usr/sbin/swagentd

3-If you are installing 10.20 TFC and the Fibre Channel Mass Storage Driver the following steps must be completed prior to installing the standard applications. If you are not installing 10.20 TFC and the Fibre Channel driver continue with step #4.

Install the “Fibre Channel Mass Storage Driver” from the applications depot. The “Match what target has” feature will NOT be used for this installation. You will need to manually mark the driver for installation. After a successful installation the system will reboot

swinstall

Now apply the Fibre Channel patches from the /var/spool/sw/fibre_channel depot. The “Match what target has” feature in swinstall will be used for this installation. After a successful installation the system will reboot

swinstall

4-Selecting appropriate applications for the installation. Application requirements may vary; the following application list will be appropriate for most installations.

MirrorDisk/UX
HP GlancePlus/UX for s800 10.20
HP OnlineJFS
HP MeasureWare Server Agent
Contributed Tools
Hewlett-Packard Jetadmin
HP 10.0 Support Tools Bundle
Lsof

Other applications that might be required include:

HP C/ANSI
HP PerfVeiw Analyzer
HP PerfView Planner
Internet Access Software
Netscape Fastrack Server

5-Use swinstall to perform these installations. The “Match what target has” feature will NOT be used for this installation. You will need to manually mark each application. The applications are located in the applications depot.

swinstall

The system will reboot as part of the application install.

6-After the system has rebooted perform the following steps to ensure all filesets were configured at boot time and observe log files for any problems that may have been introduced by the patches.

swlist –l fileset –a state | grep install

Observe the output of this command. All filesets should be in a “configured” state. If any filesets are displayed in an “installed” state you may attempt to manually configure them. Execute the following swconfig command.

*swconfig **

Upon the completion of this command use the above swlist command to verify that all filesets are now configured. If there are still “installed” filesets investigate the /var/adm/sw/swagent.log file and notify a team member.

Observe the output of *dmesg* and the log files, /etc/rc.log and /var/adm/syslog/syslog.log for any problems that may have been introduced by the applications.

Patching the New Install

- **Select Appropriate Patch Depots**
  - **Prior to Executing Swinstall**
    - **Swinstall**
  - **“Match What Target Has”**
  - **After each Depot Installation**

swlist -l fileset -a state | grep install

swconfig *

- **Cleanup Utility**

1-After a successful application installation apply the appropriate patch depots.

The patch depots are in a constant state of evolution. An installation may not utilize the same depots that were required two months ago. It is your responsibility to ask and identify the appropriate set of patch depot(s) for this installation.

List the appropriate depot(s) here:

2-Prior to the application of each patch depot stop the mwa software and kill the transaction tracker daemon (ttd).

mwa stop
ttd -k

3-Use swinstall to apply the depot(s). The “Match what target has” feature MUST be used!

swinstall

4-After each depot is installed and the system has rebooted (if necessary) execute the following swlist command to verify that all the filesets were successfully configured and observe the system log files for any problems that may have been introduced by the patches.

swlist –l fileset –a state | grep install

Observe the output of this command. All filesets should be in a “configured” state. If any filesets are displayed in an “installed” state you may attempt to manually configure them. Execute the following swconfig command.

*swconfig **

Upon the completion of this command use the above swlist command to verify that all filesets are now configured. If there are still “installed” filesets investigate the /var/adm/sw/swagent.log file and notify a team member.

Observe the output of *dmesg* and the log files /etc/rc.log and /var/adm/syslog/syslog.log for any problems that may have been introduced by the patches.

5-Repeat steps 2, 3 and 4 until all required patch depots have been applied.

6-Depending on the number of patch depots applied during the installation it may be prudent to run the cleanup utility. If 2 or fewer patch depots were installed do not execute the cleanup command, if more than 2 patch depots were installed the cleanup utility should be executed.

With no arguments this command will remove all patch history up to 1 level back and it will trim the SD logs to the last 5 entries.

cleanup

You will be asked “Would you like the logs trimmed to the last 5 entries?” enter y and hit return.

Various Installation Tasks

Standardization

- **Sequence is Significant**
- **Automated vs Manual Installation**
- **Standards Documented**

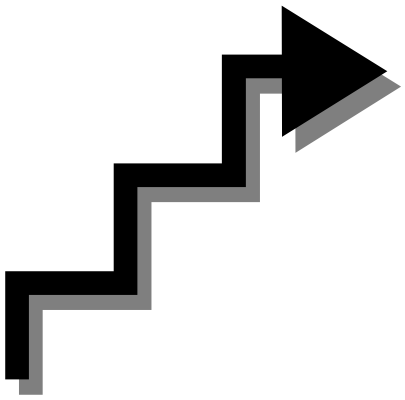

Various Installation Tasks Con't

/overlord
Local Account
nettune -s tcp_random_seq 2
secure_it RC set
sys_config script
NTP
Inetd Logging
MWA
/etc/default/fs
/sbin/ioinitrc
catman -w
sudo
Predictive/UX

Various Installation Tasks Con't

/etc/shells

SNMP

mkboot -a "hpux -lq" /dev/rdsk/....

/etc/fstab.aftersplit

tcp_wrappers

chmod -x /usr/bin/quota

monitor_setuid

Double Mirroring vg00

Creating Secondary and Tertiary Boot Devices

- **Add Devices to vg00**
- **No Quorum**
- **Update LIF Header with Diagnostics**
- **Create Mirror**
- **Alter Mirror Write Cache and Consistency for Primary Swap**

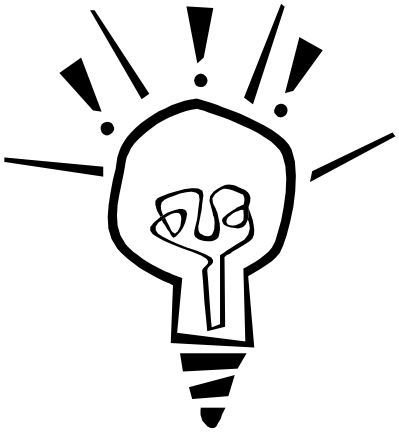

Note: This procedure contains steps that are specific to a cold installed HP-UX 10.20 system.

The devices you will use for the alternate and tertiary boot devices were selected at the beginning of this install process (see “Selecting Appropriate Devices for Installation” section). You are now ready to configure these devices and to add them to vg00. These instructions assume you are building a double mirror configuration.

	BUS X	BUS Y
POWER SUPPLY	FAN	FAN
	(6)PRI. BOOT	(6)SEC. BOOT
POWER SUPPLY	(5)	(5)
	(4)	(4)
	(3)	(3)TERT. BOOT

Jamaica Enclosure, 4 GB Devices.

1-Create the boot header on the secondary and tertiary boot devices.

pvcreate -B /dev/rdsk/c?t?d? *Secondary*
pvcreate -B /dev/rdsk/c?t?d? *Tertiary*

Where c?t?d? is the device file of the secondary and tertiary boot devices. If you are unsure of the correct device names they can be determined form the ioscan –fn command.

2-Add these devices to vg00

vgextend /dev/vg00 /dev/dsk/c?t?d? /dev/dsk/c?t?d?

3-Create the boot LIF headers on these devices, modify quorum and update the LIF headers with diagnostics.

mkboot /dev/rdsk/c?t?d? *Secondary*
mkboot /dev/rdsk/c?t?d? *Tertiary*

mkboot -a "hpux -lq" /dev/rdsk/c?t?d? *Secondary*
mkboot -a "hpux -lq" /dev/rdsk/c?t?d? *Tertiary*

Verify with lifcp /dev/rdsk/c?t?d?:AUTO -

Update boot headers with diagnostics.

cd /usr/sbin/diag/lif
mkboot –b updatediaglif –p ISL –p AUTO –p HPUX –p LABEL /dev/rdsk/c?t?d *Secondary*
mkboot –b updatediaglif –p ISL –p AUTO –p HPUX –p LABEL /dev/rdsk/c?t?d? *Tertiary*

Verify with lifls /dev/rdsk/c?t?d?.

4-Mirror the vg00 logical volumes to these devices. The order in which the mirror devices are passed to the following lvextend commands and the order in which they are issued is significant.

	Secondary	Tertiary
lvextend -m 2 /dev/vg00/lvol1	*/dev/dsk/c?t?d?*	*/dev/dsk/c?t?d?*
lvextend -m 2 /dev/vg00/lvol2	*/dev/dsk/c?t?d?*	*/dev/dsk/c?t?d?*
lvextend -m 2 /dev/vg00/lvol3	*/dev/dsk/c?t?d?*	*/dev/dsk/c?t?d?*
lvextend -m 2 /dev/vg00/lvol4	*/dev/dsk/c?t?d?*	*/dev/dsk/c?t?d?*
lvextend -m 2 /dev/vg00/lvol5	*/dev/dsk/c?t?d?*	*/dev/dsk/c?t?d?*
lvextend -m 2 /dev/vg00/lvol6	*/dev/dsk/c?t?d?*	*/dev/dsk/c?t?d?*
lvextend -m 2 /dev/vg00/lvol7	*/dev/dsk/c?t?d?*	*/dev/dsk/c?t?d?*

5-Update the LVM configuration with this new boot information.

lvlnboot -b /dev/vg00/lvol1
lvlnboot -r /dev/vg00/lvol3
lvlnboot -s /dev/vg00/lvol2

6-Boot to LVM maintenance mode and modify the mirror consistency and write cache parameters for primary swap (lvol2).

s*hutdown –r now*

Interrupt the autoboot process and the continue to boot from the primary boot device.

Interact with IPL.

From the ISL prompt enter the following to boot to LVM maintenance mode.

hpux –lm

7-Once you are booted to LVM maintenance mode enter the following commands to modify the mirror consistency and write cache parameters for primary swap (lovl2).

vgchange -a y /dev/vg00
lvchange -M n /dev/vg00/lvol2
lvchange -c n /dev/vg00/lvol2
vgchange –a n /dev/vg00

8-Reboot the system.

reboot

Testing Secondary and Tertiary Boot Devices

- **Mirror Merged**
- **Setboot**
- **Single User Mode**
- **sea IPL (tertiary device)**

Note: This procedure contains steps that are specific to a cold installed HP-UX 10.20 system.

1-Use the setboot command to verify the secondary boot path and AUTOBOOT/AUTOSEARCH flags in stable storage.

setboot

Observe the output of this command and verify that the primary and secondary boot paths are correct and that Autosearch and Autoboot are both ON or Enabled.

The following is examples of correct setboot output for each class of system in our environment:

K Class

Hostname
Primary bootpath : 10/0.6.0
Alternate bootpath : 10/8.6.0
Autoboot is ON (enabled)
Autosearch is ON (Enabled)

T600

Hostname
Primary bootpath : 0/28/52.6.0
Alternate bootpath : 0/28/36.6.0
Autoboot is ON (enabled)
Autosearch is ON (Enabled)

T520 and T500

Hostname
Primary bootpath : 0/52.6.0
Alternate bootpath : 0/36.6.0
Autoboot is ON (enabled)
Autosearch is ON (Enabled)

2-If the stable storage boot values are not correct use the setboot command to modify them. See man setboot for specific instructions. The Autosearch and Autoboot flags need only be set once not for each boot device.

3-Test booting from the secondary and tertiary boot devices. These boot tests require that the double boot mirror is synced (active).

