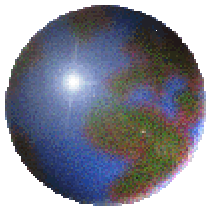

FAQ's and answers for HP-UX SysAdmins

Bill Hassell Consulting, Inc.

Agenda

- Setting up new systems
- Standardizing for easy administration
- Performance and Kernel Tuning
- Memory Management
- Disk space management
- Logfile management
- Useful tools
- Additional resources

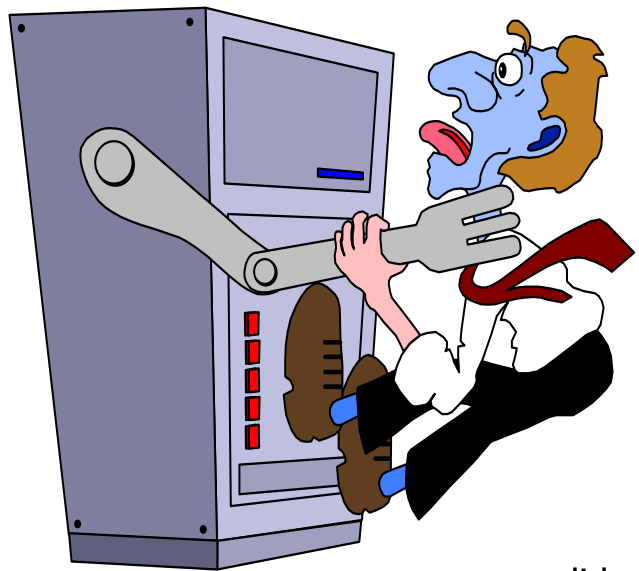

Setting up new systems

- Disk layout
  - Internal/external issues
  - VG00 design
  - Massive lvols, flat directories
- Networking
  - SAM, netconf, set_parms
  - 100BaseT issues
- Cron jobs
  - Logfiles, cleanup, backup

Setting up new systems

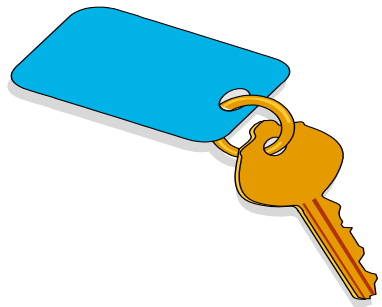

- Root's shell
  - /sbin/sh
  - POSIX /usr/bin/sh vs. Bourne
  - /usr/bin/sh vs. Korn shell
  - Users can change: chsh login-name /usr/bin/some_shell
- Login tty settings:
  - # is erase 1 char (backspace)
  - Fix all logins with a one-time command:

```
stty erase '^H' < /dev/ttyconf
```

  - Add others such as:

```
stty erase "^H" kill "^U" intr "^C" eof "^D" -parity ixoff
```

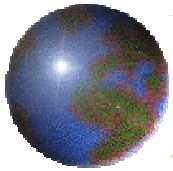

Setting up new systems

- Basic security fixes
  - `find /usr/local -type d -print -exec chmod 755 {} \;`
  - umask is missing in /etc/profile and /etc/csh.login
    - `umask 022`
    - `umask 077`
  - Search for world-write permissions in important directories:
    - `find /etc /usr /opt /dev /sbin /stand -perm -002`
    - Add: `\( -type d -o -type f \)` to look at just files and directories
    - Add secure database mountpoints, etc

Setting up new systems

- Basic security fixes (continued)
  - `/etc/path`
    - World-writable directories
    - `pwd` (`::` or `:.:` or `:.`)
    - Symlinks
    - Duplicate paths
    - Non-existent directories
  - Password and group file checks
    - `pwck`
    - `grpck`

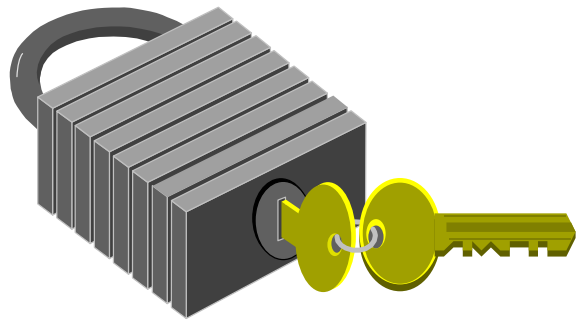

Setting up new systems

- Basic security fixes (continued)
  - `/etc/hosts.equiv .rhosts`
    - World-writable?
    - Valid user?
    - Root?
  - suid scripts/programs
    - No scripts allowed
    - Purpose
    - Owner
    - Location
    - `/etc/fstab`: `nosuid`

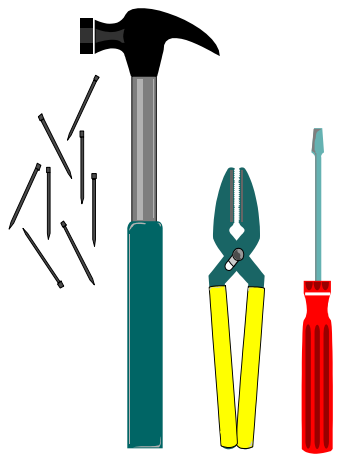

Standardizing for easy administration

- Login methods
  - telnet
  - remsh/rlogin
  - Xwindows
- Shells
  - POSIX (/usr/bin/sh)
  - Bourne (/usr/old/bin/sh)
  - Ksh (/use/bin/ksh)
  - Csh
  - Bash, tcsh, etc

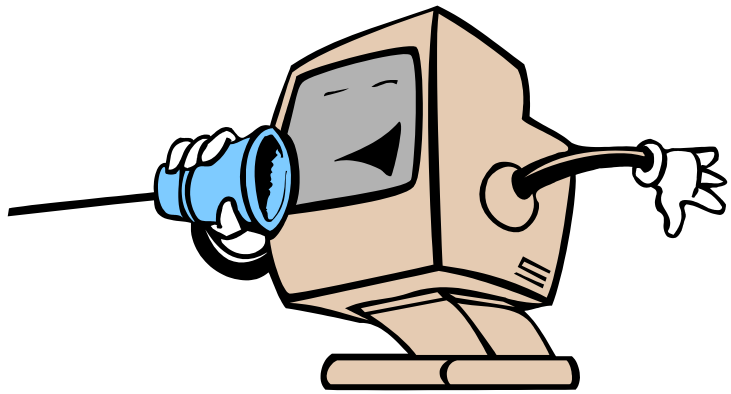

Standardizing for easy administration

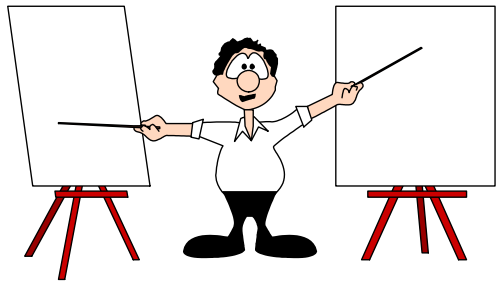

- Standard profiles
  - `set -u`
  - ulimit settings:

```
ulimit -a
time(seconds)        unlimited
file(blocks)         unlimited
data(kbytes)         65536
stack(kbytes)        8192
memory(kbytes)       unlimited
coredump(blocks)     0
nofiles(descriptors) 60
```

  - Eliminate core files: `ulimit -Sc 0`

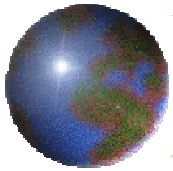

Standardizing for easy administration

- Standard profiles
  - Interactive versus batch
    - Errors: `stty: not a typewriter`
    - Separate the interactive commands:
      - Multiple `if interactive` statements
      - Or separate the interactive portion in a separate script
    - Interactive commands: `stty tput ttytype`
    - Interactive settings: `$TZ $DISPLAY`
  - Login controls: `/etc/nologin`
    - 11.0 with security patches, standard on 11i
    - Add: `NOLOGIN=1` to /etc/default/security

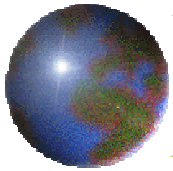

Standardizing for easy administration

- Standard profiles
  - `$HOME/.profile`
  - `/etc/skel` to standardize new users, or replace
  - Can't keep from being changed by users (`$HOME` must be writable which means no file is safe)
  - Same interactive issues with batch logins
  - May want to have `$HOME/.localprofile` for user mods
- Download sample profiles:

ftp://contrib:9unsupp8@hprc.external.hp.com/sysadmin/profiles

Standardizing for easy administration

- Other /etc/skel files:
  - `.exrc` (vi defaults)
    - `nows` (don't wrap searches around)
    - `ic` (ignore case on searches)
    - `ai` (follow prev left indent)
    - `aw` (update source at shell escapes: !)
    - `wm` (wrap to next line, =0 no wrap)
    - `report` (shows # of lines changed)
    - `showmode` (INPUT or REPLACE MODE tag)

```
set nows ic autoindent autowrite wrapmargin=2
  …
```

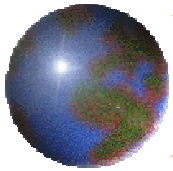

Standardizing for easy administration

- Restricted access methods
  - Restricted shells: rsh, rksh, rcsh
    - $HOME is / for user
    - NO access to /usr or any other directrories
    - Must create a local bin with binaries and set $PATH
- Backup procedures
  - Classic tools: tar cpio pax dump
    - No index, no search, no error recovery, no changer support, no largefiles
  - fbackup/frecover
    - Index on every tape, high speed search, error recovery, change hooks, largefile-capable

Standardizing for easy administration

- Backup procedures (cont'd)
  - Features:
    - Largefiles, error recovery, changer support, parallel tape drives, network backup/restore, multi-platform, centralized indexing
  - Commercial backup programs
    - HiBack
    - OmniBack
    - Veritas
- Disaster recovery (loss of boot disk)
  - Mirroring
  - Ignite/UX

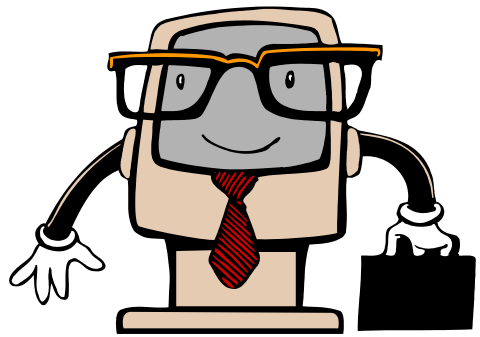

Standardizing for easy administration

- `print_manifest` (from Ignite/UX)
- from the contributed software archive:
  `ftp://contrib:9unsupp8@hprc.external.hp.com`
  - `gatherer.sh`
  - `nightowl`
  - `nickel` (HTML generator)

- Webmin: `www.webmin.com` (Perl-based)
  - Included with HP's Apache

Standardizing for easy administration

- `print_manifest` (from Ignite/UX)

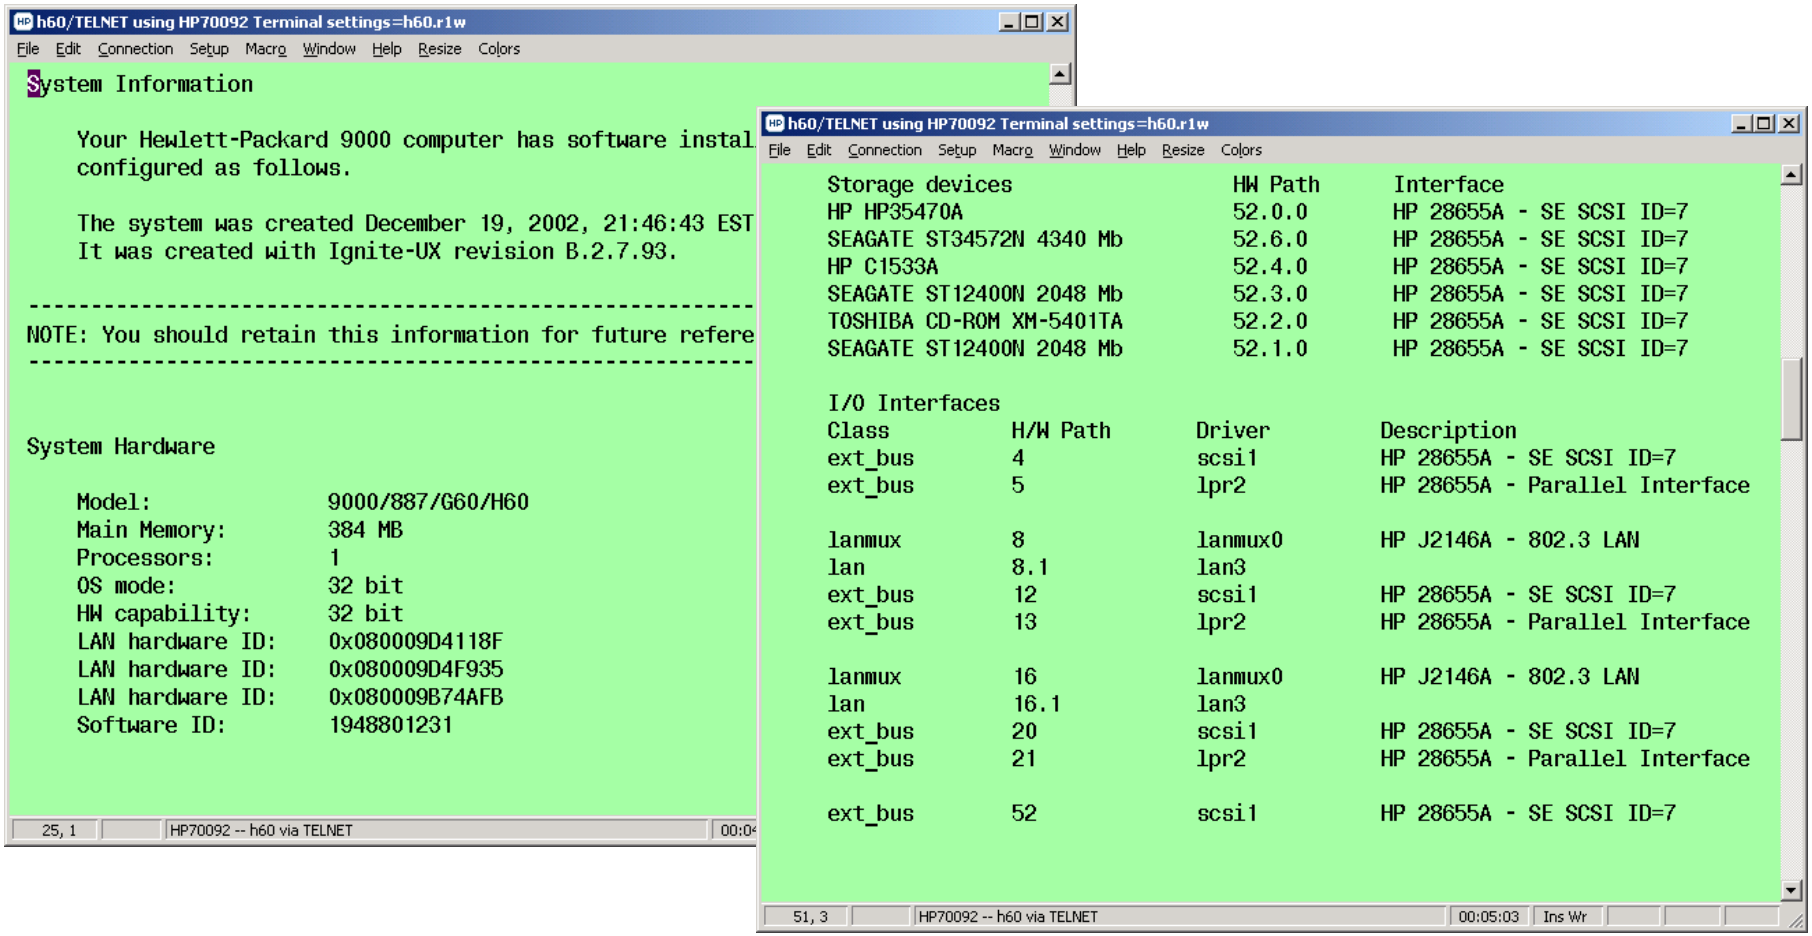

Standardizing for easy administration

- Centralized patching
  - Pick a central server
  - Create a hierarchical structure (test, pre-production, production)
  - Use swcopy to add patches and patch depots
  - Use swreg to make the depot visible on the network (don't use NFS)
  - On clients, use swlist to view the remote depot(s)

    swlist –l depot @ patch_servername
  - Network installs:

    swinstall –s depot_name@patch_servername file_set
    (or *****)

Standardizing for easy administration

- Spooler management
  - Connections
    - Parallel, serial, SCSI, Network
      - Remote (Windows, Linux, solaris, AIX)
        - RFC 1179
        - Control files and options
        - Printer scripts and filters
      - Network (HP JetDirect)
        - Port 9100
        - Jetadmin now HP Printer Installer
        - addqueue, removequeue, transferqueue
  - Troubleshooting
    - Spoolkick procedure

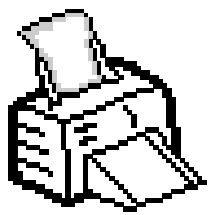

Performance and Kernel Tuning

- What can be changed
- Performance measurement
  - Built-in tools
  - SarCheck
  - MetaView, Performance Gallery
  - Glance/gpm, Measureware
- Kernel parameters
  - Filesystems
  - Processes
  - RAM and virtual memory
  - Network

Performance and Kernel Tuning

- What can be changed
  - Most kernel params size tables, provide limits or set behavior
  - CPU bound
    - Multi-CPU features
  - I/O bound
    - Disk
    - Swap (memory limited)
    - LAN

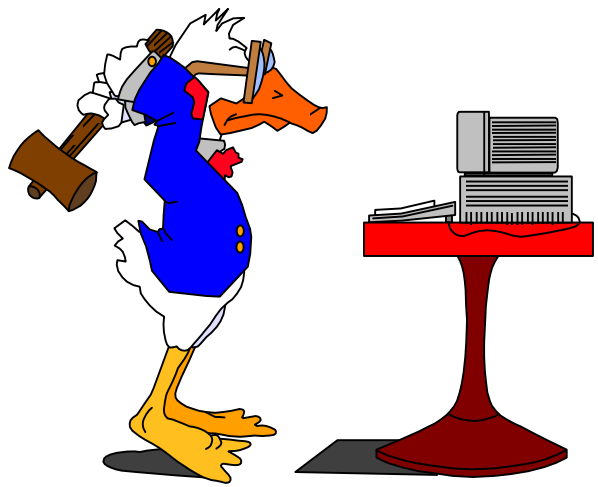

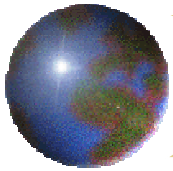

Performance and Kernel Tuning

Filesystem parameters

- **nfile:** 14min, no limit
  - Every opened file (including multiple opens)
  - Formula scales with maxusers
  - Every process has a minimum of 3
- **nflocks**: 2min, 200def, no limit
  - Maximum number of open file-locks
  - One file may have several locks
  - Application dependent
  - Databases may need hundreds

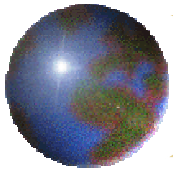

Performance and Kernel Tuning

Filesystem parameters

- **ninode:** 14min, no limit
  - In-memory cache of ***unique*** current and recent HFS file locations
  - Speeds re-open, multi-process file access
  - Indirectly controls the size of the DNLC, ncsize, ncdnode and vx_ninode.
  - 1000 to 8192, lower if ncsize is adjusted
  - Formula in SAM is not useful for large systems
  - Best recommendations are in NFS Performance book by Dave Olker

Performance and Kernel Tuning

Filesystem parameters

- **ncdnode:** 14min, no limit
  - In-memory cache of ***unique*** current and recent CDFS file locations (CDROM)
  - Speeds re-open, multi-process file access
  - Usually OK unless multiple user acce needed for CDROM or multiple CDROM drives

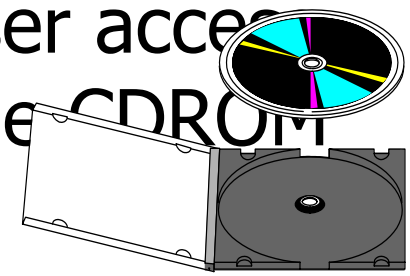

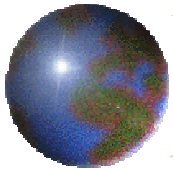

Performance and Kernel Tuning

Filesystem parameters

- **maxfiles:** 60 default
  - Maximum number of files opened by a single process
  - Used to control runaway processes
  - Override with setrlimit(2) system call or **ulimit –n**
  - Hard ceiling is **maxfiles_lim**
  - Commonly recommended too high (4096) by database vendors
  - Commonly recommended **maxfiles** = **maxfiles_lim** but not a good idea (use ulimit as needed)

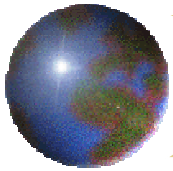

Performance and Kernel Tuning

Filesystem parameters

- **fs_async:**
  - 0 = synchronous writes to directory structures
  - 1 = async
    - Apx. 20-30% faster write speed (no read change)
    - Very high probability of data loss with a powerfail or system panic (fsck fails to fix)
- **Default_disk_ir:**
  - 0 = no immediate reporting (waits for writes to complete
  - 1 = immediate reporting (disk buffers writes)
    - Applies to all disks including raw devices

Performance and Kernel Tuning

Filesystem parameters

- **disksort_seconds:**
  - HP-UX gives priority to serial rd/wt queues
  - Intense serial I/O slows random I/O
  - Value in seconds to wait before changing priority

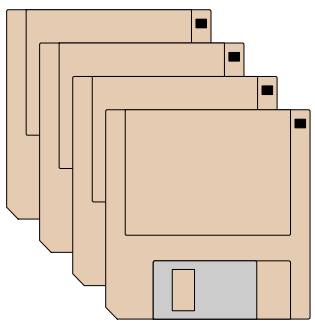
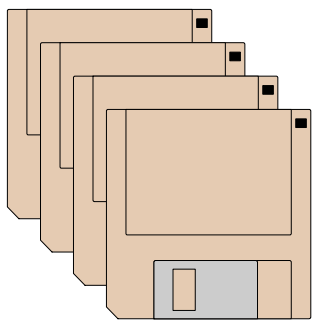
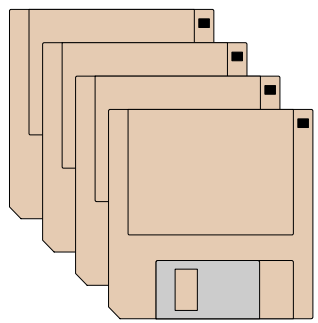

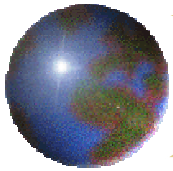

Performance and Kernel Tuning

Process parameters

- **nproc:**
  - Maximum processes to run at same time
  - Often way too low in servers
- **maxuprc:**
  - Maximum processes owned by a single UID
  - Collective, not per login
  - Generic logins can require hundreds or more

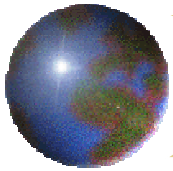

Performance and Kernel Tuning

Process parameters

- **maxdsiz:** 64 megs default
  - 32bit programs only
  - Max can be 2Gb but mapping limits are 960 and 1750 depending on compiler (pr chatr) options.
  - **ulimit –d** to create a lower limit
- **maxdsiz_64:** 64 megs default
  - 64bit programs only
  - Max is 4000 Gb

Mem_mgt + proc_mgt note (/usr/share/doc – not 11i)

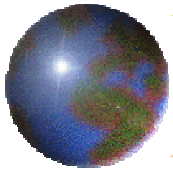

Performance and Kernel Tuning

Process parameters

- **maxssiz:** 8 megs default
  - 32bit programs only
  - Seldom needs changing
  - Exception for specialized programs or poor design (pass by data not address/pointer)
- **maxtsiz:** 64megs default
  - Maximum for unchanging instructions
  - Directly related to executable's file size
  - 2Gb max (32bit) or 4096Gb (64bit)

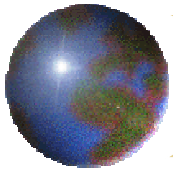

Performance and Kernel Tuning

Process parameters

- **maxswapchunks:**
  - Used to size maximum swap (32Gb)
  - Useable swap is:
    - maxswapchunks * swchunk * dev_bsize where: swchunk=2048 and dev_bsize=1024
    - Leave swchunk as default
    - Formula simplifies to: maxswapchunks=DESIRED-SWAP / 2097152
- **nswapdev, nswapfs:**
  - Maximum swap devices and filesystems

Performance and Kernel Tuning

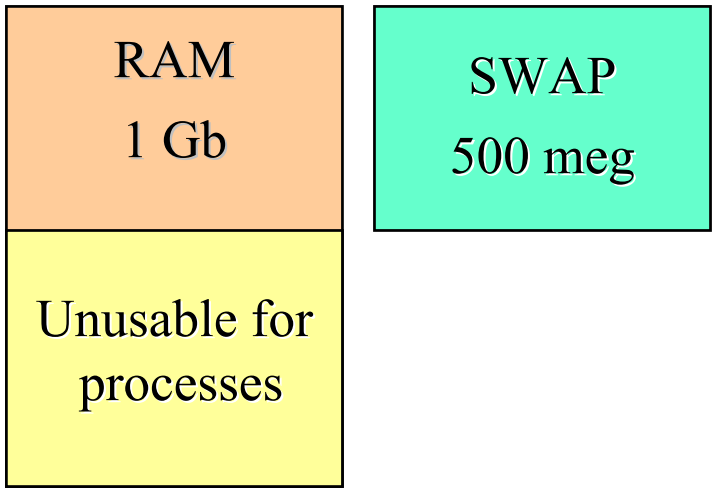

swapmem_on=0

In this example, only 500 megs is usable for processes since Virtual Memory is only 500 megs.

Performance and Kernel Tuning

swapmem_on = 0

In this example, 1 Gb is usable for processes since Virtual Memory is 1 Gb too...but no paging will take place

Performance and Kernel Tuning

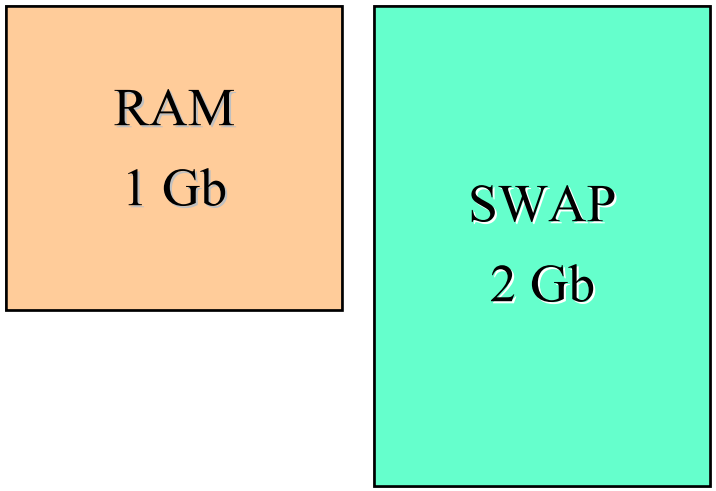

swapmem_on = 0

In this example, 2 Gb is usable for processes since Virtual Memory is 2 Gb too...some paging may be take place

Performance and Kernel Tuning

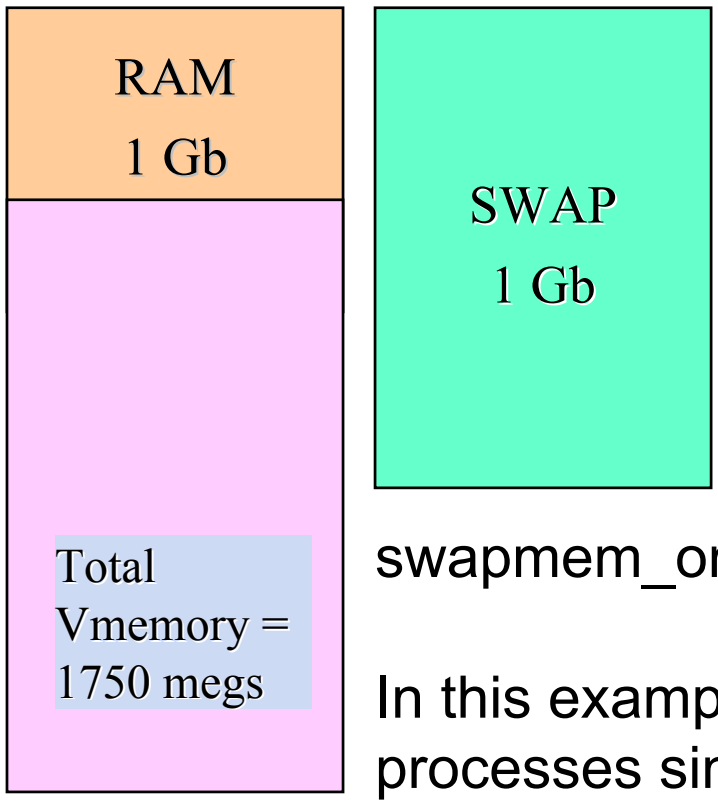

swapmem_on = 1

In this example, 1.75 Gb is usable for processes since Virtual Memory is 1 Gb plus .75 Gb in RAM which is not pageable

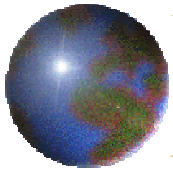

Performance and Kernel Tuning

Process parameters

- **timezone, dst:**
  - Default value for environments without $TZ (daemons)
  - **dst** controls a limited number of daylight saving rules
- **npty nstrpty,nstrtel:**
  - Controls maximum interactive network sessions
  - Must match device files (insf)
  - Use SAM to handle device files automatically
- **maxusers:**
  - Not a parameter but a formula adjustment
  - No relation to user licenses or session limits

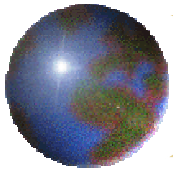

Performance and Kernel Tuning

Performance tuning

- Vocabulary
  - I/O Bound
  - Disk Thrashing
  - Swap Thrashing
  - Resource Limits

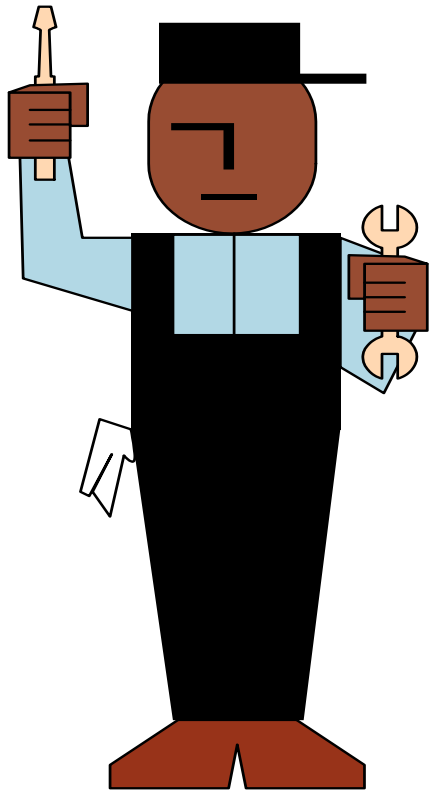

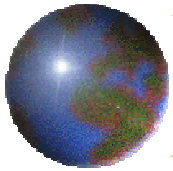

Performance and Kernel Tuning

Performance tuning

- Measurement:
  - **uptime**:
    - User count (standard logins)
    - Load average (**runqueue**)

```
uptime
     6:37pm  up 81 days, 23:43,  103 users,
     load average: 0.07, 0.08, 0.07

  uptime
     8:11pm  up 181 days, 23:43,  907 users,
     load average: 43.07, 22.08, 16.07
```

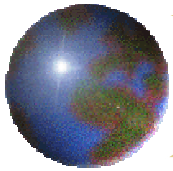

Performance and Kernel Tuning

Performance tuning

- **sar** (system activity reporter)
  - **-b** buffer cache activity
  - **-c** System Calls
  - **-d** Block Device Activity
  - **-u** CPU Utilization

