

2294
**Tru64 UNIX and Linux:
A Side-by-Side
Comparison**

Bradford Nichols

Technical Consultant
HP BCS

Transition Engineering and Consulting Group



Abstract

- This session compares and contrasts the features of Tru64 UNIX to those of Linux in the areas of general system administration and installation, configuration for users and accounts, storage and file system management, process and job control as well as clustering and resource management capabilities. The talk is intended to leave the audience with a understanding of the relative strengths of the two operating environments.
- As Tru64 UNIX customers evaluate future replacement environments to their existing Tru64 installations over the next few years. Many will be evaluating a possible role for Linux as well the HP-UX. In general, customers relying on the advanced enterprise computing features of Tru64 UNIX will be more naturally drawn to HP-UX as the replacement, particularly after TruCluster and AdvFS file systems have been ported to HP-UX. On the other hand, the popularity, flexibility and alternative cost model of Linux makes it a viable alternative for Tru64 customers in the areas of application development, web interfaces, and increasingly as time goes on, database servers.

Topics

■ Topics

- Introduction
- Brief Histories
- UNIX Look and Feel
- Storage Management Components
- Scaling and Advanced Enterprise Features

■ *Non-Topics*

- *Open Source Development and Distribution Models*
- *ISV availability*
- *Performance via Alpha, Linux on I32*
- *TOC or Economics*
- *Service and Support Models*
- *I32 implementation*
- *HPTC*

Tru64 UNIX and Linux

- Tru64 UNIX® ~88
 - starting point – 1988 Open Software Foundation’s (OSF) OSF/1
 - low level kernel services – CMU Mach
 - higher level kernel services – BSD UNIX
 - compatibility APIs, libs and tools for SysV UNIX
 - (replaced Digital’s earlier BSD based UNIX – Ultrix)
 - continued Digital/Compaq enterprise enhancements
 - SMP,
 - Partitioning
 - Multipathing,
 - Proprietary Unix
 - Mature
 - Full set of enterprise features including scaling for large servers
 - Established traditional support model

- Linux ~94
 - started point - 1991 a project by then university student Linus Torvalds,
 - inspired by partly Minix
 - goal of a simple UNIX clone to run on 386-based PCs, independent of proprietary code
 - early on utilized GNU software and other UNIX code already covered by GPL
 - Kernel – Linux
 - Compilers, Tools and Basic User-mode Environment – GNU
 - Additional Components – Desktop, Installation, GUI Admin – Other Open Source Projects or Distributor
 - made available on the internet under GNU GPL license
 - Spawned a “Open Source Phenomenon” of collaborative work by programmers all over the world
 - Distributed under GPL by vendors like Red Hat and SuSe
 - Open Source Unix
 - “Free” – low cost
 - Flexible, go your own way
 - No lock in

which Linux /platform to compare to Tru64 UNIX?



- for the purpose of this presentation;
 - Distribution: Red Hat Enterprise Linux 2.1 AS
 - Platform: HP Itanium Processor based Servers
- one of two distributions currently supported or certified for HP IPF servers
 - other is SuSe 8.0 based on United Linux UL 1.0
 - HP large contributor to Linux on IPF effort through Trillion project
 - assumption - official distributions on powerful platforms will be of most interest to customers when considering replacements for Tru64/Alpha in enterprise environments

historical timelines

- Tru64 UNIX
 - ~85 Mach 2.5, OSF/1 – Micro-Kernel concepts, Kernel Threads, memory mapped files
 - ~88 DEC OSF/1 – LVM, SysV compatibility, Shared Libraries,
 - ~93 DEC OSF/1 1.2 - Alpha, 64-bit native OS, Unified Buffer Cache
 - 04/94 DEC OSF/1 2.0a – ASE 1.0 Cluster
 - 08/94 Digital UNIX 3.0 - SMP, AdvFS, LSM replaces LVM
 - 06/96 Tru64 UNIX 4.0 - CDE, sysman tools, UNIX 95
 - 05/99 Tru64 UNIX 4.0f – Fibre Channel support
 - 08/99 Tru64 UNIX 5.0 – TruCluster (TCR) Single-Image Clustering, device location independent naming, built-in self configuring storage Multi-Pathing
 - 08/00 Tru64 UNIX 5.1 - NUMA (Global Switch), 32 CPUs, 128 GB physical memory
 - 09/02 Tru64 UNIX 5.1B - NUMA (Fabric) , Big Pages, 256GB physical memory
 - Q3/03 Tru64 UNIX “Vail”, - 64 CPU, migration tools to HP-UX and common system management
- Linux Kernel / Red Hat
 - 10/91 Linux 0.02
 - 93 ACC Corporation founded,
 - 03/94 Linux 1.0 – i386, uni-processor
 - 94 Marc Ewing creates first “Red hat” release
 - 03/95 Linux 1.2 – support for multiple CPU architectures (Alpha, MIPS,..)
 - 95 ACC buys out Ewing and creates Red Hat Software
 - 95 Red Hat Linux 2.0, includes RPM
 - 06/96 Linux 2.0 – more architectures, SMP
 - 7/97, Red Hat 4.2
 - 7/98 Red Hat - Oracle announces support
 - 01/99 Linux 2.2 – better SMP, wider hardware options
 - 04/99 Red Hat 6.0
 - 10/99 Red Hat 6.1
 - 01/01 Linux 2.4 – more SMP scaling, better desktop USB, PCMCIA, P&P
 - 01 7.1 includes 2.4 kernel, introduces software manager for red hat network – automatic bug notifications, software updates etc.
 - 01/10 Red Hat 7.2
 - 03/02 Linux Advanced Server ; Oracle, Veritas announce support
 - 09/02 Red Hat 8.0
 - 03/03 Enterprise family of operating systems products launched
 - 04/03 Red Hat 9
 - Next - Linux 2.6 16+ CPUs, True Asynchronous I/O, Better user-threads