Verify the state of the mirrors

vgdisplay –v /dev/vg00 | more

If the mirrors are synced (lvol1 through lvol7 reside on 3 physical volumes each) continue with step #4. If the mirrors are not synced, sync them before continuing.

cd /usr/local/bin
./lv_mergeit

4-Reboot the system.

shutdown –r now

Note: This boot process is specific to a K class system.

Interrupt the autoboot process.

Boot from the secondary or alternate device.

bo alt

Interact with IPL and boot to single user mode.

From the ISL prompt.

hpux –is

5-Assuming a successful boot on the secondary boot device you are now ready to test the tertiary device.

6-Reboot the system.

reboot

Note: This boot process is specific to a K class system.

Interrupt the autoboot process.

Use the search command to identify the tertiary boot device and assign a boot value to it.

s*ea ipl*

The sea command will assign p*X* values to each bootable device found.

Issue the boot command for the tertiary device.

bo pX

Interact with IPL and boot to single user mode.

From the ISL prompt.

hpux –is

7-Assuming a successful boot to the tertiary device reboot the system and allow it to come up from the primary device.

reboot

Adding Remaining Devices to vg00

- All Devices in Jamaica Enclosure
- 32 GB for vg00

All “UNUSED” devices in the Jamaica will belong to vg00.

BUS X		BUS Y
POWER SUPPLY	FAN	FAN
	(6)PRI. BOOT	(6)SEC. BOOT
POWER SUPPLY	(5)UNUSED	(5)UNUSED
	(4)UNUSED	(4)UNUSED
	(3)UNUSED	(3)TERT. BOOT

Jamaica Enclosure, 4 GB Devices.

1-For each device to be added to vg00.

pvcreate /dev/rdsk/c?t?d?

2-Add the devices to vg00. All devices may be passed to a single vgextend command.

vgextend /dev/vg00 /dev/dsk/c?t?d? dev/dsk/c?t?d? ...

Configuring Swap and Dump

- **No File System Swap**
- **Primary Swap**

512 MB
Priority = 1
Never Dump

- **Secondary Swap**

2 Logical Volumes
Equal in Size
Target 5 Disks
2x Physical Memory +
Priority = 0
Always Dump

When configuring swap and dump the following rules will always be true.

We do not use file system swap.

Primary swap will never be configured as a dump device.

The minimum amount of secondary swap created will be equal to or greater than the total amount of physical memory.

The priority settings of swap devices will be:

Primary swap will be set to a priority of 1. Regardless of the amount of secondary swap created, secondary swap will always consist of two logical volumes that are equal in size and reside on each target 5 disk in vg00. Each of these swap devices will have a priority of 0.

	BUS X	BUS Y
POWER SUPPLY	FAN	FAN
	(6)PRI. BOOT	(6)SEC. BOOT
POWER SUPPLY	(5)SEC. SWAP, PRI=0	(5)SEC. SWAP, PRI=0
	(4)	(4)
	(3)	(3)TERT. BOOT

Jamaica Enclosure, 4 GB Devices.

Configuring Swap

The minimum amount of secondary swap created will be just greater than the total amount of physical memory. The two logical volumes you create will need to be equal in size and they will reside on the target 5 devices in vg00.

1-Create two logical volumes.

lvcreate -C y -r n /dev/vg00
lvcreate -C y -r n /dev/vg00

If these installation instructions have been performed in the order they are presented these logical volumes will be lvol8 and lvol9.

2-Extend these logical volumes onto the appropriate devices.

lvextend -L xxxx /dev/vg00/lvol8/dev/dsk/c?t5d0
lvextend -L xxxx /dev/vg00/lvol9/dev/dsk/c?t5d0

3-Use sam to: define these logical volumes as device swap, set their priorities = 0, update /etc/fstab and create a new kernel (maxswapchunks). A reboot will be necessary.

s*am*

After the system has rebooted the swapinfo command can be used to verify swap.

4-Edit the /etc/fstab.aftersplit file and add the secondary swap entries that were just created in /etc/fstab.

For example add the following lines to the end of the /etc/fstab.aftersplit file.

/dev/vg00/lvol8 … swap pri=0 0 0
/dev/vg00/lvol9 … swap pri=0 0 0

Configuring Dump

The secondary swap devices we just created will now be configured as dump devices.

Primary swap will be de-configured as a dump device.

1-Issue a swapinfo command and determine the secondary swap logical volume(s).

swapinfo

2-Issue the following command for each secondary swap logical volume.

lvlnboot -d lvol? /dev/vg00

3-Remove primary swap from the dump configuration.

lvrmboot -d lvol2 /dev/vg00

4-Update the LVM configuration.

lvlnboot –R

5-Verify the configuration.

lvlnboot -v

Configuring Crash

- **Target 3 Device**
- **4 GB**

The /var/adm/crash file system is created to hold core dumps should they occur. We typically use the entire target 3 device on the X side. Therefore we end up with a 4 GB /var/adm/crash file system.

	BUS X	BUS Y
POWER SUPPLY	FAN	FAN
	(6)PRI. BOOT	(6)SEC. BOOT
POWER SUPPLY	(5)SEC. SWAP, PRI=0	(5)SEC. SWAP, PRI=0
	(4)	(4)
	(3)CRASH	(3)TERT. BOOT

Jamaica Enclosure, 4 GB Devices.

1-Create the “CRASH” logical volume.

lvcreate /dev/vg00

2-Extend this logical volume onto the appropriate device with the required size

lvextend -L ???? /dev/vg00/lvol? /dev/dsk/c?t?d?

Note the name of the logical volume created.

3-Create a file system on this logical volume.

newfs /dev/vg00/rlvol?

4-Edit /etc/fstab and add a line for this logical volume mounted as /var/adm/crash.

/dev/vg00/lvol?? /var/adm/crash vxfs delaylog 0 2

5-Issue the following mount command to read the /etc/fstab file and mount this file system.

mount –a

6-A successful mount may be verified with the bdf command.

bdf

Configuring 4th Boot Disk

- **Hot Swappable Disk Pays Off**
- **Disaster Recovery**
- **Labeled**
- **Offsite Storage**

- **Stale Logical Volume(s)**
- **monitor_config create_standard**

Note: This procedure contains steps that are specific to a cold installed HP-UX 10.20 system.

Before beginning this process verify that the logical volumes on the tertiary boot disk are synced with the mirror and note the device name of the tertiary disk. If the mirrors are not synced issue the following command to sync them.

cd /usr/local/bin
./lv_mergeit

See appendix B for the contents of lv_mergeit.

1-Execute the sys_config script to create a current online backup of this configuration.

cd /usr/local/bin
./sys_config

See appendix A for the contents of sys_config.

2-Boot the system to single user mode on the primary boot disk.

shutdown –r now

Note: This boot process is specific to a K class system.

Interrupt the autoboot process and continue to boot from the primary device.

Interact with IPL. From the ISL prompt boot to single user mode.

hpux -is

3-Mount all vg00 logical volumes

mount –a

4-Verify the mounts.

bdf

5-Remove the tertiary boot disk (see table below) from the Jamaica enclosure.. DO NOT pull the disk all the way out of the enclosure right away. Pull the disk out about 1 inch and allow it to sit for 30 seconds, then remove it from the enclosure and place it in a static bag.

	BUS X	BUS Y
POWER SUPPLY	FAN	FAN
	(6)	(6)
POWER SUPPLY	(5)	(5)
	(4)	(4)
	(3)	(3)TERT. BOOT

Jamaica Enclosure, 4 GB Devices.

6-It will be necessary to force at least one of the vg00 mirrors to become “stale”. Performing IO to any of these logical volumes will accomplish this.

cd /usr
ll

Use vgdisplay to verify the state of at least one of the mirrors is “stale”.

vgdisplay –v /dev/vg00 | more

Do not continue this procedure until you have verified that at least one of the vg00 mirrors (lvol1 through lvol7) is in a “stale” state.

7-Unmount all vg00 logical volumes (In single user mode you cannot unmount /).

umount –a

8-Boot to LVM maintenance mode.

reboot

Note: This boot process is specific to a K class system.

Interrupt the autoboot process and continue booting from the primary device.

Interact with IPL. From the ISL prompt boot to LVM maintenance mode.

hpux -lm

8-Insert the 4th boot disk into the Jamaica enclosure and allow it to spin up.

9-vgcfgrestore the device.

vgcfgrestore -n /dev/vg00 /dev/rdsk/c?t?d?

Where c?t?d? is the device name of the tertiary device.

If the vgcfgrestore command fails with a “cannot open device” message do the following.

Ioscan –fn

Now try the vgcfgrestore again.

10-After a successful vgcfgrestore boot the system and allow it to come up to multi user mode on the primary boot device.

reboot

When vg00 is activated the system will automatically sync the mirrors. This sync can be detected by the activity lights on the boot devices and the existence of an “/sbin/vgsync /dev/vg00” process.
If the system does not sync vg00 automatically issue the following vgsync command.

vgsync /dev/vg00

11-When the sync is completed issue the following mkboot commands to repopulate the boot header on the tertiary boot device.

mkboot /dev/rdsk/c?t?d?
mkboot -a "hpux -lq" /dev/rdsk/c?t?d?

Where c?t?d? is the device name of the tertiary device.

Propagate the diagnostic utilities to the tertiary boot device.

cd /usr/sbin/diag/lif
mkboot –b updatediaglif –p ISL –p AUTO –p HPUX –p LABEL /dev/rdsk/c?t?d?.

Where c?t?d? is the device name of the tertiary device.

12-Test booting from the new tertiary boot disk.

Reboot the system.

shutdown –r now

Note: This boot process is specific to a K class system.

Interrupt the autoboot process.

Use the search command to identify the tertiary boot device and assign a boot value to it.

s*ea IPL*

The sea command will assign p*X* values to each bootable device found.

Issue the boot command for the tertiary device.

bo pX

Interact with IPL and boot to single user mode. From the ISL prompt.

hpux –is

Assuming a successful boot to the tertiary device reboot the system and allow it to come up from the primary device.

reboot

13-Earlier in the install process tasks were performed to prepare for the possibility that we would need to boot from a split mirror boot disk. Verify that these files/directories exist and are correct.

The directory /split_root (root:sys 700) should exist and it should be empty.

The file /etc/fstab.aftersplit should exist and it should contain the following.

```
/dev/vg00/lvol3b   / vxfs delaylog 0 1
/dev/vg00/lvol1b   /stand hfs defaults 0 1
/dev/vg00/lvol4b   /opt vxfs delaylog 0 2
/dev/vg00/lvol5b   /tmp vxfs delaylog 0 2
/dev/vg00/lvol6b   /usr vxfs delaylog 0 2
/dev/vg00/lvol7b   /var vxfs delaylog 0 2
/dev/vg00/lvol8 … swap pri=0 0 0
/dev/vg00/lvol9 … swap pri=0 0 0
```

14-Create a copy of the /etc/fstab to account for any changes made during the install.

cp /etc/fstab /etc/fstab.b4split

15-Split the mirrors (tertiary from primary/secondary). Execute the following command file.

cd /usr/local/bin/
./lv_splitit

See appendix B for the contents of lv_splitit

Note: This is the production run state of the boot mirrors (tertiary split from primary/secondary).