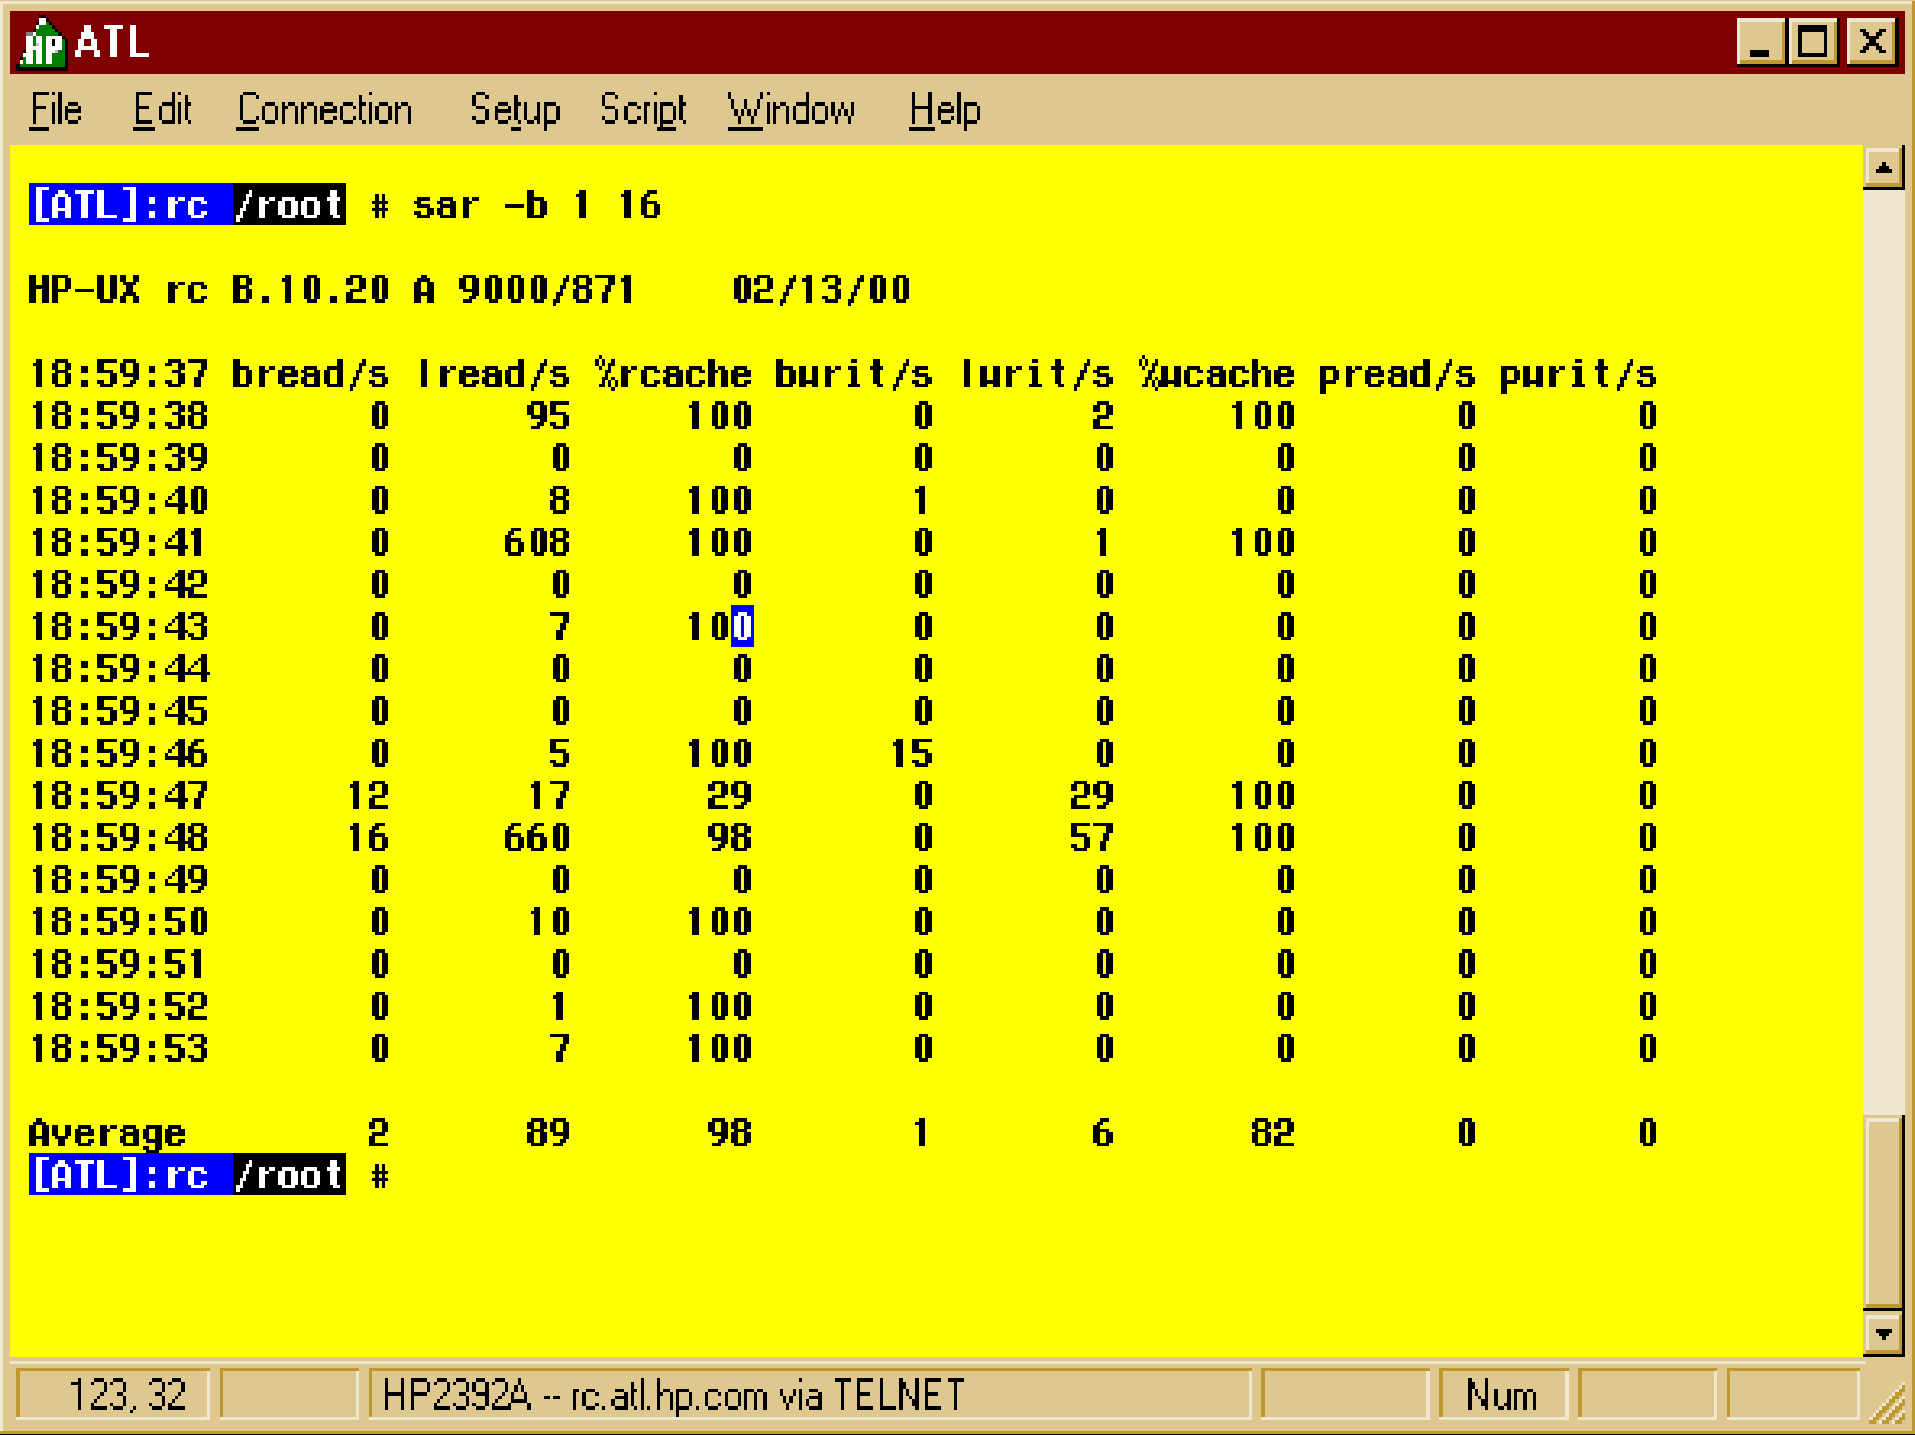
ATL
File Edit Connection Setup Script Window Help
[ATL]:rc /root # sar -b 1 16
HP-UX rc B.10.20 A 9000/871 02/13/00
18:59:37 bread/s lread/s %rcache bwrit/s lwrit/s %wcache pread/s pwrit/s
18:59:38 0 95 100 0 2 100 0 0
18:59:39 0 0 0 0 0 0 0 0
18:59:40 0 8 100 1 0 0 0 0
18:59:41 0 608 100 0 1 100 0 0
18:59:42 0 0 0 0 0 0 0 0
18:59:43 0 7 100 0 0 0 0 0
18:59:44 0 0 0 0 0 0 0 0
18:59:45 0 0 0 0 0 0 0 0
18:59:46 0 5 100 15 0 0 0 0
18:59:47 12 17 29 0 29 100 0 0
18:59:48 16 660 98 0 57 100 0 0
18:59:49 0 0 0 0 0 0 0 0
18:59:50 0 10 100 0 0 0 0 0
18:59:51 0 0 0 0 0 0 0 0
18:59:52 0 1 100 0 0 0 0 0
18:59:53 0 7 100 0 0 0 0 0
Average 2 89 98 1 6 82 0 0
[ATL]:rc /root #
123, 32
HP2392A -- rc.atl.hp.com via TELNET
Num

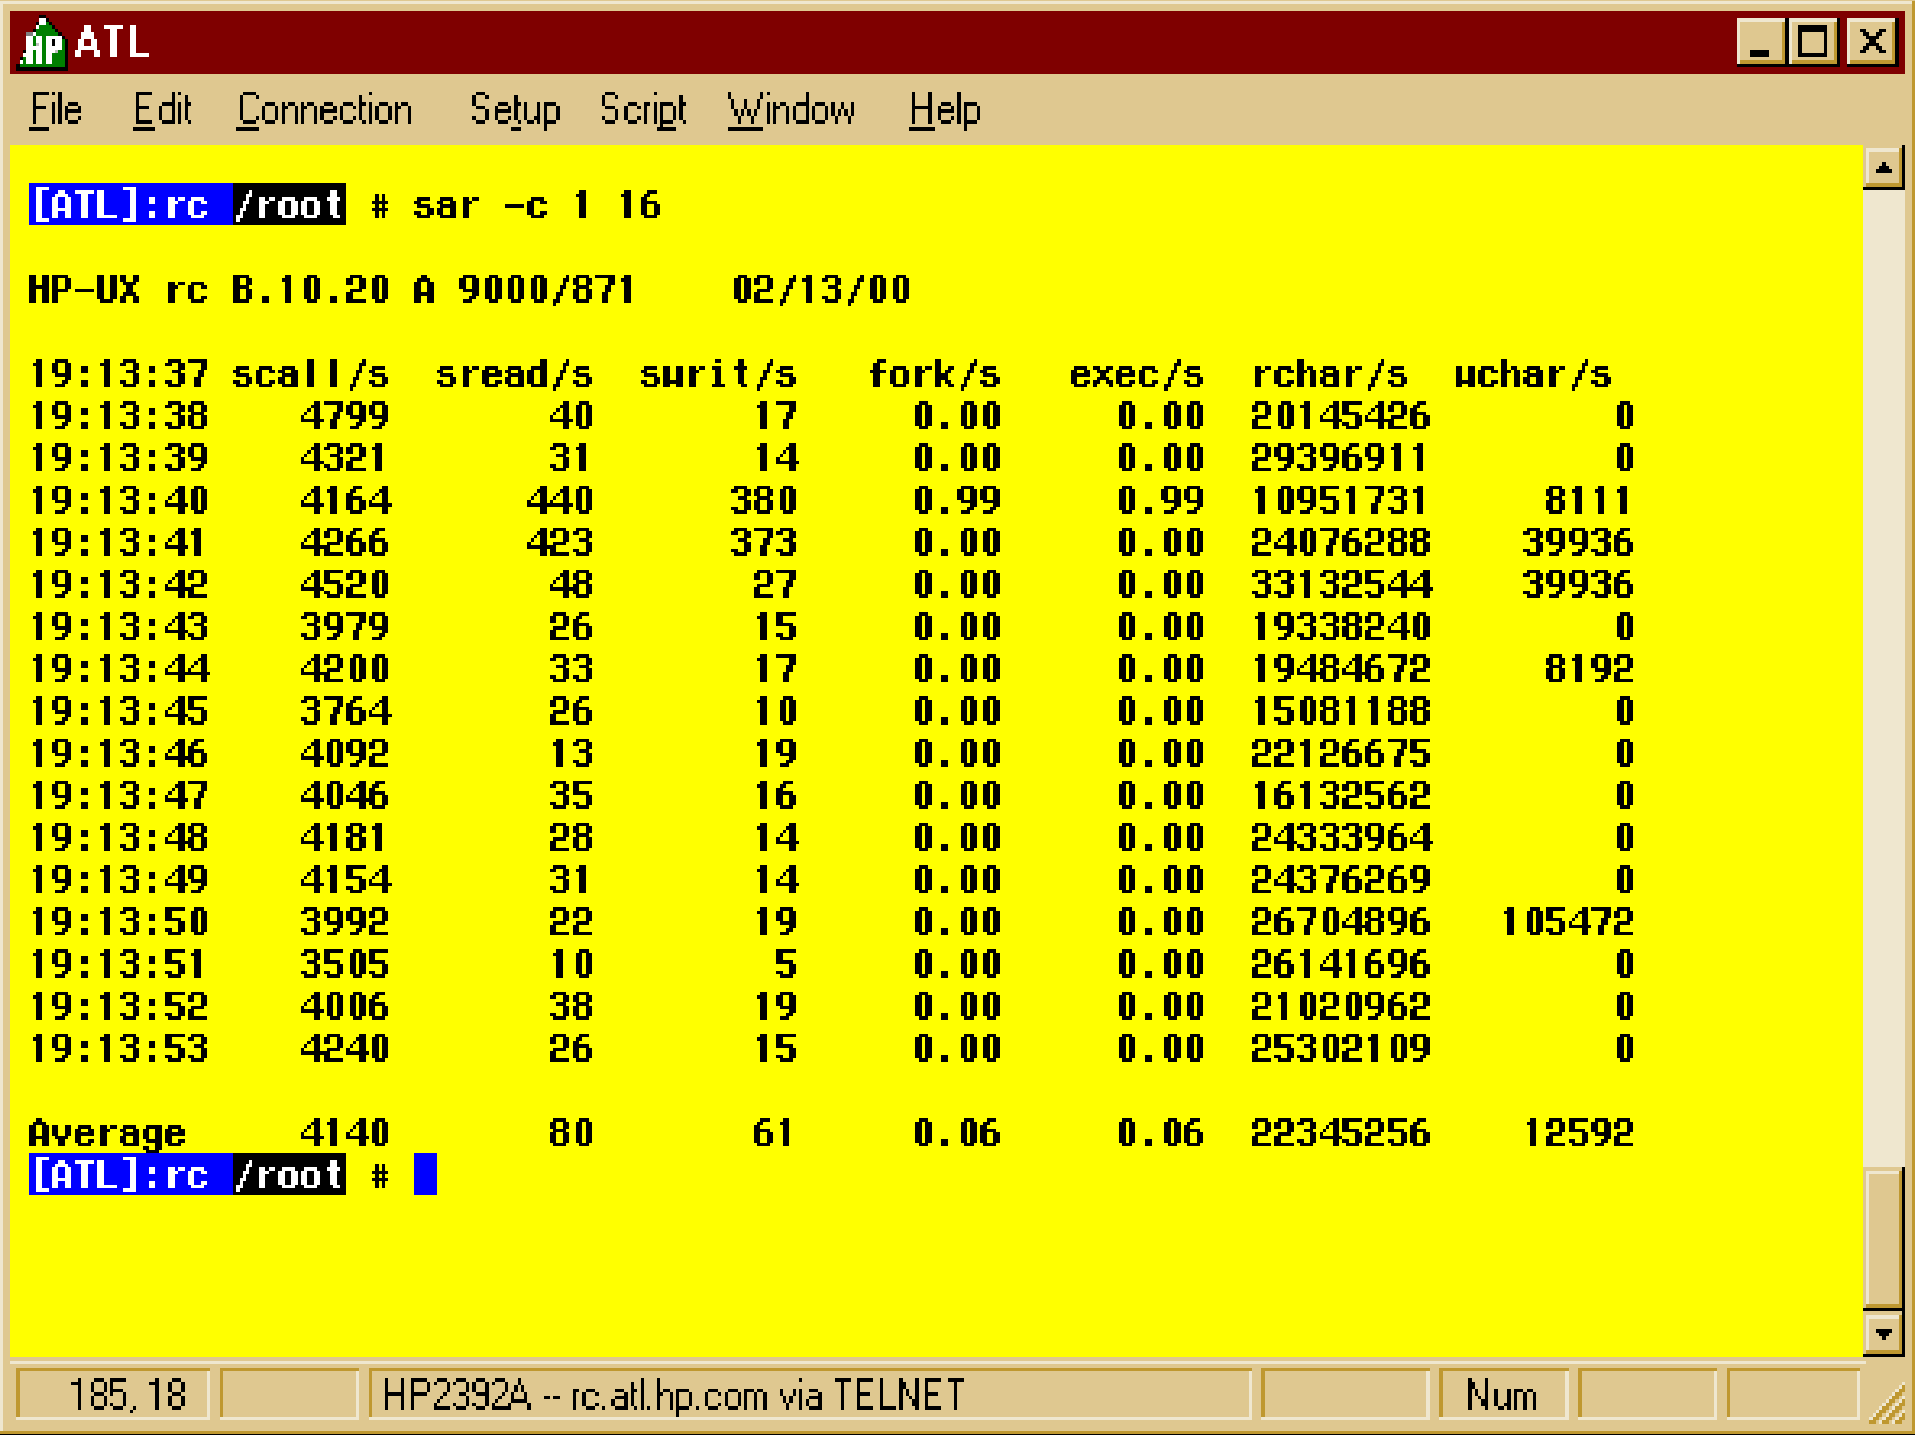
ATL
File Edit Connection Setup Script Window Help
[ATL]:rc /root # sar -c 1 16
HP-UX rc B.10.20 A 9000/871 02/13/00
19:13:37 scall/s sread/s swrit/s fork/s exec/s rchar/s wchar/s
19:13:38 4799 40 17 0.00 0.00 20145426 0
19:13:39 4321 31 14 0.00 0.00 29396911 0
19:13:40 4164 440 380 0.99 0.99 10951731 8111
19:13:41 4266 423 373 0.00 0.00 24076288 39936
19:13:42 4520 48 27 0.00 0.00 33132544 39936
19:13:43 3979 26 15 0.00 0.00 19338240 0
19:13:44 4200 33 17 0.00 0.00 19484672 8192
19:13:45 3764 26 10 0.00 0.00 15081188 0
19:13:46 4092 13 19 0.00 0.00 22126675 0
19:13:47 4046 35 16 0.00 0.00 16132562 0
19:13:48 4181 28 14 0.00 0.00 24333964 0
19:13:49 4154 31 14 0.00 0.00 24376269 0
19:13:50 3992 22 19 0.00 0.00 26704896 105472
19:13:51 3505 10 5 0.00 0.00 26141696 0
19:13:52 4006 38 19 0.00 0.00 21020962 0
19:13:53 4240 26 15 0.00 0.00 25302109 0
Average 4140 80 61 0.06 0.06 22345256 12592
[ATL]:rc /root #
185, 18
HP2392A -- rc.atl.hp.com via TELNET
Num

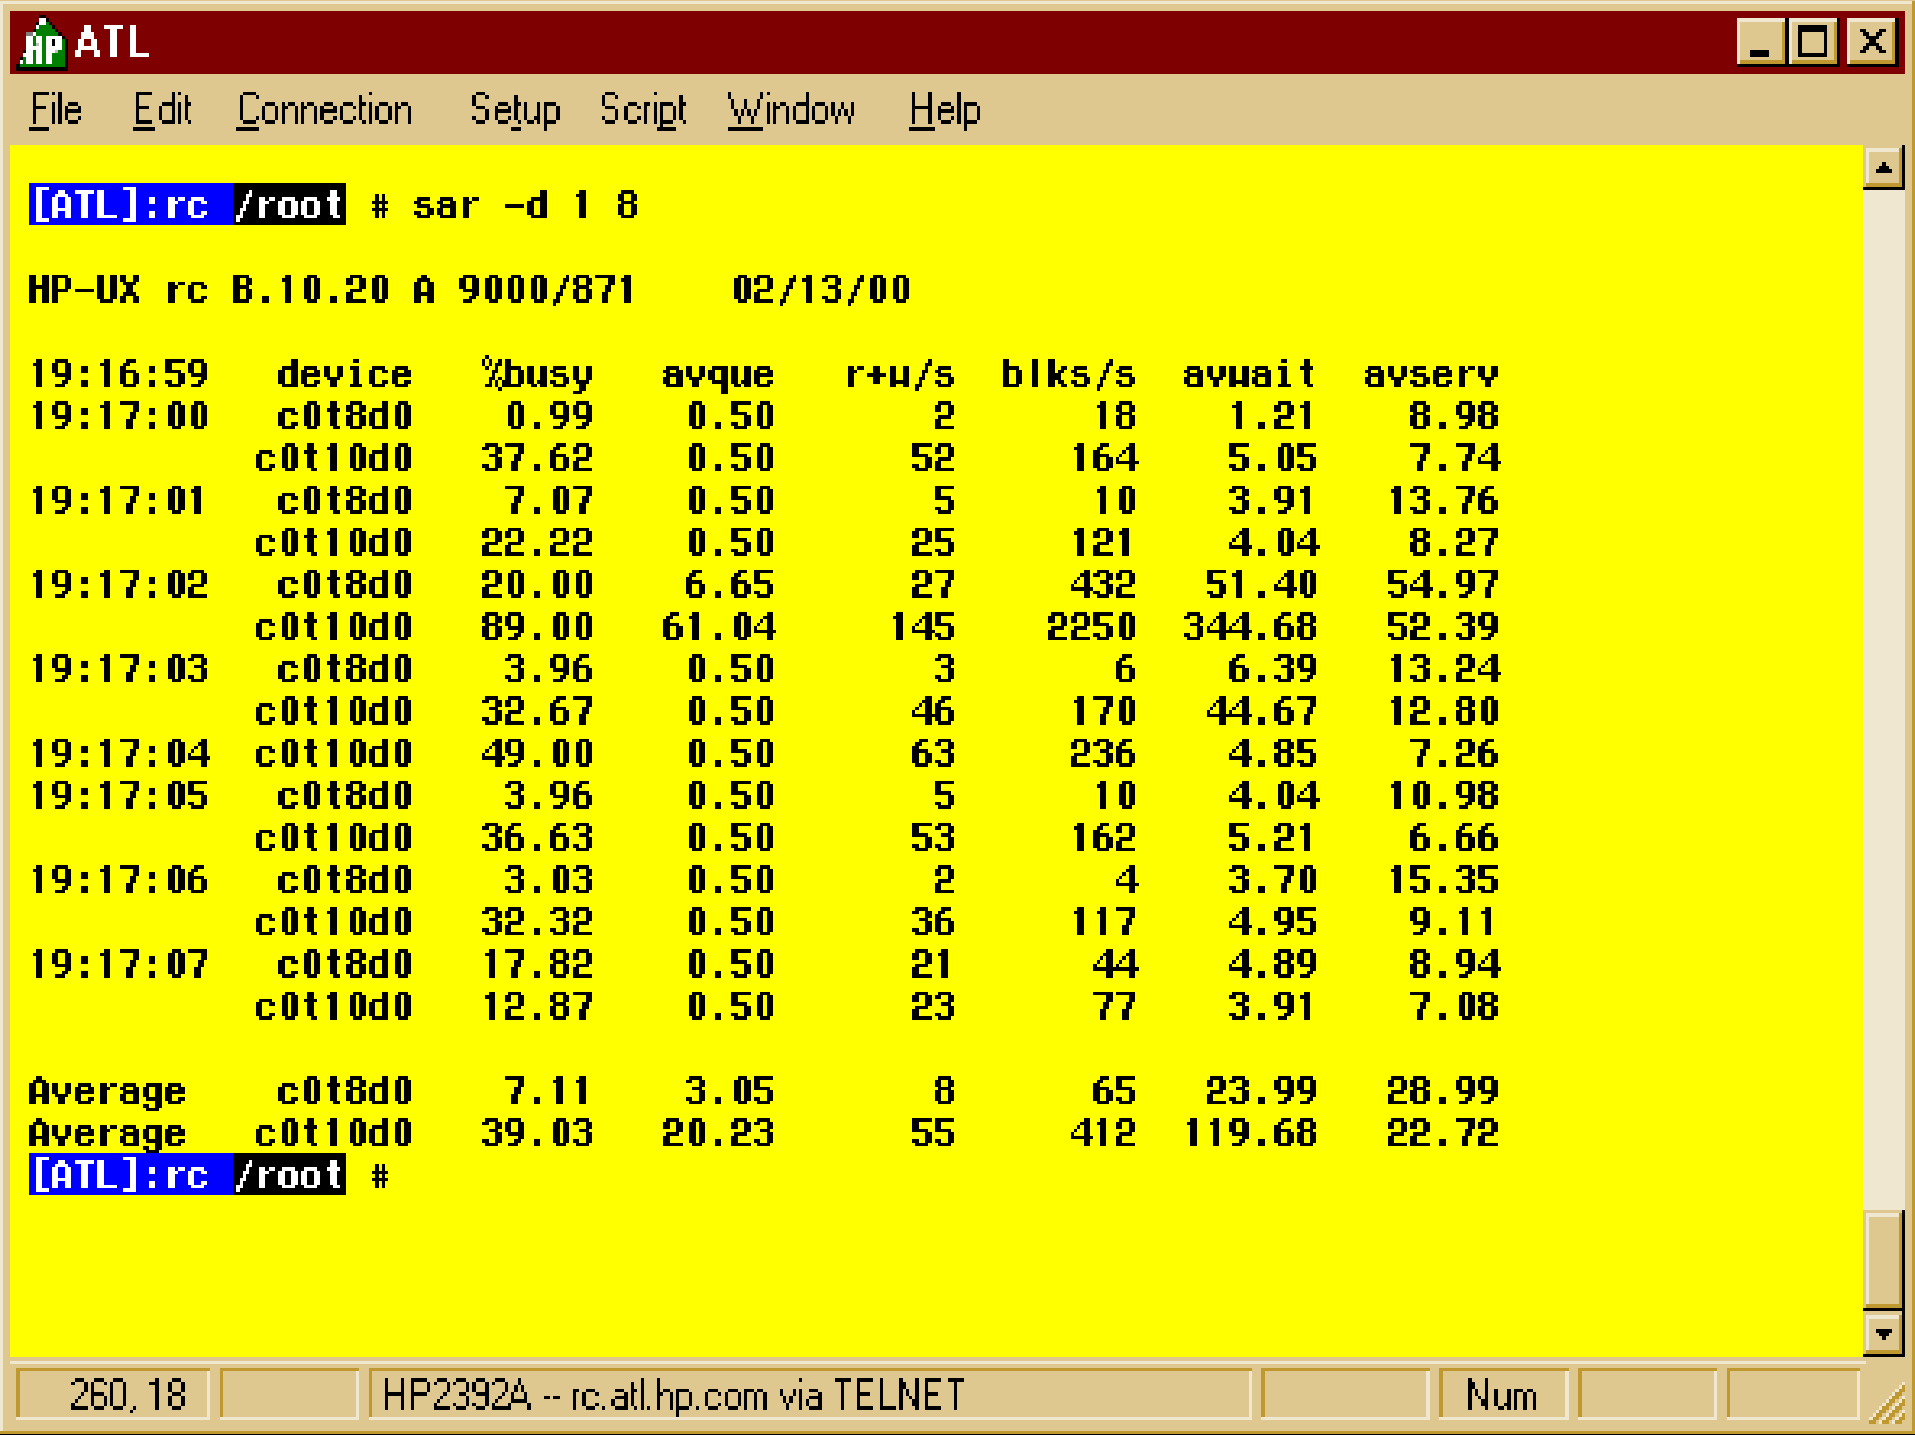
ATL
File Edit Connection Setup Script Window Help
[ATL]:rc /root # sar -d 1 8
HP-UX rc B.10.20 A 9000/871 02/13/00
19:16:59 device %busy avque r+w/s blks/s avwait avserv
19:17:00 c0t8d0 0.99 0.50 2 18 1.21 8.98
c0t10d0 37.62 0.50 52 164 5.05 7.74
19:17:01 c0t8d0 7.07 0.50 5 10 3.91 13.76
c0t10d0 22.22 0.50 25 121 4.04 8.27
19:17:02 c0t8d0 20.00 6.65 27 432 51.40 54.97
c0t10d0 89.00 61.04 145 2250 344.68 52.39
19:17:03 c0t8d0 3.96 0.50 3 6 6.39 13.24
c0t10d0 32.67 0.50 46 170 44.67 12.80
19:17:04 c0t10d0 49.00 0.50 63 236 4.85 7.26
19:17:05 c0t8d0 3.96 0.50 5 10 4.04 10.98
c0t10d0 36.63 0.50 53 162 5.21 6.66
19:17:06 c0t8d0 3.03 0.50 2 4 3.70 15.35
c0t10d0 32.32 0.50 36 117 4.95 9.11
19:17:07 c0t8d0 17.82 0.50 21 44 4.89 8.94
c0t10d0 12.87 0.50 23 77 3.91 7.08
Average c0t8d0 7.11 3.05 8 65 23.99 28.99
Average c0t10d0 39.03 20.23 55 412 119.68 22.72
[ATL]:rc /root #
260, 18
HP2392A -- rc.atl.hp.com via TELNET
Num

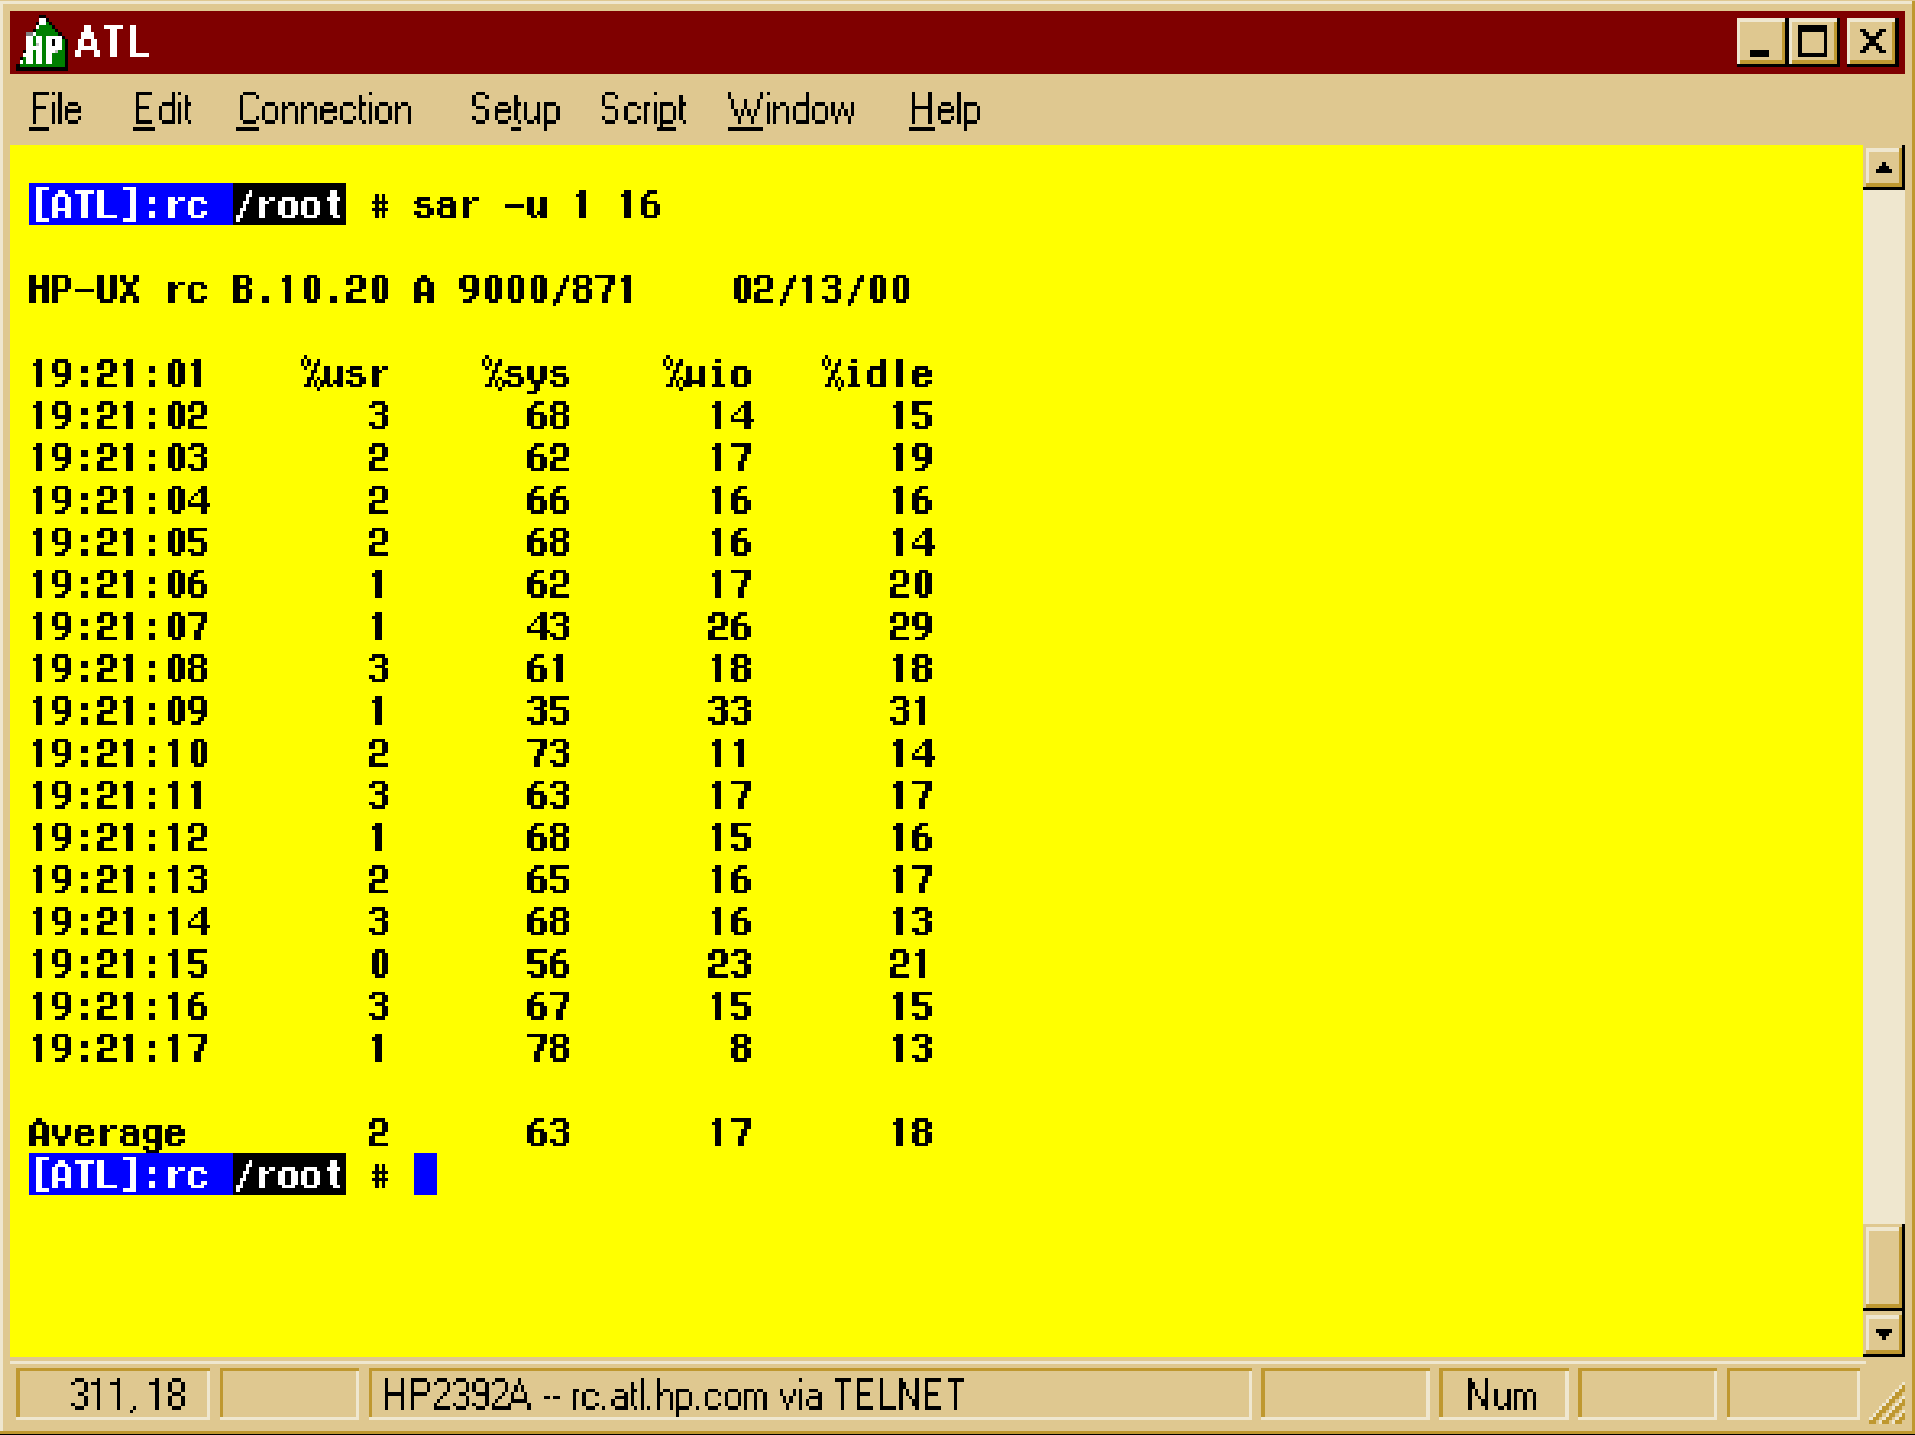
ATL
File Edit Connection Setup Script Window Help
[ATL]:rc /root # sar -u 1 16
HP-UX rc B.10.20 A 9000/871 02/13/00
19:21:01 %usr %sys %wio %idle
19:21:02 3 68 14 15
19:21:03 2 62 17 19
19:21:04 2 66 16 16
19:21:05 2 68 16 14
19:21:06 1 62 17 20
19:21:07 1 43 26 29
19:21:08 3 61 18 18
19:21:09 1 35 33 31
19:21:10 2 73 11 14
19:21:11 3 63 17 17
19:21:12 1 68 15 16
19:21:13 2 65 16 17
19:21:14 3 68 16 13
19:21:15 0 56 23 21
19:21:16 3 67 15 15
19:21:17 1 78 8 13
Average 2 63 17 18
[ATL]:rc /root #
311, 18
HP2392A -- rc.atl.hp.com via TELNET
Num

