Linux

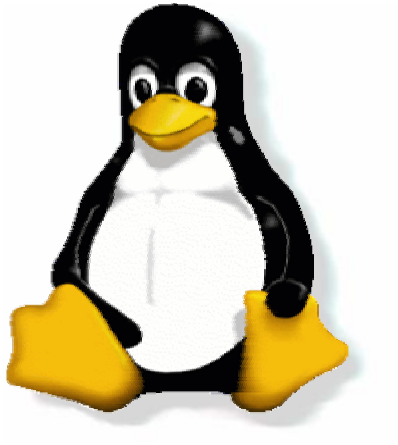

Monitoring and Maintaining

Linux Haiku Error Messages

- http://www.salonmag.com/21st/chal/1998/02/10chal2.html

A crash reduces
your expensive computer
to a simple stone.

-- James Lopez

To have no errors
Would be life without meaning
No struggle, no joy

-- Brian M. Porter

I'm sorry, there's -- um --
insufficient -- what's-it-called?
The term eludes me ...

-- Owen Mathews

The code was willing,
It considered your request,
But the chips were weak.

-- Barry L. Brumitt

There is a chasm
of carbon and silicon
the software can't bridge

-- Rahul Sonnad

Serious error.
All shortcuts have disappeared.
Screen. Mind. Both are blank.

-- Ian Hughes

Linux syslog

- **support for system logging and kernel message trapping**
  - many modern programs use this facility to provide a standardized log facility
  - the kernel, device drivers and other core software also use syslog (klogd on Linux)
  - every logged message contains at least a time and a hostname field
    - *normally a program name field, too*
  - messages are structured according to:
    - *facility*
    - *importance*
  - quite configurable via /etc/syslog.conf
    - *once edited, notify syslogd via*
      - kill -HUP `cat /var/run/syslogd.pid`

Linux Logs and Distributed Logging

- syslog logs to /var/log/messages
  - listens to a socket (/dev/log) and the writes this file
  - (klogd listens to a 4k cyclic buffer in memory)

```
# tail /var/log/messages
Dec  8 20:29:57 redhat PAM_pwdb[339]: (login) session opened for user root by (uid=0)
Dec  8 20:29:57 redhat login[339]: ROOT LOGIN ON tty1
Dec  8 20:29:57 redhat PAM_pwdb[339]: (login) session closed for user root
Dec  8 22:10:06 redhat PAM_pwdb[420]: (su) session closed for user root
Dec  9 04:02:04 redhat PAM_pwdb[1039]: (su) session opened for user nobody by (uid=99)
Dec  9 04:03:53 redhat PAM_pwdb[1039]: (su) session closed for user nobody
Dec  9 06:32:46 redhat PAM_pwdb[1085]: (login) session opened for user bob by (uid=0)
Dec  9 06:32:46 redhat login[1085]: LOGIN ON ttyp0 BY bob FROM aunty
Dec  9 06:32:46 redhat PAM_pwdb[1085]: (login) session closed for user bob
Dec  9 06:45:48 redhat PAM_pwdb[1137]: (su) session opened for user root by bob(uid=0)
```

- syslog can be configured to listen to messages sent over the network
  - provides a centralized logging facility

```
# Sample syslogd configuration file to forward all
# messages to a remote host.
*.*            @hostname
```

  - use the -r switch to syslogd
  - to have this work correctly, /etc/services must contain the following entry:

```
syslog          514/udp
```

Monitoring the System

- **Linux makes it easy to watch what is going on in the system…**
  - …but doesn't really provide the tools to tune things…
    - *the typical solution is to recompile the kernel*
    - *compare with a typical mainframe*
      - or (gasp!) Windows NT
- **an intricate subject**
  - Schrödinger's cat…
- **what can be monitored**
  - CPU, disk space, memory (real and virtual)

Linux Basic Monitoring Tools

- CPU
  - uptime/w
  - ps
  - pstree
  - top
- disk
  - du/df
  - find
- most tools now examine /proc