Verify the split.

vgdisplay -v /dev/vg00

16-Run fsck on each lvolxb to ensure there has not been any corruption across the split.

cd /usr/local/bin
./fsck_it

17-Verify that the /dev/vg00 directory on the split (tertiary) disk was updated with the correct device files when the lv_splitit script was executed.

Mount lvol3b and take a look at /dev/vg00.

mount /dev/vg00/lvol3b /split_root
cd /split_root/dev/vg00
ll

The output of the ll command should show that each split logical volume is represented by a lvolxb and an rlvolxb device file. For example in addition to the device files lvol1 and rlvol1 we should now have:

lvol1b
rlvol1b

If the correct device files do not exist in /split_root/dev/vg00 directory you will need to create them using the mknod command. If you are not absolutely sure of the appropriate parameters to pass to the mknod command do not continue, ask a team member for assistance.

18-Now umount /dev/vg00/lvol3b.

cd
umount /split_root

19-Label the boot disk; with date, time and state of system; that is now out of the Jamaica enclosure and store it offsite

20-The procedure for creating the fourth boot disk is now complete. This is a good time to create the system configuration standards files to use for comparisons. This script creates standards for configuration aspects such as lvlnboot, swap and vg00.

cd /usr/local/bin
./monitor_config create_standard

See appendix E for the contents of monitor_config script

Hardware Architectures

Standardization

- **vg00**
- **Speed of Device - Speed of Connectivity**
  Hardware priority
  HP-PB, HSC, Fibre Channel
- **PV Links**
- **RAID Levels**

Application Architectures

No
user executables, data or configuration files will reside in any vg00 file systems unless specifically approved.

Don't Cry!

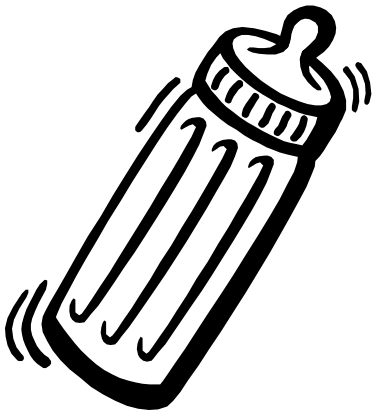

The administration team and the application owner(s) will work together to identify the best possible application architectures. The administration team must agree with the client's plans, if disagreement exists the two groups must work to solve the problem(s) and agree on a solution that will ensure the success of each application and adherence to defined standards

The following standard has been defined for this environment:

> NO application, database or user executables/data will reside in vg00.

In other words all file systems created for applications/clients will be outside of vg00 and all directory structures created for applications will also be outside of vg00. Many applications will by default put configuration files in the vg00 file systems. At times this will be tolerated however you should not assume this is the case. Soft links or additional file systems can be created to manage this situation.

This standard benefits the environment in the following ways:

> 1-If the operating system and application data are two separate physical entities troubleshooting of either one is greatly simplified.
>
> 2-Since the operating system is a separate entity it can easily be reinstalled if necessary without affecting application data/configurations.
>
> 3-Since the operating system is a separate entity it can be easily upgraded without affecting application data/configurations.

Client/DBA Configuration Requirements

Working with client determine configuration issues that are specific to the applications/databases the server will host.

Ensure that client needs are meet and standardization is maintained.

All configuration requests must be presented to the administration team in writing.

The following is a list if configuration issues that may or may not need to be addressed for a particular installation. This list may not cover all relevant issues. It is the responsibility of the client and administration team to determine the necessary configuration requirements. All configuration requests/requirements must be provided to the financial administration team in the form of a PCR.

- RAID levels/requirements
- LVM, raw device and file system
- Printers

- User accounts and groups
- Kernel parameters/drivers
- Application installation(s) and configurations
- Application start/stop configuration(s)
- Backup requirements i.e. frequency, retention etc.
- LVM, raw device and file systems
- Sendmail configuration
- Cron

Installation Complete!

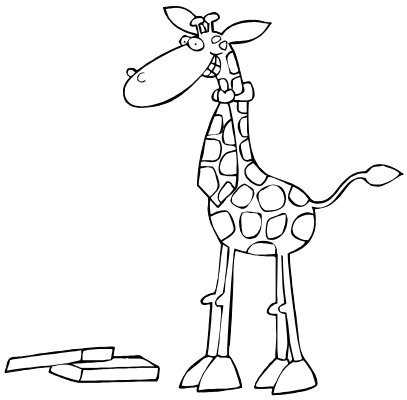

HP-UX 10.20
Administration

Some Philosophy Too!

Vg00 Double Mirror and 4th Boot Disk

Management Overview

Availability, Performance and Recovery

- **Primary Boot Disk Failure**
- **Boot Options**
- **Run State of Mirror**
  - **Protection**
  - **Performance**
- **Maintenance Events**
  - **lv_mergeit, lv_splitit then reboot**
  - **Always have a way back**

The next few topics in this section deal with specific procedures required to manage the double boot mirror and the 4th boot disk. To ensure an understanding of the boot mirrors role in this environment a general discussion on this configuration is in order.

In short this mirror configuration provides each instance of the operating system with increased availability, performance and recovery options.

The primary and secondary boot disks remain in a merged or synced mirror at all times. Should one of these devices fail the system will continue to run from the kernel that is loaded in memory. The failed device can then be replaced at some scheduled downtime in the short term. In addition to providing increased availability and data protection this mirror, as all LVM mirrors, offers increased read performance.

	BUS X	BUS Y
POWER SUPPLY	FAN	FAN
	(6)MERGED	(6)MERGED
POWER SUPPLY	(5)	(5)
	(4)	(4)
	(3)	(3)SPLIT

Jamaica Enclosure, 4 GB Devices.

The tertiary boot disk remains split from the primary/secondary mirror the majority of the time. This disk will only be merged with the mirror for brief periods of time just prior to performing system maintenance or patching. Prior to each maintenance event the tertiary disk is merged and immediately split. Should

maintenance or patching destroy the primary/secondary mirror the system can be booted from the split tertiary disk. The image on this disk can then be used to rebuild the primary and secondary devices if necessary. The fact that the tertiary disk remains split also decreases IO on the secondary/tertiary bus(bus Y in the diagram above).

> Note: When a mirror is split the system will track all changes made to the primary side for the duration of time the mirror remains split or until the system is rebooted. By tracking these changes the time required to merge a mirror is greatly reduced. Since we typically run in a split configuration for 1 to 2 months at a time, we do not want to utilize system resources tracking changes made to the primary copy of vg00. Therefore, after each merge and split of the vg00 mirror the system will be rebooted. This reboot will either be the result of maintenance performed or it will be specifically be performed as a result of the split.

If the tertiary mirror is used to overwrite the primary/secondary devices it will be critical to identify changes that have been made to the operating system i.e. /etc/fstab since the last merge of the tertiary disk. This information can be obtained from the status diary and the output of the sys_config script. Since the operating system and the applications/data on each system are kept separate no changes to applications/data should need to be recovered and the process of recovering recent changes to the operating system will be a fairly simple task

The fourth boot disk will hopefully never be used. In the event that a system is destroyed, upon replacement, this boot disk can be used to recover the operating system in the amount of time it takes to boot.

The tertiary boot disk and the 4th boot disk will be swapped out periodically. For our environment a 4 month cycle is acceptable.

Boot Device Rotation Procedure

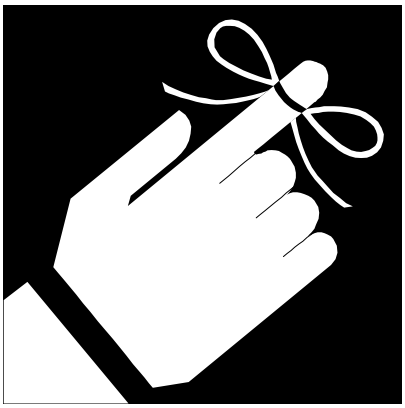

Perform quarterly just prior to a maintenance event. Label the disk, note date, time and the state of the system. Store the disk offsite.

Note: This procedure contains steps that are specific to a cold installed HP-UX 10.20 system.

This procedure will typically be performed on all systems on a quarterly basis.

1-Before beginning this process verify that the logical volumes on the tertiary boot disk are not merged with the mirror and note the device name of the tertiary disk. Issue the following command to merge the mirrors.

cd /usr/local/bin
./lv_mergeit

See appendix B for the contents of lv_mergeit.

2-Execute the sys_config script to create a current online backup of this configuration.

cd /usr/local/bin
./sys_config

See appendix A for the contents of sys_config.

3-Boot the system to single user mode on the primary boot disk.

shutdown –r now

Note: This boot process is specific to a K class system.

Interrupt the autoboot process and continue to boot from the primary device.

Interact with IPL. From the ISL prompt boot to single user mode.

hpux -is

4-Mount all vg00 logical volumes

mount –a

5-Verify the mounts.

bdf

6-Remove the tertiary boot disk (see table below) from the Jamaica enclosure. DO NOT pull the disk all the way out of the enclosure right away. Pull the disk out about 1 inch and allow it to sit for 30 seconds, then remove it from the enclosure and place it in a static bag.

	BUS X	BUS Y
POWER SUPPLY	FAN	FAN
	(6)	(6)
POWER SUPPLY	(5)	(5)
	(4)	(4)
	(3)	(3)TERT. BOOT

Jamaica Enclosure, 4 GB Devices.

7-It will be necessary to force at least one of the vg00 mirrors to become “stale”. Performing IO to any of these logical volumes will accomplish this.

cd /usr
ll

Use vgdisplay to verify the state of at least one of the mirrored logical volumes is “stale”.

vgdisplay –v /dev/vg00

Do not continue this procedure until you have verified that at least one of the vg00 mirrors (lvol1 through lvol7) is in a “stale” state.

8-Unmount all vg00 logical volumes (in single user mode you cannot unmount /).

umount –a

9-Boot to LVM maintenance mode.

reboot

Note: This boot process is specific to a K class system.

Interrupt the autoboot process and continue booting from the primary device.

Interact with IPL. From the ISL prompt boot to LVM maintenance mode.

hpux -lm

10-Insert the 4th boot disk into the Jamaica enclosure and allow it to spin up.

11-vgcfgrestore the device.

vgcfgrestore -n /dev/vg00 /dev/rdsk/c?t?d?

Where c?t?d? is the device name of the tertiary device.

If the vgcfgrestore command fails with a "cannot open device" message do the following.

Ioscan –fn

Now try the vgcfgrestore again.

12-After a successful vgcfgrestore boot the system and allow it to come up to multi user mode on the primary boot device.

reboot

When vg00 is activated the system will automatically sync the mirrors. This sync can be detected by the activity lights on the boot devices and the existence of an "/sbin/vgsync /dev/vg00" process.
If the system does not sync vg00 automatically issue the following vgsync command manually.

vgsync /dev/vg00

13-When the sync is completed issue the following mkboot commands to repopulate the boot header on the tertiary boot device.

mkboot /dev/rdsk/c?t?d?
mkboot -a "hpux -lq" /dev/rdsk/c?t?d?

Where c?t?d? is the device name of the tertiary device.

Propagate the diagnostic utilities to the tertiary boot device.

cd /usr/sbin/diag/lif
mkboot –b updatediaglif –p ISL –p AUTO –p HPUX –p LABEL /dev/rdsk/c?t?d?