basic UNIX look and feel

- Linux was intended as a UNIX clone
 - “aims to be compliant with the IEEE POSIX standard”
- So many areas of basic user and admin environments are very similar or identical
 - user/group accounts
 - shells
 - file and directory commands
 - rooted tree and basic file system layout
 - process and job control including cron and at
 - system ascii logs
 - system start, run levels and shutdown
 - basic IP network interfaces and configuration

users, groups and shells

	Tru64 UNIX® V5.1B	Red Hat 2.1 AS IPF
user and group files	<code>/etc/passwd</code> <code>/etc/group</code>	<code>/etc/passwd</code> <code>/etc/group</code>
default user account definition	<code>/usr/skel</code>	<code>/etc/skel</code>
command-line utilities to add a user	<code>useradd, userdel</code> <code>adduser, removeuser</code>	<code>useradd, userdel,</code> <code>adduser, removeuser</code>
system-wide shell startup file	<code>/etc/profile</code>	<code>/etc/profile</code>
shell information	<code>/etc/shells</code>	<code>/etc/shells</code>
Bourne shell	<code>[/usr]/bin/sh</code>	<code>/bin/bash (sym links</code> <code>/bin/sh, /bin/bash2)</code>
Korn shell	<code>[/usr]/bin/ksh</code>	
POSIX shell	<code>[/usr]/bin/posix/sh</code>	
C shell	<code>[/usr]/bin/csh</code>	<code>/bin/csh, /bin/tcsh</code>

manipulating files and file systems



Tru64 UNIX® V5.1B

Red Hat AS 2.1 IPF

user file and dir
commands

`ls, cd, find,
more ...`

`ls, cd, find,
more, less...`

mounting/unmounting
file systems

`mount, umount`

`mount, umount`

boot time mounted
file systems

`/etc/fstab
/sbin/bcheckrc`

`/etc/fstab
/etc/rc.d/rc.sysinit
/etc/init.d/netfs`

list mounted file
systems

`df`

`df`

file system directory hierarchy

Tru64 UNIX® V5.1B

/
/dev, /devices
/etc
/home, /usr/users
/lost+found (UFS)
/mnt
/opt, /usr/opt, /var/opt
/sbin,
/shlib
/sbin, /usr/bin, /usr/sbin
(/bin sym link to /usr/bin)
/vmunix
/subsys, /sys
/tmp
/usr
/usr/lib, /usr/shlib
/var
NA

Red Hat AS 2.1 IPF

/
/dev
/etc
/home
/lost+found
/mnt, /mnt/cdrom
/var/opt
/sbin, /bin
/lib
/bin, /sbin, /usr/bin, /usr/
sbin
/boot/vmlinuz-X.X
/usr/src/linux-X.X
/tmp
/usr
/usr/lib
/var
/usr/src

device special files

configuration files

default user home dirs

fsck output area

temporary mount

optional software

single user-mode binaries

single user-mode shared libs

binaries

bootstrap, kernel

kernel config and build

libraries

source area

man hier

```
[tru64] # man hier
hier(5)
  hier(5)
```

NAME

hier - Standard file system hierarchy

DESCRIPTION

A Tru64 UNIX operating system has a standard file system hierarchy. So does the X11 Window System. The operating system has a unique root directory, the identity of which is compiled into the operating system kernel and is activated when the operating system is bootstrapped. The X11 Window System is a file system within the operating system hierarchy.

...

Base System Hierarchy	Description of Directory or File
/	The root directory for the root file system of the operating system
/cluster/ members/	Directory for a cluster of which this system could be a member Root directory for cluster member0, this system, whether or not it is in a cluster
/dev/	Block and character device special files

...

```
[redhat] # man hier
HIER(7)          Linux Programmer's Manual
  HIER(7)
```

NAME

hier - Description of the file system hierarchy

DESCRIPTION

A typical Linux system has, among others, the following directories:

- / This is the root directory. This is where the whole tree starts.
- /bin This directory contains executable programs which are needed in single user mode and to bring the system up or repair it.
- /boot Contains static files for the boot loader. This directory only holds the files which are needed during the boot process. The map installer and configuration files should go to /sbin and /etc.
- /dev Special or device files, which refer to physical devices. See mknod(1).

...

basic processes and jobs

Tru64 UNIX® V5.1B

Red Hat AS 2.1 IPF

process
control

`ps, kill, nice, renice`

`ps, kill, nice, renice`

cron, at,
batch

`/usr/sbin/cron`

`/usr/sbin/crond`

`/var/adm/cron/cron.allow,
cron.deny`

`/etc/cron.allow, cron.deny`

`/var/adm/cron/log`

`/var/log/cron`

`/var/spool/cron/crontabs/
USER`

`/var/spool/cron/USER,
/etc/cron.d/XXX`

`/etc/crontab (system)`

`/var/spool/cron/atjobs`

`/var/spool/at`

(handled as part of cron)

`/var/spool/at/spool`

`/etc/at.allow`

`/etc/at.deny`

ps

```
[tru64] # ps -ef
UID          PID    PPID      C  STIME     TTY          TIME CMD
root         524288      0    1.3   Jul 17   ??          06:36:05 [kernel idle]
root         524289  524288    0.0   Jul 17   ??          0:00.24 /sbin/init -a
root         524290  524288    0.0   Jul 17   ??          0:00.00 [kproc_creator_da]
```

...

```
[redhat] # ps -ef
UID          PID    PPID      C  STIME     TTY          TIME CMD
root          1      0      0  Aug06   ?           00:00:05 init
root          2      1      0  Aug06   ?           00:00:00 [migration_CPU0]
root          3      1      0  Aug06   ?           00:00:00 [migration_CPU1]
```