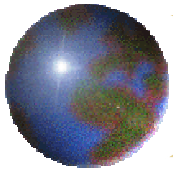

Performance and Kernel Tuning

Performance tuning

- **sar** (system activity reporter)
  - **-a** File Access
  - **-v** Kernel Parameters

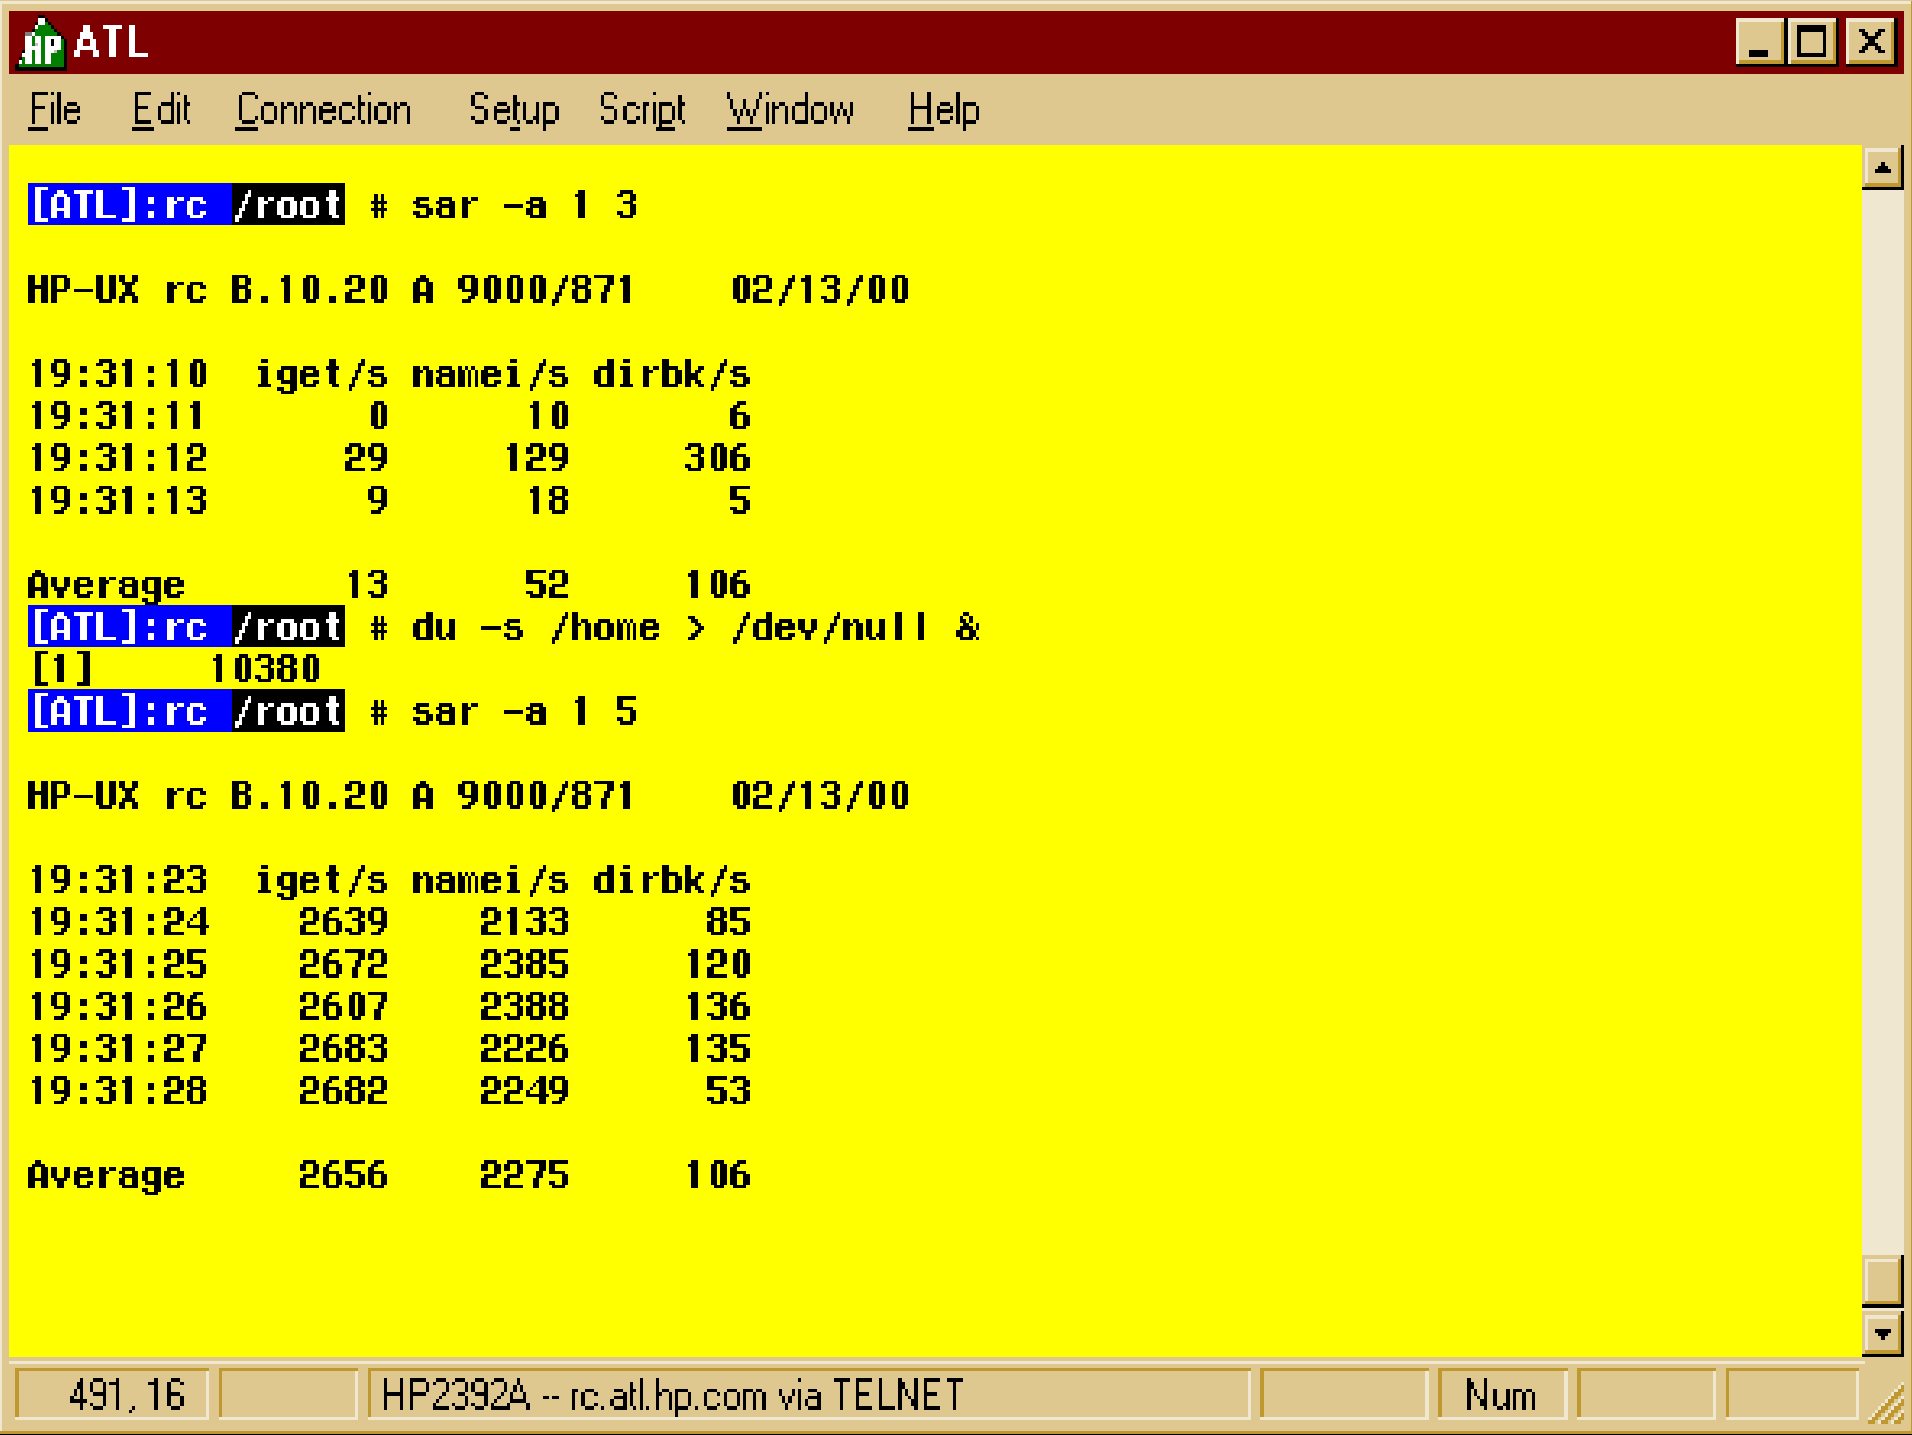
ATL
File Edit Connection Setup Script Window Help
[ATL]:rc /root # sar -a 1 3
HP-UX rc B.10.20 A 9000/871 02/13/00
19:31:10 iget/s namei/s dirbk/s
19:31:11 0 10 6
19:31:12 29 129 306
19:31:13 9 18 5
Average 13 52 106
[ATL]:rc /root # du -s /home > /dev/null &
[1] 10380
[ATL]:rc /root # sar -a 1 5
HP-UX rc B.10.20 A 9000/871 02/13/00
19:31:23 iget/s namei/s dirbk/s
19:31:24 2639 2133 85
19:31:25 2672 2385 120
19:31:26 2607 2388 136
19:31:27 2683 2226 135
19:31:28 2682 2249 53
Average 2656 2275 106
491, 16
HP2392A -- rc.atl.hp.com via TELNET
Num

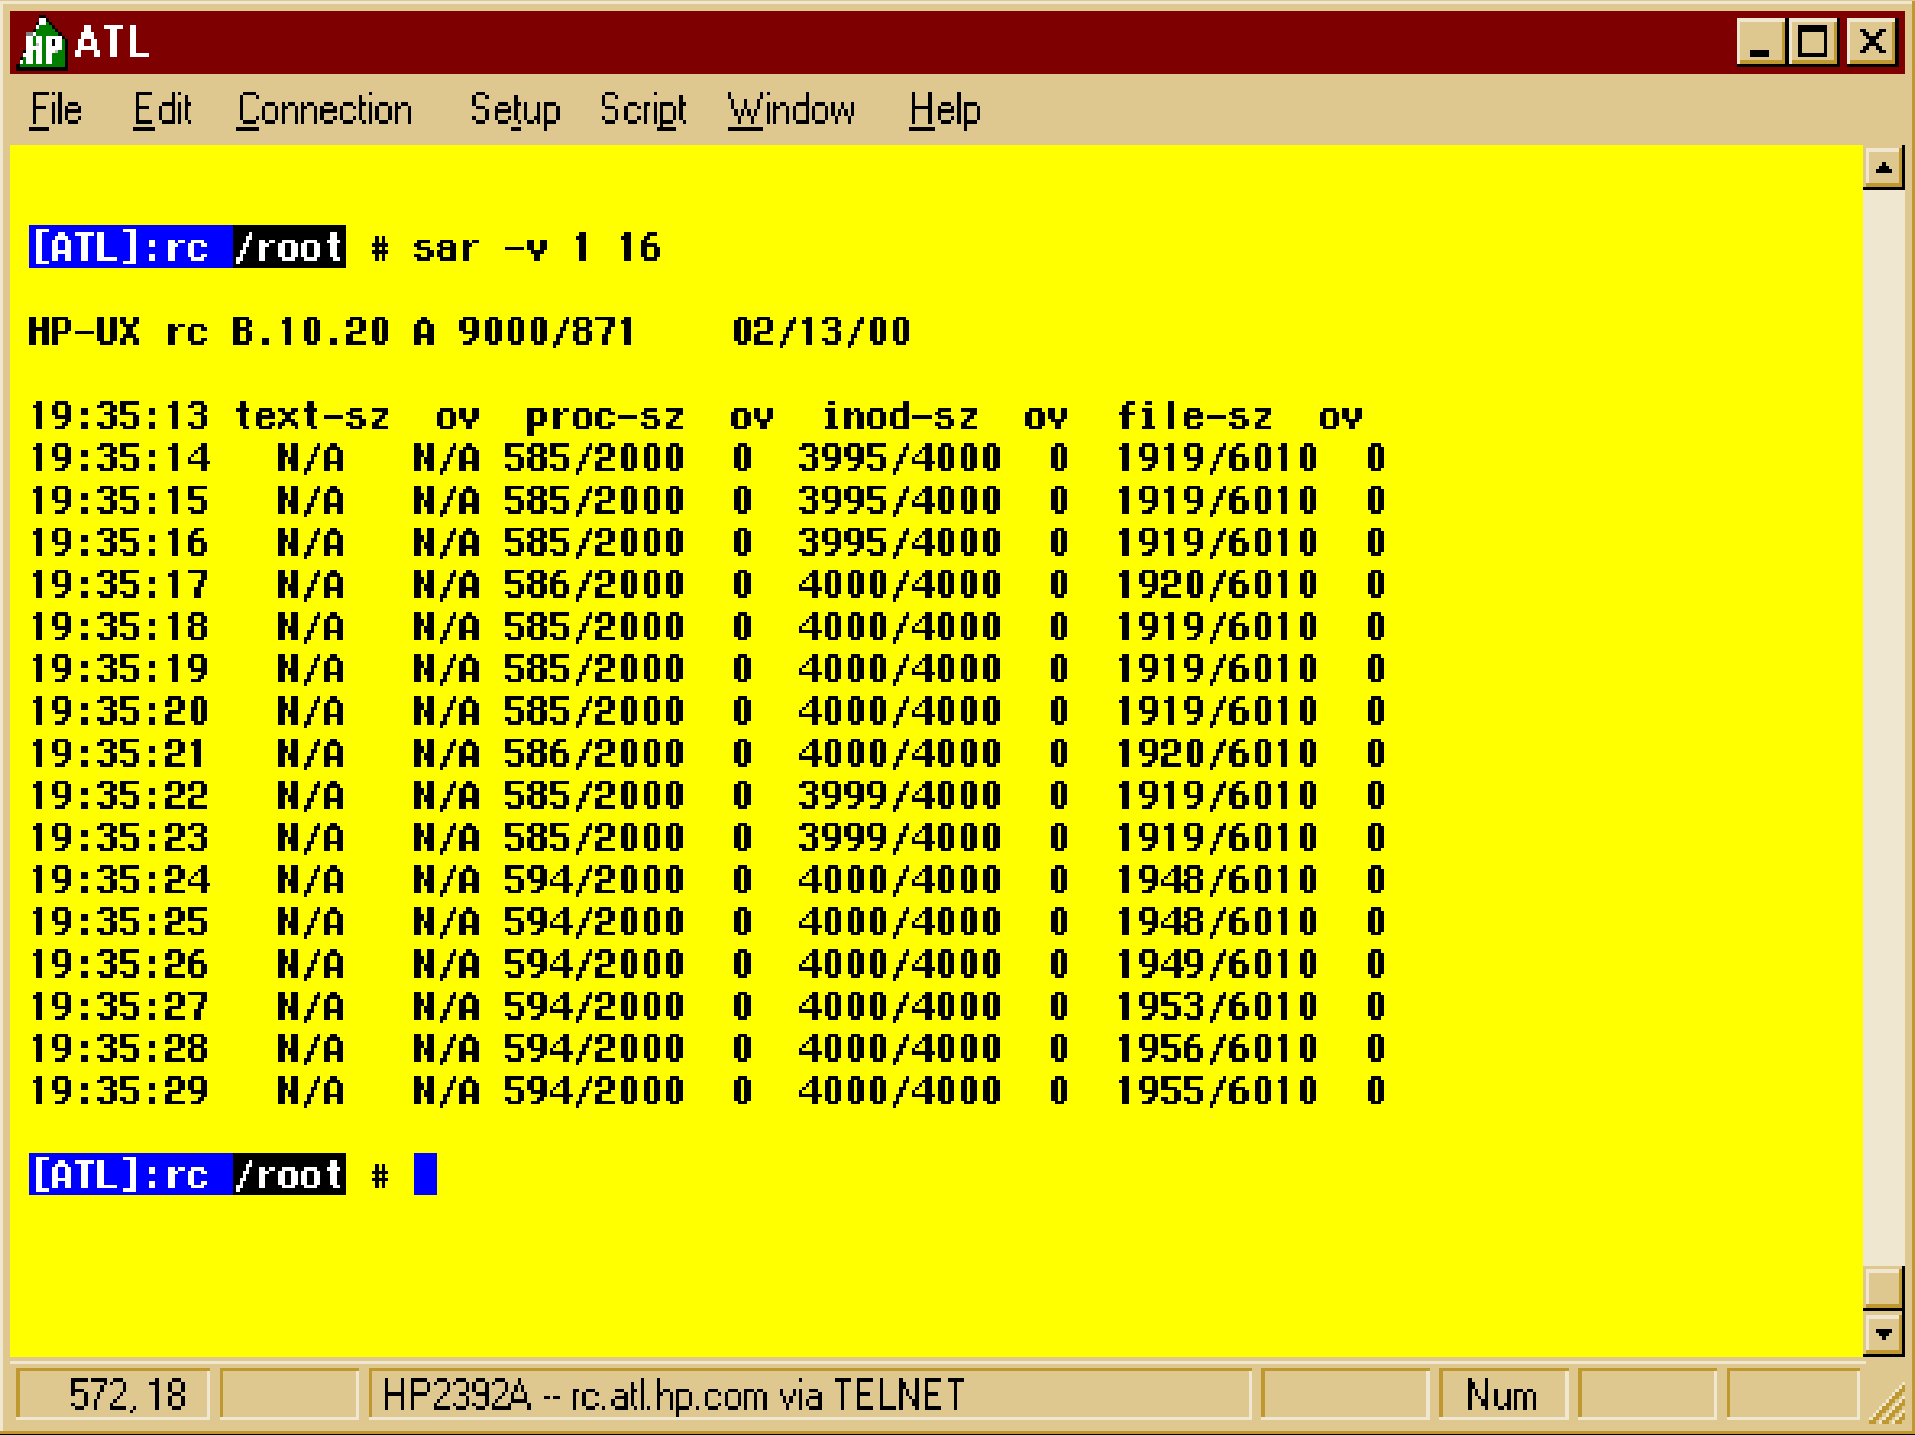
ATL
File Edit Connection Setup Script Window Help
[ATL]:rc /root # sar -v 1 16
HP-UX rc B.10.20 A 9000/871 02/13/00
19:35:13 text-sz ov proc-sz ov inod-sz ov file-sz ov
19:35:14 N/A N/A 585/2000 0 3995/4000 0 1919/6010 0
19:35:15 N/A N/A 585/2000 0 3995/4000 0 1919/6010 0
19:35:16 N/A N/A 585/2000 0 3995/4000 0 1919/6010 0
19:35:17 N/A N/A 586/2000 0 4000/4000 0 1920/6010 0
19:35:18 N/A N/A 585/2000 0 4000/4000 0 1919/6010 0
19:35:19 N/A N/A 585/2000 0 4000/4000 0 1919/6010 0
19:35:20 N/A N/A 585/2000 0 4000/4000 0 1919/6010 0
19:35:21 N/A N/A 586/2000 0 4000/4000 0 1920/6010 0
19:35:22 N/A N/A 585/2000 0 3999/4000 0 1919/6010 0
19:35:23 N/A N/A 585/2000 0 3999/4000 0 1919/6010 0
19:35:24 N/A N/A 594/2000 0 4000/4000 0 1948/6010 0
19:35:25 N/A N/A 594/2000 0 4000/4000 0 1948/6010 0
19:35:26 N/A N/A 594/2000 0 4000/4000 0 1949/6010 0
19:35:27 N/A N/A 594/2000 0 4000/4000 0 1953/6010 0
19:35:28 N/A N/A 594/2000 0 4000/4000 0 1956/6010 0
19:35:29 N/A N/A 594/2000 0 4000/4000 0 1955/6010 0
[ATL]:rc /root #
572, 18
HP2392A -- rc.atl.hp.com via TELNET
Num

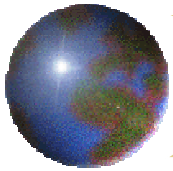

Performance and Kernel Tuning

Performance tuning

- **iostat** - report I/O Statistics
  - Number of seeks per second
  - Kbytes transferred per second
  - Milliseconds per average seek

ATL

File Edit Connection Setup Script Window Help

```
[ATL]:rc /root # iostat 2 3

  device     bps      sps     msps

  c0t9d0       0      0.0      1.0
 c0t11d0       0      0.0      1.0
  c0t8d0       0      0.0      1.0
 c0t10d0       0      0.0      1.0
  c0t5d0       0      0.0      1.0

  c0t9d0       0      0.4      1.0
 c0t11d0    5112     10.2       1.0
  c0t8d0       0      0.0      1.0
  c0t5d0    5111      9.8       1.0

  c0t9d0       0      0.5      1.0
 c0t11d0    5554     10.8       1.0
  c0t8d0       1      0.0      1.0
  c0t5d0    5556     10.8       1.0

[ATL]:rc /root #
```

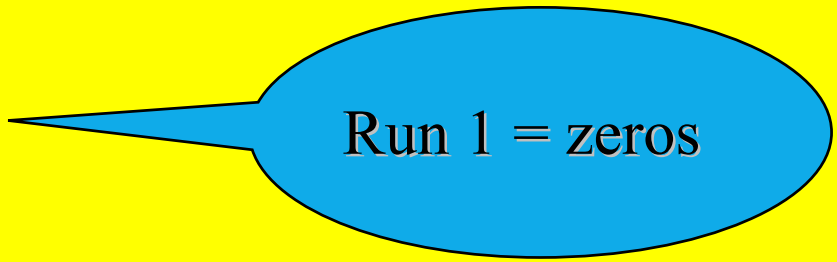

933, 18 | HP2392A -- rc.atl.hp.com via TELNET | Num

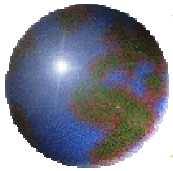

Performance and Kernel Tuning

Performance tuning

- **vmstat**
  - **vmstat [-dnS] [interval [count]]**
  - **vmstat -f | -s | -z**
- Virtual Memory Statistics
  - **-d** Adds disk transfers per second
  - **-n** Format for 80 columns
  - **-S** Processes swapped rather than paging

ATL

File Edit Connection Setup Script Window Help

```
[ATL]:rc /root # vmstat 1 18
         procs           memory                   page                              faults       cpu
    r     b     w      avm    free   re  at    pi  po    fr   de    sr     in     sy    cs    us sy id
    0     0     0    13331   67976    0   0     0   0     0    0     0      0      0     0    14 10 76
    1     0     0    13643   67946   14   5     0   0     0    0     0    347    660   100     3 -2 99
    1     0     0    13643   67946   11   4     0   0     0    0     0    347    602   109    11 -1 90
    1     0     0    13643   67946    9   3     0   0     0    0     0    350    572   113     5  0 95
    1     0     0    13643   67946    7   2     0   0     0    0     0    353    536   112    -3  1 102
    1     0     0    13643   67815   18   2     0   0     0    0     0    374   1065   139     0 91  9
    5     0     0    14100   67811   15   1     0   0     0    0     0    411   1527   188    11 83  7
    5     0     0    14100   67805   12   0     0   0     0    0     0    416   1996   207     6 87  6
    5     0     0    14100   67805    9   0     0   0     0    0     0    417   2375   226     1 92  7
    5     0     0    14100   67804    7   0     0   0     0    0     0    426   2520   239    11 82  7
    5     0     0    14100   67802    9   1     1   0     0    0     0    432   2722   251     6 89  5
    4     0     0    14492   67800    8   0     0   0     0    0     0    431   2672   256     0 94  6
    4     0     0    14492   67800    6   0     0   0     0    0     0    437   2688   261     2 93  6
    4     0     0    14492   67793    4   0     0   0     0    0     0    431   2670   262     5 89  6
    4     0     0    14492   67793    3   0     0   0     0    0     0    433   2808   266    -1 95  6
    4     0     0    14492   67790    2   0     0   0     0    0     0    427   2798   262     1 93  6
    5     0     0    14737   67790    1   0     0   0     0    0     0    427   2744   260     1 94  5
    5     0     0    14737   67790    0   0     0   0     0    0     0    441   2786   279     2 93  6
[ATL]:rc /root #
```

276, 18 | HP70092 -- rc.atl.hp.com via TELNET | Num

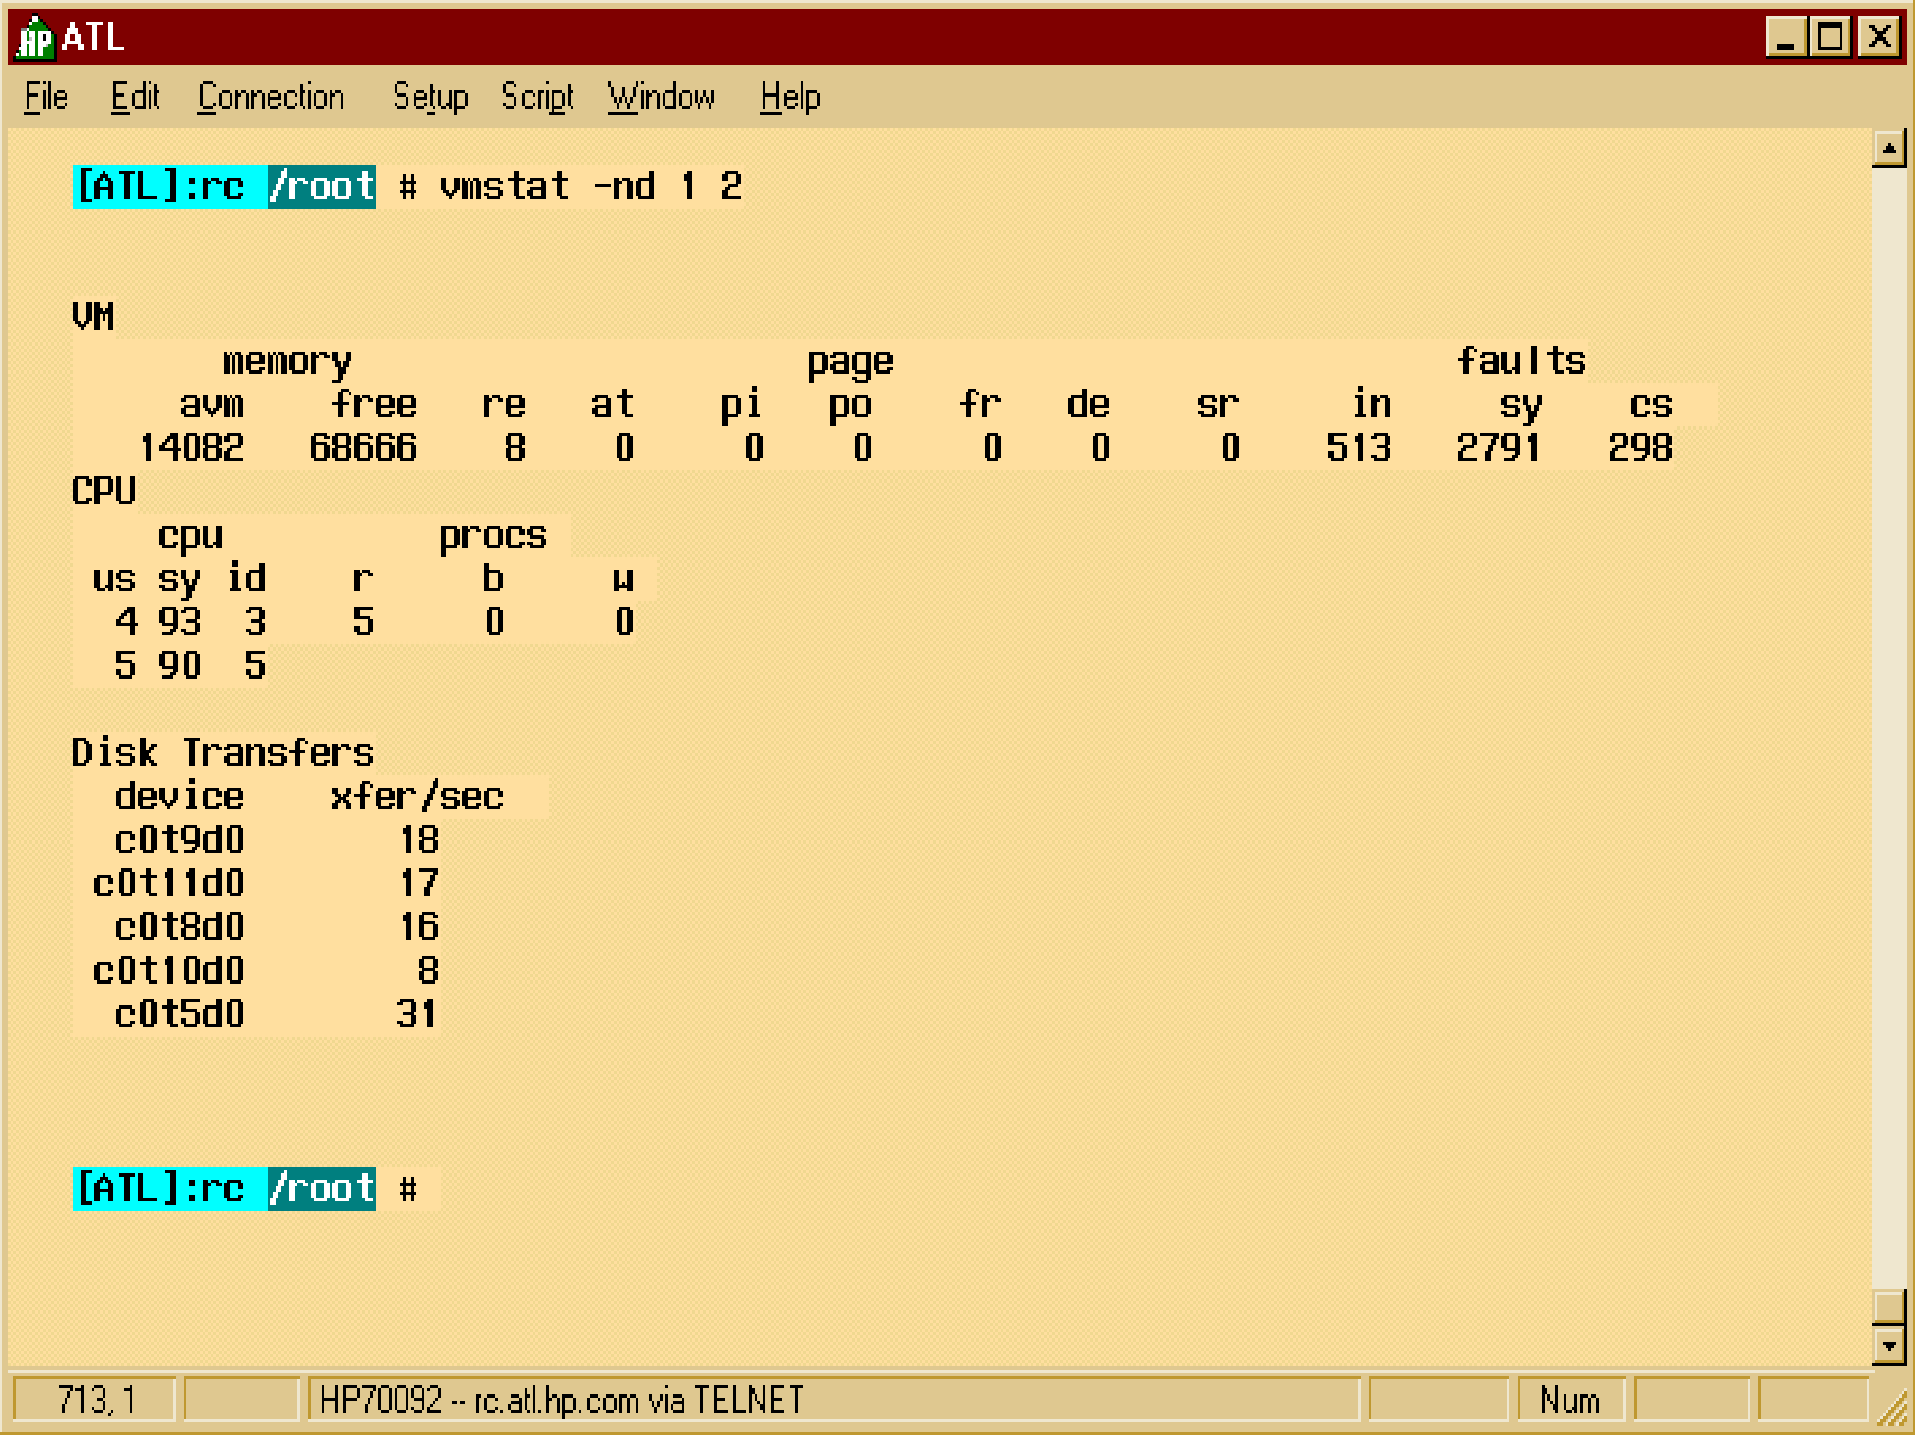
ATL
File Edit Connection Setup Script Window Help
[ATL]:rc /root # vmstat -nd 1 2
VM
memory page faults
avm free re at pi po fr de sr in sy cs
14082 68666 8 0 0 0 0 0 0 513 2791 298
CPU
cpu procs
us sy id r b w
4 93 3 5 0 0
5 90 5
Disk Transfers
device xfer/sec
c0t9d0 18
c0t11d0 17
c0t8d0 16
c0t10d0 8
c0t5d0 31
[ATL]:rc /root #
713,1
HP70092 -- rc.atl.hp.com via TELNET
Num