Where c?t?d? is the device name of the tertiary device.

14-Test booting from the new tertiary boot disk.

Reboot the system.

shutdown –r now

Note: This boot process is specific to a K class system.

Interrupt the autoboot process.

Use the search command to identify the tertiary boot device and assign a boot value to it.

s*ea IPL*

The sea command will assign p*X* values to each bootable device found.

Issue the boot command for the tertiary device.

bo pX

Interact with IPL and boot to single user mode. From the ISL prompt.

hpux –is

Assuming a successful boot to the tertiary device reboot the system and allow it to come up to multi user mode from the primary device.

reboot

15-During the install process of each system tasks were performed to prepare for the possibility that we would need to boot from a split mirror boot disk. Verify that these files/directories exist and are correct.

The directory /split_root should exist and it should be empty.

The file /etc/fstab.aftersplit should exist and it should contain the following.

```
/dev/vg00/lvol3b   / vxfs delaylog 0 1
/dev/vg00/lvol1b   /stand hfs defaults 0 1
/dev/vg00/lvol4b   /opt vxfs delaylog 0 2
/dev/vg00/lvol5b   /tmp vxfs delaylog 0 2
/dev/vg00/lvol6b   /usr vxfs delaylog 0 2
/dev/vg00/lvol7b   /var vxfs delaylog 0 2
/dev/vg00/lvol8 … swap pri=0 0 0
/dev/vg00/lvol9 … swap pri=0 0 0
```

16-Assuming the /etc/fstab file has changed since the last boot device rotation make a copy of it.

cp /etc/fstab /etc/fstab.b4split

17-Split the mirrors (tertiary from primary/secondary). Execute the following command file.

cd /usr/local/bin/
/lv_splitit

Note: This is the production run state of the boot mirrors (tertiary split from primary/secondary).

Verify the split.

vgdisplay -v /dev/vg00

18-Run fsck on each lvolxb to ensure there has not been any corruption across the split.

cd /usr/local/bin
./fsck_it

19-Verify that the /dev/vg00 directory on the split (tertiary) disk was updated with the correct device files when the lv_splitit script was executed.

Mount lvol3b and take a look at /dev/vg00.

mount /dev/vg00/lvol3b /split_root
cd /split_root/dev/vg00
ll

The output of the ll command should show that each split logical volume is represented by and lvolxb and an rlvolxb device file. For example in addition to the device files lvol1 and rlvol1 we should now have:

lvol1b
rlvol1b

If the correct device files do not exist in /split_root/dev/vg00 directory you will need to create them using the mknod command. If you are not absolutely sure of the appropriate parameters to pass to the mknod command do not continue.

20-Now umount /dev/vg00/lvol3b.

cd
umount /split_root

21-Label the disk; with date, time and state of system; that is now out of the Jamaica and store it offsite.

Booting From A “Split” Tertiary Device

- **Why?**
- **Be Careful!**
- **Preparations Have Been Made**

- **Run Level 3 - The system will run fine in this state but this is an unsupported configuration.**

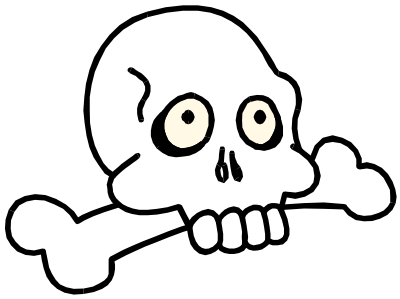

Note: This procedure contains steps that are specific to a cold installed HP-UX 10.20 system.

Like many of the procedures in this document this one is dangerous. If you are booting from a split tertiary boot disk you are most likely trying to repair a damaged system and you are under pressure to fix the problem as fast as possible. Do not let this pressure cloud your judgement. It is imperative that the steps outlined here are executed in the order they are presented. If at any point during this procedure you observe errors or unexpected results from a command STOP, do not continue until you can explain the cause of the problem and understand the steps to correct it. Better to ask what you might think is a dumb question than it is to spend the night at the datacenter rebuilding a system.

1-Reboot the system to LVM maintenance mode on the split tertiary disk.

shutdown -r now

Note: This boot process is specific to a K class system.

2-Interrupt the autoboot process.

From the boot admin menu you may either set the primary boot path to the tertiary disk or you may use sea to boot off the tertiary disk. We will assume your goal is to overwrite the primary mirror and therefore sea is you best option.

3-From the boot admin menu, search for bootable devices.

sea IPL

Identify the correct pX value for the tertiary device and boot from it.

bo pX

Interact with IPL and boot to LVM maintenance mode. From the ISL prompt.

hpux -lm

Note: When booted to LVM maintenance mode any changes you make, such as editing files are not propagated to the boot mirrors even if the mirrors are synced.

4-Once the system is booted to LVM maintenance mode, activate vg00.

vgchange -a y /dev/vg00

5-Mount the split logical volume for usr to /usr. This will give access to the cp command.

mount /dev/vg00/lvol6b /usr

6-Create a backup copy of the current /etc/fstab file and copy /etc/fstab.aftersplit to /etc/fstab. The /etc/fstab.aftersplit file was created at the time the system was installed it should contain "lvolXb" entries for vg00 lvol1b through lvol7b and the secondary swap entries. No data file systems are included in this file.

cp /etc/fstab /etc/fstab.b4split
cp /etc/fstab.aftersplit /etc/fstab

6-Modify lvlnboot to reflect the split logical volumes and correct dump devices.

Note: In LVM maintenance mode the lvrmboot and lvlnboot commands do update the mirror.

Delete the current lvlnboot definitions:

lvrmboot -r /dev/vg00

Update lvlnboot with the split mirror logical volumes for boot, root and swap.

lvlnboot -b /dev/vg00/lvol1b
lvlnboot -r /dev/vg00/lvol3b
lvlnboot -s /dev/vg00/lvol2b

Update lvlnboot with the correct dump device information. To determine what the dump logical volumes are issue the following command.

grep swap /etc/fstab

The output will be the entries for secondary swap/dump logical volumes. For example:

/dev/vg00/lvol8 … swap pri=1 0 0
/dev/vg00/lvol10 … swap pri=1 0 0

Update lvlnboot with the dump information using the correct logical volume names. Remember that these are not split or “b” logical volumes..

lvlnboot -d /dev/vg00/lvolx
lvlnboot -d /dev/vg00/lvolx

Verify the LVM configuration.

lvlnboot -v

7-Unmount /usr

umount /usr

8-Deactivate vg00.

vgchange -a n /dev/vg00

9-Reboot the system, and come up in single user mode on the tertiary drive.

reboot

Note: This boot process is specific to a K class system.

Interrupt the autoboot process. If the primary boot path is set to the tertiary disk simply boot from it.

bo pri

If the primary boot path has not been altered perform an sea.

sea IPL

Identify the correct pX value for the tertiary device and boot from it.

bo pX

Interact with IPL, from the ISL prompt boot to single user mode.

hpux -is

10-Mount all vg00 logical volumes.

mount -a

11-Ensure all logical volumes mounted properly.

bdf

The output should show /, /opt, /var, /tmp, /usr and /stand are mounted on “b” logical volumes.

If / is mounted to /dev/root, remove /etc/mnttab, and execute another *mount -a* to recreate /etc/mnttab.

12-Verify that lvlnboot was updated correctly.

lvlnboot -v

13-Verify that swap is configured correctly.

swapinfo

Note: In single user mode only primary swap is enabled. Therefore you should see one “dev” line and the logical volume associated with it will be lvol2b.

14-At this point you could boot to multi-user mode (*init 3)* and the system should run fine. However you are be booted off the tertiary device on “b” logical volumes and you have no recovery options should the tertiary boot disk fail.

Overwriting Primary Mirror With Tertiary Mirror

- **Assumes all steps in previous section have been completed successfully.**

Note: This procedure contains steps that are specific to a cold installed HP-UX 10.20 system.

Note: This procedure assumes all steps in “Booting From A “SPLIT” Tertiary Boot Device” have been successfully completed and the system is booted from the tertiary device.

If you are booted from the tertiary boot mirror disk and are in single user mode go to step 3. If you are booted to multi user mode on the tertiary disk start with step 1.

1-Reboot the system and boot to single user mode on the tertiary disk.

shutdown -r now *If you are currently in multi-user mode.*

Note: This boot process is specific to a K class system.

2-Interrupt the autoboot process. If the primary boot path is set to the tertiary disk simply boot from it.

bo pri

If the primary boot path has not been altered perform an sea.

sea IPL

Identify the correct pX value for the tertiary device and boot from it.

bo pX

Interact with IPL, from the ISL prompt boot to single user mode.

hpux -is

3-Issue the following lvmerge commands to overwrite vg00 lvol1 through lvol7 with lvol1b through lvol7b. These commands will also delete the logical volumes lvol1 through lvol7.

!!! ARE YOU SURE YOU WANT TO DO THIS !!!

lvmerge /dev/vg00/lvol1 /dev/vg00/lvol1b
lvmerge /dev/vg00/lvol2 /dev/vg00/lvol2b
lvmerge /dev/vg00/lvol3 /dev/vg00/lvol3b
lvmerge /dev/vg00/lvol4 /dev/vg00/lvol4b
lvmerge /dev/vg00/lvol5 /dev/vg00/lvol5b
lvmerge /dev/vg00/lvol6 /dev/vg00/lvol6b
lvmerge /dev/vg00/lvol7 /dev/vg00/lvol7b

After the merges are complete a *vgdisplay -v /dev/vg00* will show that vg00 lvol1 through lvol7 no longer exist and that lvol1b through lvol7b are synced and each logical volume is using 3 physical devices.

4-At this point, you could boot to multi-user mode (*init 3)* and the system should behave normally. You are booted off of the tertiary device using “b” logical volumes and the triple mirror is active. This is however not a configuration we support, this option should only be used in emergency situations and only for a brief period of time.

If you will be using the system in this configuration you will need to edit the /etc/fstab (add data file system entries and verify required vg00 entries are “b” logical volumes, simply copying the /etc/fstab.beforesplit file will not work) file and boot to multi user mode (init 3).

If the goal is to return to the standard configuration continue this procedure.

5-Unmount any file systems that may be mounted.

cd
umount -a

6-Perform the following mv's to replace all lvol*x*b and rlvol*x*b with lvol*x* and rlvol*x*.

Note: Since /dev/vg00/lvol3b is currently mounted as / the mv will switch the mount to /dev/root.
Note: The following mv's will change the minor numbers of each logical volume, this is not a problem. If the situation warranted mknod could be used to create the non "b" logical volume entries and the "b" logical volumes could be deleted.

mv /dev/vg00/lvol1b /dev/vg00/lvol1
mv /dev/vg00/lvol2b /dev/vg00/lvol2
….…
mv /dev/vg00/lvol7b /dev/vg00/lvol7

mv /dev/vg00/rlvol1b /dev/vg00/rlvol1
mv /dev/vg00/rlvol2b /dev/vg00/rlvol2
….…
mv /dev/vg00/rlvol7b /dev/vg00/rlvol7

7-Verify these moves with a vgdisplay.

v*gdisplay -v /dev/vg00*

The output should display vg00 lvol1 through lvol7 as synced using 3 physical disks each. The "b" logical volumes should be gone.