...

```
[tru64] # ps aux
USER          PID  %CPU  %MEM    VSZ   RSS TTY      S   STARTED          TIME COMMAND
root         524288    0.8   4.1  5.05G  169M ??        R <   Jul 17          06:36:04 [kernel idle]
root         954155    0.1   0.0   2.52M  360K console  S     Aug 06           0:00.04 -ksh (ksh)
root         524564    0.1   0.0   3.45M  480K ??        S     Jul 17           0:12.10 /usr/sbin/evmlogger -o /
```

...

```
[redhat]# ps aux
USER          PID  %CPU  %MEM    VSZ   RSS TTY      STAT START   TIME COMMAND
root          1    0.0   0.0   2864  1296 ?        S     Aug06    0:05 init
root          2    0.0   0.0     0     0 ?        SW    Aug06    0:00 [migration_CPU0]
root          3    0.0   0.0     0     0 ?        SW    Aug06    0:00 [migration_CPU1]
```

...

system logs

Tru64 UNIX® V5.1B

traditional
ASCII logs

`/etc/syslog.conf`

`syslogd`

`/var/adm/syslog.dated/`

`XXX/ [kern|daemon,] .log`

`/var/adm/messages`
(kern.debug only)

kernel logs

binary

`/etc/binlog.conf`

`binlogd`

`/dev/kbinlog`

`/var/adm/binary.errlog`

Red Hat AS 2.1 IPF

`/etc/syslog.conf`

`syslogd`

`/var/log/messages.X`

ascii

embedded launch from
`/sbin/init.d/syslog`

klogd

`/proc/kmsgs`

`/var/log/messages via`
syslogd

boot trace in messages

```
[tru64] # more /var/adm/messages
```

```
...
Jun 24 19:29:30 localhost vmunix: Alpha boot: available memory from 0x3bbe000 to 0xffff4000
Jun 24 19:29:30 localhost vmunix: Compaq Tru64 UNIX P5.1B (Rev. 173); Tue Dec 17 15:49:27 EST 20
Jun 24 19:29:30 localhost vmunix: physical memory = 4096.00 megabytes.
Jun 24 19:29:30 localhost vmunix: available memory = 3965.19 megabytes.
Jun 24 19:29:30 localhost vmunix: using 15646 buffers containing 122.23 megabytes of memory
Jun 24 19:29:30 localhost vmunix: Master cpu at slot 0
Jun 24 19:29:30 localhost vmunix: Starting secondary cpu 1
Jun 24 19:29:30 localhost vmunix: Starting secondary cpu 2
Jun 24 19:29:30 localhost vmunix: Starting secondary cpu 3
```

```
[redhat] # more /var/adm/messages
```

```
...
Aug 6 21:04:25 localhost kernel: Inspecting /boot/System.map-2.4.18-e.31smp
...
Aug 6 21:04:25 localhost kernel: Loaded 16789 symbols from /boot/System.map-2.4.18-e.31smp.
Aug 6 21:04:25 localhost kernel: Symbols match kernel version 2.4.18.
Aug 6 21:04:25 localhost kernel: Error seeking in /dev/kmem
Aug 6 21:04:25 localhost kernel: Symbol #eepro100, value 00240000
Aug 6 21:04:25 localhost kernel: Error adding kernel module table entry.
Aug 6 21:04:25 localhost kernel: Linux version 2.4.18-e.31smp (bhcompile@natasha.devel.redhat.c
Aug 6 21:04:25 localhost kernel: EFI v1.02 by INTEL: SALsystab=0x7ff24e40 ACPI=0x7 ffda0c0 ACPI
Aug 6 21:04:25 localhost kernel: CPU 0: mapping PAL code [0x7ff40000-0x7ff7b000) into [0xe00000
```

startup and shutdown

Tru64 UNIX® V5.1B

Red Hat AS 2.1 IPF

startup process `init`
/command

`init`

`/etc/inittab`
`(initdefault 3)`

`(telinit)`

`/etc/inittab`

rc script(s) `/sbin/rc[2|3]`

`(initdefault 5)`

`/sbin/rc.d/rc [0..6]`

rc directories `/sbin/rc[0|2|3].d`

`/sbin/rc[0..6].d`

subsystem start script `/sbin/init.d/XXX`

`/etc/init.d/XXX`

subsystem settings `/etc/rc.config`

`/etc/sysconfig/XX`

Shutdown `shutdown, reboot`

`shutdown, reboot`

run levels

	Tru64 UNIX® V5.1B	Red Hat AS 2.1 IPF
Halt	0	0
Single User	S	1
Multi-User (gettys), fstab, no networking	2	-
Multi-User, networking interfaces	-	2
Mutli-User, full networking, fstab	-	3
Multi-user, full networking, X	3	5
Unused	-	4,6