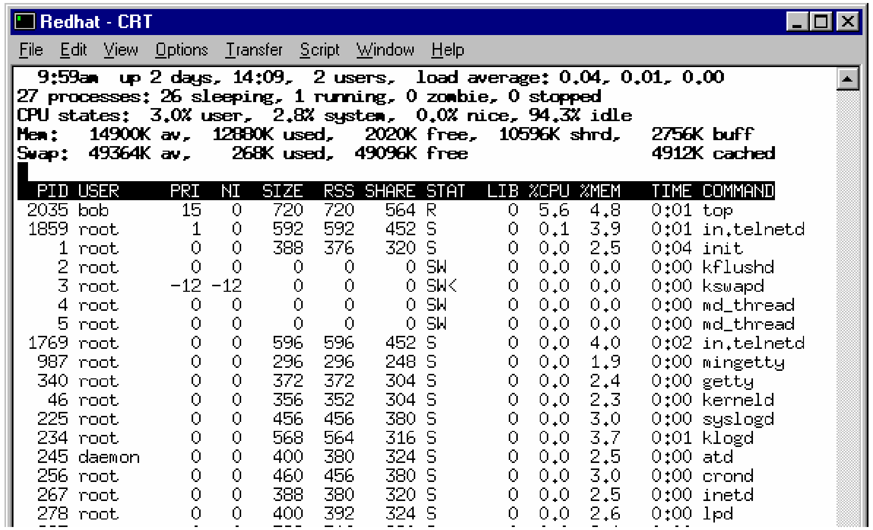

```
  9:59am  up 2 days, 14:09,  2 users,  load average: 0.04, 0.01, 0.00
27 processes: 26 sleeping, 1 running, 0 zombie, 0 stopped
CPU states:  3.0% user,  2.8% system,  0.0% nice, 94.3% idle
Mem:   14900K av,  12880K used,   2020K free,  10596K shrd,   2756K buff
Swap:  49364K av,    268K used,  49096K free                  4912K cached

  PID USER     PRI  NI  SIZE  RSS SHARE STAT  LIB %CPU %MEM   TIME COMMAND
 2035 bob       15   0   720  720   564 R       0  5.6  4.8   0:01 top
 1859 root       1   0   592  592   452 S       0  0.1  3.9   0:01 in.telnetd
    1 root       0   0   388  376   320 S       0  0.0  2.5   0:04 init
    2 root       0   0     0    0     0 SW      0  0.0  0.0   0:00 kflushd
    3 root     -12 -12     0    0     0 SW<     0  0.0  0.0   0:00 kswapd
    4 root       0   0     0    0     0 SW      0  0.0  0.0   0:00 md_thread
    5 root       0   0     0    0     0 SW      0  0.0  0.0   0:00 md_thread
 1769 root       0   0   596  596   452 S       0  0.0  4.0   0:02 in.telnetd
  987 root       0   0   296  296   248 S       0  0.0  1.9   0:00 mingetty
  340 root       0   0   372  372   304 S       0  0.0  2.4   0:00 getty
   46 root       0   0   356  352   304 S       0  0.0  2.3   0:00 kerneld
  225 root       0   0   456  456   380 S       0  0.0  3.0   0:00 syslogd
  234 root       0   0   568  564   316 S       0  0.0  3.7   0:01 klogd
  245 daemon     0   0   400  380   324 S       0  0.0  2.5   0:00 atd
  256 root       0   0   460  456   380 S       0  0.0  3.0   0:00 crond
  267 root       0   0   388  380   320 S       0  0.0  2.5   0:00 inetd
  278 root       0   0   400  392   324 S       0  0.0  2.6   0:00 lpd
```

```
# cat /proc/meminfo
        total:    used:    free:  shared: buffers:  cached:
Mem:  15257600 12050432  3207168 10539008  1638400  5320704
Swap: 50548736   274432 50274304
MemTotal:      14900 kB
MemFree:        3132 kB
MemShared:     10292 kB
Buffers:        1600 kB
Cached:         5196 kB
SwapTotal:     49364 kB
SwapFree:      49096 kB
```

```
# pstree
init-+-atd
     |-crond
     |-getty
     |-gpm
     |-httpd---2*[httpd]
     |-inetd-+-in.telnetd---tcsh---pstree
     |       `-in.telnetd---tcsh---man---sh-+-gunzip
     |                                      `-less
     |-kerneld
     |-kflushd
     |-klogd
     |-kswapd
     |-lpd
     |-2*[md_thread]
     |-2*[mingetty]
     |-nmbd
     |-smbd
     |-syslogd
     `-update
```

```
# w
  9:58am  up 2 days, 14:09,  2 users,  load average: 0.08, 0.02, 0.01
USER     TTY      FROM              LOGIN@   IDLE   JCPU   PCPU  WHAT
bob      ttyp0    aunty             8:44am  0.00s  3.42s  0.15s  w
bob      ttyp1    aunty             9:33am  2:15   1.81s  1.81s  -tcsh
```

```
# du -s /home/bob
32949   /home/bob
```