8-Mount /usr

mount /dev/vg00/lvol6 /usr

9-Copy the original /etc/fstab into place.

c*p /etc/fstab.b4split /etc/fstab*

10-Mount all vg00 logical volumes and verify that vg00 lvol1 through lvol7 are used.

mount -a
bdf

11-Reboot the system to LVM maintenance mode on the tertiary disk.

reboot

Note: This boot process is specific to a K class system.

Interrupt the autoboot process. If the primary boot path is set to the tertiary disk simply boot from it.

bo pri

If the primary boot path has not been altered perform an sea.

sea IPL

Identify the correct pX value for the tertiary device and boot from it.

bo pX

Interact with IPL, from the ISL prompt boot to LVM maintenance mode.

hpux -lm

12-Activate vg00

vgchange -a y /dev/vg00

13-Modify lvlnboot to reflect the original logical volumes and correct dump devices.

Delete the current lvlnboot definitions:

lvrmboot -r /dev/vg00

Update lvlnboot with the primary logical volumes for boot, root and swap.

lvlnboot -b /dev/vg00/lvol1
lvlnboot -r /dev/vg00/lvol3
lvlnboot -s /dev/vg00/lvol2

You will also need to update lvlnboot with the correct dump device information. To determine what the dump logical volumes are issue the following command.

grep swap /etc/fstab

The output will be the entries for secondary swap/dump logical volumes. For example:

/dev/vg00/lvol8 … swap pri=1 0 0
/dev/vg00/lvol10 … swap pri=1 0 0

Update lvlnboot with the dump information using the correct logical volume numbers for the system you are on.

lvlnboot -d /dev/vg00/lvolx
lvlnboot -d /dev/vg00/lvolx

14-Verify the LVM configuration

lvlnboot -v

15-Deactivate vg00

vgchange -a n /dev/vg00

16-Reboot the system

reboot

Note: This boot process is specific to a K class system.

Interrupt the autoboot process.

If the primary boot path was modified earlier in this procedure change it back now. If the primary boot path is set to the primary disk simply boot from it.

bo pri

Interact with IPL, from the ISL prompt boot to single user mode.

hpux -is

17-0Mount all vg00 file systems.

mount -a

18-Ensure all logical volumes mounted properly.

bdf

The output should show that /, /usr, /opt, /var, /tmp and /stand are NOT mounted on the “b” logical volumes.

If / is mounted to /dev/root, remove /etc/mnttab, and execute another *mount -a* to recreate /etc/mnttab.

19-Verify that lvlnboot was updated correctly.

lvlnboot -v

20-Verify that swap is configured correctly.

swapinfo

Note: In single user mode only primary swap is enabled. Therefore you should see one “dev” line and the logical volume associated with it will be lvol2.

21-If the system looks good, boot to multi user mode.

init 3

Obtaining Support

Some support personnel are better than others, the same is absolutely true about customers.

- **WORKING Support Calls**
- **Providing Information (email/ftp)**
- **Executing Provided Scripts**
- **Record All Information**
- **Keep an Open Mind**
- **Persistence**
- **Escalations**
- **Response Time (call back/on site)**

Often when dealing with support organizations you are trying to resolve a problem that management and the user community are painfully aware of. This is not a fun position to be in and if you are staring at a down system, it's even worse. Do not allow the stress of the situation distract you from the problem at hand.

The fact is, some support personnel are better than others, the same is absolutely true about customers.

When you are WORKING a call with a support organization, and I do mean WORKING, keep the following in mind.

- Save yourself a little time when logging the call. Typically the person who answers the phone does not need detail they're not going to understand anyway. They do need a topic (command, subsystem, application, hardware) to log the call under and connect you with the correct support personnel.
- When each call is logged be very specific about the severity or urgency of the problem.
- Record all call ID's and document call specifics.
- You will take responsibility for each call. You are required to drive each call to a solution. If you are not getting what you need in a timely manner call again, escalate the call, contact your lead, contact the response center advocate and or contact local HP resource(s).
- If you do not provide accurate information to the response center engineer you can bet that the solution you are provided with will be incorrect. If you're not sure the picture of a problem you have is complete include a team member or your lead; they will be more than happy to assist you.

- If the response center engineer requests you to email/ftp information regarding a problem caution should be exercised. For example:
  -Only information that is relevant to the problem will be provided. We have received requests in the past to run collect scripts that gather system wide information and email the output; this is not acceptable.
  -We do not email entire log files such as syslog.log. This file contains sensitive information regarding the entire environment. If a support engineer needs information from syslog.log email only the specific lines pertaining to the problem.
  -If you are unsure about what is appropriate to email/ftp and what is not ask.
- If the response center engineer asks you to execute a script on a system YOU are responsible for the result of the script. Each script should be investigated prior to execution. Obviously many support issues will be time sensitive: if you need assistance with an investigation ask.
- If a response center engineer requests dial in access to our servers the answer is NO. Company policy is that no dial in lines exist in the datacenter. THIS IS NOT UP FOR DISCUSSION. Any information the engineer needs regarding a call will be provided by you. If you are not comfortable gathering the information requested ask for assistance from a team member.
- Provided requests for information are applicable to the call and do not represent a security violation, information will be provided in a timely manner.
- Stay focused on the problem at hand. It is extremely easy to be distracted by unrelated issues, if you allow yourself to be distracted then the response center engineer is distracted as well.
- Keep an open mind. If from the beginning of a problem you are convinced the issue is patching, for example, and you lead the HP engineer down that path you may end up spending your weekend working because you didn't consider other possibilities and you convinced the HP engineer to do the same.
- Remember when a vendor provides you with a procedure you are responsible for the result(s). If you are not comfortable with the solution provided, be persistent. If you need assistance from a team member do not hesitate to contact them.
- Displaying frustration and anger will most likely be counter productive, however there is nothing wrong with being persistent.
- If you reach a point during any call where you do not know how to move forward, you are not comfortable with a provided solution, you are not getting what you need from the support engineer, etc. page your lead immediately.
- Communicate all information regarding each call to team lead.

HP Support Contracts

- **Accuracy of Contracts**
- **2 User vs Unlimited User**
- **Aging Hardware Maintenance Costs**
  - **New Purchases**
  - **Move to Dev/Test Systems**
- **Reduced Support Costs By $$$ Despite Growth**

HP-UX Patching Philosophy

- **To Patch or Not to Patch?**
- **Average Production Availability From 09/97 to 04/99 = 99.88%**

To patch, or not to patch? For the majority of the IT organizations I have been involved with the answer to this question is “not to patch”. The idea that changing a configuration i.e. updating the operating system can result in problems is the driving reason for this decision. I couldn’t disagree more, patching is a necessity. If you don’t patch as a planned maintenance event and you think you’re reducing the risk of downtime you’re wrong. A system will go down as a result of a problem a patch would have corrected or you will patch as a result of the daily troubleshooting process. Administrators are painfully aware of the fact that when you open a call with the response center 9 times out of 10 their first question is have you applied the following patches. So instead of patching on a per problem basis why not apply patches as a planned downtime event. At least then if the patching does cause a problem you will have a downtime window to address the issue. I also enjoy answering “yes” when HP asks, “ Have you applied the following patches?”

We have applied a new patch depot to our systems every month since 09/97 and despite all of this change and potential for problems we have achieved an average 99.88% availability on production systems (average production availability calculated from 09/97 to 04/99)! Why? There are a number of reasons for this:

- We have well defined patch management procedures.
- All of our systems are standardized. Applying a patch depot to 20 systems that are all configured differently is a hell of a lot more risky than applying a depot to 20 systems that are essentially identical.
- We always have a way back. In the event we do have a problem that destroys all or part of the operating system we have an image of the operating system that was created just prior to patching, the split tertiary boot disk. It’s is worth noting that to date it has not been necessary to utilize this recovery option as a result of a patching event
- When creating patch depots we do not include patches that are less than 20 days old.
- The procedures we use to perform patch analysis/depot creation are very thorough. As much as possible is done to eliminate problems prior to ever applying a specific depot.
- All patch depots are applied to all development and test machines first, then at the next production downtime (typically 1-month later) they are applied to all production machines.

Patch Management

- **Standardized Procedures**
- **Standardized Systems**
- **Always Have A Way Back**

- **Patch Analysis/Depot Creation**
  Analyze against 2 systems
  No patches < 20 days old
  Dev/Test then production
  Historical depots

The TCI HP administration team performs a patch analysis on a monthly basis. This analysis uses the cpm_collect.sh script (/usr/local/bin/cpm_collect.sh) and the HP web site http://us-support.external.hp.com. The collect script is typically run against 1 production machine and 1 development machine. The machines chosen are those that best represent the entire environment. The resulting depots are the merged into one and all patches that are less than 20 days old are eliminated.

The depots created form these analysis are kept in /var/spool/sw and they follow the naming convention 10.20_800_patches_15, where the number 15 changes accordingly. The historical depots are maintained for a period of time. Periodically (4 to 6 months) we will perform an analysis on a "clean" system and use this depot as a starting point for new installs. All depots are first applied to the development systems and then at the next scheduled downtime they are applied to the production systems. Therefore our production environment is always 1 to 2 months behind our development environment.

When applying patch depots created by the analysis process described above you can assume that all patch dependencies, special instructions etc. have been investigated and resolved. If however you are installing patches that are not a part of the depots we create on a monthly basis it is your responsibility to resolve dependency issues and to address any special instructions that may be included in the patches you are installing.

Patching Procedures

- **Merge and Split Mirrors**
- **cleanup**
- **fsck_it**
- **mwa stop, ttd -k**
- **swinstall "Match What Target Has"**
- **rc.log, syslog.log**
- **swlist -l fileset -a state | grep install**
- **swconfig ***
- **monitor_setuid create_standard**

1-Just prior to patching or performing any downtime maintenance merge and split the tertiary mirror.

cd /usr/local/bin
./lv_mergeit

After the merge has completed verify it with vgdisplay –v then split the mirrors.

cd /usr/local/bin
./lv_splitit

After the split has completed verify it with vgdisplay –v.

3-Execute the cleanup utility.

cleanup

You will be asked "Would you like the logs trimmed to the last 5 entries?" enter y and hit return.

4-Verify the consistency of the split logical volumes.

cd /usr/local/bin
./fsck_it

5-Stop the MeasureWare daemons and ttd.

mwa stop
ttd -k

6-Use swinstall to apply the current patch depot. The “Match What Target Has” feature in swinstall MUST be used for all patch applications.

swinstall

7-After the system has rebooted (if necessary) view the following log files and use swlist to ensure all filesets were configured.

View the following log files and search for any problems.

/var/adm/syslog/syslog.log
/var/adm/shutdownlog
/etc/rc.log
execute dmesg

Use swlist to verify fileset configuration.

swlist –l fileset –a state | grep install

If any filesets are in an “installed” state use swconfig to manually configure them.

*swconfig **

View the log file /var/adm/sw/swagent.log for any swconfig errors/problems that may have occurred. Execute the swlist command to re-verify fileset configuration. If any configurations are unsuccessful notify a team member.