inittab

```
[tru64] # more /etc/inittab
...
is:3:initdefault:
ss:Ss:wait:/sbin/rc0 shutdown < /dev/console >
s0:0:wait:/sbin/rc0 off < /dev/console > /dev/console 2>&1
fs:23:wait:/sbin/bcheckrc < /dev/console > /dev/console 2>&1
esm_init:23:wait:/sbin/init.d/esm init </dev/null >...
esmd:23:respawn:/usr/sbin/esmd </dev/null >/dev/null 2>&1
kls:Ss:sysinit:/sbin/kloadsrv < /dev/console > ...
hsd:Ss:sysinit:/sbin/hotswapd < /dev/console > ...
sysconfig:23:wait:/sbin/init.d/autosysconfig start < ...
update:23:wait:/sbin/update > /dev/console 2>&1
smsync:23:wait:/sbin/sysconfig -r vfs smoothsync-age=30 ...
smsyncS:Ss:wait:/sbin/sysconfig -r vfs smoothsync-age=0 ...
it:23:wait:/sbin/it < /dev/console > /dev/console 2>&1
kmk:3:wait:/sbin/kmknod > /dev/console 2>&1
s2:23:wait:/sbin/rc2 < /dev/console > /dev/console 2>&1
s3:3:wait:/sbin/rc3 < /dev/console > /dev/console 2>&1
cons:1234:respawn:/usr/sbin/getty console console ...
cms:s:sysinit:/sbin/sysconfig -o cms 100 > ...
```

```
[redhat] # more /etc/inittab
...
id:5:initdefault:
...
si::sysinit:/etc/rc.d/rc.sysinit
10:0:wait:/etc/rc.d/rc 0
11:1:wait:/etc/rc.d/rc 1
12:2:wait:/etc/rc.d/rc 2
13:3:wait:/etc/rc.d/rc 3
14:4:wait:/etc/rc.d/rc 4
15:5:wait:/etc/rc.d/rc 5
16:6:wait:/etc/rc.d/rc 6
...
ud::once:/sbin/update
...
ca::ctrlaltdel:/sbin/shutdown -t3 -r now
...
pf::powerfail:/sbin/shutdown -f -h +2 "Power Failure;
System Shutting Down"
...
pr:12345:powerokwait:/sbin/shutdown -c "Power
Restored; Shutdown Cancelled"
...
co:012345:respawn:/sbin/agetty ttyS0 115200 vt100
1:2345:respawn:/sbin/mingetty tty1
2:2345:respawn:/sbin/mingetty tty2
3:2345:respawn:/sbin/mingetty tty3
4:2345:respawn:/sbin/mingetty tty4
5:2345:respawn:/sbin/mingetty tty5
6:2345:respawn:/sbin/mingetty tty6
...
x:5:respawn:/usr/bin/kdm -nodaemon
```

example nfs startup - scripts

```
[tru64] # find /sbin/rc* -name "*nfs*" -
print
/sbin/rc0.d/K30nfs
/sbin/rc0.d/K35nfsmount
/sbin/rc2.d/K35nfs
/sbin/rc2.d/K40nfsmount
/sbin/rc3.d/S19nfs
/sbin/rc3.d/S20nfsmount
[tru64] # ls /sbin/init.d/*nfs*
/sbin/init.d/nfs
/sbin/init.d/nfsmount
```

```
[redhat]# find /etc/rc.d -name "*nfs*" -
print
/etc/rc.d/init.d/nfs
/etc/rc.d/init.d/nfslock
/etc/rc.d/rc0.d/K20nfs
/etc/rc.d/rc0.d/K86nfslock
/etc/rc.d/rc1.d/K20nfs
/etc/rc.d/rc1.d/K86nfslock
/etc/rc.d/rc2.d/K20nfs
/etc/rc.d/rc2.d/K86nfslock
/etc/rc.d/rc3.d/K20nfs
/etc/rc.d/rc3.d/S14nfslock
/etc/rc.d/rc4.d/K20nfs
/etc/rc.d/rc4.d/S14nfslock
/etc/rc.d/rc5.d/S60nfs
/etc/rc.d/rc5.d/S14nfslock
/etc/rc.d/rc6.d/K20nfs
/etc/rc.d/rc6.d/K86nfslock
[redhat]# ls /etc/init.d/*nfs*
/etc/init.d/nfs /etc/init.d/nfslock
[redhat]# ls /var/lock/subsys/*nfs*
/var/lock/subsys/nfs
/var/lock/subsys/nfslock
```

example NFS startup - settings

```
[tru64] # more /etc/rc.config
#!/bin/sh
...
#
. /etc/rc.config.common
DISPLAYTYPE=
MAX_NETDEVS=
NETDEV_2=
...
[tru64] # more /etc/rc.config.common
#!/bin/sh
...
NUM_NFSIOD="7"
export NUM_NFSIOD
AUTOMOUNT="1"
export AUTOMOUNT
AUTOMOUNT_ARGS="-v -D MACH=alpha -D OS=osf1 -D NET=f"
export AUTOMOUNT_ARGS
NFS_CONFIGURED="1"
export NFS_CONFIGURED
NFSSERVING="1"
export NFSSERVING
NFSLOCKING="1"
export NFSLOCKING
MOUNTOPTS="-i"
export MOUNTOPTS
NONROOTMOUNTS="0"
export NONROOTMOUNTS
...

```

```
[redhat] more /etc/init.d/nfs
...
nfs:RPCNFSDCOUNT=8
```

- Start links created and removed by GUI services tools as services are enabled and disabled. Less need for a NFSSERVING? Type flag
- autofs/automount flags are inbedding in /etc/init.d/autofs or are in /etc/auto.master.