Linux vmstat

- **virtual machine statistics:**
  - procs
    - *r: processes waiting for run time*
    - *b: processes in uninterruptable sleep*
    - *w: processes swapped out but otherwise runnable*
  - memory (kB)
    - *swpd: virtual memory used*
    - *free: idle memory*
    - *buff: memory used as buffers*
  - swap (kB/s)
    - *si: memory swapped in from disk*
    - *so: memory swapped to disk*
  - IO (blocks/s)
    - *bi: Blocks sent to a block device*
    - *bo: Blocks received from a block device*
  - system
    - *in: interrupts per second, including the clock*
    - *cs: The number of context switches per second*
  - CPU (%)
    - *us: user time*
    - *sy: system time*
    - *id: idle time*

```
# vmstat 5
 procs                  memory    swap        io    system         cpu
 r b w  swpd  free  buff cache  si  so   bi   bo   in   cs  us  sy  id
 0 0 0   268  1732  3012  4952   0   0    0    0  102    4   0   0 100
 0 0 0   268  1800  3012  4952   0   0    0    0  103   10   0   2  98
 0 0 0   268  1800  3012  4952   0   0    0    1  107    4   1   1  98
 0 0 0   268  1800  3012  4952   0   0    0    0  103    4   1   1  98
```

Linux free

- **free**
  - a little simpler to understand than vmstat
    - *but only examines memory*

```
# free -s 5
             total       used       free     shared    buffers     cached
Mem:         14900      13068       1832      10316       3012       4944
-/+ buffers/cache:       5112       9788
Swap:        49364        268      49096

             total       used       free     shared    buffers     cached
Mem:         14900      13072       1828      10352       3012       4944
-/+ buffers/cache:       5116       9784
Swap:        49364        268      49096

             total       used       free     shared    buffers     cached
Mem:         14900      13164       1736      10684       3012       4952
-/+ buffers/cache:       5200       9700
Swap:        49364        268      49096
```

Linux GKrellM

- **themed stack of system monitoring tools**
  - lots of tools…
    - *many also have associated configurable alarm conditions*
  - reads /proc

http://sfstation.members.easyspace.com/fbcpict.htm

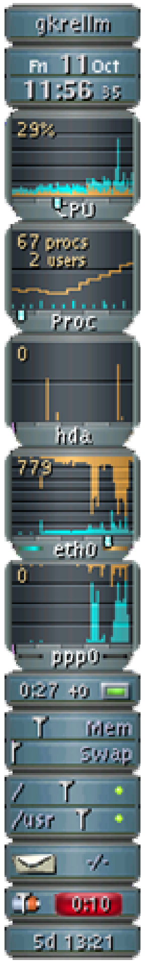

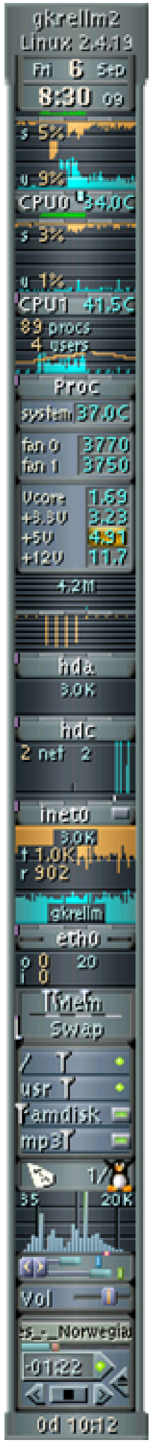

"If you have seen the movie Forbidden Planet, you might recall the Krell had a room with wall to wall meters for monitoring their power systems, and that is what I was thinking of when I came up with the GKrellM name."