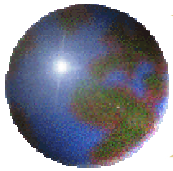

Performance and Kernel Tuning

Performance tuning

- **vmstat** (cont)
  - **-f** fork and page summary:

    vmstat -f

    9043939 forks, 448300514 pages, average= 49.57
  - **-s** Summary
  - **-z** Zero kernel sums


ATL
File Edit Connection Setup Script Window Help
[ATL]:rc /root # vmstat -s
0 swap ins
0 swap outs
0 pages swapped in
0 pages swapped out
38679 total address trans. faults taken
13410 page ins
0 page outs
180 pages paged in
0 pages paged out
12738 reclaims from free list
13231 total page reclaims
0 intransit blocking page faults
24941 zero fill pages created
10379 zero fill page faults
9358 executable fill pages created
0 executable fill page faults
0 swap text pages found in free list
5636 inode text pages found in free list
0 revolutions of the clock hand
0 pages scanned for page out
0 pages freed by the clock daemon
210553 cpu context switches
809959 device interrupts
55205 traps
1572971 system calls
[ATL]:rc /root #
54, 18
HP70092 -- rc.atl.hp.com via TELNET
Num

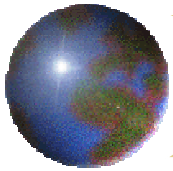

Performance and Kernel Tuning

Performance tuning

- **top**:
  - System data - Summary
    - name and time
    - Load Average 1,5,15 minutes
    - % time in user,nice,system,idle,etc
  - Memory:Virtual and Real
  - Individual Processes

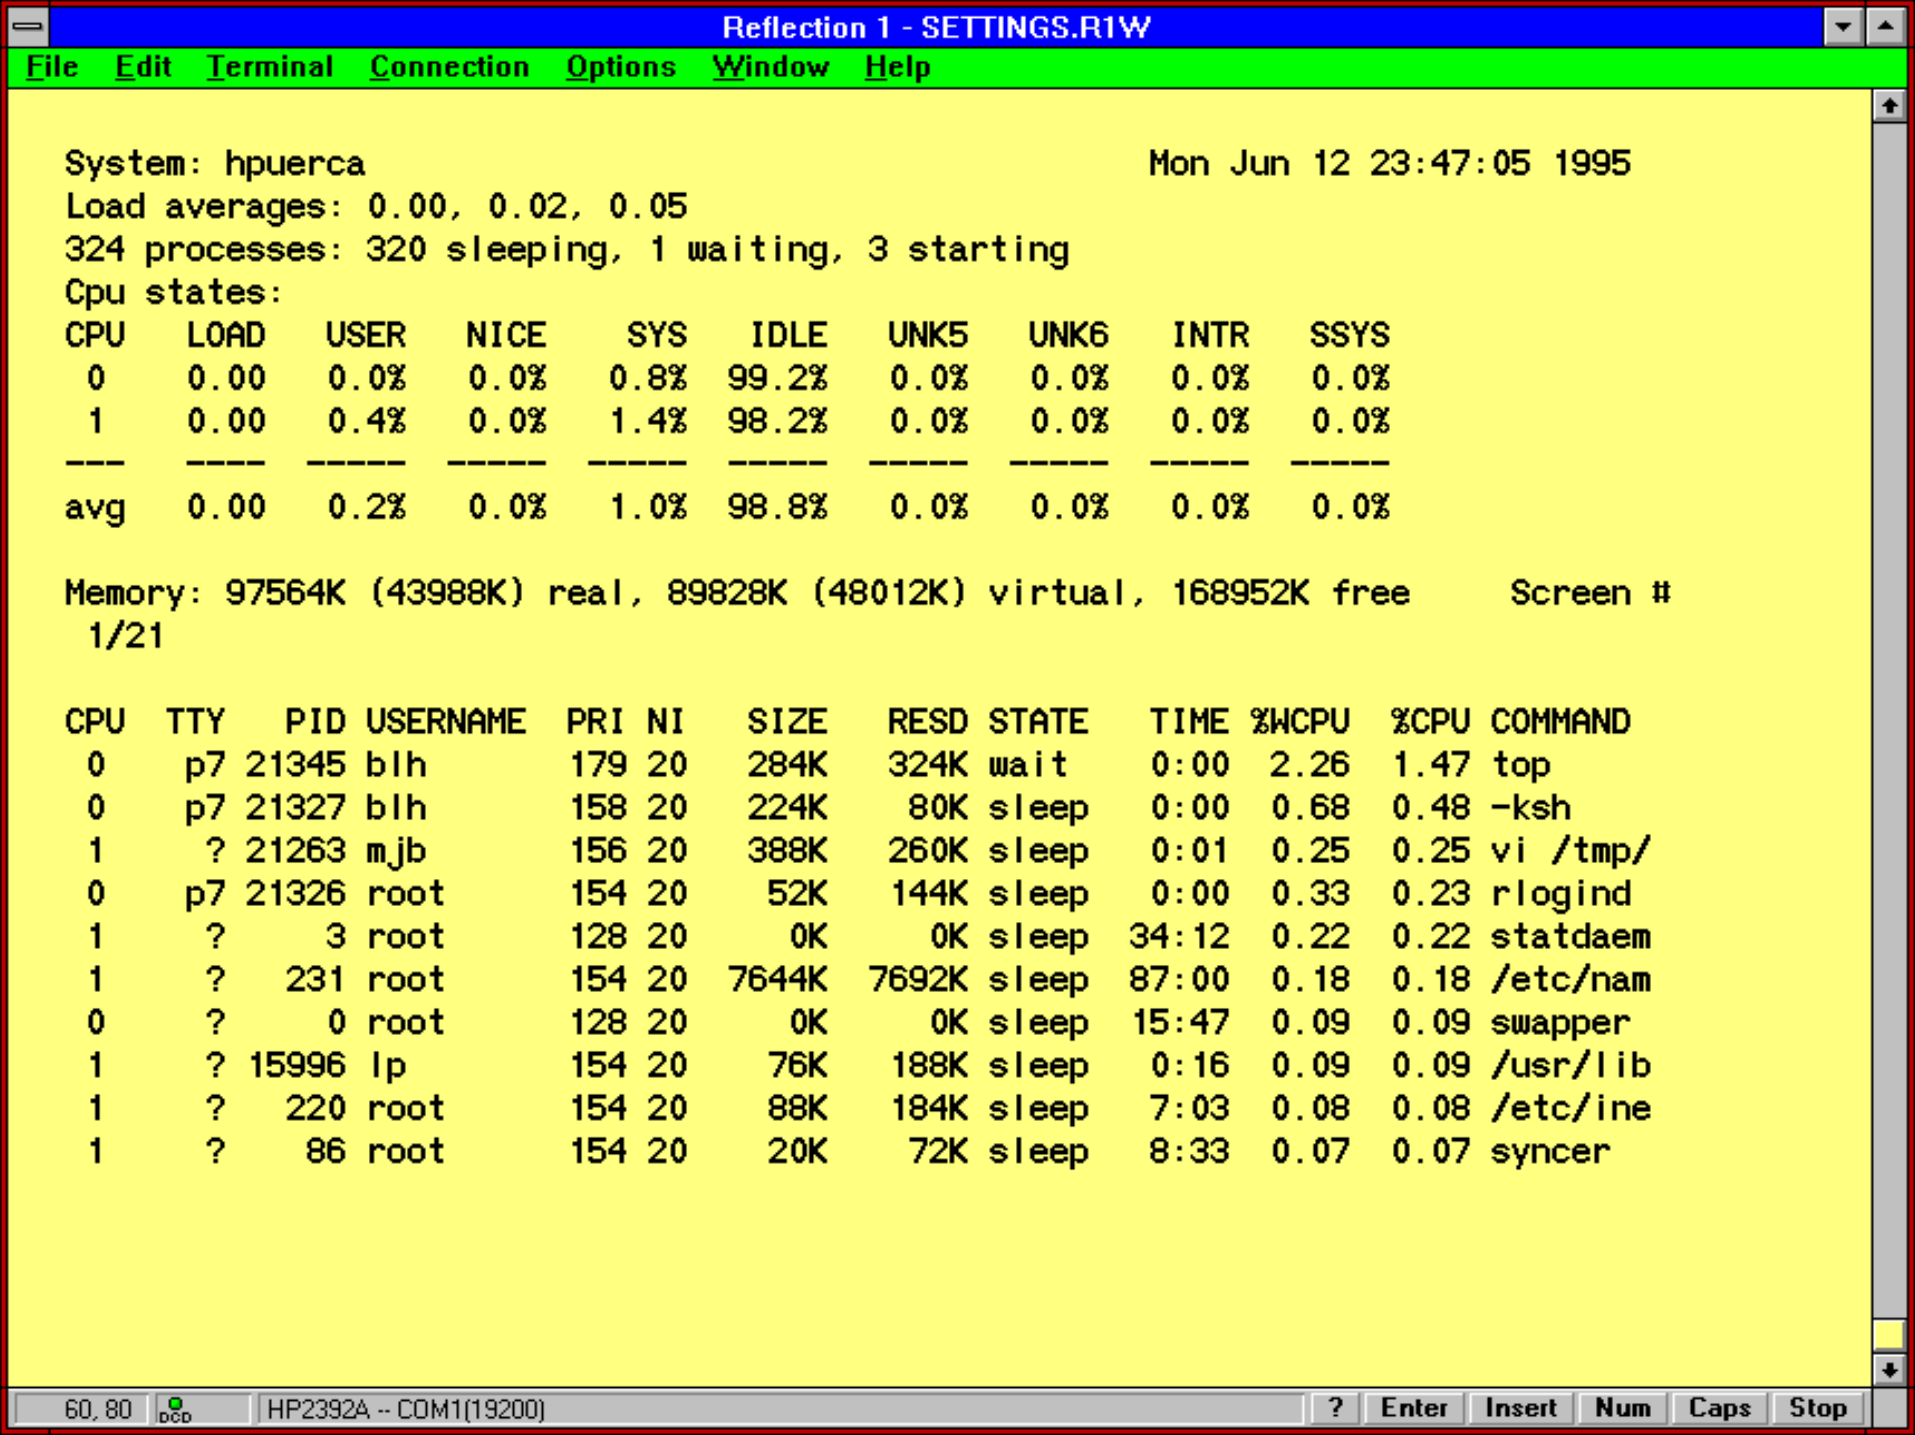
Reflection 1 - SETTINGS.R1W
File Edit Terminal Connection Options Window Help
System: hpuerca Mon Jun 12 23:47:05 1995
Load averages: 0.00, 0.02, 0.05
324 processes: 320 sleeping, 1 waiting, 3 starting
Cpu states:
CPU LOAD USER NICE SYS IDLE UNK5 UNK6 INTR SSYS
0 0.00 0.0% 0.0% 0.8% 99.2% 0.0% 0.0% 0.0% 0.0%
1 0.00 0.4% 0.0% 1.4% 98.2% 0.0% 0.0% 0.0% 0.0%
avg 0.00 0.2% 0.0% 1.0% 98.8% 0.0% 0.0% 0.0% 0.0%
Memory: 97564K (43988K) real, 89828K (48012K) virtual, 168952K free Screen #
1/21
CPU TTY PID USERNAME PRI NI SIZE RESD STATE TIME %WCPU %CPU COMMAND
0 p7 21345 blh 179 20 284K 324K wait 0:00 2.26 1.47 top
0 p7 21327 blh 158 20 224K 80K sleep 0:00 0.68 0.48 -ksh
1 ? 21263 mjb 156 20 388K 260K sleep 0:01 0.25 0.25 vi /tmp/
0 p7 21326 root 154 20 52K 144K sleep 0:00 0.33 0.23 rlogind
1 ? 3 root 128 20 0K 0K sleep 34:12 0.22 0.22 statdaem
1 ? 231 root 154 20 7644K 7692K sleep 87:00 0.18 0.18 /etc/nam
0 ? 0 root 128 20 0K 0K sleep 15:47 0.09 0.09 swapper
1 ? 15996 lp 154 20 76K 188K sleep 0:16 0.09 0.09 /usr/lib
1 ? 220 root 154 20 88K 184K sleep 7:03 0.08 0.08 /etc/ine
1 ? 86 root 154 20 20K 72K sleep 8:33 0.07 0.07 syncer
60, 80 DCD HP2392A -- COM1(19200) ? Enter Insert Num Caps Stop

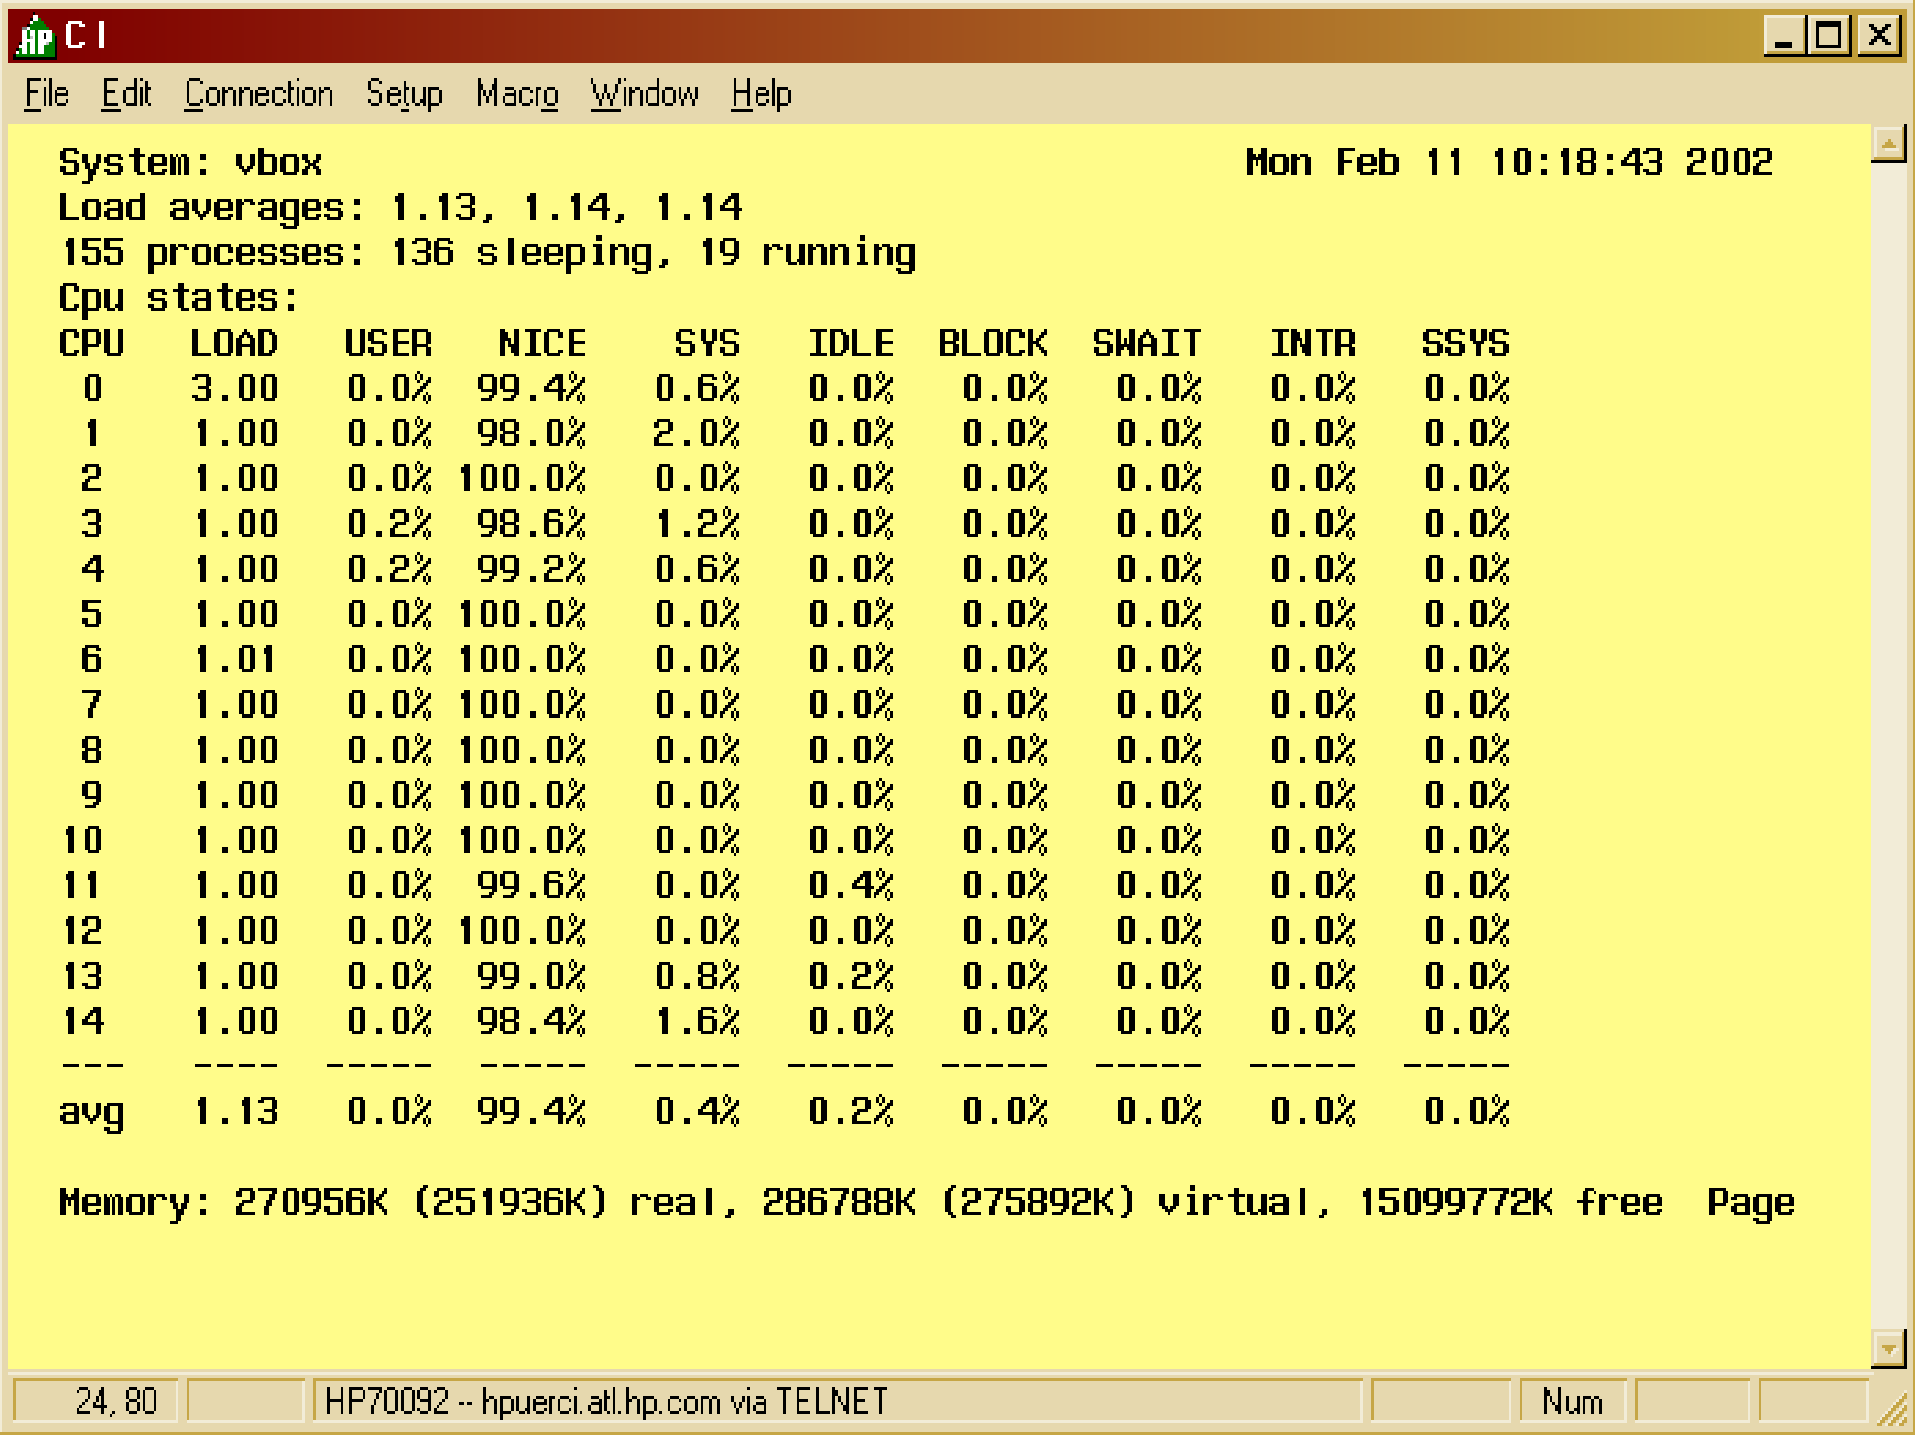
CI
File Edit Connection Setup Macro Window Help
System: vbox Mon Feb 11 10:18:43 2002
Load averages: 1.13, 1.14, 1.14
155 processes: 136 sleeping, 19 running
Cpu states:
CPU LOAD USER NICE SYS IDLE BLOCK SWAIT INTR SSYS
0 3.00 0.0% 99.4% 0.6% 0.0% 0.0% 0.0% 0.0% 0.0%
1 1.00 0.0% 98.0% 2.0% 0.0% 0.0% 0.0% 0.0% 0.0%
2 1.00 0.0% 100.0% 0.0% 0.0% 0.0% 0.0% 0.0% 0.0%
3 1.00 0.2% 98.6% 1.2% 0.0% 0.0% 0.0% 0.0% 0.0%
4 1.00 0.2% 99.2% 0.6% 0.0% 0.0% 0.0% 0.0% 0.0%
5 1.00 0.0% 100.0% 0.0% 0.0% 0.0% 0.0% 0.0% 0.0%
6 1.01 0.0% 100.0% 0.0% 0.0% 0.0% 0.0% 0.0% 0.0%
7 1.00 0.0% 100.0% 0.0% 0.0% 0.0% 0.0% 0.0% 0.0%
8 1.00 0.0% 100.0% 0.0% 0.0% 0.0% 0.0% 0.0% 0.0%
9 1.00 0.0% 100.0% 0.0% 0.0% 0.0% 0.0% 0.0% 0.0%
10 1.00 0.0% 100.0% 0.0% 0.0% 0.0% 0.0% 0.0% 0.0%
11 1.00 0.0% 99.6% 0.0% 0.4% 0.0% 0.0% 0.0% 0.0%
12 1.00 0.0% 100.0% 0.0% 0.0% 0.0% 0.0% 0.0% 0.0%
13 1.00 0.0% 99.0% 0.8% 0.2% 0.0% 0.0% 0.0% 0.0%
14 1.00 0.0% 98.4% 1.6% 0.0% 0.0% 0.0% 0.0% 0.0%
avg 1.13 0.0% 99.4% 0.4% 0.2% 0.0% 0.0% 0.0% 0.0%
Memory: 270956K (251936K) real, 286788K (275892K) virtual, 15099772K free Page
24, 80
HP70092 -- hpuerci.atl.hp.com via TELNET
Num

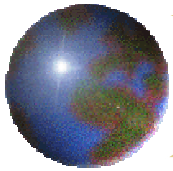

Performance and Kernel Tuning

Performance tuning

- **Glance/Plus**
  - HP measurement package
  - midaemon for better metrics
  - character mode interface
  - **gpm** for Xwindcows
  - Extensive measurement options
  - Alarms

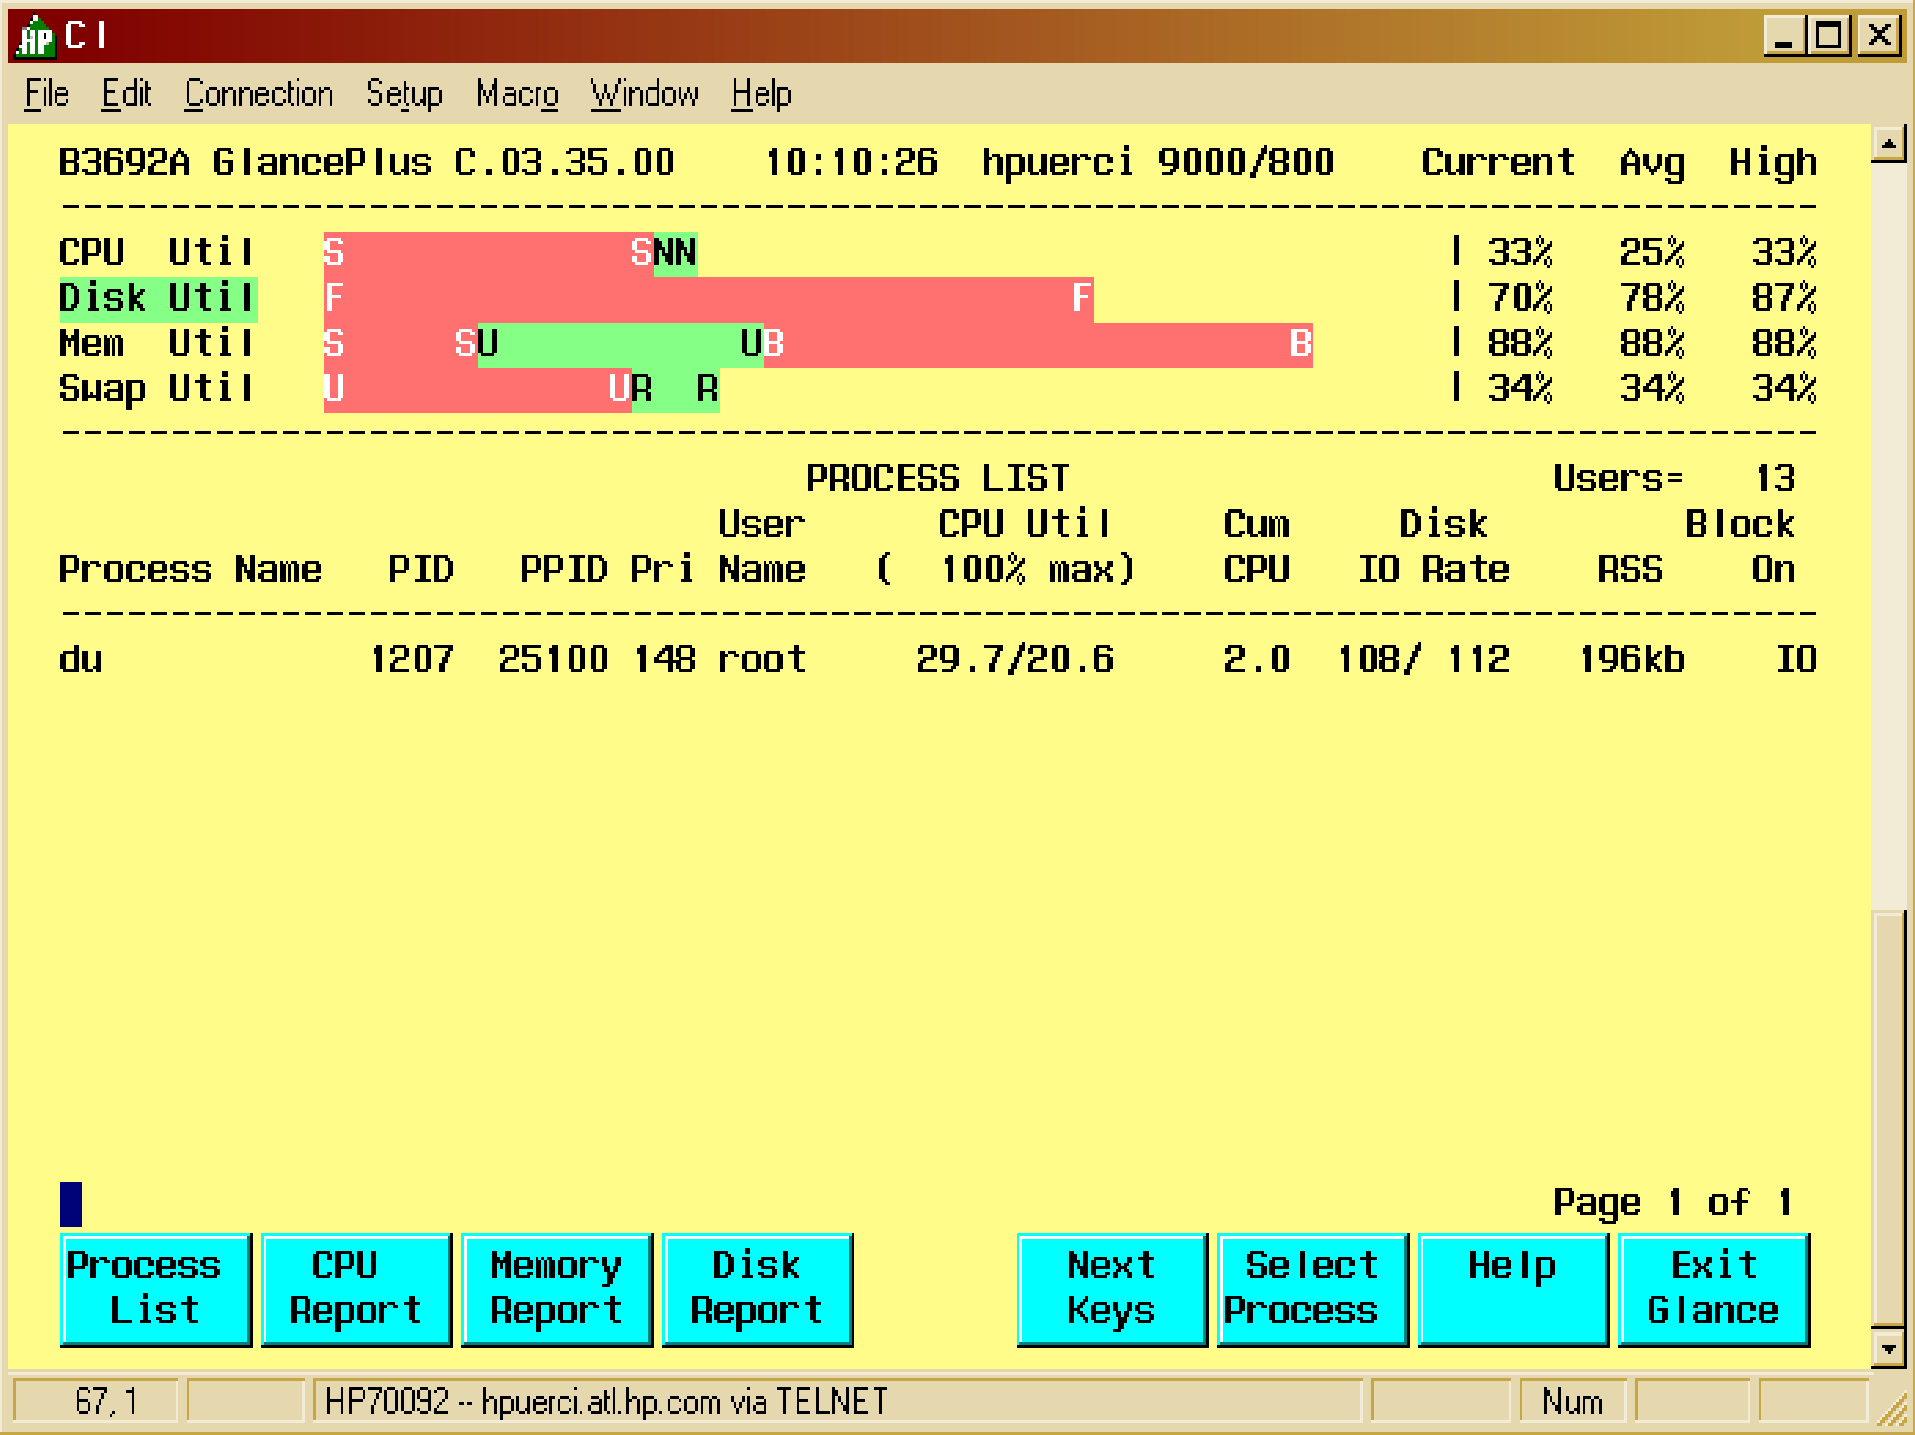
CI
File Edit Connection Setup Macro Window Help
B3692A GlancePlus C.03.35.00 10:10:26 hpuerci 9000/800 Current Avg High
CPU Util | 33% 25% 33%
Disk Util | 70% 78% 87%
Mem Util | 88% 88% 88%
Swap Util | 34% 34% 34%
PROCESS LIST Users= 13
Process Name PID PPID Pri User Name CPU Util (100% max) Cum CPU Disk IO Rate RSS Block On
du 1207 25100 148 root 29.7/20.6 2.0 108/ 112 196kb IO
Page 1 of 1
Process List
CPU Report
Memory Report
Disk Report
Next Keys
Select Process
Help
Exit Glance
67,1
HP70092 -- hpuerci.atl.hp.com via TELNET
Num