8-Recreate the standard file for setuid, setgid and sticky bit monitoring. Applying patches to this system will potentially modify, add or delete existing setuid, setgid and sticky bit files and directories therefore now is a good time to recreate the standards we monitor against for the existence of these files.

cd /usr/local/bin
./monitor_setuid create_standard

See appendix D for the contents of monitor_setuid

System Standardization

- **Each System is Configured Identically: Specific to Hardware and Operating System.**

- **Documentation is Mandatory**
- **Automated Monitoring of Standards is Mandatory**
- **Education is Mandatory**
- **Change Management Scheme is Mandatory**

Benefits Of Standardization

- **Simplifies and Facilitates:**
  Troubleshooting
  Documentation
  Change Management
  Operating System Upgrades
- **Decreases Possibility of Mistakes**
- **Lays the Foundation for a Successful Security Implementation**
- **Reduces Risks Associated with Patching and Maintenance Events**

An environment is standardized when each system in the environment is configured identically per the standards defined. Documentation, automated monitoring, education of the administrators and a change management scheme designed to maintain current standards and implement new standards are required to achieve and maintain standardization. For our purposes this standardization applies strictly to the operating system and hardware.

New standards are subject to approval. Prior to the introduction of new standards they will be thoroughly tested and documented.

This system standardization is the groundwork that provides for the success of all other aspects of this environment. I cannot stress enough the importance of designing, documenting and maintaining operating system and hardware standards.

Benefits of a standardized configuration include:

- Simplification and facilitation of troubleshooting
- Simplification and facilitation of documentation
- Simplification and facilitation of change management
- Decreased possibility of mistakes
- Lays the ground work for a successful security plan
- Reduced risk, associated with patching
- Facilitation of OS upgrades

Change Management

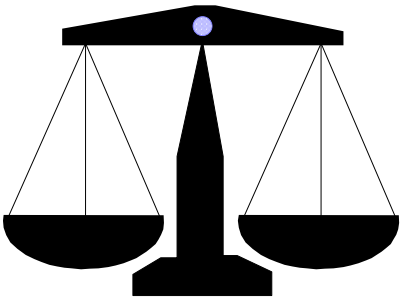

- **Processes created to control types, amounts and rates of change.**
- **Standardization <> Change Management**
- **Concept suggests you have something in a known and consistent state and you want to maintain it. This known state equates to standardization.**

A **Change Management** scheme is made up of the documentation and processes that are created to control the amount and rate of change allowed in computing resources, including but not limited to the hardware, applications, databases, operating systems, accounts, security, etc. Ultimately change management should minimize the negative impacts change can bring by maintaining the standardization of the resources it is applied to.

Change management facilitates system standardization. The opposite is also true. Standardization facilitates change management.

The concept of change management suggests that you have something in a known and consistent state and you want to maintain it, if changes are introduced you want to document them and to control how they are deployed. This "known state" equates to standardization.

A Scenario

- **Change Management is not worth much if computing resources it is applied to are not standardized.**
- **Change management is often implemented for the wrong reasons and at the wrong time.**
- **Authority over change, where should it reside?**
- **Vendors often blamed for problems**
- **Low expectations of computing resources as a whole.**

So all computing environments that have change management applied to them are standardized and documented, right? Unfortunately the answer to this question is most often no. Let's create a scenario and investigate this a little further.

> Let's say there is a computing environment consisting of 20 HP-UX servers and various Oracle database instances. Several different administrators installed the operating systems. Those administrators were provided with no standards to work from. They simply configured the systems to the best of their ability per their varying experience. The various instances of Oracle running in the environment were installed in much the same way. Neither has been patched according to any plan and versions therefore vary. As these system and database administrators support this environment they end up fixing problems on various servers but no attempt is made to consistently implement their fixes across the environment. No discrete test and development systems exist, often development occurs on production systems. Occasionally a developer's code takes a production systems memory utilization to 100% and the system spends some time paging.
>
> This environment has seen its share of problems. System and database downtime resulting from maintenance events is commonplace and response time is typically less than desirable. Yesterday the SA's had a bad day. They were in a little trouble because they tested a change on a system (that was a test system, right?) and it worked fine but when they moved it around 2 out of 5 servers panicked. It's too bad that neither of them had any dump space configured.
>
> About this time someone who is taking a lot of heat for this situation has decided to eliminate these problems by controlling change. They are tired of those incompetent administrators screwing things

up! They might even go as far as creating a whole new department to manage change. Better yet, let's put the users in charge of it. Oh, and by the way if HP doesn't improve their performance and Oracle can't stabilize their products they're going to replace everything with Sun and Sybase.

This real world scenario demonstrates many points regarding change management.

- Change management is often implemented for the wrong reason and at the wrong time. Change management should be implemented because you never want to be in the above situation not as a result of it.
- If you don't have a handle on the current state of each system/database because you haven't standardized you're too late. Implementing change management in an environment that is in this state of disarray will not solve your problems; in fact it will most likely compound them.
- Administrators and vendors are often blamed for problems of this nature. The fact is that computing environments must be managed. A lack of direction will typically result in problems similar to those demonstrated above.

The sad fact is that many people involved in IT and many users have extremely low expectations regarding computing resources. Environments in a state similar to the above scenario are all they have ever experienced, therefore, they have come to expect problems similar to those demonstrated.

History on Environment

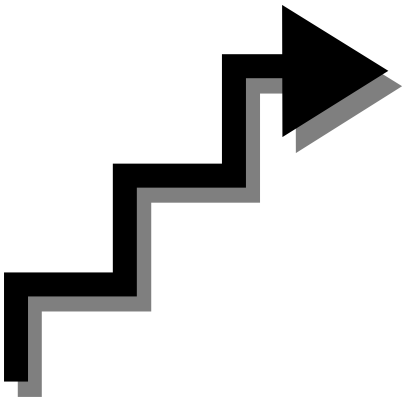

- **09/97**
  Not Standardized, No Change Management Scheme
- **11/98**
  Standards Defined and Successfully implemented
- **1999**

- **What was involved?**

When we began the process of evaluating the TCI Finance environment in 1997, no complete change management scheme existed and the systems were certainly not standardized. The environment was well established and therefore we were facing issues similar to those described in the scenario above. After creating an initial set of standards (including double boot mirror, vg00 configuration, patch management standards, etc.) we began deploying what would evolve into the standards we enjoy today. The implementation of these standards and the change management scheme to maintain them required that each existing system in the environment be completely rebuilt (cold install of HP-UX 10.20 and vgexport and vgimport of data filesystems). This process required 1 year of monthly scheduled maintenance events to complete. Today, with the exception of kernels, files system layouts, etc. every system in this environment is essentially identical.

Change Management Scheme

Specific to Hardware and Operating System
Whole Document Defines Overall Scheme

- **Change Authority Controlled by Lead**
- **Standards Documented**
- **Discrete Development, Test and Productions Systems**
- **Education**
- **vg00 vs Applications**
- **New Purchases (Hardware and Software)**

The following sums up the change management scheme for this environment. This scheme is specific to the hardware and operating systems under our control. This list may appear incomplete; many of the change management rules we use are defined elsewhere in this document. For example the patch management, security and account management all have specific change management standards associated with them. These standards are defined in the appropriate documentation.

- This document and the on-line configuration backups of each system serve as the standards that all change management is based on.
- Each new standard that is defined will be incorporated into this document by the system administration team.
- All aspects of change regarding hardware resources are controlled and subject to approval by the lead.
- All aspects of change regarding operating systems are controlled and subject to approval by the lead.
- All development, test and production efforts will occur on systems dedicated to development, test and production (applies to administrators and users alike).
- No user data (data, configuration files or executables owned by the user community) will reside in the root or vg00 file systems of any server.
- No system will be installed using any process other than those defined in this document.
- All application installations, development efforts and testing must fall within the constraints of our change management scheme. For example: If an application is to be installed but the users failed to reserve or purchase the disk resources, the application will not be installed in the vg00 file systems. Not even if it is only 1k in size.
- No new server will be purchased without the necessary components required to support standardization. For example a Jamaica enclosure and Mirror/UX are a mandatory part of any server purchase.
- No disk resources will be purchased without the necessary components to support alternate links.

Change management plans often overlook a major feature that will make the difference in success or failure. Once you have implemented a change management scheme and the standardization that it consists of how do you enable administrators to successfully manage it?

We utilize documentation, repetition, education and automated monitoring tools to ensure that each administrator is aware of our standards and maintain those standards.

- By documenting the installation process we ensure that each system is built to our standards.
- The standardized installation process educates each administrator about the various aspects of the installation and therefore the overall configuration. Administrators who understand the configuration are equipped to maintain it.
- By documenting patch management procedures and applying all patch depots in specific sequence we ensure system configurations are consistent across the environment.
- By documenting the administration of the environment we ensure the administration team as a whole is aware of each standard and is therefore equipped to maintain them.
- By implementing and documenting a change management scheme we ensure that as new configuration standards are identified they are documented and implemented across the environment.
- Automated configuration monitoring tools have been created. These tools notify administration team of inconsistencies on the servers via email.

HP's Openview ITA was purchased in January of 99'. Ultimately ITA will be utilized to provide a single point of administration for change management.

Why It Works

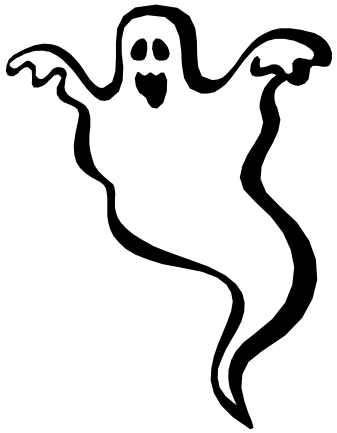

- **Documented**
- **Educated**
- **Automated**
- **Standardized**
- **Authority**

System Monitoring

- **Timely Notification**
- **What to Monitor**
- **Single Point of Administration**

The tools we currently use to monitor the environment have served us well, however they lack the timeliness and flexibility this environment demands. The main monitoring limitations we face are:

- Monitoring scheme does not provide for timely notifications.
- Monitoring scheme does not offer a single point of administration.

Openview ITO was purchased in January of 99'. It will be used to create a monitoring scheme that is appropriate for this environment. Until the ITO configuration is complete we will continue to rely on the following monitoring configuration.

1-Each morning Monday through Friday the on-call SA's first priority is to perform a thorough investigation of each system. Upon completion of this sweep a status email is sent to various clients and IT personnel.

The following logfiles, tools and commands may be viewed or utilized as a part of the system investigation.

/var/adm/syslog.log, /var/adm/shutdownlog, /var/adm/messages, /var/adm/cron/log, /var/adm/btmp (lastb), /var/tombstones, STM, /var/adm/sulog, dmesg, netstat, ioscan, psconfig, Netbackup job monitor/log files and root's email account on the administration server.

In addition to the above, the output of monitoring scripts we have created may also be viewed. This includes: monitor_setuid, monitor_config, monitor_passwd, finger_it, finger_it_all, mirror_stale and sys_config.

2-The operations staff has been provided with a system checklist. This checklist is used to monitor the physical environment of the servers. This check is performed Monday through Friday at 6:00 A.M. and 6:00 P.M..

If a problem is identified the operators will page the on-call pager.

3-Our current on-call paging system is simplistic but reliable.

The paging script /usr/local/bin/ping_script root:sys 700 executes via cron on multiple servers. This script pings all the servers in the environment every 5-min. If a system is not responding the on-call pager is notified.

The paging script /opt/oracle/notifier/ora_config.dat executes on multiple servers. This program monitors system log files for specific events such as bus resets and LVM powerfail messages. If errors are discovered the on-call pager is notified.