Linux Control Tools

- **limited and primitive**
  - nice/renice
    - *a process's requested priority*
    - *lower gets more CPU attention*
    - *users can be 'nice' to other users and mark a process as less important by setting a high nice number*
    - *only the super user can set a low nice number to give priority to a process*

```
% nice +5 my_long_job
% renice 0 3486
```

  - swapon
    - *used to specify devices on which paging and swapping are to take place*
    - *usually executed during system boot*
  - kill and killall
  - kernel configuration
  - buy more and bigger…
    - *CPU, Disk, RAM, etc.*

Linux Limits

- **limit/ulimit**
  - csh/bash builtin command
  - can be set by administrator
    - “Now for the bad news. Current UNIX resource limits are completely useless … for several reasons. First, the hard limits are often hard-wired into the kernel and cannot be changed by the system administrator. Second, users can always change their own soft limits. All an administrator can do is place the desired commands into users’ .profile or .cshrc files and hope. Third, the limits are on a per-process basis. Unfortunately, many real jobs consist of may processes, not just one. … Finally, in many cases, limits are not even enforced; this is probably most often true of the ones you probably care about the most: CPU time and memory use.”

```
% limit -h
cputime         unlimited
filesize        unlimited
datasize        unlimited
stacksize       8192 kbytes
coredumpsize    unlimited
memoryuse       unlimited
descriptors     256
memorylocked    unlimited
maxproc         256
openfiles       256
```

```
$ ulimit -a
core file size (blocks)    1000000
data seg size (kbytes)     unlimited
file size (blocks)         unlimited
max memory size (kbytes)   unlimited
stack size (kbytes)        8192
cpu time (seconds)         unlimited
max user processes         256
pipe size (512 bytes)      8
open files                 256
virtual memory (kbytes)    2105343
```

Linux /proc Filesystem

"The /proc filesystem is a direct reflection of the system kept in memory and represented in a hierarchal manner."

- A relatively recent introduction
- Virtual filesystem
- Provides dynamic information about the system in an easily accessible manner instead of having to invoke difficult to understand system calls
  - Readable *and* writeable
    - *Show and change system-level information*

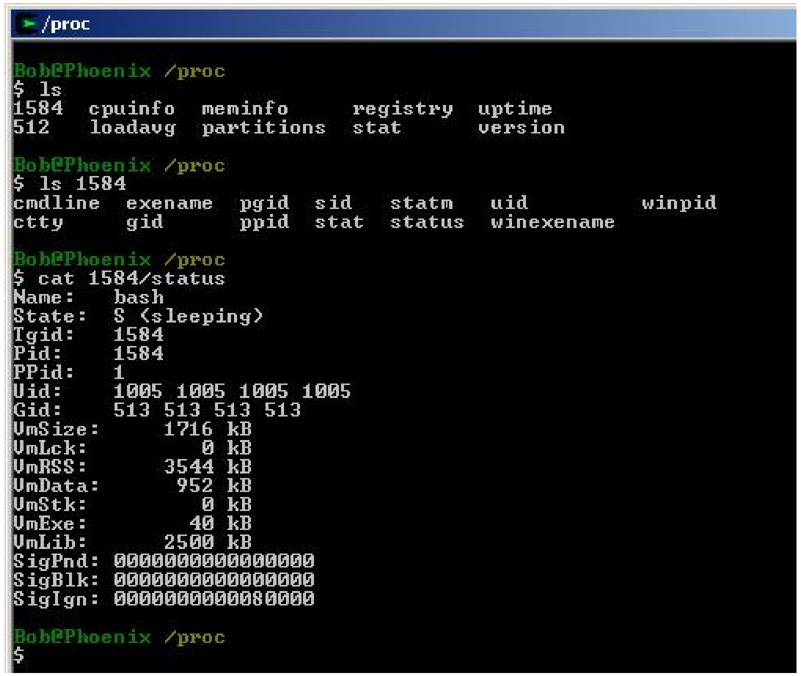

```
Bob@Phoenix /proc
$ ls
1584  cpuinfo  meminfo     registry  uptime
512   loadavg  partitions  stat      version

Bob@Phoenix /proc
$ ls 1584
cmdline  exename  pgid  sid   statm   uid         winpid
ctty     gid      ppid  stat  status  winexename

Bob@Phoenix /proc
$ cat 1584/status
Name:   bash
State:  S (sleeping)
Tgid:   1584
Pid:    1584
PPid:   1
Uid:    1005 1005 1005 1005
Gid:    513 513 513 513
VmSize:     1716 kB
VmLck:         0 kB
VmRSS:      3544 kB
VmData:      952 kB
VmStk:         0 kB
VmExe:        40 kB
VmLib:      2500 kB
SigPnd: 0000000000000000
SigBlk: 0000000000000000
SigIgn: 0000000000080000

Bob@Phoenix /proc
$
```

```
# increase the system limit on open files…
echo 32768 > /proc/sys/fs/file-max
```