network interfaces and services

	Tru64 UNIX® V5.1 B	Red Hat AS 2.1 IPF
interface names	<code>lnX, eeX</code>	<code>ethX</code>
interface settings	<code>/etc/rc.config</code>	<code>/etc/sysconfig/network</code> <code>/etc/sysconfig/network-scripts/ifcfg-ethX</code>
show configured interfaces	<code>ifconfig -a</code>	<code>ifconfig -a</code>
network services daemon	<code>/usr/sbin/inetd</code>	<code>/usr/sbin/xinetd</code>
network services daemon config file	<code>/etc/inetd.conf</code>	<code>/etc/xinetd.conf</code> <code>/etc/xinetd.d/XXX</code>
network services config file	<code>/etc/services</code>	<code>/etc/services</code>
failover between physical NICs	NetRAIN	Channel Bonding
aggregation between physical NICs	Link Aggregation (LAG)	Channel Bonding

ifconfig -a

```
[tru64] # ifconfig -a
ee0: flags=1000c63<UP,BROADCAST,NOTRAILERS,RUNNING,MULTICAST,SIMPLEX,CLUIF>
    inet 10.1.0.1 netmask ffffffff broadcast 10.1.0.255 ipmtu 1500

ee1: flags=c63<UP,BROADCAST,NOTRAILERS,RUNNING,MULTICAST,SIMPLEX>
...
lo0: flags=100c89<UP,LOOPBACK,NOARP,MULTICAST,SIMPLEX,NOCHECKSUM>
    inet 127.0.0.1 netmask ff000000 ipmtu 4096
...
tu0: flags=c63<UP,BROADCAST,NOTRAILERS,RUNNING,MULTICAST,SIMPLEX>
    inet 16.29.128.102 netmask ffffffff broadcast 16.29.128.255 ipmtu 1500
```

```
[redhat]# ifconfig -a
eth0      Link encap:Ethernet  HWaddr 00:03:47:D0:C8:C6
          inet addr:16.141.8.56  Bcast:16.141.8.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:2545655 errors:0 dropped:0 overruns:0 frame:0
          TX packets:504 errors:0 dropped:0 overruns:0 carrier:0
          collisions:4 txqueuelen:100
          RX bytes:351953576 (335.6 Mb)  TX bytes:45451 (44.3 Kb)
          Interrupt:49 Base address:0xf000

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:480 errors:0 dropped:0 overruns:0 frame:0
          TX packets:480 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:31846 (31.0 Kb)  TX bytes:31846 (31.0 Kb)
```

interface settings

```
[tru64] # more /etc/rc.config
NETDEV_1=
NETDEV_2=
NETDEV_3=
NETDEV_4=
NETDEV_5=
NETDEV_7=
IFCONFIG_1=
IFCONFIG_2=
IFCONFIG_3=
IFCONFIG_4=
IFCONFIG_5=
IFCONFIG_7=
...
NETDEV_6="tu0"
export NETDEV_6
HOSTNAME="localhost.zko.dec.com"
export HOSTNAME
...
NETDEV_0="ics0"
export NETDEV_0
IFCONFIG_0="10.0.0.1 netmask 255.255.255.0 filter"
export IFCONFIG_0
IFCONFIG_6="16.29.128.102 netmask 255.255.255.0 filter"
export IFCONFIG_6
NUM_NETCONFIG="2"
export NUM_NETCONFIG
...
```

```
[redhat]# more /etc/sysconfig/network
NETWORKING=yes
HOSTNAME=yusuke
GATEWAY=16.141.8.1

[redhat]# ls /etc/sysconfig/network-scripts/ifcfg*
/etc/sysconfig/network-scripts/ifcfg-eth0
/etc/sysconfig/network-scripts/ifcfg-lo

[redhat]# more /etc/sysconfig/network-scripts/ifcfg*
:::::::::::::
/etc/sysconfig/network-scripts/ifcfg-eth0
:::::::::::::
DEVICE=eth0
BOOTPROTO=static
BROADCAST=16.141.8.255
IPADDR=16.141.8.56
NETMASK=255.255.255.0
NETWORK=16.141.8.0
ONBOOT=yes
:::::::::::::
/etc/sysconfig/network-scripts/ifcfg-lo
:::::::::::::
DEVICE=lo
IPADDR=127.0.0.1
NETMASK=255.0.0.0
NETWORK=127.0.0.0
# If you're having problems with gated making 127.0.0.0/8 a
  martian,
# you can change this to something else (255.255.255.255, for
  example)
BROADCAST=127.255.255.255
ONBOOT=yes
NAME=loopback
```

kernel builds and configuration

Tru64 UNIX® V5.1B

Red Hat AS 2.1 IPF

largely a re-link of
subsystem modules

largely a recompile from source
files as well as linking

location of kernel

`/vmunix`

`/boot/vmlinuz-X.X` (EFI
partition)

kernel build area

`/sys/HOST`

`/usr/src/linux-X.X/`

build definition file

`/sys/conf/HOST`

`/usr/src/linux-X.X/.config`

tools

`doconfig`

`make config, make
menuconfig, make xconfig
make clean
make bzImage
make modules
make install`

kernel build config file

```
[tru64] # more /sys/conf/localhost
ident          "localhost"

options        UERF
options        OSF
options        _LMF_
options        BIN_COMPAT
options        COMPAT_43
options        MACH
options        MACH_IPC_TCACHE
options        MACH_IPC_WWA
options        MACH_IPC_XXXHACK
options        BUFCACHE_STATS
...
```

```
[redhat]# make config
rm -f include/asm
( cd include ; ln -sf asm-i386 asm)
/bin/sh scripts/Configure arch/i386/config.in
#
# Using defaults found in configs/kernel-2.4.20-i686.config
#
*
* Code maturity level options
*
Prompt for development and/or incomplete code/drivers
(CONFIG_EXPERIMENTAL) [Y/n/?] n
*
* Loadable module support
*
Enable loadable module support (CONFIG_MODULES) [Y/n/?]
Set version information on all module symbols
(CONFIG_MODVERSIONS) [Y/n/?]
Kernel module loader (CONFIG_KMOD) [Y/n/?]
...
more [redhat]# more .config
#
# Automatically generated make config: don't edit
#
CONFIG_X86=y
# CONFIG_SBUS is not set
CONFIG_UID16=y

#
# Code maturity level options
#
# CONFIG_EXPERIMENTAL is not set

#
# Loadable module support
#
CONFIG_MODULES=y
CONFIG_MODVERSIONS=y
CONFIG_KMOD=y
```