The script /usr/local/bin/mirror_stale root:sys 700 runs from cron 6 times a day. This script monitors all LVM mirrors, if a mirror is in a “stale” state an email is sent to the on-call pager and the admin staff.

The arraymond daemon monitors all RAID devices on each system. This daemon has been configured to email the on-call pager if any failures are detected.

4-Predictive/UX is currently configured on multiple systems. Predictive/UX runs each night; performing a complete system check. If any problems are identified predictive calls the HP response center and logs a support event, the response center then pages the on-call pager. Predictive is also configured to email the admin staff regarding logged events.

5-The script /usr/local/bin/monitor_setuid is used to track the existence of setuid, setgid and sticky bit files and directories. This script compares the current setuid, setgid and sticky files and directories to a standard list of these files that is generated after each patching event. This script executes via cron each night. It notifies the administration team of changes, additions or deletions of setuid, setgid and sticky bit files and directories via email. To view the contents of this script see Appendix D – Monitor_setuid.

6-The finger_it script monitors all systems for root logins. The output of this script is used to monitor where root logins originate.

7-The script /usr/local/bin/monitor_config is used to monitor aspects of the configuration such as lvlnboot, swap, cron, setboot, vg00 and the contents of boot LIF headers. This script compares the current configurations to standards for each configuration created after the last known configuration change(s). This script executes via cron each night. It notifies the administration team of configuration changes via email. To view the contents of this script see Appendix E.

8-The script /usr/local/bin/monitor_passwd is used to monitor each /etc/passwd file and the NIS password map for the existence of null password values. This script executes via cron each night. It notifies the administration team of configuration changes via email. To view the contents of this script see Appendix F.

Various Administration Procedures

- **Deleting Unused Disk Device Files**
- **Modifying /etc/passwd and /etc/group**
- **Unmounting Busy File Systems**
- **Lsof**
- **Using Kill**
- **File System Full**
- **Vxfs Largefiles**
- **Extending and Reducing File Systems OnLine**
- **Copying File Systems and Directory Structures**
- **Cleanup**

The procedures contained here are by no means a complete list of the tasks administrators are required to perform; they are simply the procedures I have had time to document and hope are useful. These procedures are presented in no particular order.

As always if any of these procedures are not clear consult the appropriate man pages or ask a team member before continuing.

Deleting Unused Disk Device Files

In the event that disk resources are moved to different buses on a system or removed all together the original device files will be hanging around with no devices to reference and the device definitions may still be listed in the ioconfig map of the system. The existence of these device files in /dev and the existence of devices in a NO_HW (no hardware) state in the output of ioscan can be confusing and downright messy. The following procedure should be used to clean up disk device files i.e. those located in /dev/dsk and /dev/rdsk (possibly /dev/floppy and /dev/rfloppy) as well as ioconfig map entries.

To identify unused disk device files perform the following.

cd /dev/dsk
*lssf * | grep ???*

The output of the lssf command will be device files that are no longer associated with a hardware path. hence the ??? where the hardware path should be displayed. For example.

Disc3 card instance 6 SCSI target 1 SCSI LUN 3 section 0 at address ??? c6t3d0

If you issue the following rmsf command for a given device the –a should tell rmsf to remove all associated device files for the device.

rmsf –a /dev/dsk/disk_device_file_name

The rmsf man page is confusing. As stated, the –a should remove all associated device files for a given device, it does not. If you cd into /dev/rdsk and look for the device file that you just deleted you will find that it still exists (it may also still exist in /dev/floppy and /dev/rfloppy depending on the original driver the device was associated with). The above command will simply remove the device file in /dev/dsk and the definition for this /dev/dsk/ device file from the ioconfig map.

If you perform an ioscan –fn you may see that the device definition still exists in a NO_HW state and that it is associated with a device file in /dev/rdsk/ (depending on system reboots this will not always be the case).

To ensure that all associated device files are removed and that the ioconfig map definition is deleted perform a rmsf for each device file that exists.

rmsf –a /dev/dsk/device_file_name
rmsf –a /dev/rdsk/device_file_name
rmsf –a /dev/floppy/device_file_name (if necessary)
rmsf –a /dev/rfloppy/device_file_name (if necessary)

Modifying /etc/passwd and /etc/group

-Modifying /etc/passwd:

Anytime the /etc/passwd file is edited the vipw command will be used to perform the edit. The vipw command is beneficial for the following reasons.

1-The vipw command places a lock on the /etc/passwd file, therefore no one else will be allowed to edit it the same time you are.
2-The vipw command performs a consistency check on the root entry; it will not put (overwrite) the current passwd file if the root entry is not formatted correctly.

After editing the /etc/passwd file use pwck to verify it’s consistency.

pwck

-Modifying /etc/group:

After editing the /etc/group file grpck should be used to verify it’s consistency.

grpck

Modifying /etc/hosts

We have tcp_wrappers configured on all systems in this environment. The result of this is that all IP's must resolve to *hostname*.tci.com (whether files or DNS is used) or connections from the device may be refused depending on the service requested. Therefore all entries in the /etc/hosts file must follow the format:

xxx.xxx.xxx.xxx *hostname*.tci.com *hostname*

In an attempt to limit the number of DNS requests from the servers each server on the finance network will be listed in the /etc/hosts file and the hosts entry in the nsswitch.conf file will be set to files, dns, nis. The default host file is maintained on the administration system and put in place on each system at installation. If a new system is added to the finance network it will be necessary to modify the default hosts file. It would be ideal to extend this configuration to all devices that access this environment, we are however not provided with the information necessary.

Unmounting Busy File Systems

We have all issued a umount /some_filesystem and seen.

umount: cannot unmount /some_filesystem : Device busy

Use the following procedure to identify the processes that are referencing the file system in question.

Use bdf to identify the name of the logical volume the filesystem is on.

bdf | grep file_system_name

Once you have determined the name of the logical volume run the fuser command on it to determine what processes are referencing the file system in question.

fuser –u /dev/vgxx/lvolx

The output of the fuser command will look something like.

/dev/vgxx/lvolx: 6595o(root) 2348co(userid)

This output identifies the process ID and the owner of the process. Use the ps command to identify additional information regarding each process and/or userid. Once you have identified the processes that are referencing the file system you may either ask the user to terminate the process or you may use the kill command (see Using Kill in this section). It is up to you to determine if the processes in question can be killed. If you are unsure consult with a team member.

Once all the referencing processes have been terminated the file system may be umounted.

LSOF (list open files)

The lsof command is configured on each system at installation. Lsof is worth mentioning here because it is such a powerful and flexible tool. The lsof man page offers a very detailed description of the command and it's myriad of options so I will not attempt to re-invent the wheel here. I simply want each administrator to be aware of this command and to suggest a thorough review of its man page.

Using Kill

Ever issued a "kill –l" (that's an l as in lard). If you issued this command you see that there are 33 kill values besides the "-9" you see so often. The "–9" or SIGKILL option will almost always kill the process it is sent to. The problem is that SIGKILL does not allow processes to clean up after themselves and in some cases can result in the creation of a zombie process(es). When you use kill to eliminate processes the "kill –9" should be the last option you use, try the following kill options first.

kill –2 *process_id*(SIGINT)
kill –15 *process_id* (SIGTERM)
kill –3 *process_id*(SIGQUIT)

If the process(es) cannot be successfully killed using on of the above options; the –9 should be used.

Another interesting option to kill is "-6" or SIGABRT. This kill option will cause the process to core dump. In the event that you have a process(es) that frequently get out of hand this option can be used to create a core file that can then be analyzed and hopefully solve the problem once and for all.

For a listing of the different kill options and a brief description of each issue the following command.

grep "define _SIG" /usr/include/sys/signal.h

A man on kill and signal may also be used to investigate these options further.

File System Full

This can be a time consuming and frustrating problem. Unfortunately there is not one way of approaching it that will always lead to a quick resolution. Each of these approaches to the file system full problem has it's own weaknesses or flaws.

- I usually start by executing a du –s –k * in the mount point of the file system that is full. This will show the size of all files and directories in kilobytes. It's usually necessary to issue this command in several directory levels before identifying the culprit. The problem with this solution is unless you are very familiar with the typical sizes of files and directories in a given mount point its hard to identify what is larger than normal.

- Execute "find . –size +*xxxx*c" (where *xxxx* is a size in bytes) in the mount point of the file system that is full. This will find all files greater in size than the value you entered for *xxxx*. Again if you are not familiar with the contents of this file system this may not lead you to any conclusions.

- Execute "find . –mtime –1" in the mount point of the file system that is full. This will display all files that have been modified in the last day. Depending on the contents and purpose of the file system in question this may be a complete waste of time, then again it may help.

- Begin by freeing up enough space in the file system in question to allow you to touch a new file, "touch /filesystem_mount_point/touchfile". Execute "find filesystem_mount_point –newer /filesystem_mount_point/touchfile". This will display all files that are newer (or been modified) since touchfile was created. This is great except:
  -The process(es) that filled the file system in the first place may no longer be running and therefore the files they were responsible for are no longer being modified.
  -The file system is full and therefore very few if any files can be modified.
  -During the time period the find command was executing no modifications were made to the files that you are looking for.

- Lsof may also be used against the mount point or specific files within it to investigate who or what processes have files open in the file system in question. The problem here is that depending on the file system the lsof output can be very large.

If none of the above solutions work for you try creating a cron job that collects periodic du's or find's then use diff to compare the output of these commands based on different times. Hopefully, with a better picture of the contents and behavior of the file system in question you will be able to identify the problem.

Largefiles Option for vxfs File Systems

The default file size limit on vxfs file systems at HP-UX 10.20 is 2 GB. This default file size limit can be increased to 128 GB on a per file system basis (must be version 3 vxfs file system, man fsadm_vxfs).

Turn largefiles compatibility bit on.

fsadm –o largefiles /mount/point

Turn the largefiles compatibility bit off.

fsadm –o nolargefiles /mount/point

Extending and Reducing vxfs File Systems On-Line

Note: These procedures require the HP On-Line JFS product.

The fsadm command can be utilized to extend or reduce vxfs file systems on-line.

All references to the fsadm –b option in the man page (man fsadm_vxfs) are in sectors. The default vxfs sector is 1 kilobyte. Therefore values passed to fsadm –b can be expressed in kilobytes (unless the default sector/block size has been modified with mkfs).

All values in Kilobytes passed to fsadm –b should be a multiple of the LVM extent size, 4 MB.

Reducing vxfs file systems on-line:

The fsadm command will not reduce a file system if the sectors it is attempting to reduce are utilized. The lvreduce command will not reduce a logical volume if the extents it is attempting to reduce contain a file system. Therefore it is difficult to adversely affect the integrity of a file system when reducing it.