Linux RPM

- **RedHat Package Manager**
  - manages the maintenance of software packages
  - a package is an archive of files, and package information, including name, version, and description.
  - ten basic modes of operation
    - *install, query, verify, check package signature, uninstall, build, rebuild database, fix permissions, set owners and groups and show rc file*
  - can perform upgrades without overwriting config files, etc.
  - can do automatic dependency following
    - *if package X requires package Y, ensure that Y is installed before installing X*
  - rpm package format allows for the inclusion of digital signatures
    - *ensure that a package comes from a trusted source and hasn't been tampered with*
  - can install across an ftp link from the internet
    - *if package source is given as an ftp URL*

"RPM emulates the local council; it always tells you why you can't load a package."

Linux RPM Examples

```
# rpm -qip which-1.0-8.i386.rpm
Name       : which              Distribution : Manhattan
Version    : 1.0                      Vendor : Red Hat Software
Release    : 8                    Build Date : Tue Apr 28 02:59:13 1998
Install date                : (not installed)  Build Host    : porky.redhat.com
Group      : Utilities/File       Source RPM : which-1.0-8.src.rpm
Size       : 7227                    License : distributable
Packager   : Red Hat Software <bugs@redhat.com>
Summary    : Finds a program 'which' is in one of the directories on your PATH
Description:
Give it a program name, and it tells you if it is on your 'PATH'.

For example, 'which ls' would print '/bin/ls', because the ls program,
which is in one of the directories listed in your PATH environment
variable, is located in the /bin directory.
```

```
# rpm -ivh dump-0.3-13.i386.rpm
dump                        ##################################################
```

```
# rpm -qf /usr/bin/which
which-1.0-8
```

```
# rpm -qlp which-1.0-8.i386.rpm
/usr/bin/which
/usr/doc/which-1.0
/usr/doc/which-1.0/Makefile
/usr/doc/which-1.0/blah
/usr/doc/which-1.0/blah/Makefile
/usr/doc/which-1.0/which.c
/usr/man/man1/which.1
```

Linux up2date

- **RedHat's simplified system for maintenance**
  - linked into, and requires registration on, the RedHat Network
    - *the marketing droids are beginning to stir, methinks!*
  - both graphical and command-line tools available