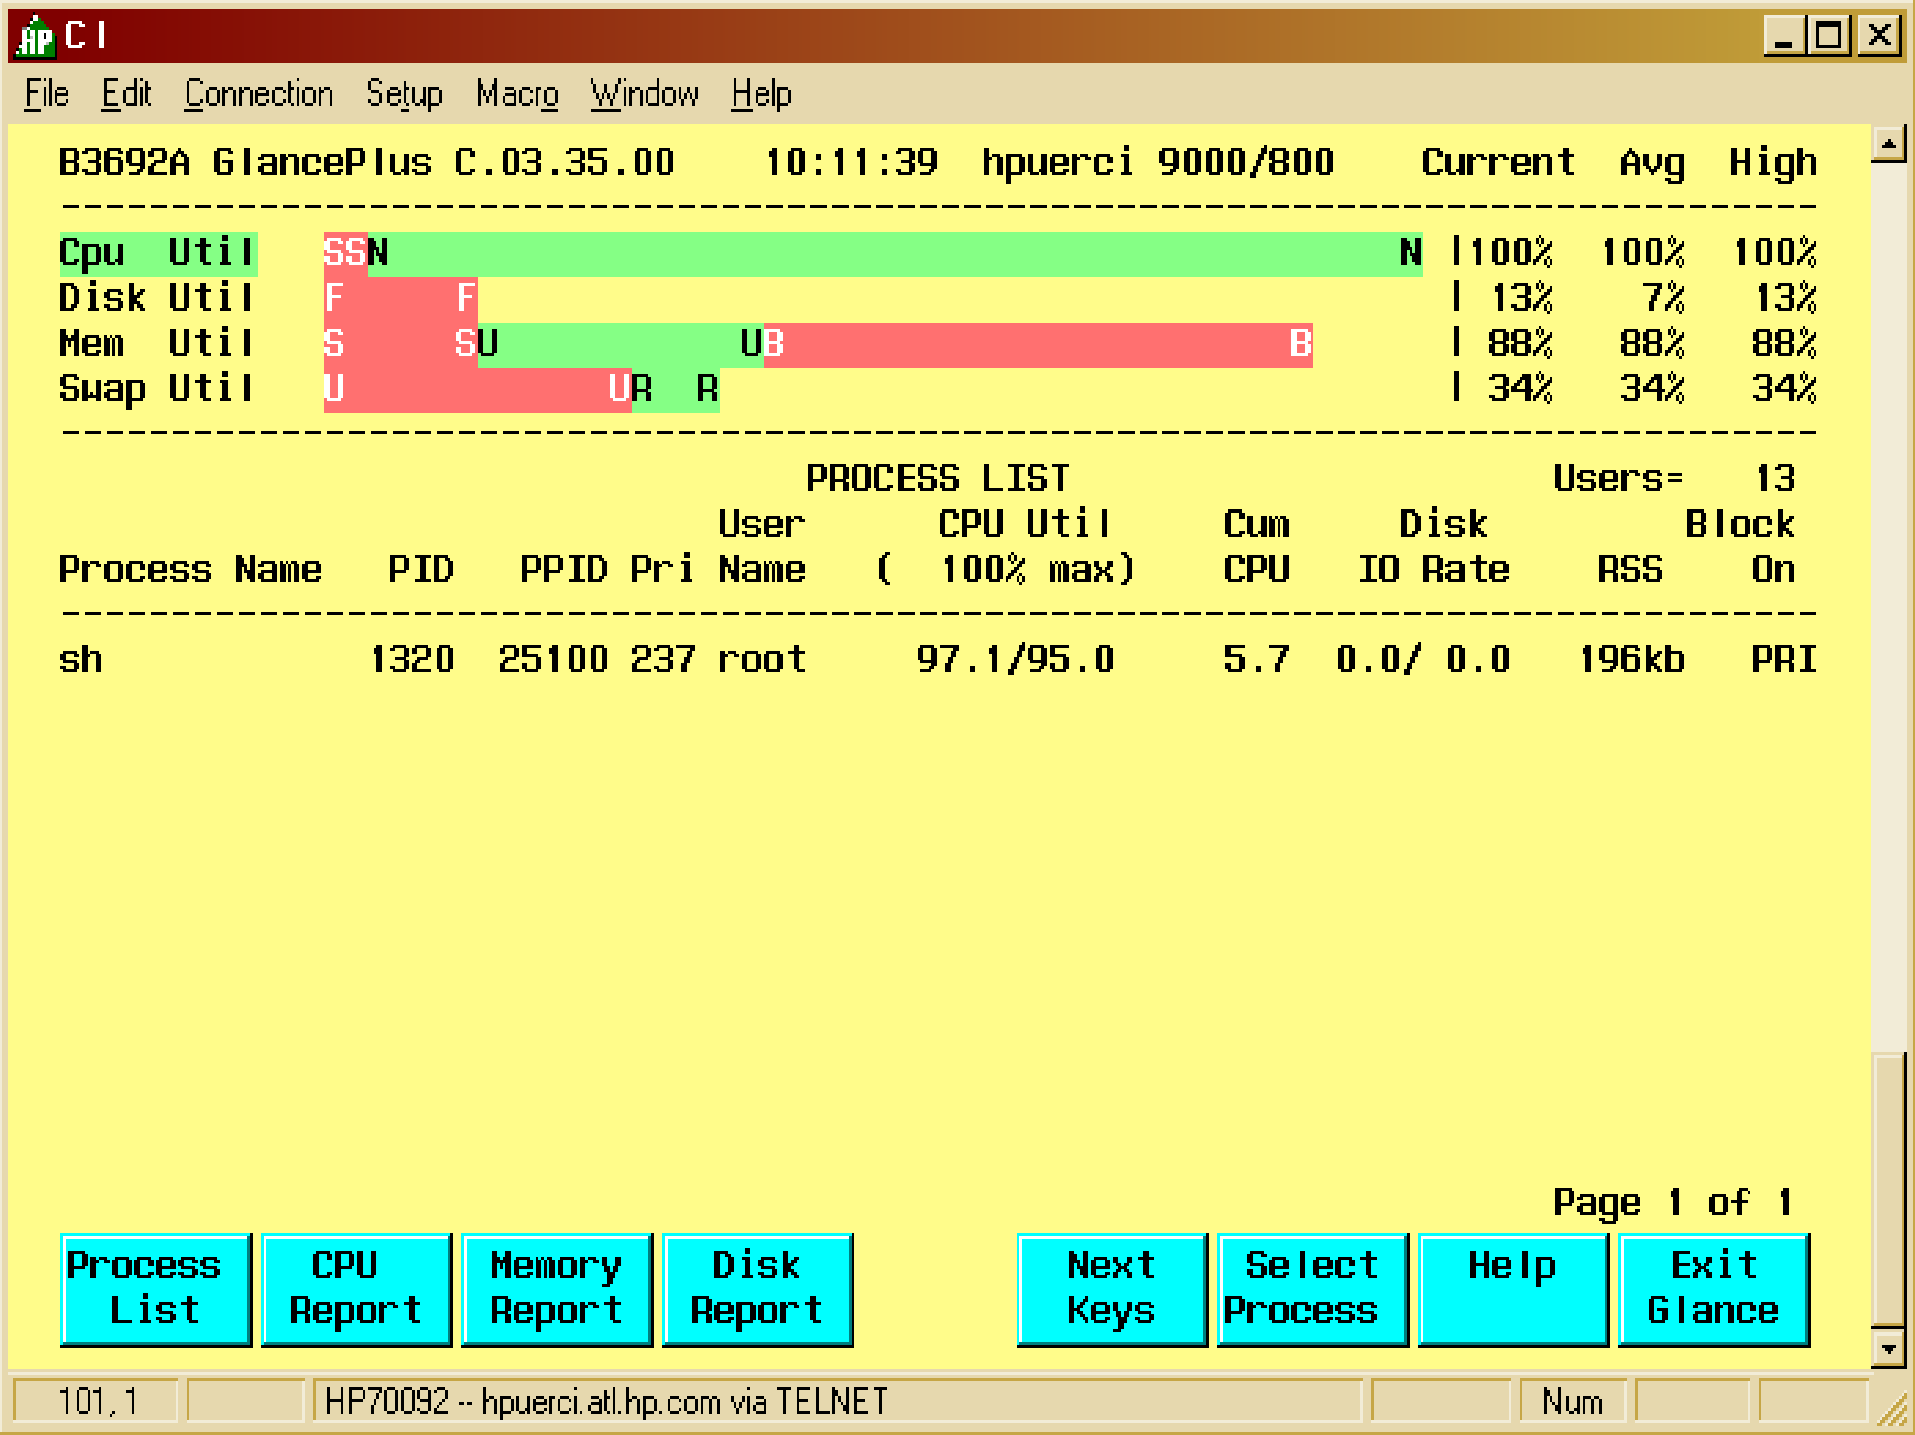
CI
File Edit Connection Setup Macro Window Help
B3692A GlancePlus C.03.35.00 10:11:39 hpuerci 9000/800 Current Avg High
Cpu Util 100% 100% 100%
Disk Util 13% 7% 13%
Mem Util 88% 88% 88%
Swap Util 34% 34% 34%
PROCESS LIST Users= 13
Process Name PID PPID Pri User Name CPU Util (100% max) Cum CPU Disk IO Rate RSS Block On
sh 1320 25100 237 root 97.1/95.0 5.7 0.0/ 0.0 196kb PRI
Page 1 of 1
Process List
CPU Report
Memory Report
Disk Report
Next Keys
Select Process
Help
Exit Glance
101,1
HP70092 -- hpuerci.atl.hp.com via TELNET
Num

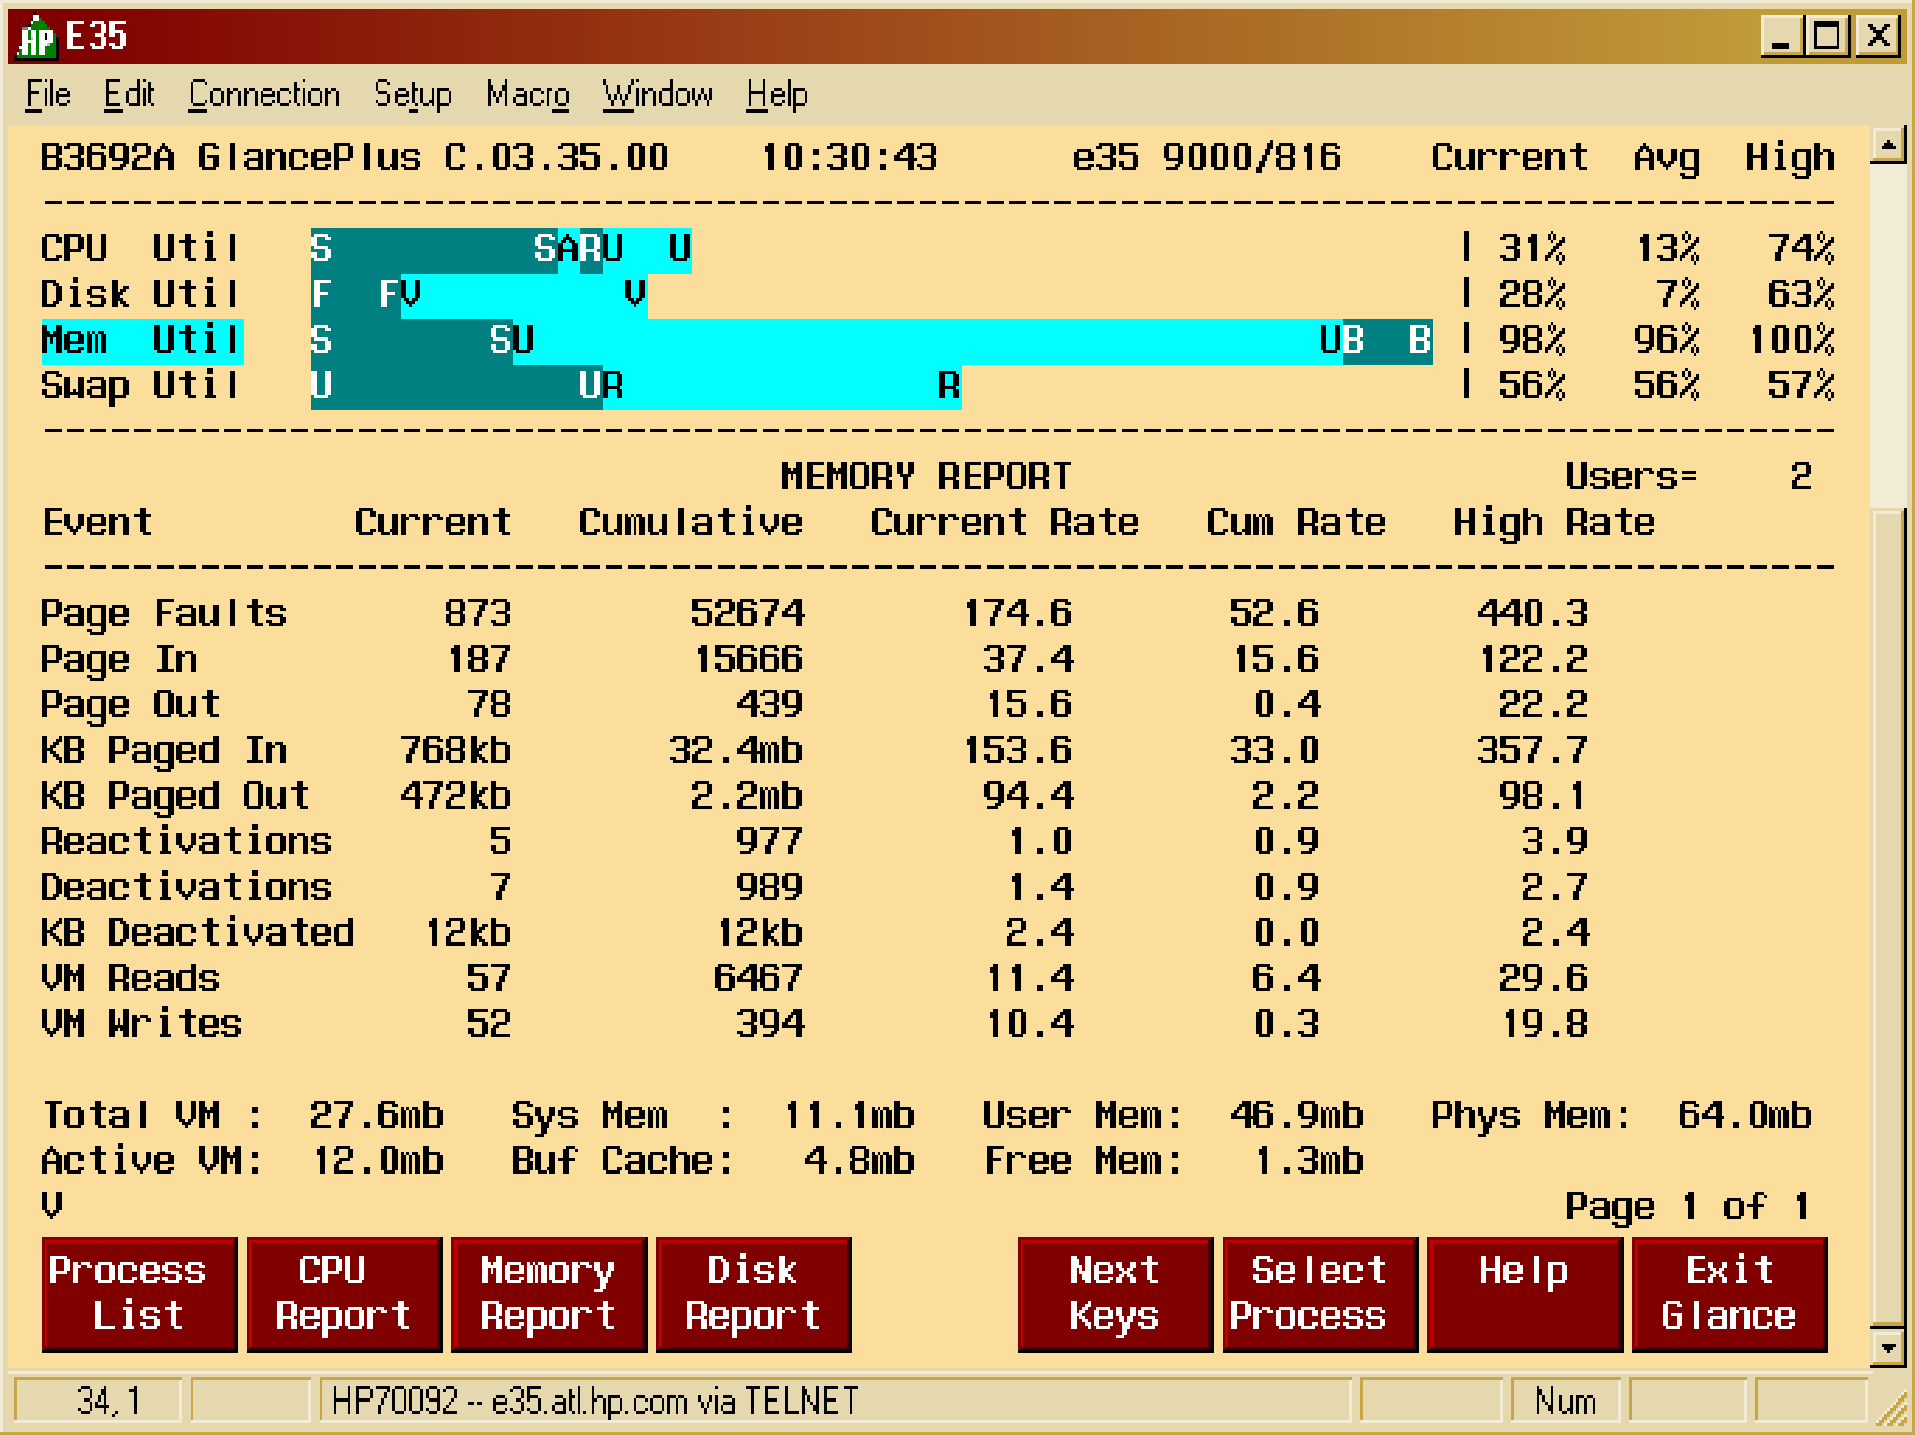
E35
File Edit Connection Setup Macro Window Help
B3692A GlancePlus C.03.35.00 10:30:43 e35 9000/816 Current Avg High
CPU Util | 31% 13% 74%
Disk Util | 28% 7% 63%
Mem Util | 98% 96% 100%
Swap Util | 56% 56% 57%
MEMORY REPORT Users= 2
Event Current Cumulative Current Rate Cum Rate High Rate
Page Faults 873 52674 174.6 52.6 440.3
Page In 187 15666 37.4 15.6 122.2
Page Out 78 439 15.6 0.4 22.2
KB Paged In 768kb 32.4mb 153.6 33.0 357.7
KB Paged Out 472kb 2.2mb 94.4 2.2 98.1
Reactivations 5 977 1.0 0.9 3.9
Deactivations 7 989 1.4 0.9 2.7
KB Deactivated 12kb 12kb 2.4 0.0 2.4
VM Reads 57 6467 11.4 6.4 29.6
VM Writes 52 394 10.4 0.3 19.8
Total VM : 27.6mb Sys Mem : 11.1mb User Mem: 46.9mb Phys Mem: 64.0mb
Active VM: 12.0mb Buf Cache: 4.8mb Free Mem: 1.3mb
Page 1 of 1
Process List
CPU Report
Memory Report
Disk Report
Next Keys
Select Process
Help
Exit Glance
34,1
HP70092 -- e35.atl.hp.com via TELNET
Num

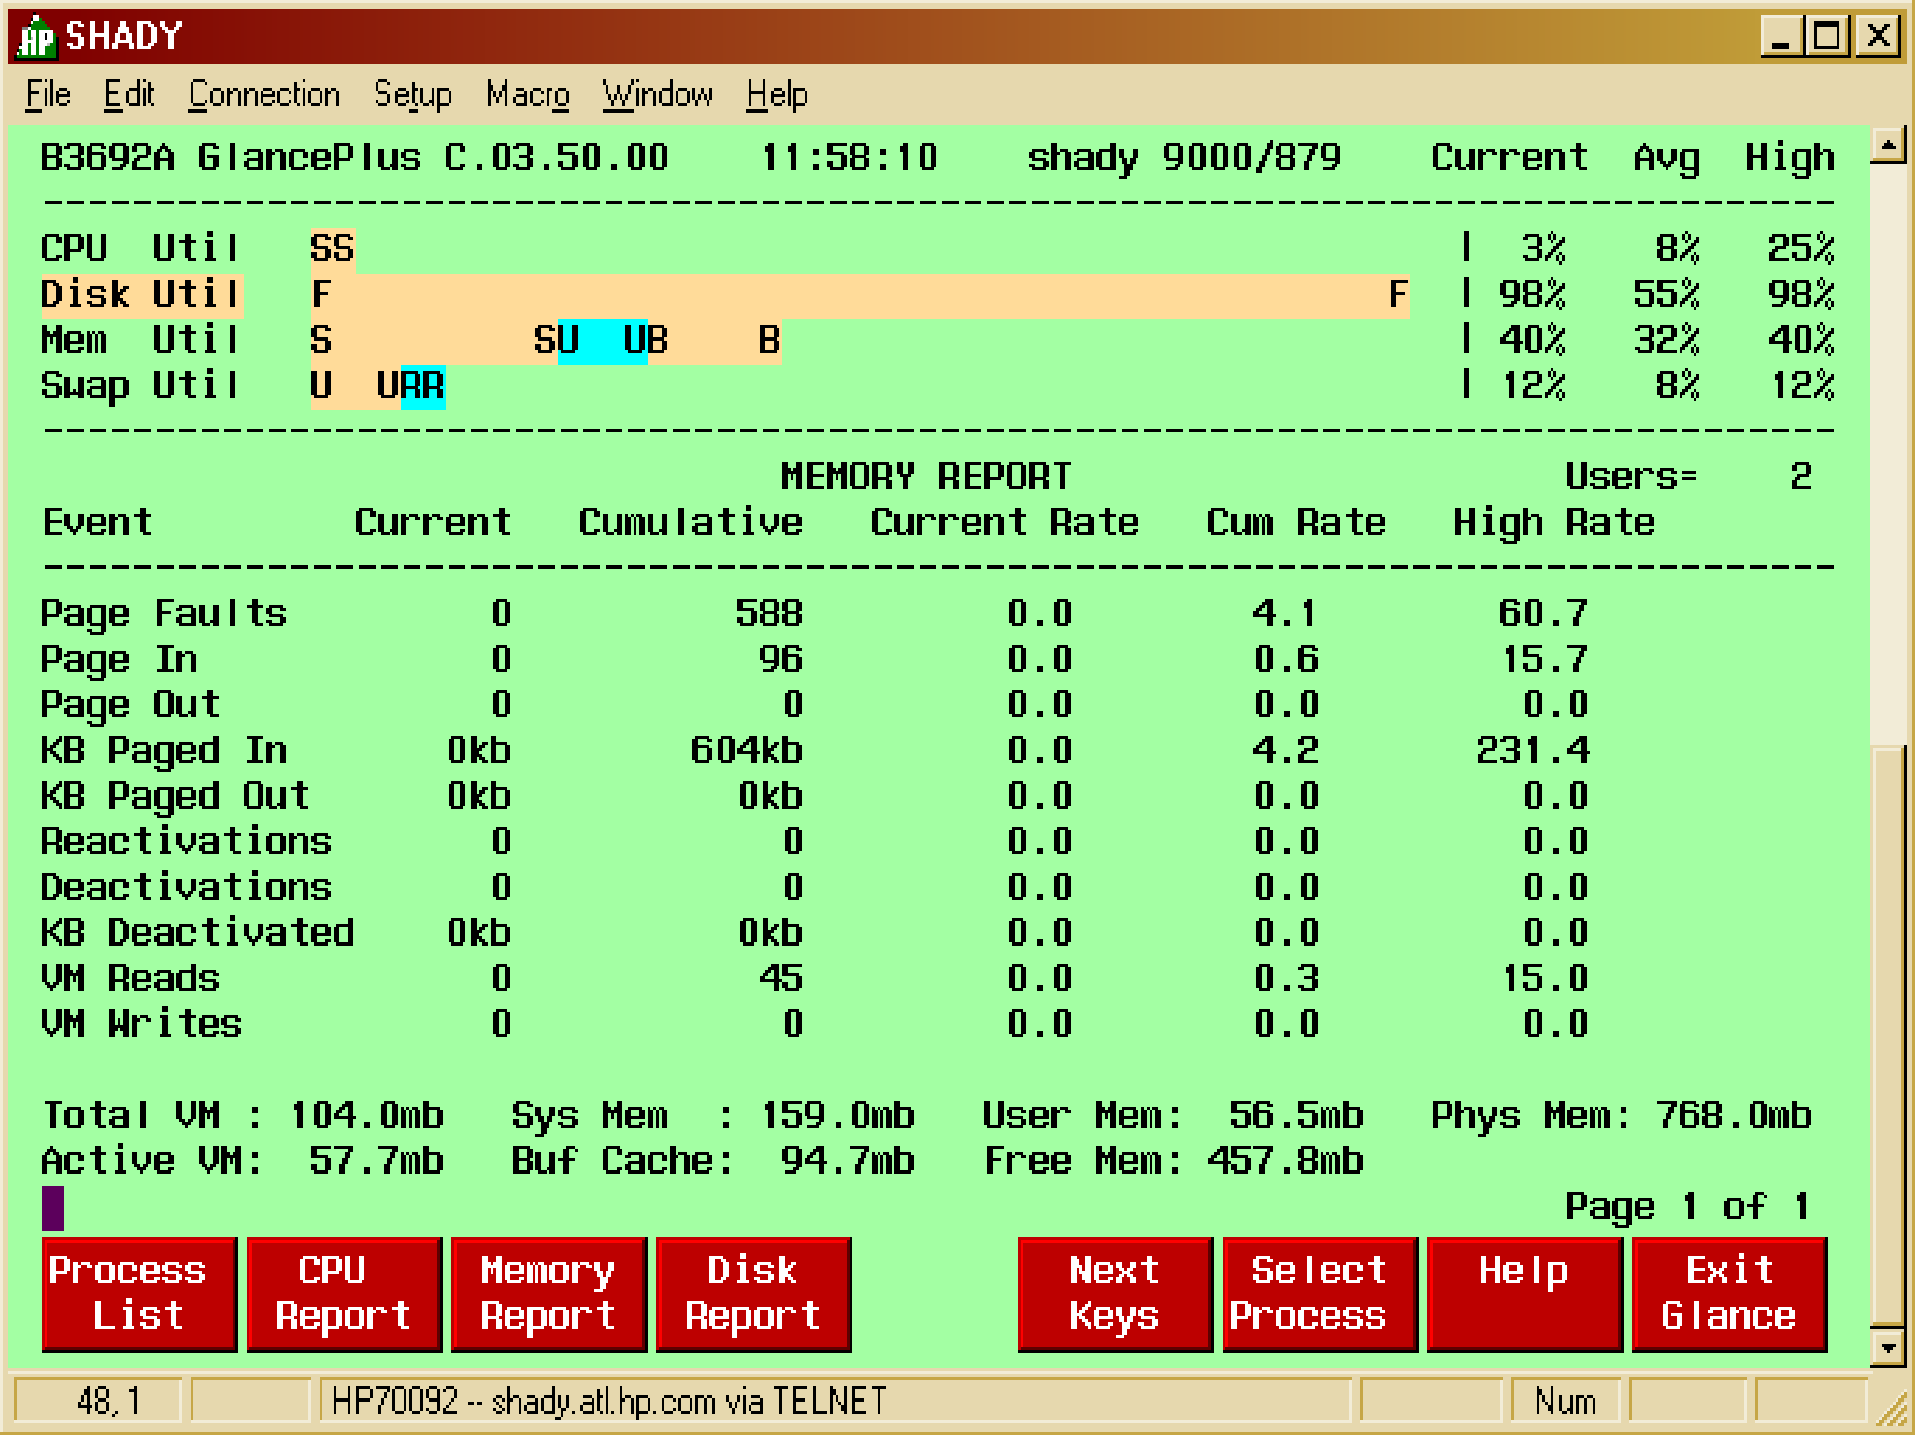
SHADY
File Edit Connection Setup Macro Window Help
B3692A GlancePlus C.03.50.00 11:58:10 shady 9000/879 Current Avg High
CPU Util SS | 3% 8% 25%
Disk Util F F | 98% 55% 98%
Mem Util S SU UB B | 40% 32% 40%
Swap Util U URR | 12% 8% 12%
MEMORY REPORT Users= 2
Event Current Cumulative Current Rate Cum Rate High Rate
Page Faults 0 588 0.0 4.1 60.7
Page In 0 96 0.0 0.6 15.7
Page Out 0 0 0.0 0.0 0.0
KB Paged In 0kb 604kb 0.0 4.2 231.4
KB Paged Out 0kb 0kb 0.0 0.0 0.0
Reactivations 0 0 0.0 0.0 0.0
Deactivations 0 0 0.0 0.0 0.0
KB Deactivated 0kb 0kb 0.0 0.0 0.0
VM Reads 0 45 0.0 0.3 15.0
VM Writes 0 0 0.0 0.0 0.0
Total VM : 104.0mb Sys Mem : 159.0mb User Mem: 56.5mb Phys Mem: 768.0mb
Active VM: 57.7mb Buf Cache: 94.7mb Free Mem: 457.8mb
Page 1 of 1
Process List
CPU Report
Memory Report
Disk Report
Next Keys
Select Process
Help
Exit Glance
48,1
HP70092 -- shady.atl.hp.com via TELNET
Num