Reduce the file system to the desired size. (If the fsadm command fails, complaining that sectors are in use the fsadm command may be used with the –d and –e options to reorganize (defragment) the file system, then attempt the reduction again. This reorganization may clear the used sectors and allow the reduction.)

fsadm –b size_in_KB /mount/point

Verify the reduction with a bdf.

bdf

Reduce the logical volume. (This step is not mandatory, if however it is not performed the portion of the logical volume that does not contain a file system will be unavailable for other uses.)

lvreduce –L size_in_MB /dev/vgxx/lvolx

Verify with an lvdisplay.

lvdisplay –v /dev/vgxx/lvolx

Extending vxfs file systems on-line:

Extend the logical volume of the mounted file system to the desired size.

lvextend –L size_in_MB /dev/vgxx/lvolx

Issue the fsadm command to extend the file system.

fsadm –b size_in_KB /mount/point

Verify the extend with a bdf.

bdf

Copying File Systems and Directory Structures

For some reason we always seem to be copying some directory structure or moving file systems to different devices. Much of the time cp –r –p will suffice but links and hidden files can create problems. To avoid these problems we created the tar_it script. Source and destination directories must be passed to the script upon execution. This script is located on each system in /usr/local/bin. See Appendix G for the contents of tar_it.

Using the Cleanup Utility

Information about each patch that is installed on a system is maintained in /var/adm/sw/patch. Included in this directory are patch descriptions, the file lists the patch effects and the remove scripts for each patch. Over time the contents of this directory can become very large, ultimately filling the /var file system. The cleanup utility will be used to clean up patch history that is no longer needed and ensure that sufficient free space is available in /var.

Caution should be exercised when using the cleanup utility. The possibility that patches may need to be removed from a system always exists. If the historical patch information in /var/adm/sw/patch is removed for a given patch, the patch can no longer be removed from the system without first re-installing it. Therefore the cleanup utility should not be run immediately after a patching event. The appropriate time to "clean up" patch history is just prior to the application of a new depot.

The amount of patch history the cleanup utility removes is controlled by the options passed to it. With no arguments this command will remove all patch history up to 1 level back and it will trim the SD logs to the last 5 entries. Typically this is acceptable for our environment.

Periodically execute the cleanup utility on all systems just prior to the application of patch depots.

cleanup

You will be asked "Would you like the logs trimmed to the last 5 entries?" enter y and hit return.

After the cleanup has completed recreate the standard files for setuid, setgid and sticky bit monitoring. Cleaning up the patch history on the system will definitely delete existing setuid, setgid and sticky bit files

and directories therefore now is a good time to recreate the standards we monitor against for the existence of these files.

cd /usr/local/bin
./monitor_setuid create_standard

Environment Status Diaries

- **What a Mess**
- **History**
- **Track Downtime**
- **Drive Environment Forward**

These documents have proven to be extremely valuable to the success of this environment; they ensure consistency in our efforts to move the environment forward. These documents also provide a diary that can be referenced to determine historical trends and events, to track downtime and to drive status meetings. At the beginning of each month the current document is renamed (status-0998) and all completed tasks are removed.

The format of these documents leaves something to be desired. They are most of the time a mess, it is difficult to sort through them and we do not keep them online which would allow them to be updated by all team members. Despite all of this they work and they require minimal time to manage.

We currently have this history on the environment since 09/97.

See Appendix H for an example status diary.

Uptime and Availability Statistics

- **Valuable Tool - Assists in keeping perceptions about environment grounded in fact.**
- **Available System Defined as:**

- **Uptime considers planned downtime and ignores unplanned downtime.**
- **Availability considers unplanned downtime and ignores planned downtime**
- **Absolute Availability**

Uptime and Availability statistics are an important tool for systems administrators and managers. This information can help to minimize misguided perceptions about computing resources and ensure that decisions regarding these resources are grounded in fact.

The uptime and availability information we track is limited to the hardware and operating systems under the direct control of the IR Finance administration team. Application and database statistics are not included. While this information would obviously be useful we are not in a position to gather valid information for the numerous database instances and applications running throughout the environment.

Each planned and unplanned downtime event is tracked by the administration team and recorded in the system status diaries. The availability/uptime document is then updated on a monthly basis. If required, detailed historical information can be viewed upon request.

For our purposes an “up” or “available” system is defined as:

> An instance of the operating system that is booted to run level 3. This run level implies networking. ALL associated daemons/services including but not limited to NIS, DNS, NFS, inetd (TCP, UDP), DCE, and RPC must be functioning as expected and available to all applications and users.
>
> All hardware resources must be “CLAIMED” (except in the case of PV Link devices) and available to all applications and users.

The **uptime** statistic accounts for planned downtime and ignores unplanned downtime. Typically this statistic is given for periods of time as small as 1 week. Uptime is usually expressed in hours and days. For example 12 hours of planned downtime in one week would be 24 X 6.5 or 24 hours by 6.5 days. Frankly an uptime statistic presented in this format doesn’t do much for me. Therefore we present the uptime statistic as a percentage referencing up to 1 year.

Uptime is a direct reflection of the amount of change an environment is experiencing and also the ability of the administrators to successfully manage planned downtime.

The **availability** statistic accounts for unplanned downtime and ignores planned downtime. Availability is expressed as a percentage referencing up to 1 year.

Availability is a direct reflection of the stability of hardware and operating systems.

Availability is usually the number that generates the most interest and it is certainly the most widely used.

While both uptime and availability statistics are useful it is also valuable to know the absolute percentage of time that systems are available for processing. The **absolute availability** statistic accounts for both planned and unplanned downtime it is expressed as a percentage referencing up to 1 year.

The following is an example of how these statistics are presented.

HOST	**Sys_a**	**Sys_b**	**Sys_c**
Model	K420/1	T600/8	K420/2
HP-UX	10.20	10.20	10.20
HP/RT	4 hours	2 hours	2 hours
Start Date	1/1/99	1/1/99	1/1/99
Available Hours Per Day	24.0	24.0	24.0
Days Tracked	90	90	90
Max. Avail. Hours To Date	2160.00	2160.00	2160.00
Uptime			
Scheduled Downtime HTD	14.00	7.50	7.50
Uptime Hours To Date	2146.00	2152.50	2152.50
% Uptime To Date	**99.35%**	**99.65%**	**99.65%**
Availability			
Unscheduled Downtime HTD	0.25	0.25	0.25
Availability Hours To Date	2159.75	2159.75	2159.75
% Available To Date	**99.99%**	**99.99%**	**99.99%**
Absolute Availability			
Absolute Downtime HTD	14.25	7.75	7.75
Absolute Availability HTD	**2145.75**	**2152.25**	**2152.25**
% Absolutely Available TD	**99.34%**	**99.64%**	**99.64%**

This document and the calculations it contains are based on hardware and operating system availability. Application availability is not considered here.

*The **Uptime** calculation does not consider unscheduled downtime; it does consider scheduled downtime.*
*The **Availability** calculation does not consider scheduled downtime; it does consider unscheduled downtime.*

*The **Absolute Availability** calculation considers both scheduled and unscheduled downtime.*

HPD	*Hours per day*
HTD	*Hours to date*
TD	*To date*
HP/RT	*Contracted HP on site response time.*

Last Updated 4/1/99

Firmware and Microcode

- **Time Intensive Task**
- **8 to 12 Month Cycle**
- **Consistent Revisions Across Each Hardware Class**

- **New Hardware**
- **Upgrades Outside of Standard Cycle**
- **HP's Customer Care Script**

Security

- **Reasonable Security: Provide defense against a hackers attack, authorized users mistakes; hardware and operating system failures.**
- **Specific to hardware, operating system and associated applications.**
- **Not limited to this section.**
- **Result in confidentiality, data integrity, availability and ease of administration.**

Overview

The goal of this security implementation is to define reasonable security standards for the TCI Information Resources HP/Financial environment. Reasonable security means that the systems in this environment should provide reasonable defenses against a hackers attack or an authorized user's mistakes while still affording the users the access and functionality required to perform their work. In the cases where functionality outweighs security, this documentation will serve as a record as to why specific security standards are not implemented.

These Security standards apply directly to the operating system configurations, user accounts, passwords, data protection, change management, networking, backups/restores and disaster recovery of the TCI Finance systems. They do not apply to application, database, intranet/internet, firewall or network security outside of specific UNIX network configurations.

This environment is in a constant state of evolution and thus the security standards must also evolve.

This security effort will result in increased confidentiality, data integrity, availability, standardization, and ease of administration.

The specific descriptions of the security standards we have implemented are not limited to this security section; this entire document describes our security configuration. This Security section serves as a centralized place to document security standards that are not already described elsewhere in this document. The security standards and outstanding security issues listed here are not presented in any particular order.

Security Standards

- /etc/motd and /etc/issue
  - monitor_setuid
  - monitor_config
  - monitor_passwd
    - finger_it
  - tcp_random_seq
  - secure_it RC set
    - ftpusers
  - tcp_wrappers
    - /overlord

Outstanding Security Issues

- Document issues that have not been resolved.

Security Knowledge

- **Security Tools**
- **Security Threats**
- **Signs of a Security Breach**
- **Recovering from a Security Breach**

Security Tools

In addition to the tools we have created in house the following third party tools can be used to identify, monitor and resolve security issues.

Cops - Monitors common procedural UNIX problems
PGP - Public and private keys
Tcpdump - Network packet information
Swatch - System activity monitor
Npassword - Replacement for passwd, incorporates password checking
Satan - Network security

Security Threats

The following terms are often used with regard to the different types of security threats. They are defined here to provide administrators with a basic understanding of potential risks.

Trojan Horses: A program that appears to have one function but performs another. For example, a simple script can be written to mimick the login program; this script will actually capture user login names and passwords.

Spoofs: A program that masquerades as another. For example, a user creates a script called su and puts that script in a directory that is included in his modified $PATH (or if root's $PATH contains a "."). The user then requests assistance from an SA. Assuming that the SA will su from the users account. Therefore executing the users su program rather than the systems.

Trap Doors: Creating root access to a system that bypasses normal security authentication. A user discovers that root has a . in its $PATH. The user creates a script called ls that copies /bin/ksh to some directory and executes a chmod 4555 on it. The user then requests assistance from an administrator in the hopes that he will perform an ls in the directory where his script resides. Since the setuid for root has been set on this shell the user can now execute it and become root.

Time Bombs: A program that has hidden features that execute when specific conditions are met.

Viruses: A program that inserts copies of itself into other executables on a system.

Worms: A program that propagates itself from computer to computer. It does not necessarily modify other programs on machines.

Signs of a Security Breach

A sudden unexpected increase in system load.
Unexpected changes in configuration files or executables.
The existence of non-standard setuid and setguid files.
Logins or login attempts from subnets outside the intranet.
Logins or login attempts from IP's unknown to DNS.
Unusual internet services activity.
Root activity/logins from unusual hosts.
Processes running on the system resulting in denial of service.
A sudden unexpected increase in network utilization.
Unusual entries in system log files.
Unusual hardware events on any consistent basis.
Unexpected console activity.
Unexpected changes in the utilization (I/O or capacity) of vg00.
An increase in the number of bad login attempts.

Recovering from a Security Breach

1-Immediately notify DBA's to halt all databases
2-Shutdown the system
3-Notify Management
4-Boot the system to single user mode, activate all volume groups and mount all file systems
5-Determine origin of attack
6-Analyze damage
7-Eliminate security weakness from environment
8-Save any evidence of attack
9-Notify management and proper authorities
10-Test system integrity
11-Monitor other systems in environment
12-Boot the system from either third or fourth boot disk and reestablish integrity of the operating system from mirror
13-Test integrity of user files, databases and applications
14-Restore user files, databases and applications from tape if necessary

Appendices

Scripts have been modified for this presentation.

Use at your own risk!