kernel build procedure

```
[tru64] # doconfig

*** KERNEL CONFIGURATION AND BUILD PROCEDURE ***

Enter a name for the kernel configuration file. [localhost]:

A configuration file with the name 'localhost' already exists.
Do you want to replace it? (y/n) [n]: y

Saving /sys/conf/localhost as /sys/conf/localhost.bck

*** KERNEL OPTION SELECTION ***

  Selection  Kernel Option
-----
  1          System V Devices
  2          NTP V3 Kernel Phase Lock Loop (NTP_TIME)
  3          Kernel Breakpoint Debugger (KDEBUG)
  4          Packetfilter driver (PACKETFILTER)
  ...
  20         All of the above
  21         None of the above
  22         Help
  23         Display all options again
-----

Enter your choices.

Choices (for example, 1 2 4-6) [21]: 20
...
Do you want to edit the configuration file? (y/n) [n]:
...
*** PERFORMING KERNEL BUILD ***

A log file listing special device files is located in
/dev/MAKEDEV.log
Working...Thu Aug 7 15:14:09 EDT 2003

The new kernel is /sys/localhost/vmunix
```

```
[redhat]# make clean
make[1]: Entering directory `/usr/src/linux-2.4.20-6/arch/i386/boot'
rm -f tools/build
rm -f setup bootsect zImage compressed/vmlinux.out
rm -f bsetup bbootsect bzImage compressed/bvmlinux.out
make[2]: Entering directory `/usr/src/linux-2.4.20-6/arch/i386/boot/compressed'
rm -f vmlinux bvmlinux _tmp_*
...

[redhat]# make bzImage
make[1]: Entering directory `/usr/src/linux-2.4.20-6/arch/i386/boot'
make[1]: Nothing to be done for `dep'.
make[1]: Leaving directory `/usr/src/linux-2.4.20-6/arch/i386/boot'
scripts/mkdep -- init/*.c > .depend
scripts/mkdep -- `find /usr/src/linux-2.4.20-6/include/asm
/usr/src/linux-2.4.20-6/include/linux /usr/src/linux-2.4.20-6/include/scsi /usr/src/linux-2.4.20-6/include/net
/usr/src/linux-2.4.20-6/include/math-emu \ ( -name SCCS -o -name
.svn \) -prune -o -follow -name \*.h ! -name modversions.h -
print > .hdepend
make _sfdep kernel _sfdep drivers _sfdep mm _sfdep fs _sfdep net
_sfdep ipc _sfdep lib _sfdep crypto _sfdep arch/i386/kernel
_sfdep arch/i386/mm _sfdep arch/i386/lib
_FASTDEP ALL SUB_DIRS="kernel drivers mm fs net ipc lib crypto
arch/i386/kernel arch/i386/mm arch/i386/lib"
...
Root device is (3, 2)
Boot sector 512 bytes.
Setup is 4846 bytes.
System is 1033 kB
warning: kernel is too big for standalone boot from floppy
make[1]: Leaving directory `/usr/src/linux-2.4.20-6/arch/i386/boot'

[redhat]# make modules
make -r -f tmp_include_depends all
make[1]: Entering directory `/usr/src/linux-2.4.20-6'
make[1]: Circular /usr/src/linux-2.4.20-6/include/asm/smplock.h <-
/usr/src/linux-2.4.20-6/include/linux/interrupt.h dependency
dropped.
...
```

kernel dynamic modules and interfaces



Tru64 UNIX® V5.1B

Red Hat AS 2.1 IPF

dynamically
loadable kernel
modules

`yes, rarely used`

`yes, commonly used`

module
configuration

`sysconfig,`

`lsmod, modprobe, depmod,
rmmod,`

boot-time loaded
modules

`/etc/sysconfigtab
/sbin/init.d/autosysc
onfig`

`/etc/modules.conf`

dynamic kernel
tuning

`sysconfig,
/etc/sysconfigtab`

`sysctl,
/proc/sys/subsys/param
/etc/sysctl.conf`

listing dynamically loaded kernel modules

```
[tru64] # grep "^SUBSYSTEM_LIST=" /sbin/init.d/autosysconfig
SUBSYSTEM_LIST="hwautoconfig"
```

```
[tru64] ls /var/subsys
lat.mod          marvel_pfm.mod  marvel_pfm.mth  pfm.mod
  pfm.mth
```

```
[tru64] # sysconfig -m | more
cm: static
hs: static
ksm: static
psm: static
generic: static
io: static
ipc: static
...
hwautoconfig: dynamic
envmmon: dynamic
lat: dynamic
```

```
[redhat] more /etc/modules.conf
alias parport_lowlevel parport_pc
alias eth0 eeepro100
alias scsi_hostadapter qla1280
alias eth1 e100
alias sound-slot-0 cs4281
post-install sound-slot-0 /bin/aumix-minimal -f /etc/.aumixrc
-L >/dev/null 2>&1
|| :
pre-remove sound-slot-0 /bin/aumix-minimal -f /etc/.aumixrc -
S >/dev/null 2>&1 |
| :
alias usb-controller usb-uhci
[root@yusuke redhat]#
```

```
[redhat]# lsmod
Module                Size  Used by    Not tainted
nfs                   221648  1 (autoclean)
nfsd                   195264  8 (autoclean)
lockd                  126480  1 (autoclean) [nfs nfsd]
sunrpc                 196544  1 (autoclean) [nfs nfsd]
  lockd]
ide-cd                  74744  0 (autoclean)
cdrom                   73464  0 (autoclean) [ide-cd]
cs4281                 122424  0 (autoclean)
soundcore              14416  3 (autoclean) [cs4281]
button                  8808  0 (unused)
autofs                 31328  0 (autoclean) (unused)
eeepro100              53488  1
ipchains               110632  13
nls_iso8859-1          6048  1 (autoclean)
nls_cp437              7728  1 (autoclean)
vfat                   31096  1 (autoclean)
fat                    88664  0 (autoclean) [vfat]
usb-uhci               70744  0 (unused)
usbcore                176888  1 [usb-uhci]
ext3                   167144  2
jbd                    126688  2 [ext3]
```

storage stack

Tru64 UNIX® V5.1B

storage device
naming

physical location
independent (dskX)

software RAID /
Volume Manager

LSM - licensed port of
Veritas VxVM

software multi-
pathing

built-in to single system and
cluster, auto-configured and
automatically used on all
storage