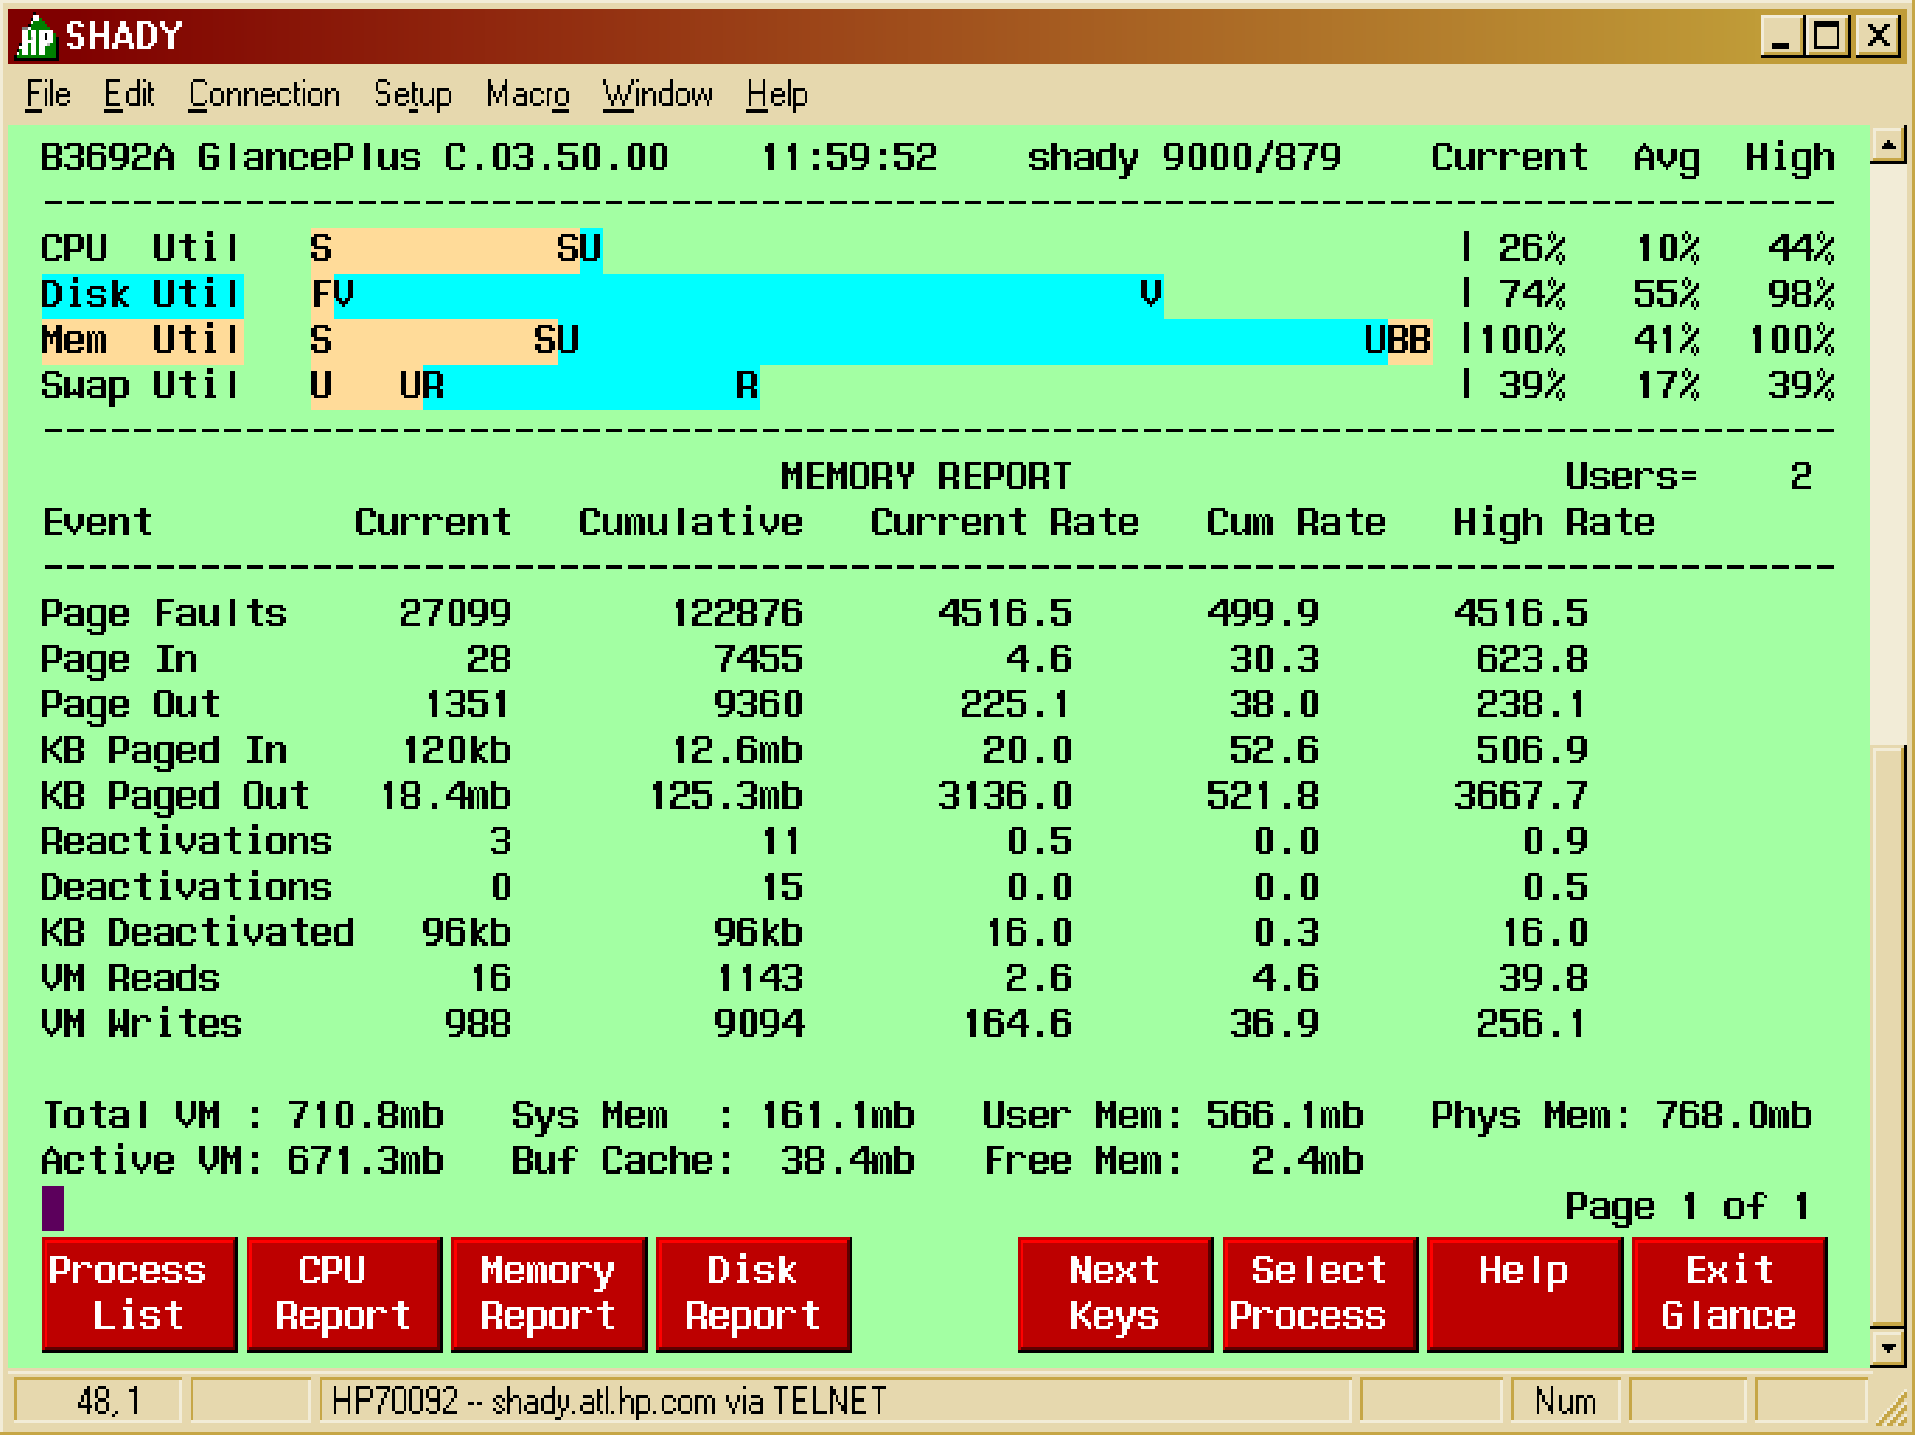
SHADY
File Edit Connection Setup Macro Window Help
B3692A GlancePlus C.03.50.00 11:59:52 shady 9000/879 Current Avg High
CPU Util S SU | 26% 10% 44%
Disk Util FV V | 74% 55% 98%
Mem Util S SU UBB |100% 41% 100%
Swap Util U UR R | 39% 17% 39%
MEMORY REPORT Users= 2
Event Current Cumulative Current Rate Cum Rate High Rate
Page Faults 27099 122876 4516.5 499.9 4516.5
Page In 28 7455 4.6 30.3 623.8
Page Out 1351 9360 225.1 38.0 238.1
KB Paged In 120kb 12.6mb 20.0 52.6 506.9
KB Paged Out 18.4mb 125.3mb 3136.0 521.8 3667.7
Reactivations 3 11 0.5 0.0 0.9
Deactivations 0 15 0.0 0.0 0.5
KB Deactivated 96kb 96kb 16.0 0.3 16.0
VM Reads 16 1143 2.6 4.6 39.8
VM Writes 988 9094 164.6 36.9 256.1
Total VM : 710.8mb Sys Mem : 161.1mb User Mem: 566.1mb Phys Mem: 768.0mb
Active VM: 671.3mb Buf Cache: 38.4mb Free Mem: 2.4mb
Page 1 of 1
Process List
CPU Report
Memory Report
Disk Report
Next Keys
Select Process
Help
Exit Glance
48,1
HP70092 -- shady.atl.hp.com via TELNET
Num

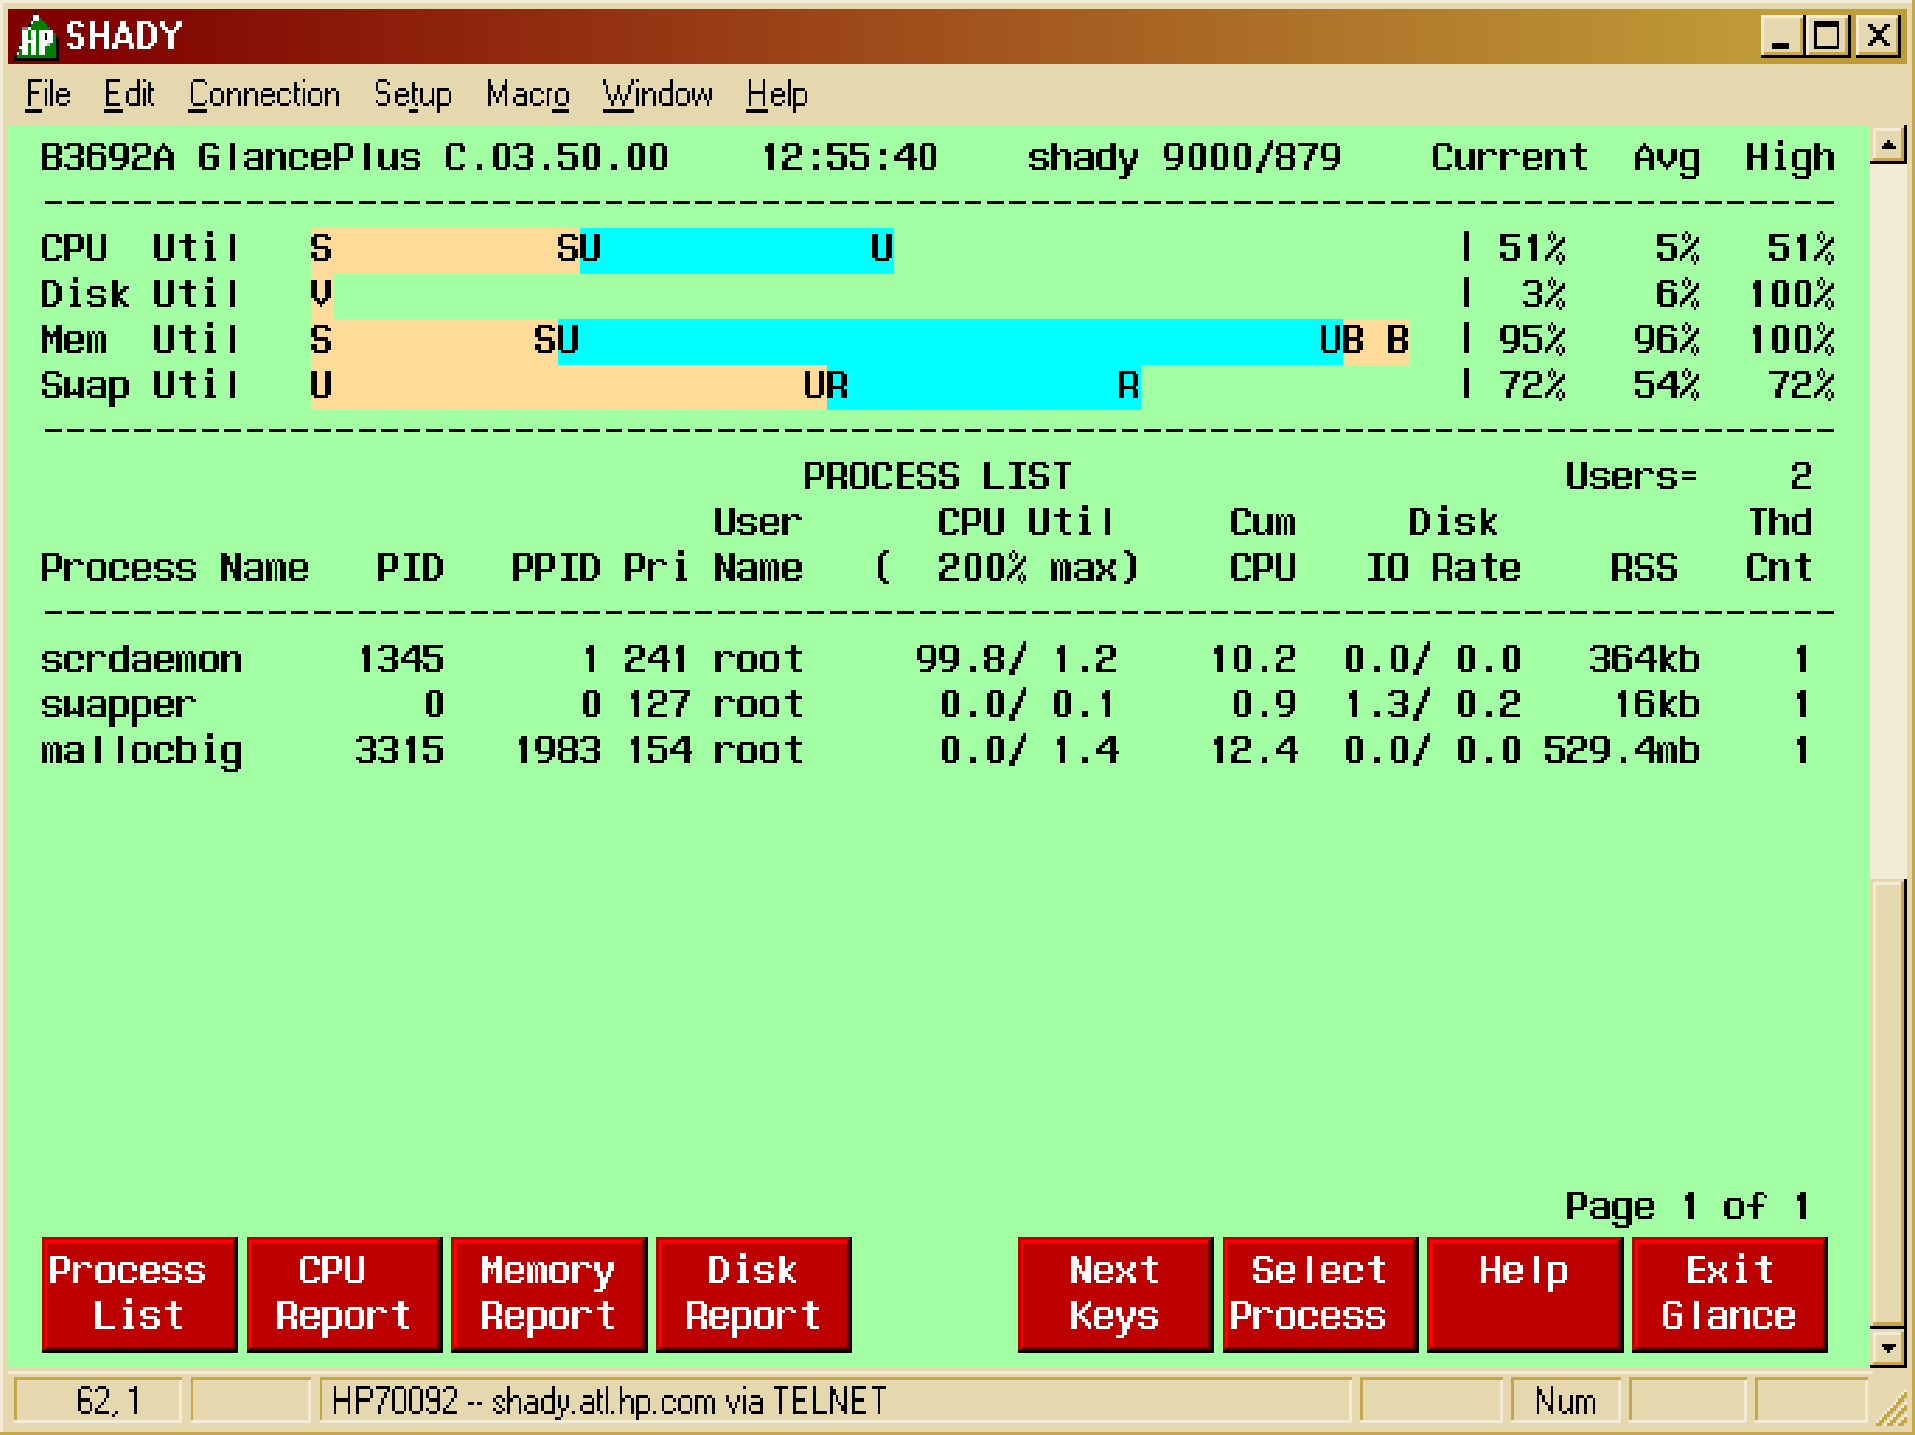
SHADY
File Edit Connection Setup Macro Window Help
B3692A GlancePlus C.03.50.00 12:55:40 shady 9000/879 Current Avg High
CPU Util S SU U | 51% 5% 51%
Disk Util V | 3% 6% 100%
Mem Util S SU UB B | 95% 96% 100%
Swap Util U UR R | 72% 54% 72%
PROCESS LIST Users= 2
Process Name PID PPID Pri User Name CPU Util (200% max) Cum CPU Disk IO Rate RSS Thd Cnt
scrdaemon 1345 1 241 root 99.8/ 1.2 10.2 0.0/ 0.0 364kb 1
swapper 0 0 127 root 0.0/ 0.1 0.9 1.3/ 0.2 16kb 1
mallocbig 3315 1983 154 root 0.0/ 1.4 12.4 0.0/ 0.0 529.4mb 1
Page 1 of 1
Process List
CPU Report
Memory Report
Disk Report
Next Keys
Select Process
Help
Exit Glance
62,1
HP70092 -- shady.atl.hp.com via TELNET
Num

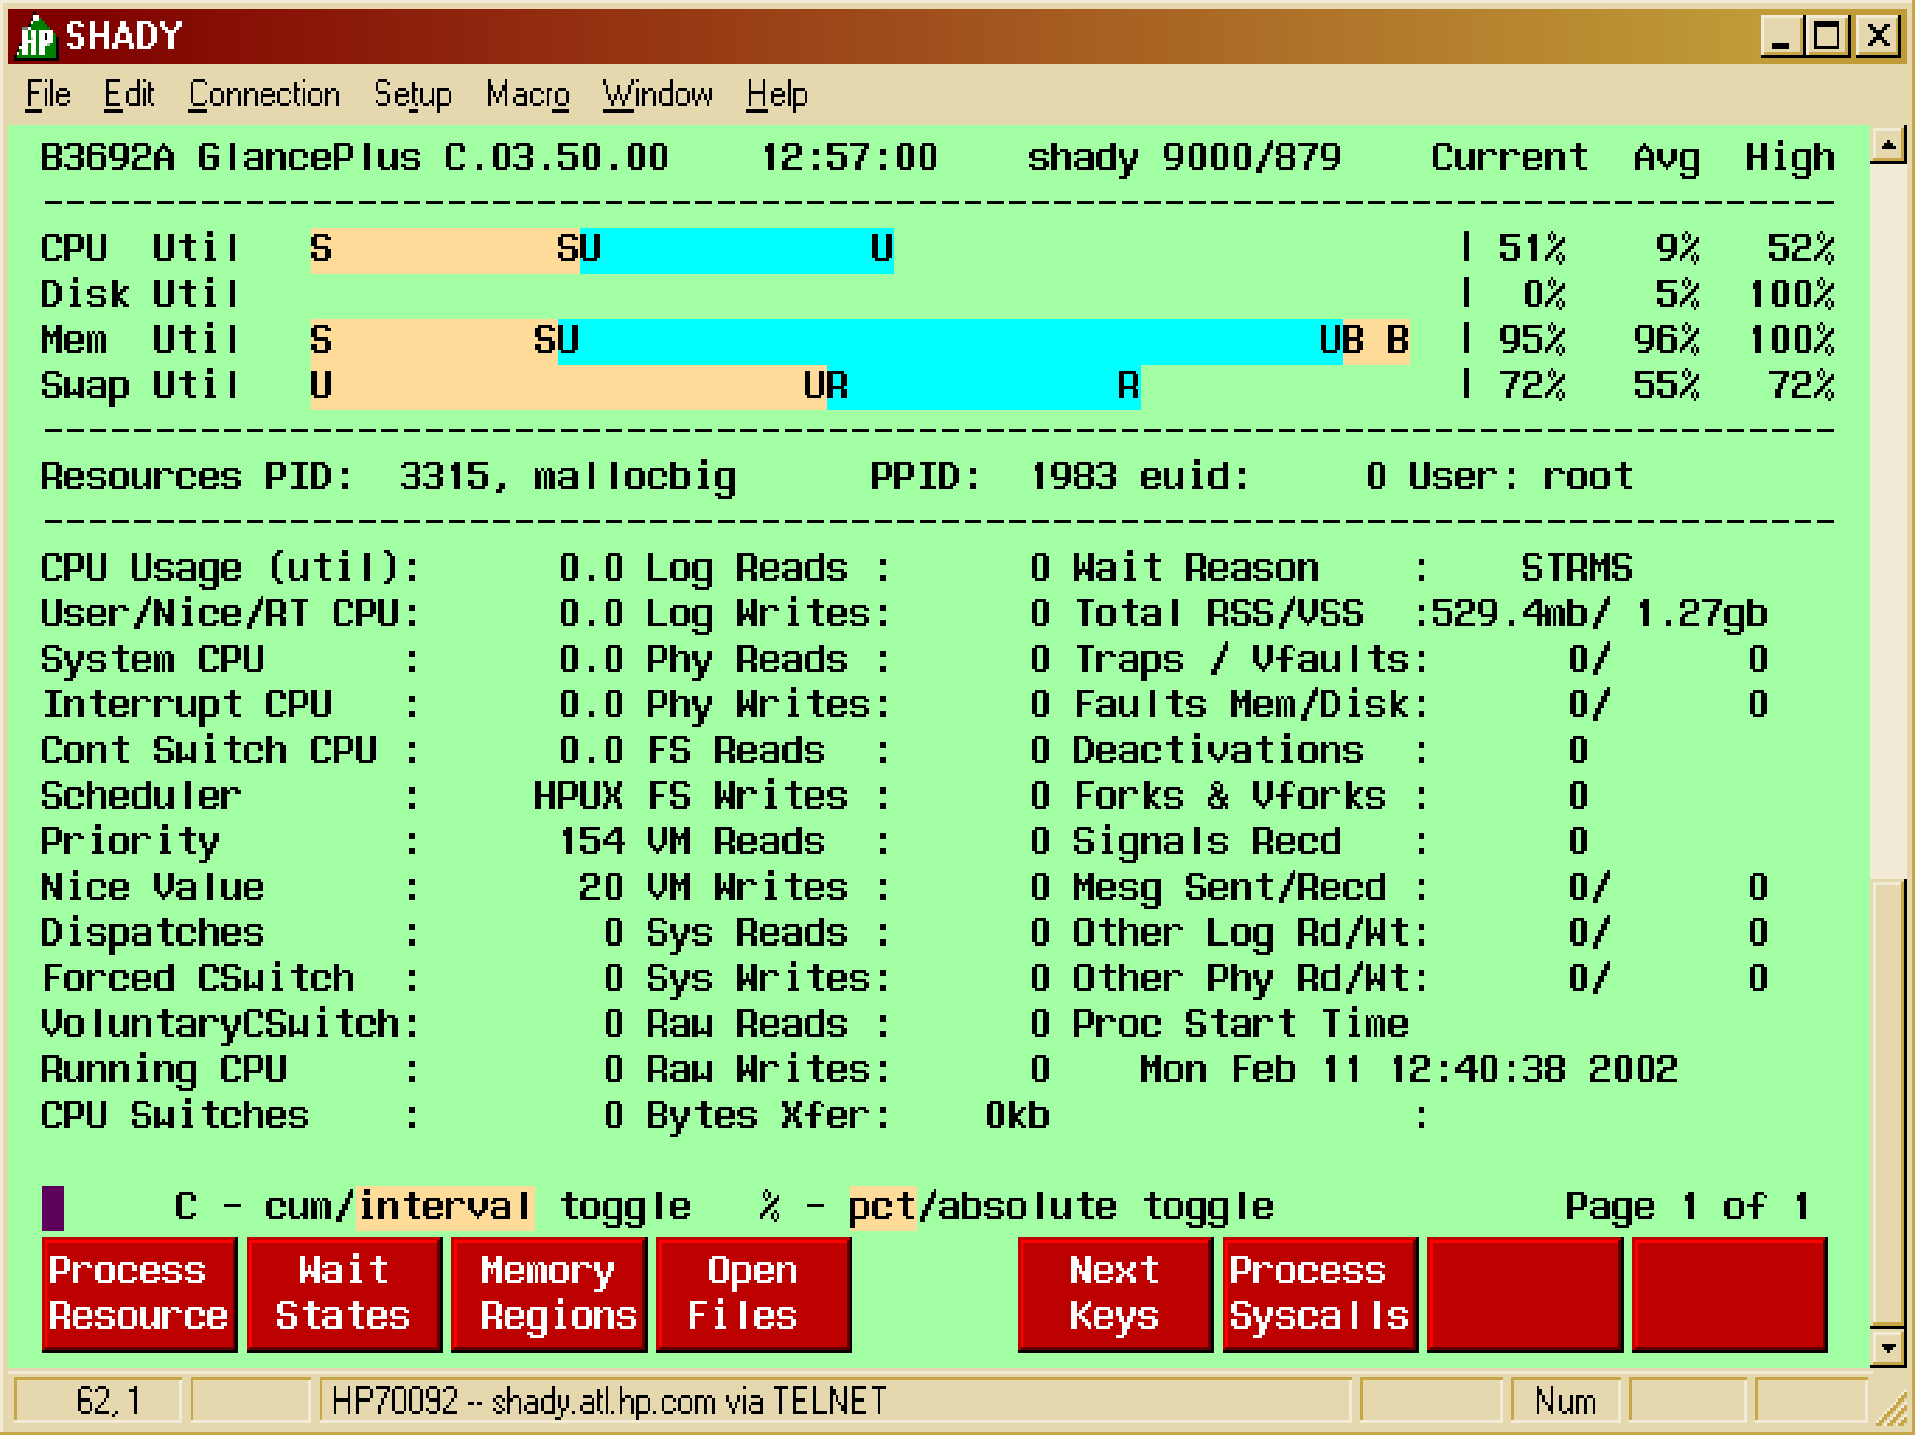
SHADY
File Edit Connection Setup Macro Window Help
B3692A GlancePlus C.03.50.00 12:57:00 shady 9000/879 Current Avg High
CPU Util S SU U | 51% 9% 52%
Disk Util | 0% 5% 100%
Mem Util S SU UB B | 95% 96% 100%
Swap Util U UR R | 72% 55% 72%
Resources PID: 3315, mallocbig PPID: 1983 euid: 0 User: root
CPU Usage (util): 0.0 Log Reads : 0 Wait Reason : STRMS
User/Nice/RT CPU: 0.0 Log Writes: 0 Total RSS/VSS :529.4mb/ 1.27gb
System CPU : 0.0 Phy Reads : 0 Traps / Vfaults: 0/ 0
Interrupt CPU : 0.0 Phy Writes: 0 Faults Mem/Disk: 0/ 0
Cont Switch CPU : 0.0 FS Reads : 0 Deactivations : 0
Scheduler : HPUX FS Writes : 0 Forks & Vforks : 0
Priority : 154 VM Reads : 0 Signals Recd : 0
Nice Value : 20 VM Writes : 0 Mesg Sent/Recd : 0/ 0
Dispatches : 0 Sys Reads : 0 Other Log Rd/Wt: 0/ 0
Forced CSwitch : 0 Sys Writes: 0 Other Phy Rd/Wt: 0/ 0
VoluntaryCSwitch: 0 Raw Reads : 0 Proc Start Time
Running CPU : 0 Raw Writes: 0 Mon Feb 11 12:40:38 2002
CPU Switches : 0 Bytes Xfer: 0kb :
C - cum/interval toggle % - pct/absolute toggle Page 1 of 1
Process Resource
Wait States
Memory Regions
Open Files
Next Keys
Process Syscalls
62,1
HP70092 -- shady.atl.hp.com via TELNET
Num

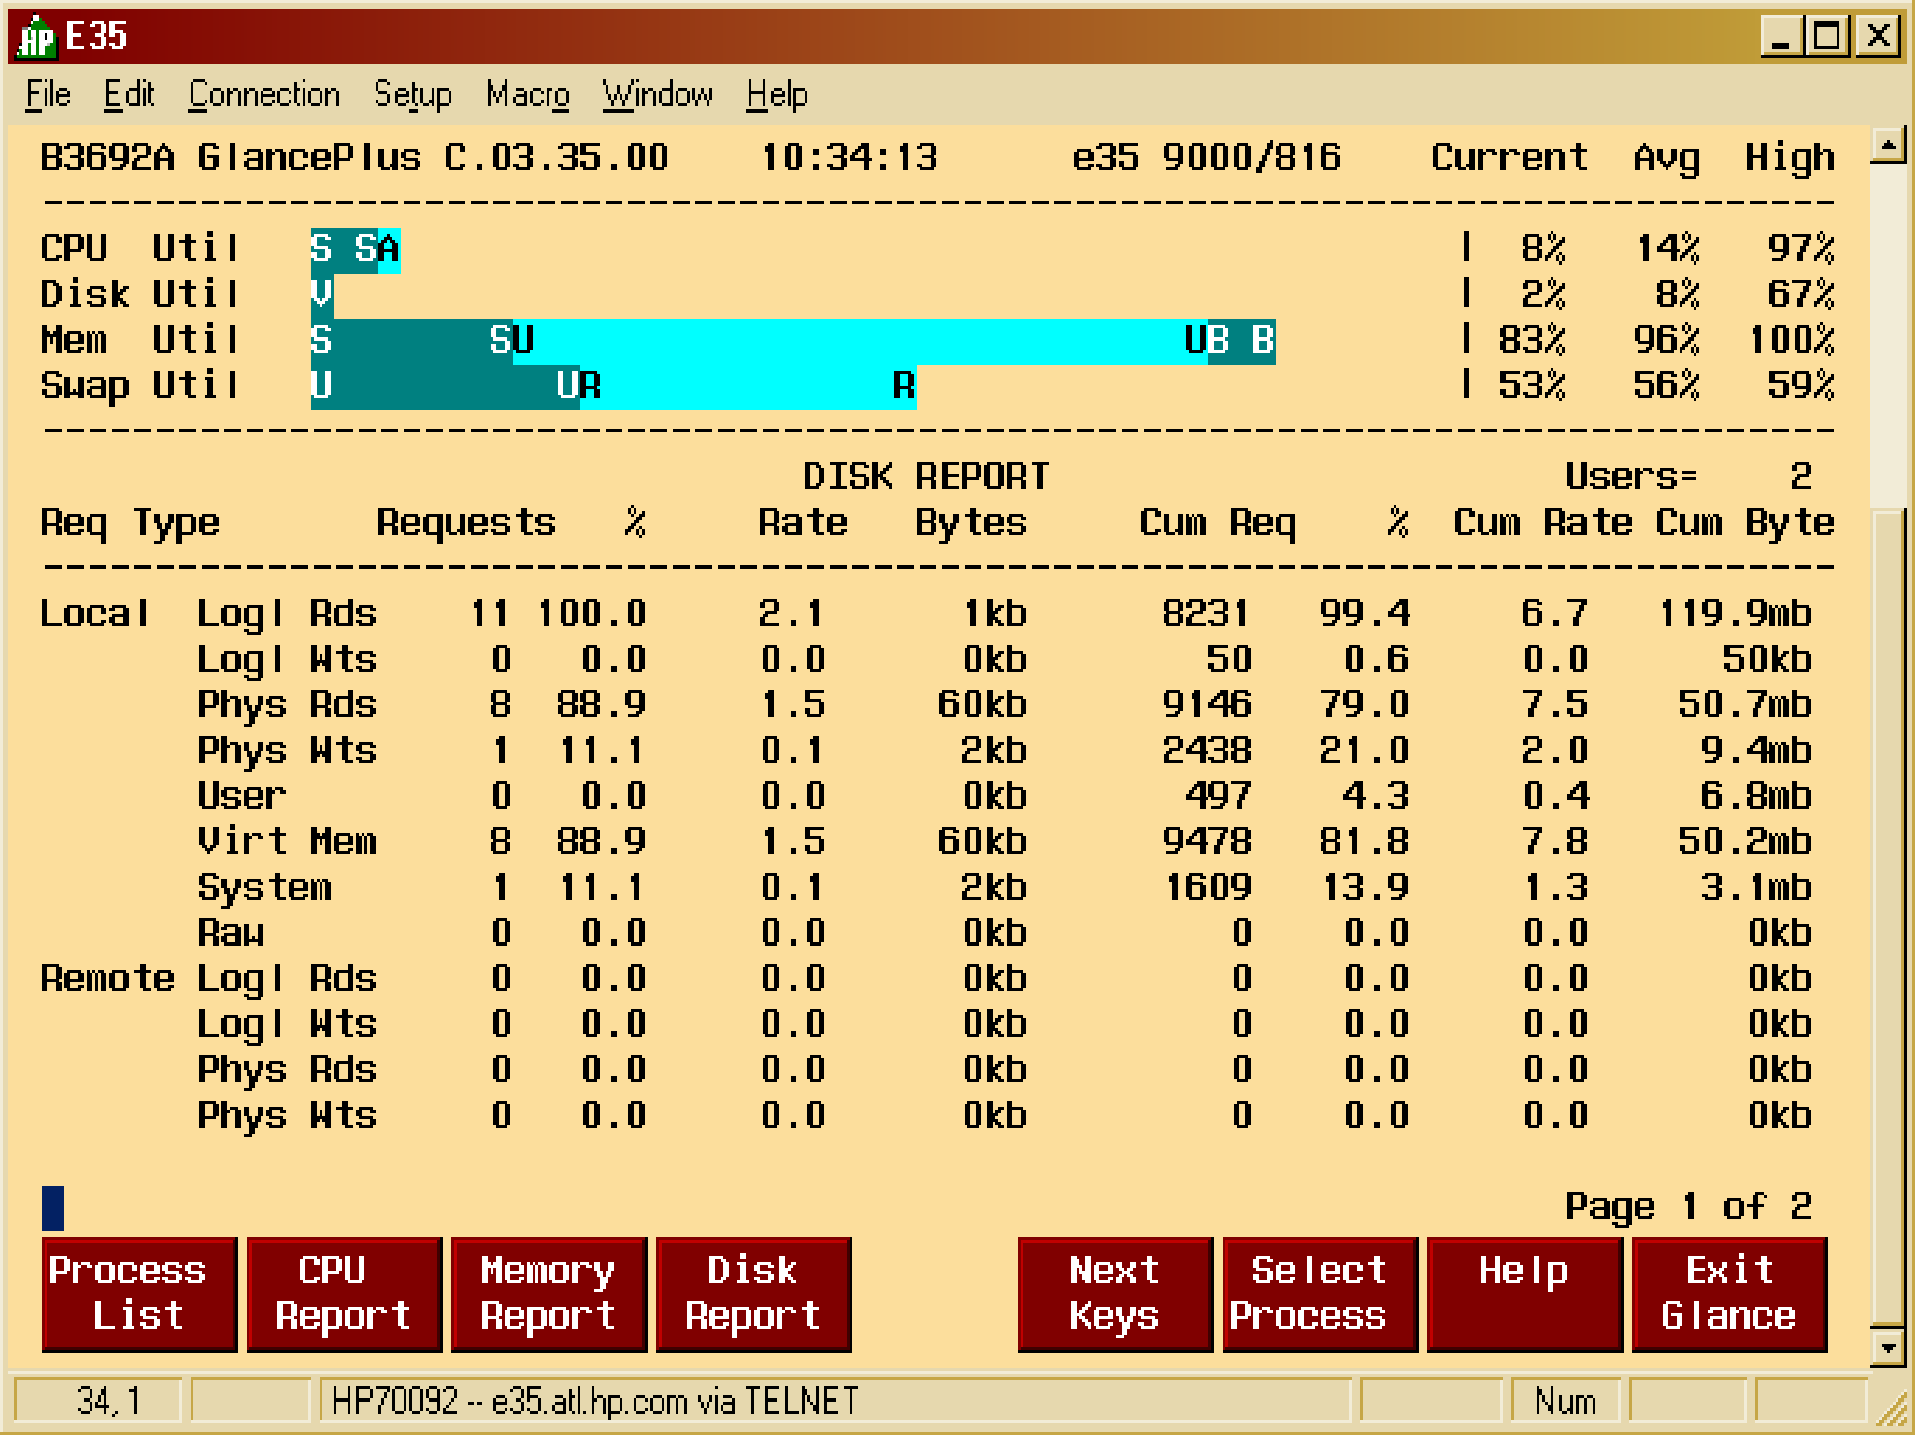
E35
File Edit Connection Setup Macro Window Help
B3692A GlancePlus C.03.35.00 10:34:13 e35 9000/816 Current Avg High
CPU Util | 8% 14% 97%
Disk Util | 2% 8% 67%
Mem Util | 83% 96% 100%
Swap Util | 53% 56% 59%
DISK REPORT
Users= 2
Req Type Requests % Rate Bytes Cum Req % Cum Rate Cum Byte
Local Logl Rds 11 100.0 2.1 1kb 8231 99.4 6.7 119.9mb
Logl Wts 0 0.0 0.0 0kb 50 0.6 0.0 50kb
Phys Rds 8 88.9 1.5 60kb 9146 79.0 7.5 50.7mb
Phys Wts 1 11.1 0.1 2kb 2438 21.0 2.0 9.4mb
User 0 0.0 0.0 0kb 497 4.3 0.4 6.8mb
Virt Mem 8 88.9 1.5 60kb 9478 81.8 7.8 50.2mb
System 1 11.1 0.1 2kb 1609 13.9 1.3 3.1mb
Raw 0 0.0 0.0 0kb 0 0.0 0.0 0kb
Remote Logl Rds 0 0.0 0.0 0kb 0 0.0 0.0 0kb
Logl Wts 0 0.0 0.0 0kb 0 0.0 0.0 0kb
Phys Rds 0 0.0 0.0 0kb 0 0.0 0.0 0kb
Phys Wts 0 0.0 0.0 0kb 0 0.0 0.0 0kb
Page 1 of 2
Process List
CPU Report
Memory Report
Disk Report
Next Keys
Select Process
Help
Exit Glance
34,1
HP70092 -- e35.atl.hp.com via TELNET
Num