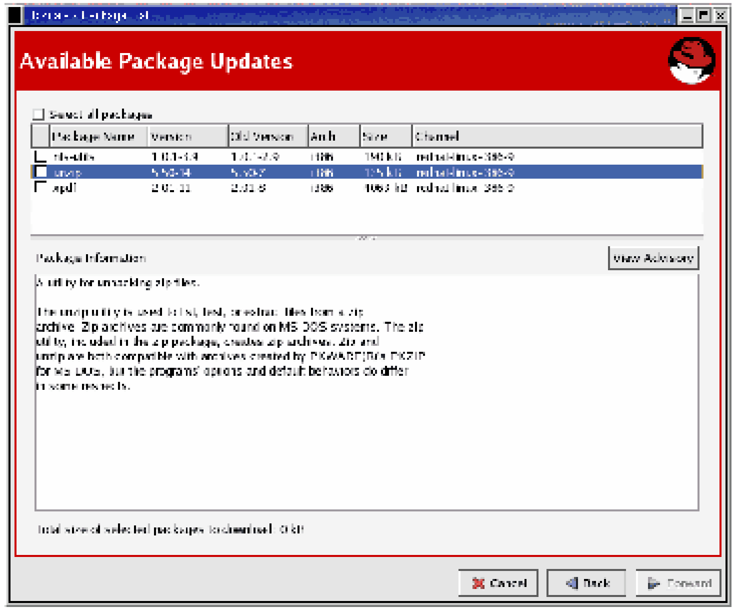

Available Package Updates

□ Select all packages

	Package Name	Version	Old Version	Arch	Size	Channel
□	[illegible]	1.0.1-4.4	1.0.1-2.9	i386	190 kB	redhat-linux-i386-9
□	unzip	5.50-34	5.50-7	i386	125 kB	redhat-linux-i386-9
□	xpdf	2.01-11	2.01-8	i386	1063 kB	redhat-linux-i386-9

Package Information

A utility for unpacking zip files.

The unzip utility is used to list, test, or extract files from a zip archive. Zip archives are commonly found on MS-DOS systems. The zip utility, included in the zip package, creates zip archives. Zip and unzip are both compatible with archives created by PKWARE(R)'s PKZIP for MS-DOS, but the programs' options and default behaviors do differ in some respects.

Total size of selected packages to download: 0 kB

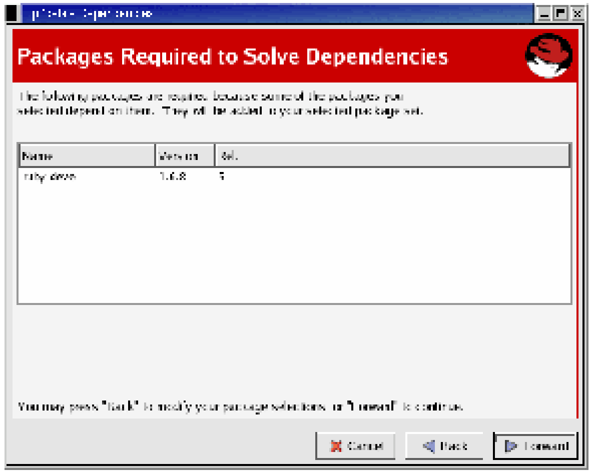

Packages Required to Solve Dependencies

The following packages are required because some of the packages you selected depend on them. They will be added to your selected package set.

Name	Version	Rel.
ruby-devel	1.6.8	9

You may press "Back" to modify your package selections or "Forward" to continue.

rdist

- **maintains identical copies of files over multiple hosts**
  - very useful for updating system configuration files
    - *can be used to distribute updated programs (and anything else...)*
  - uses rsh to make connections to remote hosts
  - tasks are driven via a 'distfile'
    - *something like a 'makefile'*
    - *provides a rich set of configuration options*
      - update iff newer, iff binary comparison fails, etc.
      - send an email notification after doing something, log to syslog, etc.
      - maintain exception lists
      - do post-installation processing
      - etc.

Linux rdist Example

```
# distfile
HOSTS = ( localhost )

FILES = ( /home/bob/distfile )

${FILES} -> ${HOSTS}
        install -ocompare /tmp/bob/distfile;

${FILES} :: /home/bob/distfile.tstamp
        notify bob@redhat ;
```

```
% rdist
/home/bob/stamp.bob: /home/bob/distfile: file is newer
/home/bob/stamp.bob: notify  ( bob@redhat )
localhost: updating host localhost
localhost: redhat: /tmp/bob/distfile: updated
localhost: updating of localhost finished
/home/bob/stamp.bob: updating of /home/bob/stamp.bob finished
```

```
% mail
Mail version 8.1 6/6/93.  Type ? for help.
"/var/spool/mail/bob": 2 messages 2 new
>N  1 rdist@redhat.skewst.  Sun Dec 20 11:39  15/490   "files updated after S"
&
Message 1:
From bob  Sun Dec 20 11:39:15 1998
Date: Sun, 20 Dec 1998 11:39:14 +1000
From: rdist@redhat.skewst.home.net.au (Remote distribution program)
To: bob@redhat.skewst.home.net.au
Subject: files updated after Sun Dec 20 11:38:23 1998



/home/bob/distfile.tstamp: /home/bob/distfile: file is newer

&
```