BSD style file
system

ufs

journal file system

AdvFS (owned and
developed by HP)

other local rw file
systems

-

Red Hat AS 2.1 IPF

physical location
dependent (hdX, sdX)

Multi Device (MD) Driver,
LVM of IBM heritage

?MD driver?

ufs

ext3, reiserfs

ext2,
msdos, vfat,

disks

Tru64 UNIX® V5.1B

Red Hat AS 2.1 IPF

what disks?

```
hwmgr -show scsi, hwmgr -
view dev
```

hwbrowser (X)

```
ls /proc/ide
```

```
more /proc/scsi
```

labels, partitions

```
disklabel
```

parted, fdisk

special file naming

```
dskXY
```

hdXY, sdXY

where:

where;

x: digit(s) which tie the name to the WWID of a disk(LUN), independent of paths or detection order in a probe

x: letter(s) indicating the disk [a...z,aa..zz] based on detection order in probe

X: letter [a..h] for partition

y: digit for a partition

listing attached disks

```
[tru64] # hwmgr -view dev
```

HWID:	Device Name	Location	Mfg	Model
6:	/dev/dmapi/dmapi			
7:	/dev/scp_scsi			
8:	/dev/kevm			
75:	/dev/disk/floppy0c		3.5in floppy	fdi...
93:	/dev/disk/dsk0c		COMPAQ	BB018222B8 bus-1-targ...
94:	/dev/disk/cdrom0c		COMPAQ	CD-224E bus-3-targ...
95:	/dev/random			
96:	/dev/urandom			
101:	/dev/cport/scp0		HSG80CCL	bus-2-targ-...
113:	/dev/disk/dsk911c		DEC	HSG80 IDENTIFIER=911
114:	/dev/disk/dsk912c		DEC	HSG80 IDENTIFIER=912
115:	/dev/disk/dsk913c		DEC	HSG80 IDENTIFIER=913
116:	/dev/disk/dsk914c		DEC	HSG80 IDENTIFIER=914
117:	/dev/disk/dsk915c		DEC	HSG80 IDENTIFIER=915
118:	/dev/disk/dsk916c		DEC	HSG80 IDENTIFIER=916
316:	/dev/disk/dsk921c		DEC	HSG80 IDENTIFIER=921
317:	/dev/disk/dsk922c		DEC	HSG80 IDENTIFIER=922
318:	/dev/disk/dsk923c		DEC	HSG80 IDENTIFIER=923
319:	/dev/disk/dsk924c		DEC	HSG80 IDENTIFIER=924
320:	/dev/disk/dsk925c		DEC	HSG80 IDENTIFIER=925
321:	/dev/disk/dsk926c		DEC	HSG80 IDENTIFIER=926
322:	/dev/disk/dsk927c		DEC	HSG80 IDENTIFIER=30

```
[redhat] # ls /proc/ide
```

```
drivers hda hdb hdc ide0 ide1 piix
```

```
[redhat] # more /proc/scsi/scsi
```

```
Attached devices:
```

```
Host: scsi0 Channel: 00 Id: 00 Lun: 00
```

```
Vendor: QUANTUM Model: ATLAS10K2-TY184L Rev: DA40
```

```
Type: Direct-Access ANSI SCSI  
revision: 03
```

```
Host: scsi0 Channel: 00 Id: 02 Lun: 00
```

```
Vendor: COMPAQPC Model: ATLAS10K2-TY184L Rev: DDC2
```

```
Type: Direct-Access ANSI SCSI  
revision: 03
```

disklabels / partitions

```
[tru64] # disklabel -r dsk925
# /dev/rdisk/dsk925c:
type: SCSI
disk: HSG80
label: clu_member2
flags:
bytes/sector: 512
sectors/track: 254
tracks/cylinder: 20
sectors/cylinder: 5080
cylinders: 1400
sectors/unit: 7109115
rpm: 3600
interleave: 1
trackskew: 7
cylinderskew: 26
headswitch: 0          # milliseconds
track-to-track seek: 0 # milliseconds
drivedata: 0

8 partitions:
#      #      size      offset  fstype  fsize  bsize  cpgr
#      # ~Cyl values
a:      #      524288      0      AdvFS
#      0 - 103*
b:      #      6582779  524288      swap
#      103*- 1399*
c:      #      7109115      0      unused      0      0
#      0 - 1399*
d:      #      0      0      unused      0      0
#      0 - 0
e:      #      0      0      unused      0      0
#      0 - 0
f:      #      0      0      unused      0      0
#      0 - 0
g:      #      3357949  393216      unused      0      0
#      77*- 738*
h:      #      2048      7107067      cnx
#      1399*- 1399*
```

```
[redhat] # parted /dev/hdc print
Disk geometry for /dev/hdc: 0.000-57220.458 megabytes
Disk label type: GPT
Minor  Start      End      Filesystem  Name
  Flags
1      0.017      100.016  FAT
boot, lba
2      100.017    2100.016 linux-swap
lba
3      2100.017   4100.016 ext3
lba
4      4100.017   8100.016 ext3
lba
[root@yusuke redhat]#
```


software RAID

Tru64 UNIX® V5.1B

LSM

Red Hat AS 2.1 IPF

LVM

MD driver

Root Support	yes	yes (/boot /efi no)	yes (/boot/efi no)
Installation support	yes	yes	yes
Multi-pathing	NA (handled in IO stack)	no*	yes
Linear/Append	yes	yes	yes
RAID 0	yes	yes	yes
RAID 1	yes	no	yes
RAID 0+1	yes	no	yes
RAID 3/5	yes	no	4/5