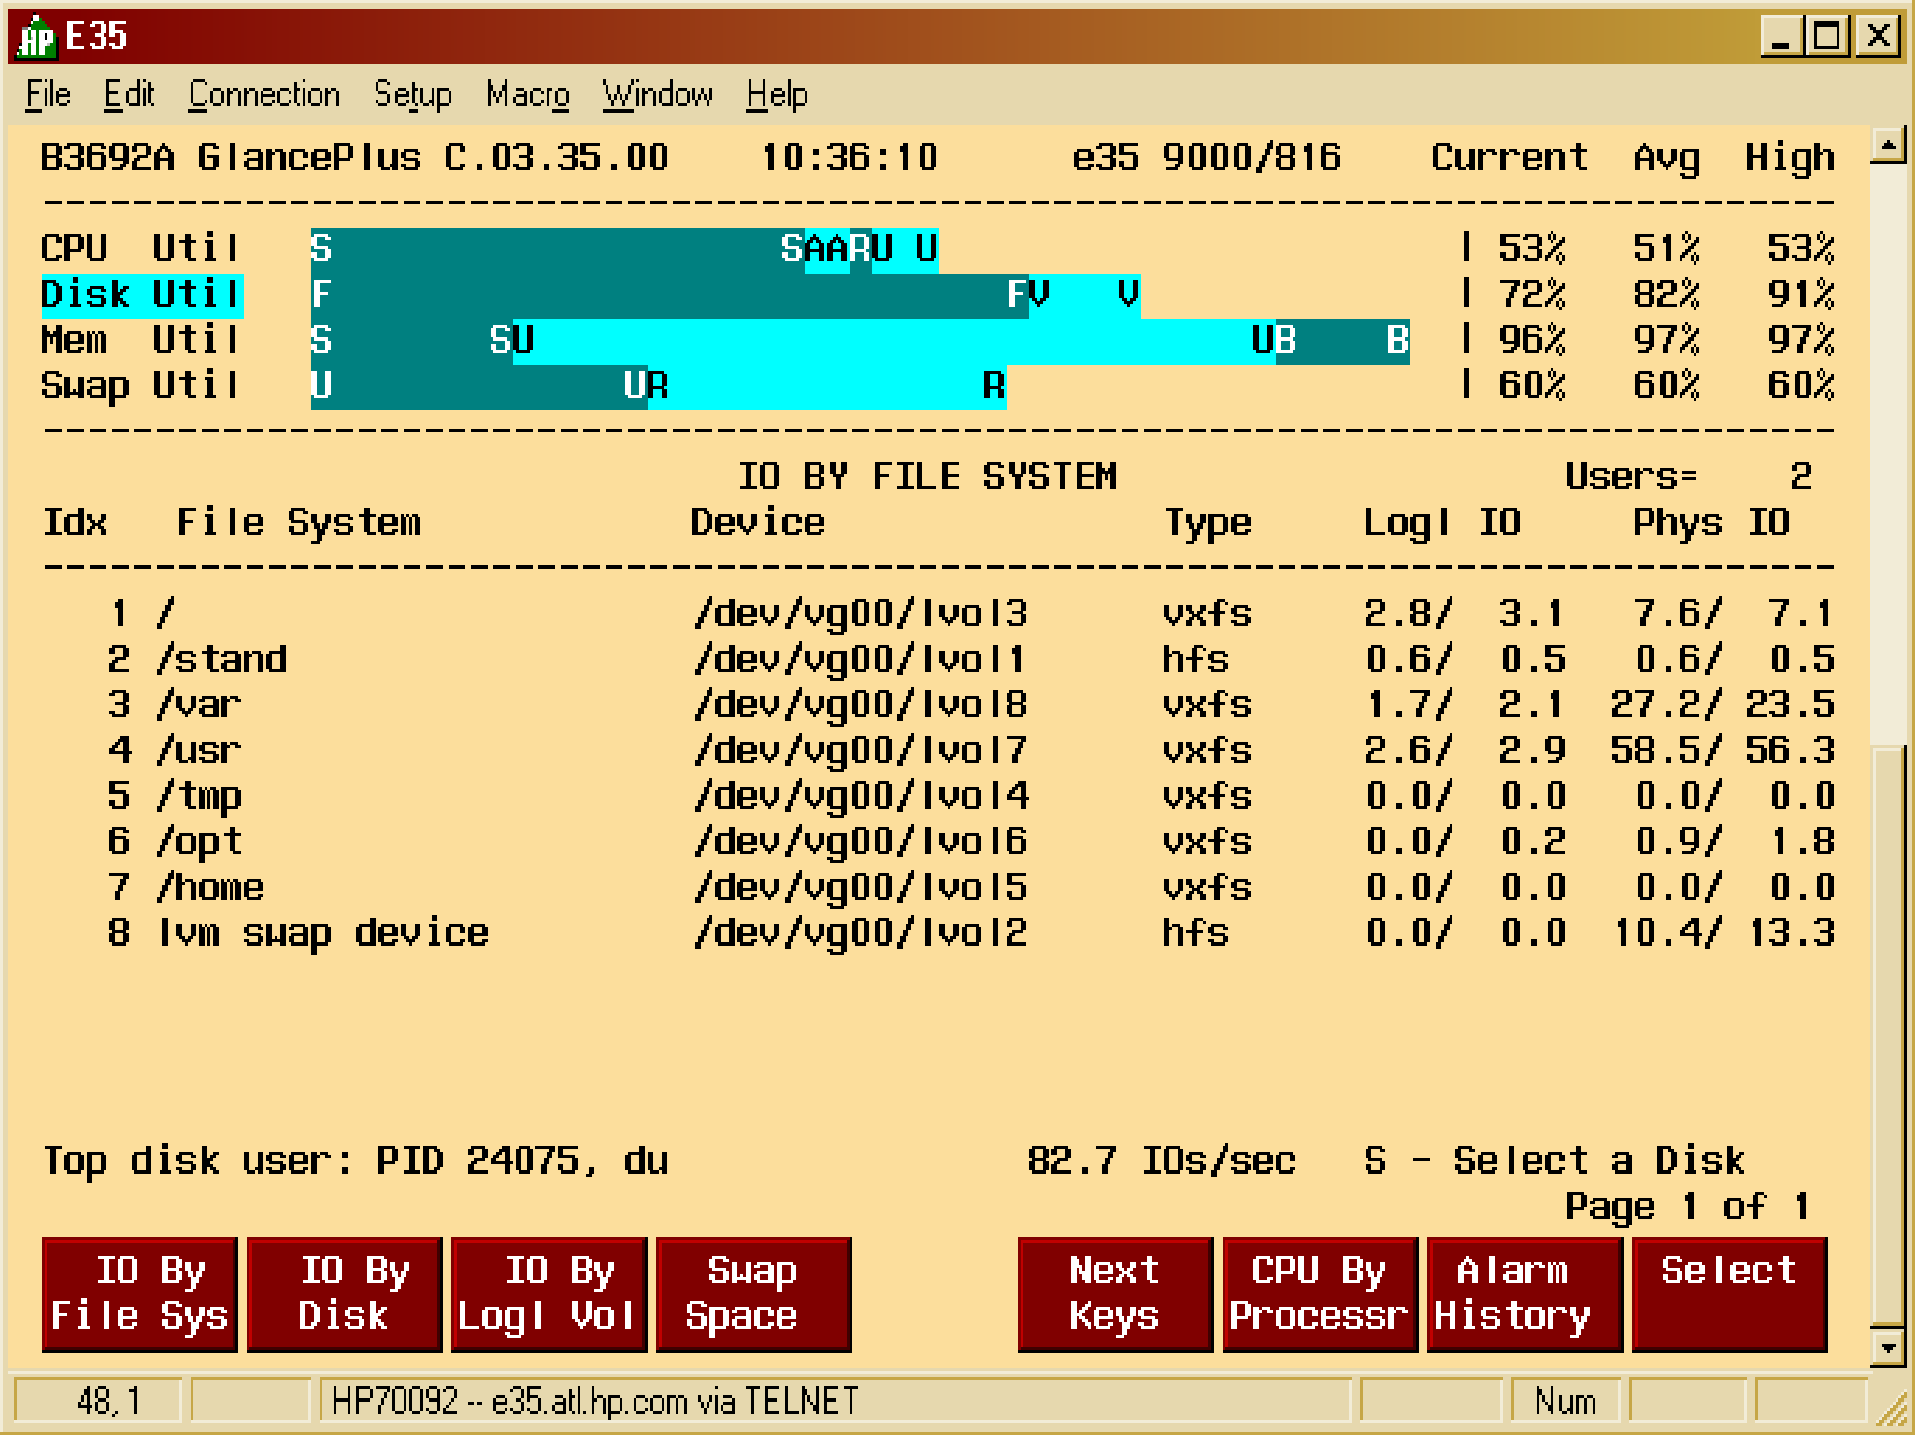
E35
File Edit Connection Setup Macro Window Help
B3692A GlancePlus C.03.35.00 10:36:10 e35 9000/816 Current Avg High
CPU Util | 53% 51% 53%
Disk Util | 72% 82% 91%
Mem Util | 96% 97% 97%
Swap Util | 60% 60% 60%
IO BY FILE SYSTEM Users= 2
Idx File System Device Type Logl IO Phys IO
1 / /dev/vg00/lvol3 vxfs 2.8/ 3.1 7.6/ 7.1
2 /stand /dev/vg00/lvol1 hfs 0.6/ 0.5 0.6/ 0.5
3 /var /dev/vg00/lvol8 vxfs 1.7/ 2.1 27.2/ 23.5
4 /usr /dev/vg00/lvol7 vxfs 2.6/ 2.9 58.5/ 56.3
5 /tmp /dev/vg00/lvol4 vxfs 0.0/ 0.0 0.0/ 0.0
6 /opt /dev/vg00/lvol6 vxfs 0.0/ 0.2 0.9/ 1.8
7 /home /dev/vg00/lvol5 vxfs 0.0/ 0.0 0.0/ 0.0
8 lvm swap device /dev/vg00/lvol2 hfs 0.0/ 0.0 10.4/ 13.3
Top disk user: PID 24075, du
82.7 IOs/sec
S - Select a Disk
Page 1 of 1
IO By File Sys
IO By Disk
IO By Logl Vol
Swap Space
Next Keys
CPU By Processr
Alarm History
Select
48,1
HP70092 -- e35.atl.hp.com via TELNET
Num

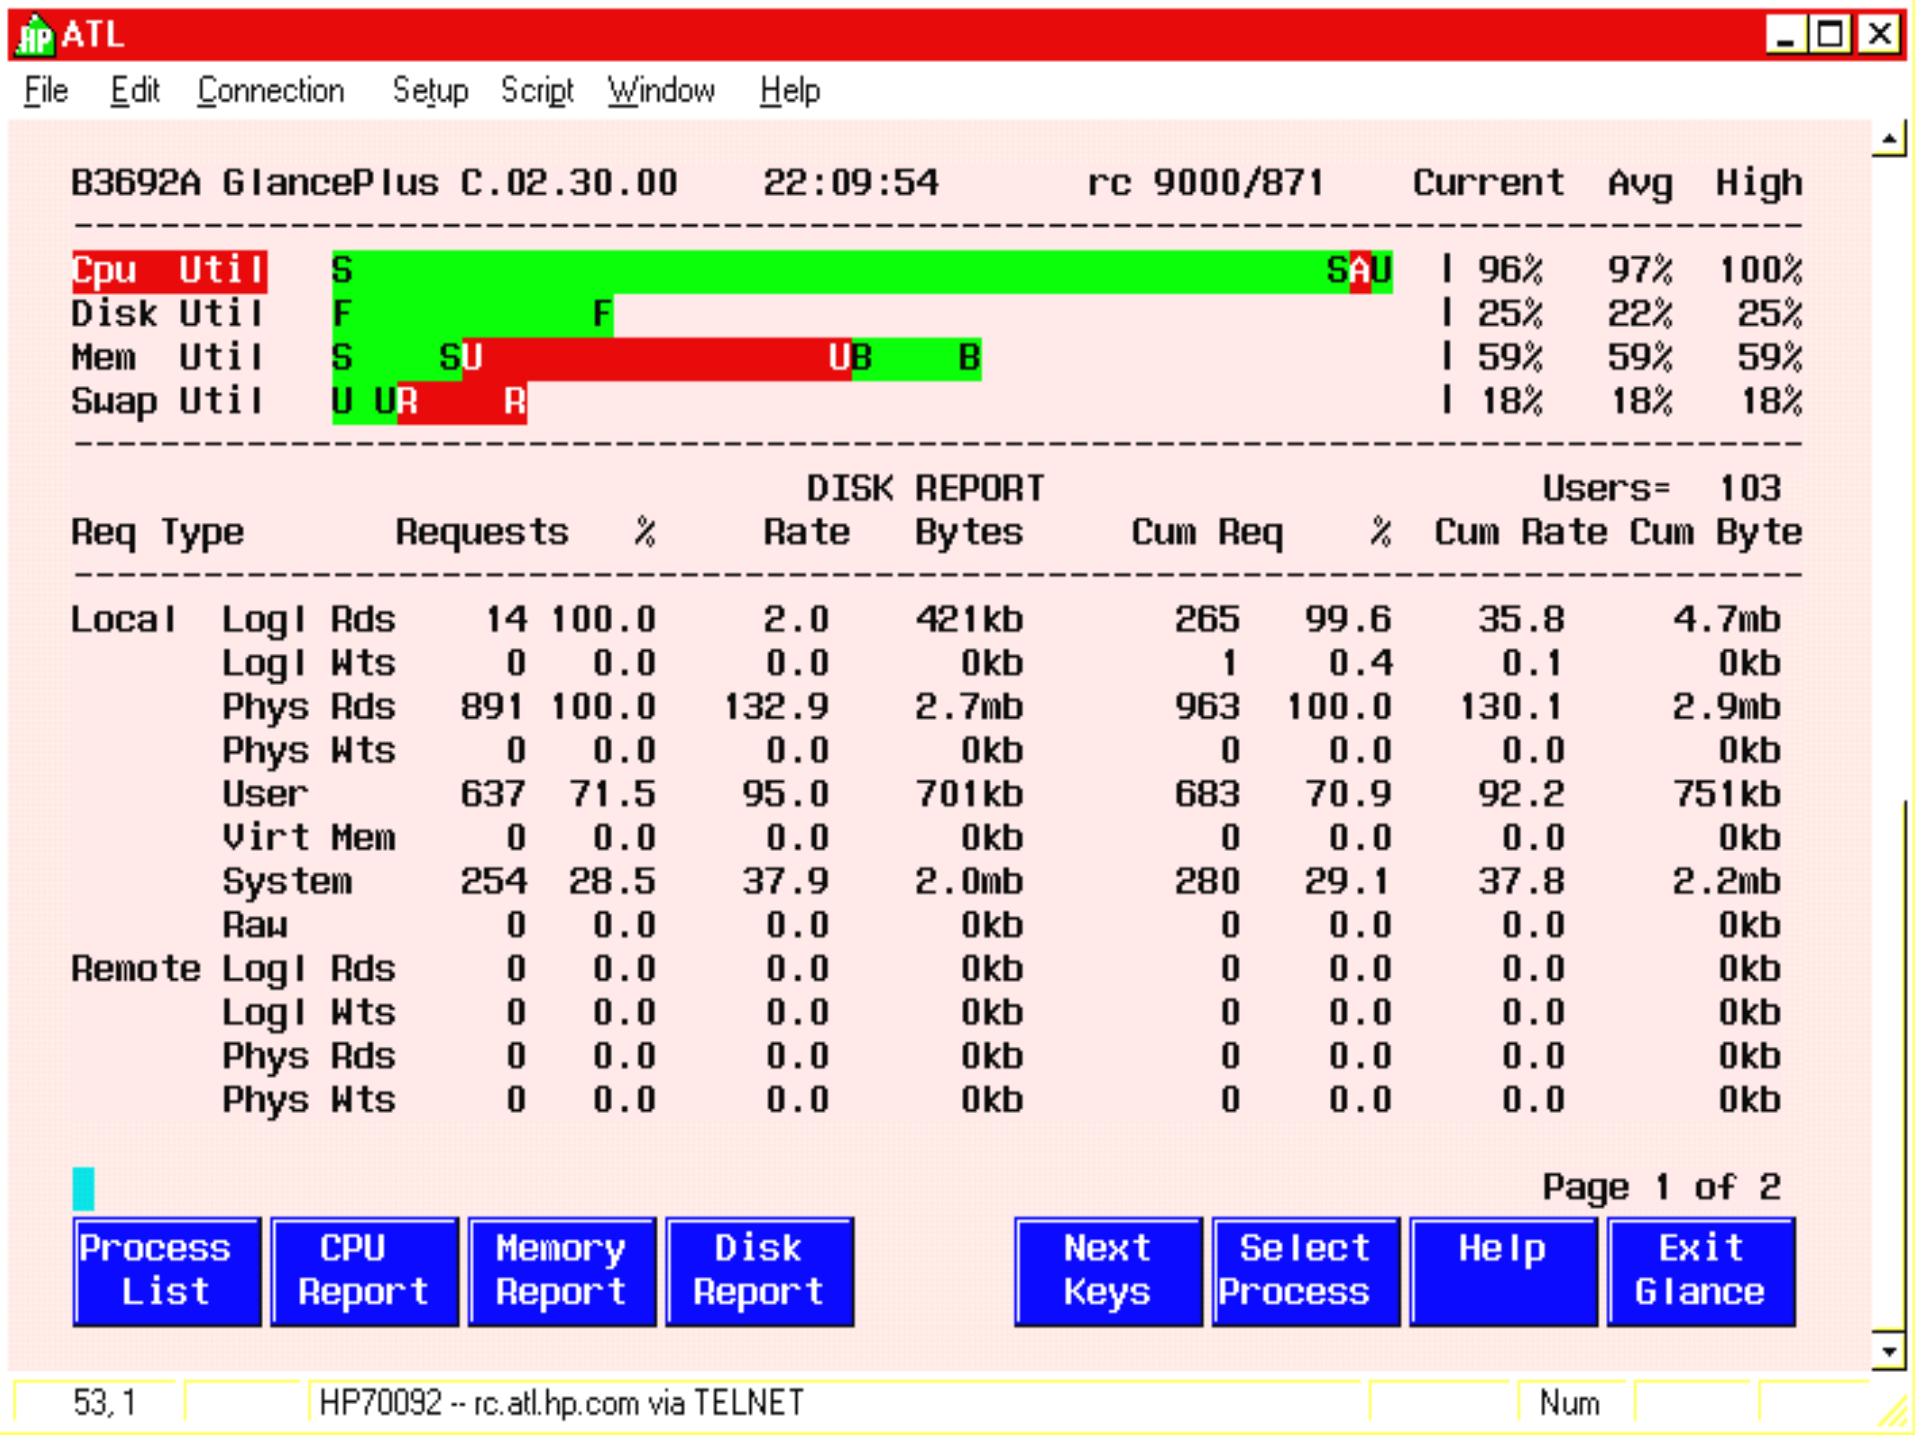
ATL
File Edit Connection Setup Script Window Help
B3692A GlancePlus C.02.30.00 22:09:54 rc 9000/871 Current Avg High
Cpu Util | 96% 97% 100%
Disk Util | 25% 22% 25%
Mem Util | 59% 59% 59%
Swap Util | 18% 18% 18%
DISK REPORT
Users= 103
Req Type Requests % Rate Bytes Cum Req % Cum Rate Cum Byte
Local Logl Rds 14 100.0 2.0 421kb 265 99.6 35.8 4.7mb
Logl Wts 0 0.0 0.0 0kb 1 0.4 0.1 0kb
Phys Rds 891 100.0 132.9 2.7mb 963 100.0 130.1 2.9mb
Phys Wts 0 0.0 0.0 0kb 0 0.0 0.0 0kb
User 637 71.5 95.0 701kb 683 70.9 92.2 751kb
Virt Mem 0 0.0 0.0 0kb 0 0.0 0.0 0kb
System 254 28.5 37.9 2.0mb 280 29.1 37.8 2.2mb
Raw 0 0.0 0.0 0kb 0 0.0 0.0 0kb
Remote Logl Rds 0 0.0 0.0 0kb 0 0.0 0.0 0kb
Logl Wts 0 0.0 0.0 0kb 0 0.0 0.0 0kb
Phys Rds 0 0.0 0.0 0kb 0 0.0 0.0 0kb
Phys Wts 0 0.0 0.0 0kb 0 0.0 0.0 0kb
Page 1 of 2
Process List
CPU Report
Memory Report
Disk Report
Next Keys
Select Process
Help
Exit Glance
53,1
HP70092 -- rc.atl.hp.com via TELNET
Num

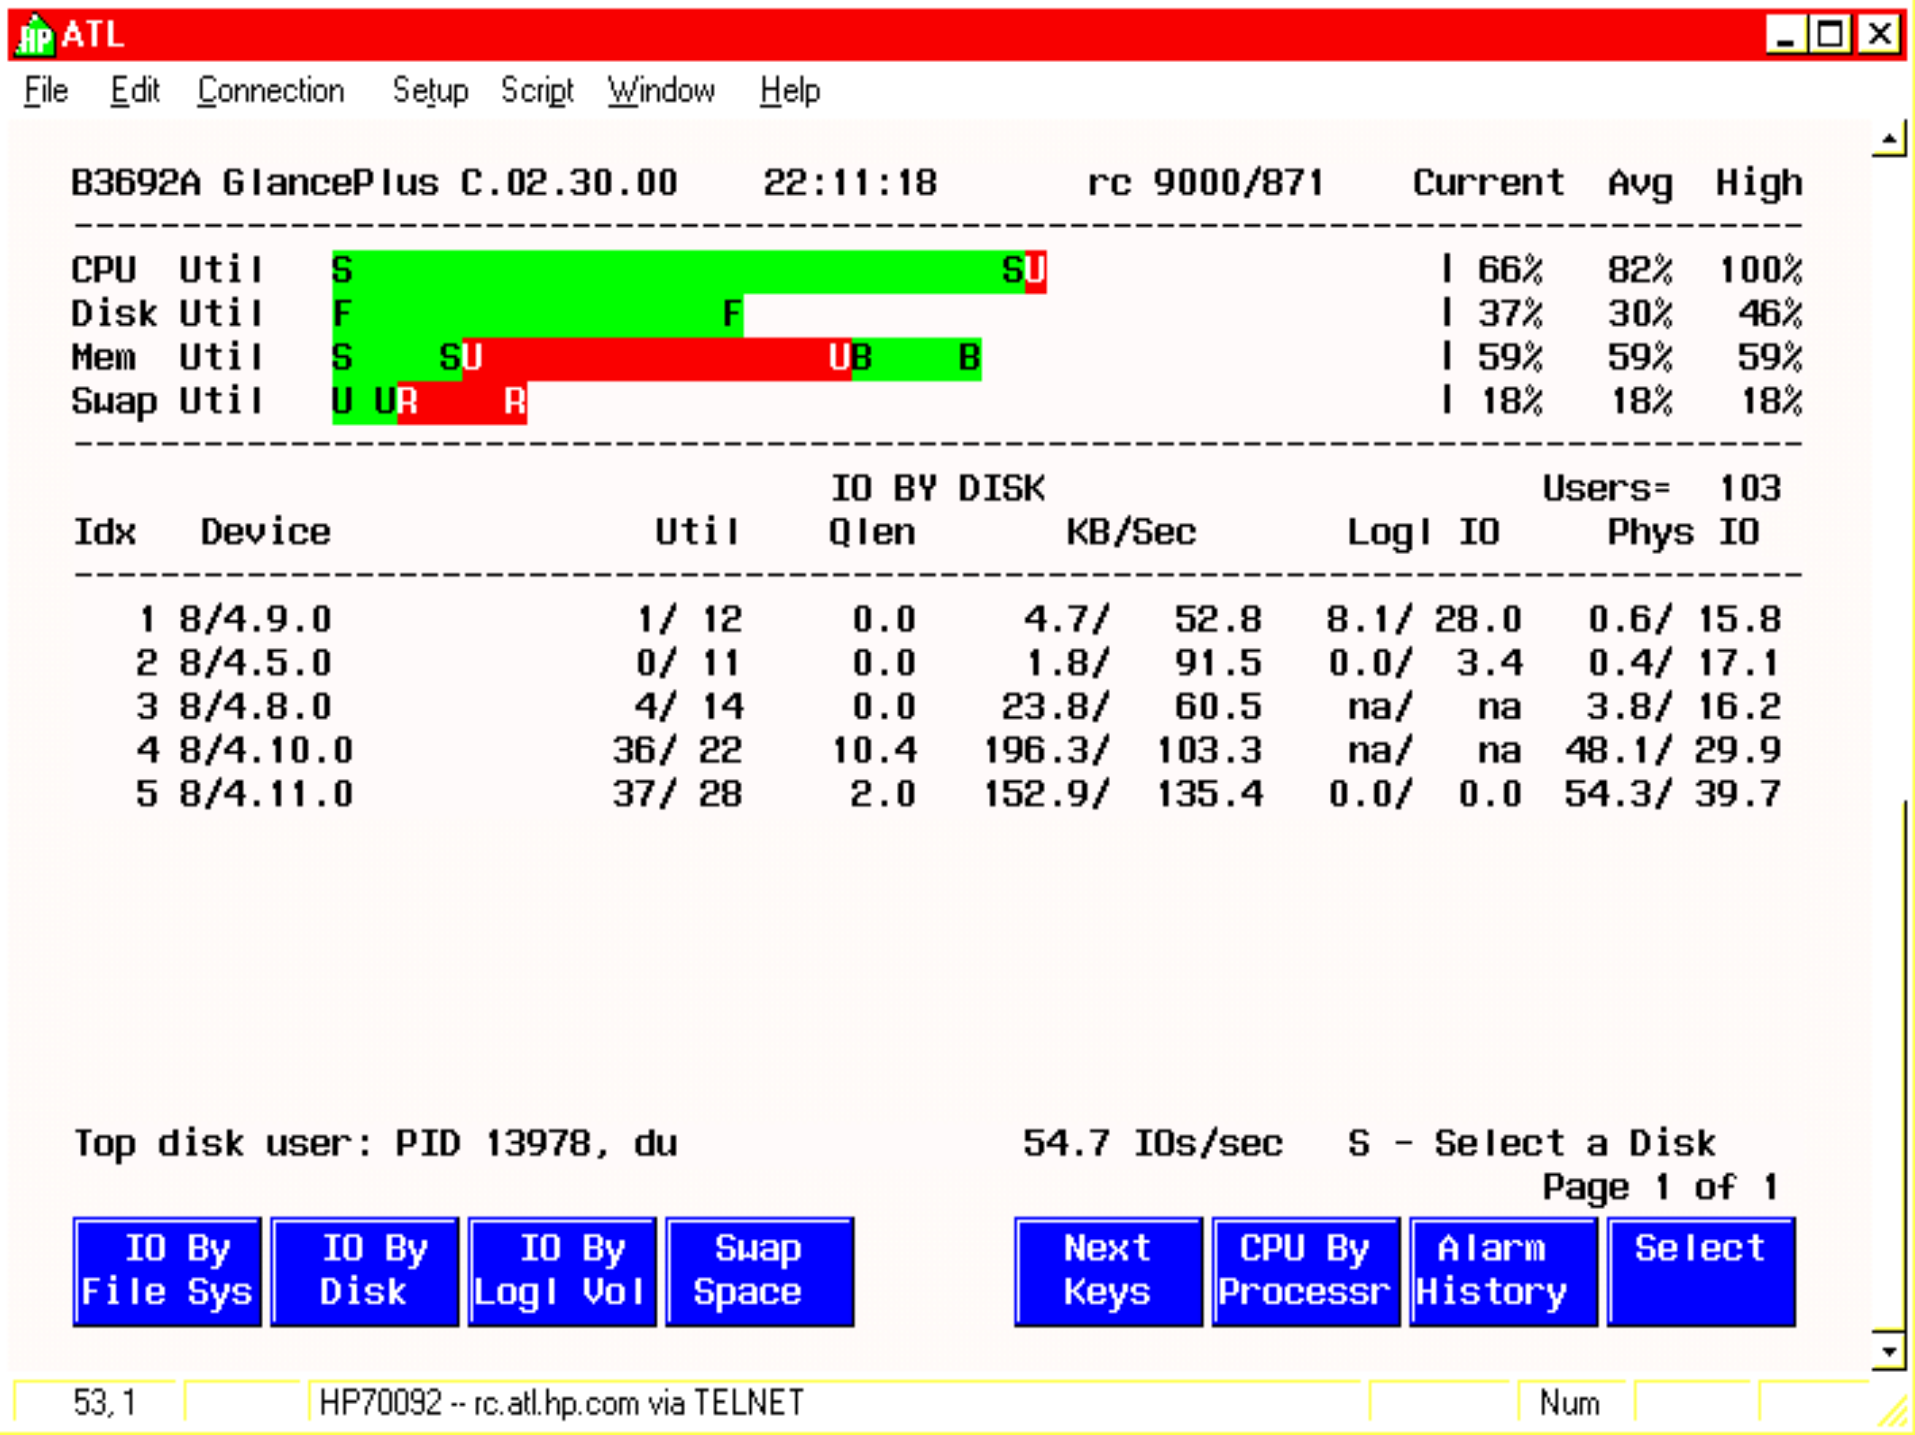
ATL
File Edit Connection Setup Script Window Help
B3692A GlancePlus C.02.30.00 22:11:18 rc 9000/871 Current Avg High
CPU Util | 66% 82% 100%
Disk Util | 37% 30% 46%
Mem Util | 59% 59% 59%
Swap Util | 18% 18% 18%
IO BY DISK
Users= 103
Idx Device Util Qlen KB/Sec Logl IO Phys IO
1 8/4.9.0 1/ 12 0.0 4.7/ 52.8 8.1/ 28.0 0.6/ 15.8
2 8/4.5.0 0/ 11 0.0 1.8/ 91.5 0.0/ 3.4 0.4/ 17.1
3 8/4.8.0 4/ 14 0.0 23.8/ 60.5 na/ na 3.8/ 16.2
4 8/4.10.0 36/ 22 10.4 196.3/ 103.3 na/ na 48.1/ 29.9
5 8/4.11.0 37/ 28 2.0 152.9/ 135.4 0.0/ 0.0 54.3/ 39.7
Top disk user: PID 13978, du
54.7 IOs/sec S - Select a Disk
Page 1 of 1
IO By File Sys
IO By Disk
IO By Logl Vol
Swap Space
Next Keys
CPU By Processr
Alarm History
Select
53,1
HP70092 -- rc.atl.hp.com via TELNET
Num

B3692A GlancePlus C.02.30.00 22:12:17 rc 9000/871

		Current	Avg	High
CPU Util	S SU	62%	72%	100%
Disk Util	F F	40%	35%	53%
Mem Util	S SU UB B	59%	59%	59%
Swap Util	U UR R	18%	18%	18%