LSM and LVM abstraction comparison

abstraction	LSM	LVM
physical disk	disk media (dm)	physical volume (pv)
logical volume	volume(v)	logical volume (lv)
administrative unit of physical disks and logical volumes	disk group (dg)	volume group (vg)
group of physical volume with different HBAs	NA	physical volume group (pvg)
contiguous extent on a physical volume	subdisk (sd)	physical extent (pe)
contiguous logical extent of storage for use in a volume	plex (p)	logical extent (le)

LSM and LVM

command equivalence

physical volumes/disk media



LSM

voldisksetup brings a disk under VxVM control
1 in the voldiskadm menu adds or initializes one or more disks

voldisk list lists information about VxVM disks

LVM

pvcreate makes a disk an LVM disk

pvdisplay displays information about physical volumes in a volume group

LSM and LVM

command equivalence

volume groups/disk groups

LSM		LVM	
voldiskadd/voldg init	creates a new disk group and/or adds disks to a disk group	vgcreate	creates a volume group
voldg list volprint	displays the contents of a disk group displays information about all objects or a subset of objects.	vgdisplay	displays information on all volume groups.
voldiskadd	adds a disk to the disk group	vgextend	extends a volume group by adding one or more disks to it
voldg deport	deports a disk group from the system.	vgexport	removes a volume group from the system
voldg import	imports a disk group.	vgimport	adds a volume group to the system by scanning physical volumes which have been exported using vgexport

LSM and LVM

command equivalence

logical volume / volume

LSM		LVM	
volassist	creates volumes with the make parameter	lvcreate	create a logical volume
volassist	increases a volume in size with the growto or growby parameter	lvextend	grow the size of a logical volume
volassist snapshot	the snapshot operation takes one of the attached temporary mirrors and creates a new volume with the temporary mirror as its one plex	-	
volcreco ver/volu me start	the vxrecover command performs resynchronize operations for the volumes, or for volumes residing on the named disks (medianame or the VxVM name for the disk)	-	

MD driver software RAID

- In kernel, configuration?

- `cat /proc/mdstat`

- config file `/etc/raidtab`

```
raiddev /dev/md0
    raid-level                0
    nr-raid-disks             2
    persistent-superblock     0
    chunk-size                8

    device                    /dev/sda1
    raid-disk                 0
    device                    /dev/sdb1
    raid-disk                 1
```

...

- raidtools

- `lsraid, mkraid, raidstart, raidstop, raidreconf`

journal file systems

	Tru64 UNIX® V5.1B	Red Hat AS 2.1 IPF
Journal File System	AdvFS	ext3, reiserfs
storage model	multi-volume	single volume
quotas	yes	yes
creating	mkfdmn, mkfset	mkfs,
	-	e2label
resize	addvol, rmvol, mount -u -o extend	resize2fs (ext2), resize_reiserfs (reiserfs)
conversion tools	-	tune2fs (ext2 to ext3)

swap

Tru64 UNIX® V5.1B

Red Hat AS 2.1 IPF

supported containers

partitions

partitions, files

configured

/etc/sysconfigtab

/etc/fstab

addition/removal

swapon -a

**mkswap, swapon,
swapoff**

status

swapon -s

**more /proc/swaps
free**

unified, graphical and distributed admin frameworks



Tru64 UNIX 5.1B

Red Hat AS 2.1 IPF

single system admin

`sysman`, `sysman station` - unified, graphical (curses, X, web)

collection of `redhat-xxxx` command line
collection of X based admin tools ...

limits/scaling

Tru64 UNIX 5.1B

Red Hat AS 2.1 IPF

CPU	32	4-8
Memory	256 GB	(132 – 16GB); 96 GB on largest currently shipping IPF server rx5670
effective file system size lowest common denominator of kernel, volume manager, file system limits	16 TB	2 TB through 2.4 kernel, 16 TB in 2.5 LBD work backported in patches to 2.4
threading	Kernel Threads	Kernel Light Weight Processes
	NXM for User Mode	Nx1
Big Pages	Yes	Yes
NUMA	2nd Generation (5.1, 5.1B)	No

SMP and resource management

Tru64 UNIX® V5.1B

Red Hat AS 2.1 IPF

SMP scheduling

soft processor affinity
with binding options

soft processor affinity

CPU binding options

runon, processor
sets, class scheduler

-

platform partitioning

HP Alpha Servers -
hard partitions

HP IPF Servers -
npartitions (hard)

resource manager

-

Future HP GWM

commercial clustering

	Tru64 UNIX® V5.1B	Red Hat 2.1 AS IPF
IP load balancing to a collection of servers	IP Routing is a component of TruCluster (See Below)	Advanced Server IP load balancing (Piranha)
failover clustering	was Available Server Environment (ASE) – now superseded by TruCluster Server	Advanced Server Cluster Manager
single image cluster failover or parallel applications (Oracle9i RAC) using single cluster wide rooted-tree filesystem	TruCluster Server	-

Note: many other non-Red Hat clustering products exist for Linux including HP ServiceGuard for Linux



HP WORLD 2003

Solutions and Technology Conference & Expo

Interex, Encompass and HP bring you a powerful new HP World.