SYSTEM TABLES REPORT Users= 103

System Table	Available	Used	Utilization	High(%)
Proc Table (nproc)	2000	590	30	30
File Table (nfile)	6010	1940	32	33
Shared Mem Table (shmmni)	200	9	5	5
Message Table (msgmni)	50	2	4	4
Semaphore Table (semmni)	64	8	13	13
File Locks (nflocks)	500	65	13	13
Pseudo Terminals (npty)	1400	148	11	11
Buffer Headers (nbuf)	23556	23556	100	100

Global Waits | Global Syscalls | DCE Global | System Tables | Next Keys | Netuk By Intrface | NFS Global | NFS By System

Performance and Kernel Tuning

Performance tuning

- gpm: Xwindows interface

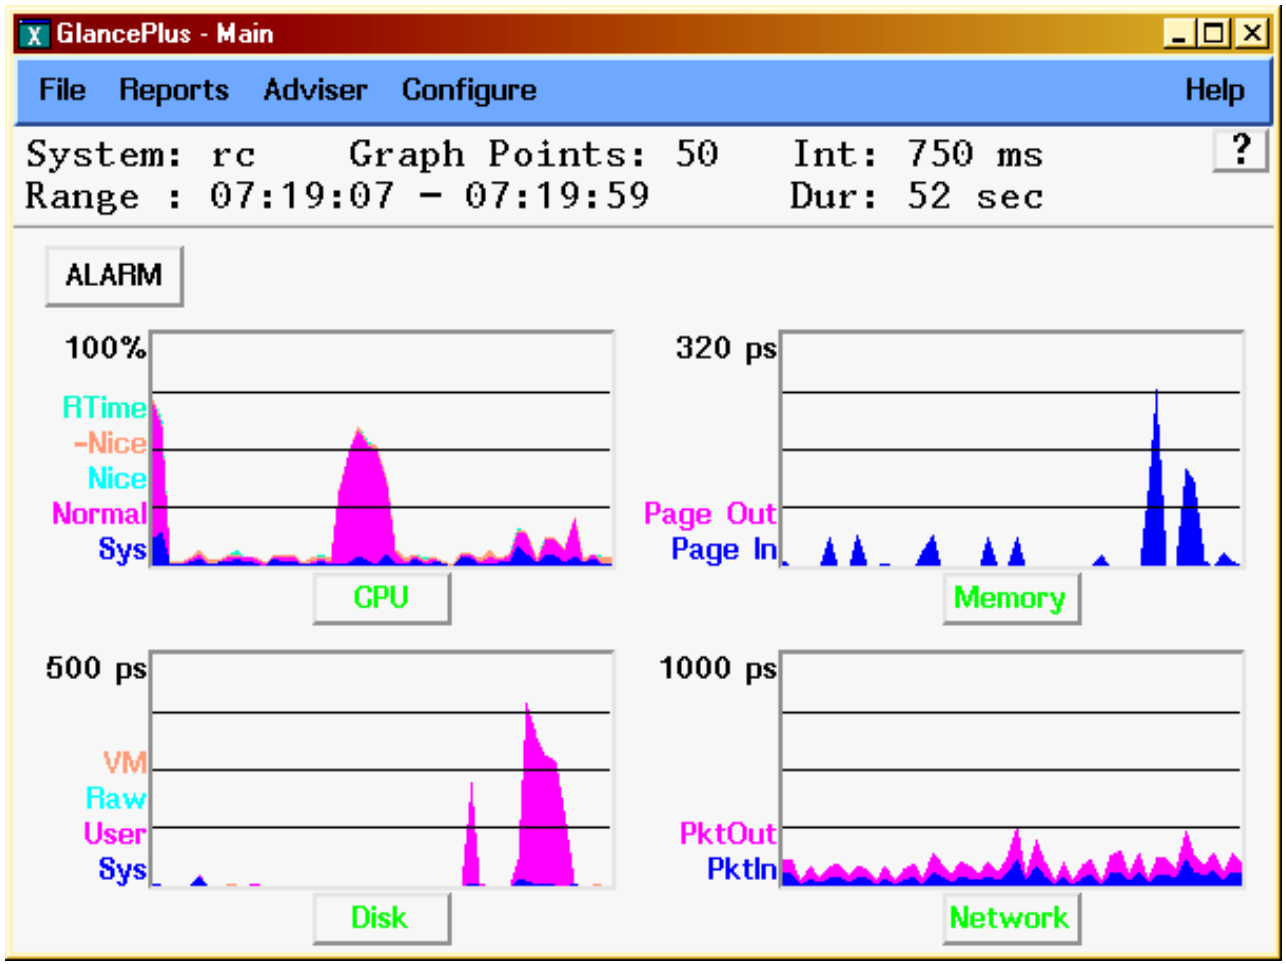

Performance and Kernel Tuning

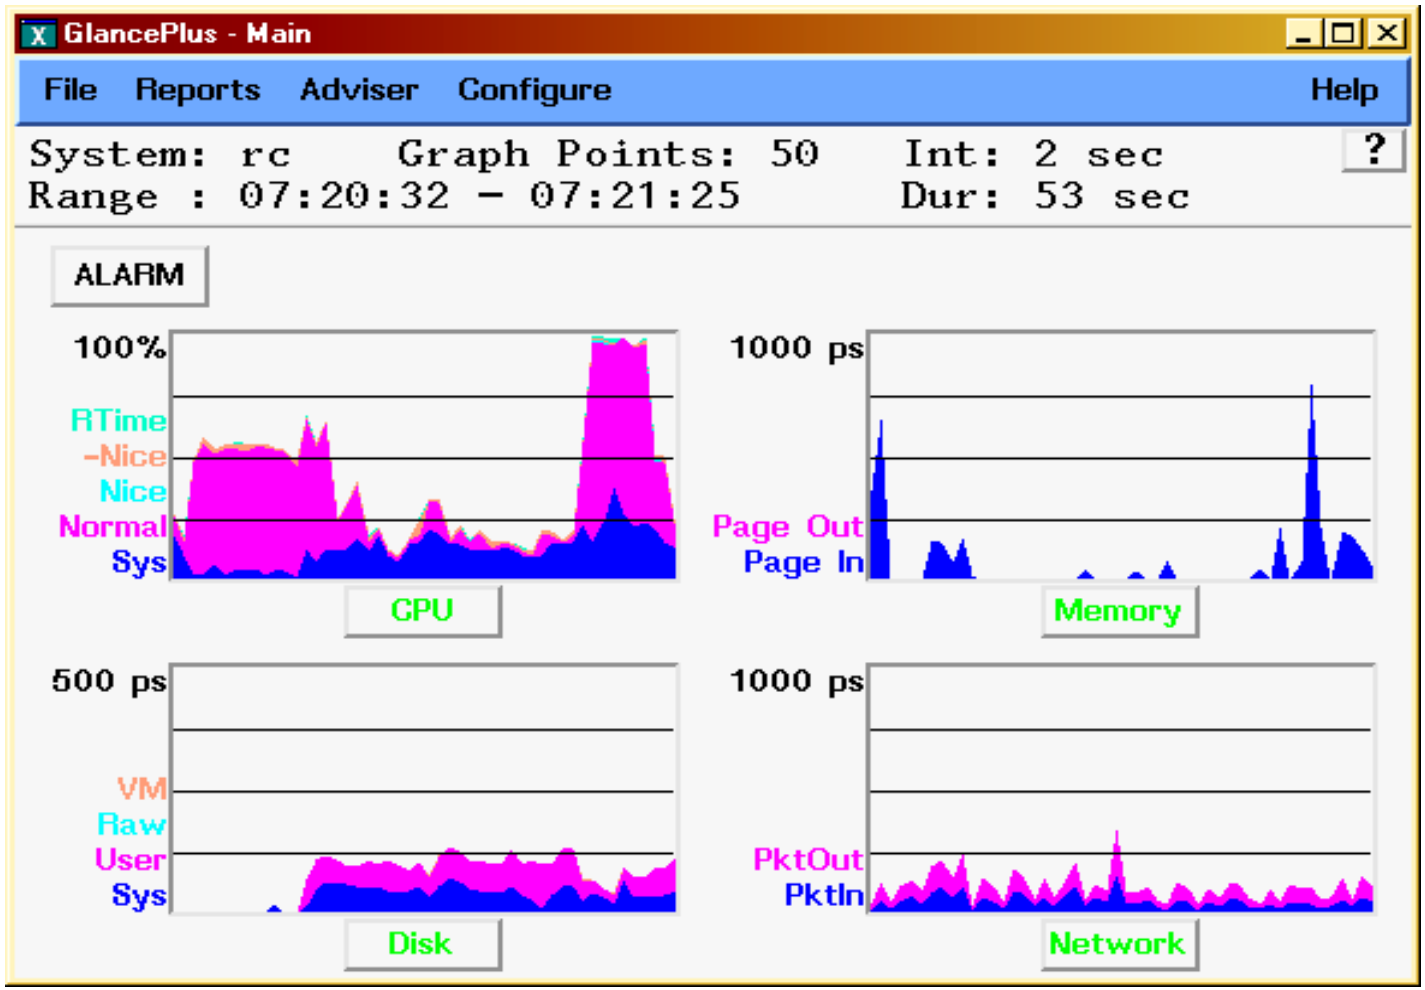

Performance and Kernel Tuning

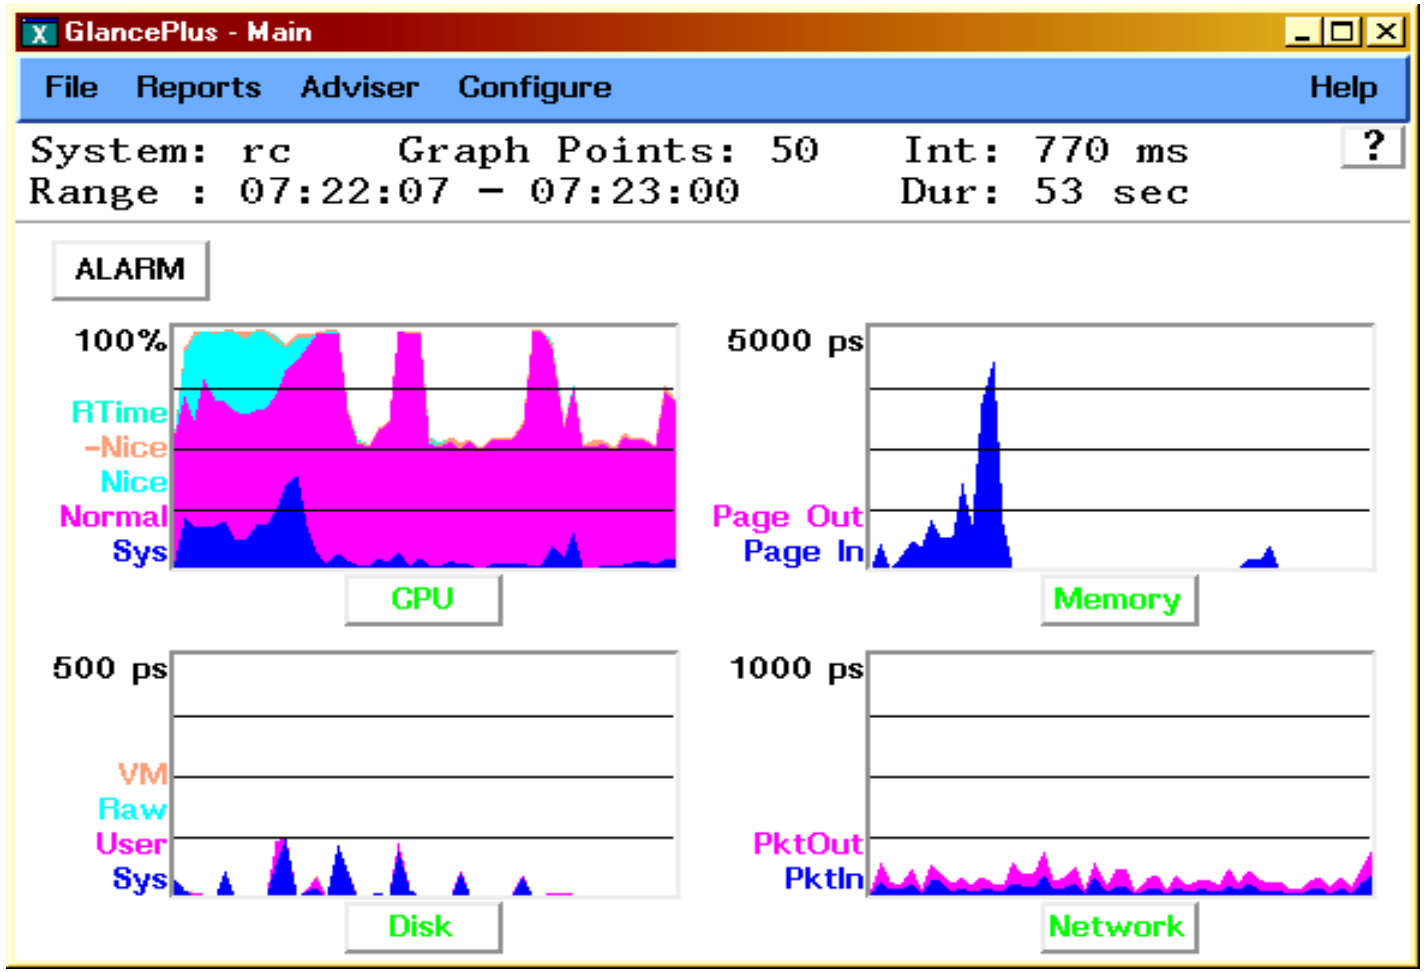

Performance and Kernel Tuning

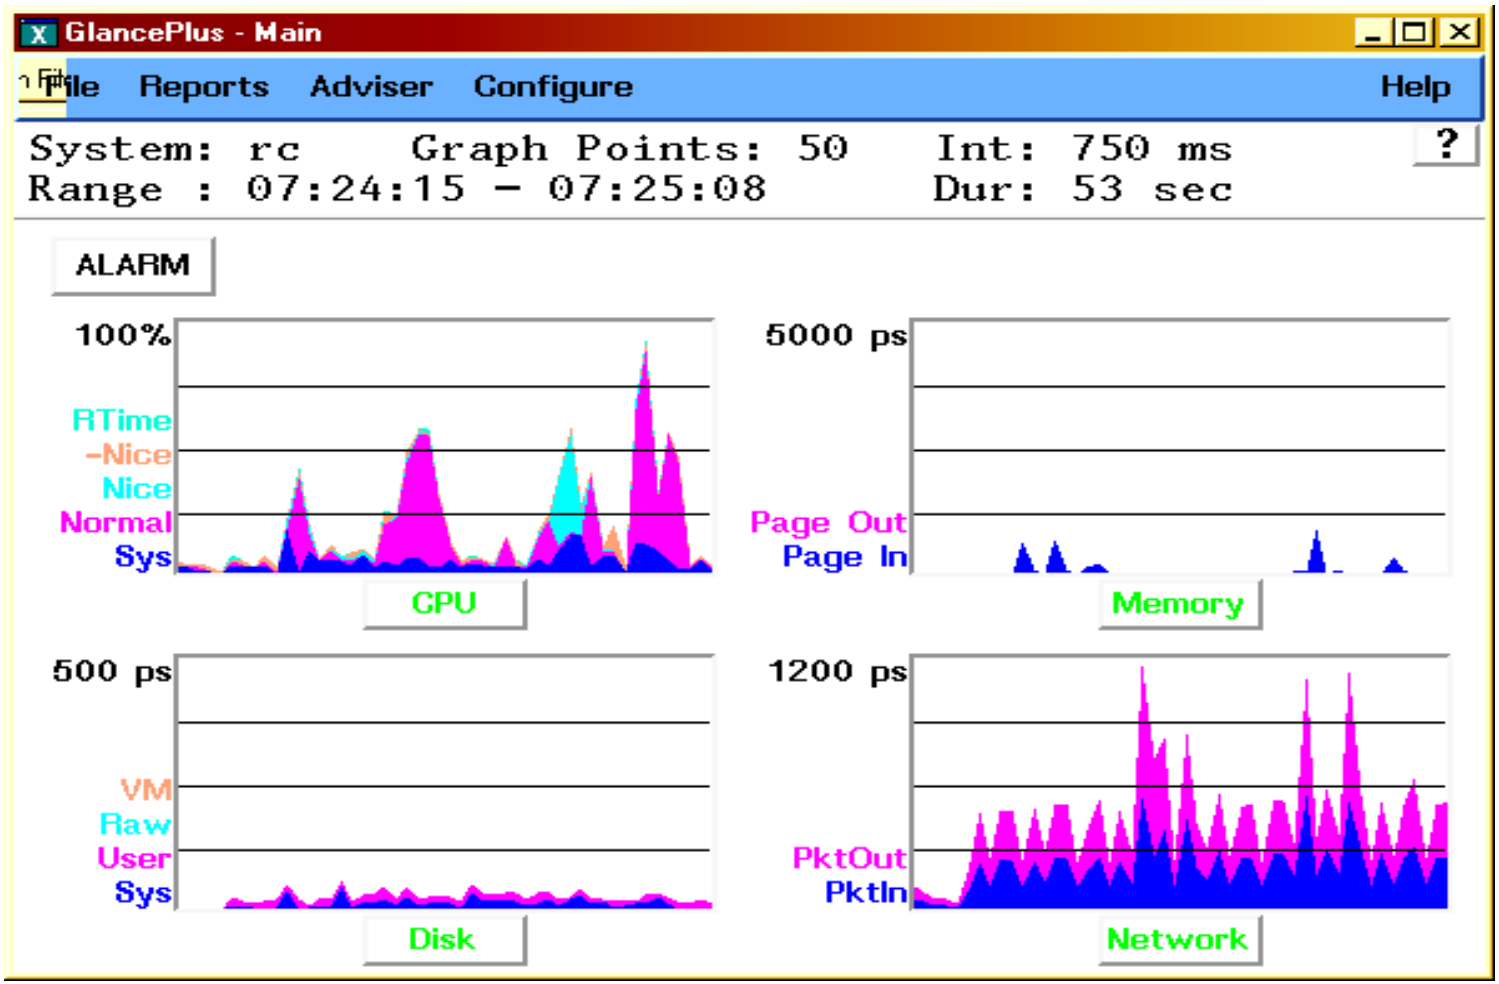

Performance and Kernel Tuning

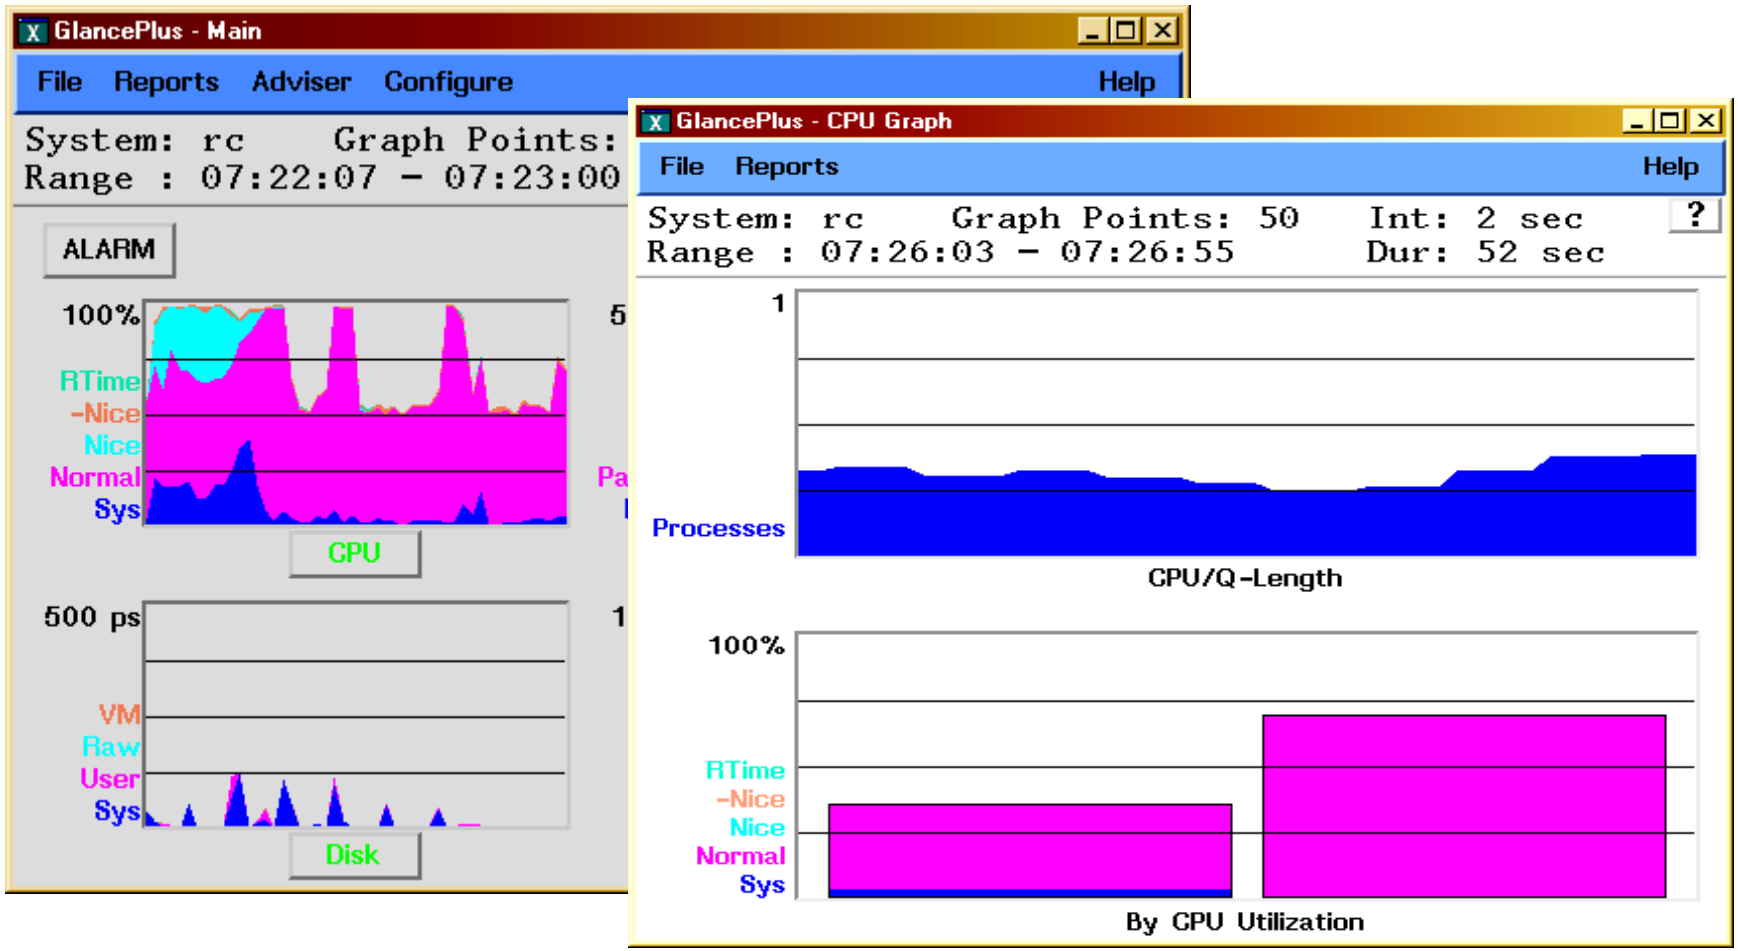

Performance and Kernel Tuning

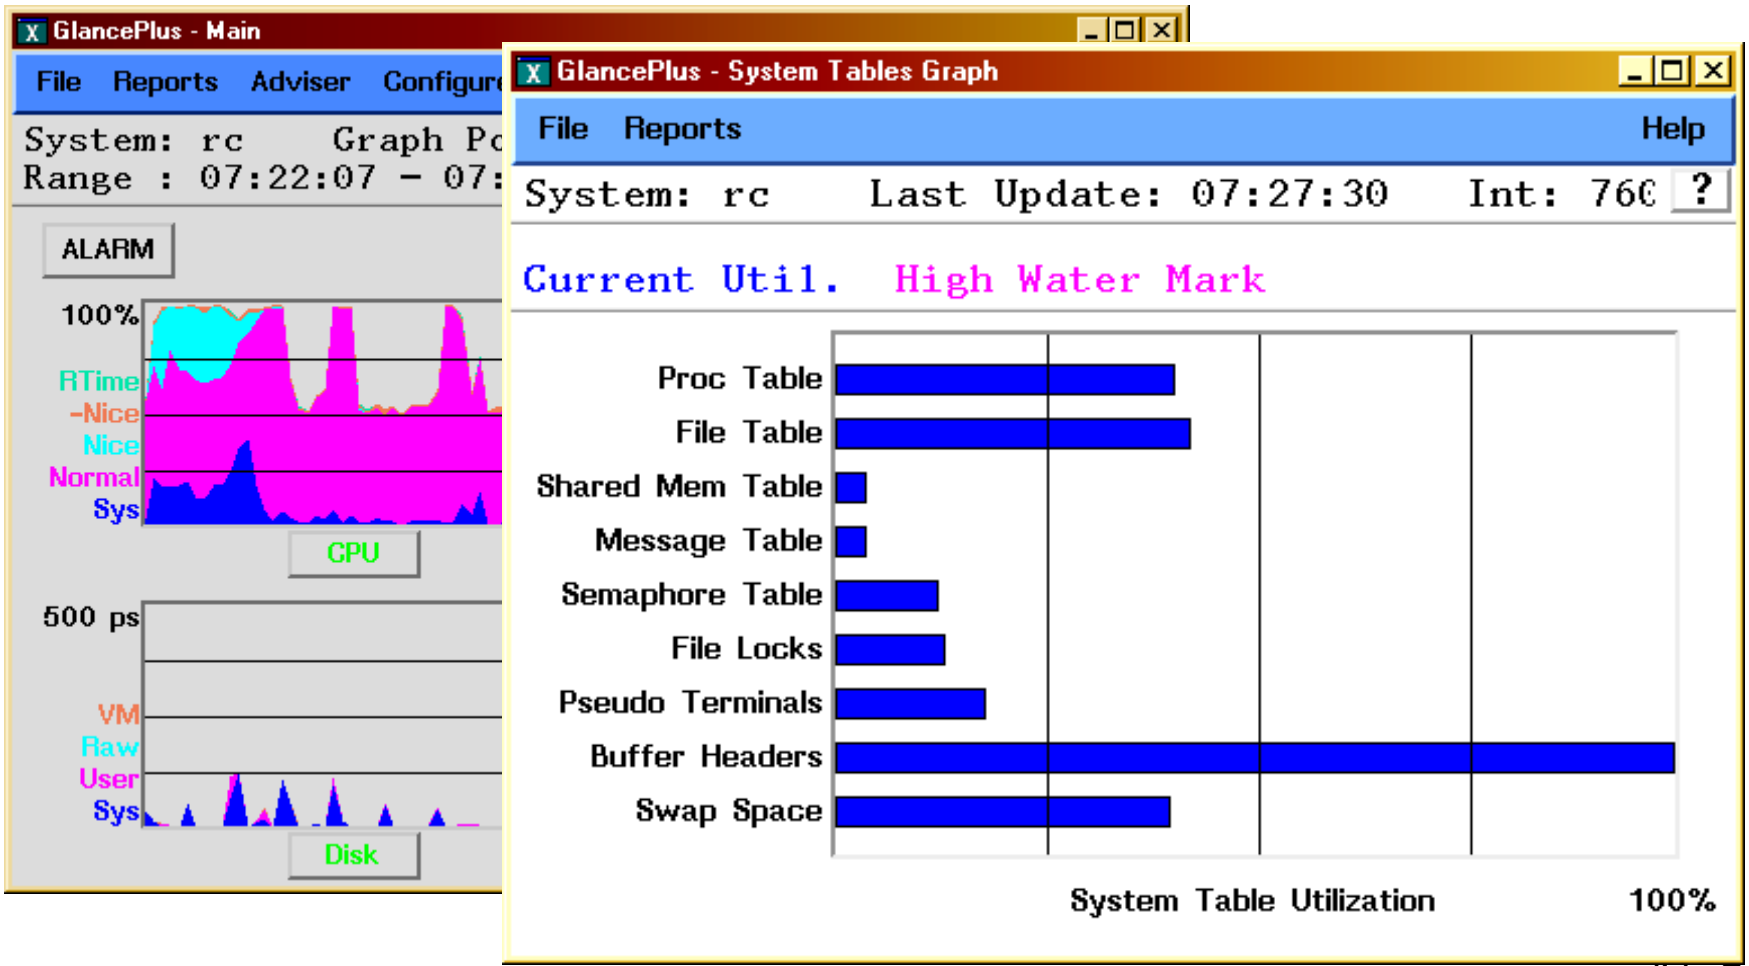

Performance and Kernel Tuning

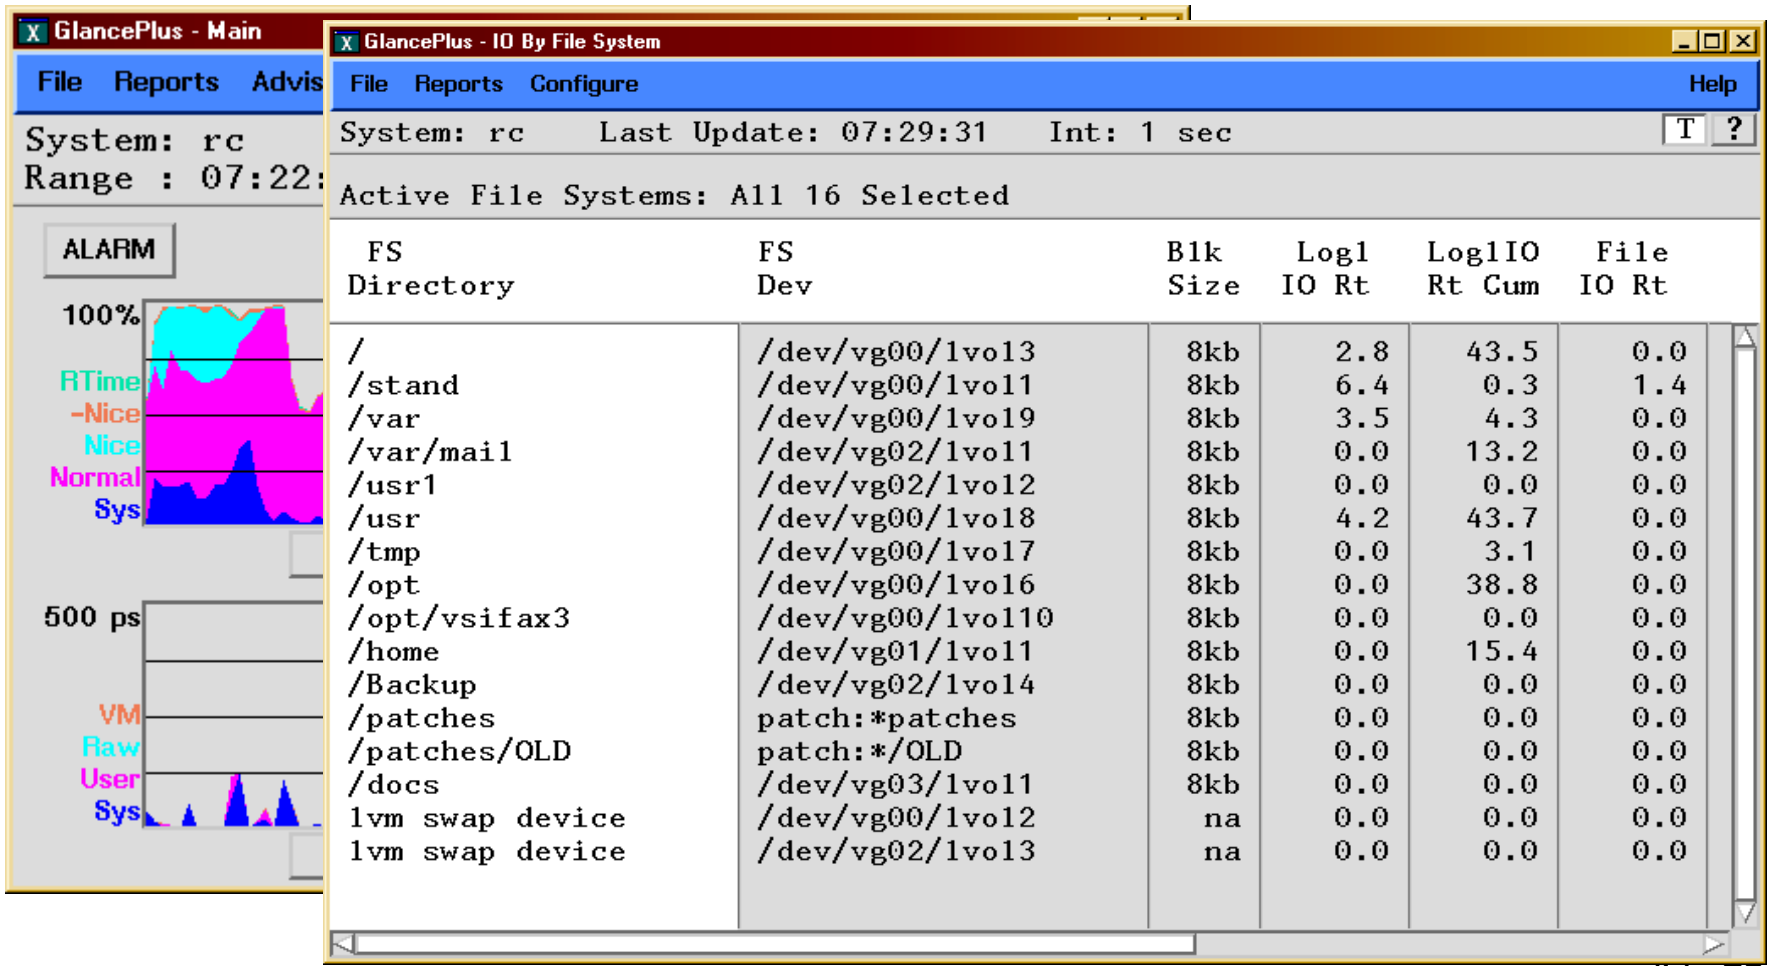

FS Directory	FS Dev	Blk Size	Logl IO Rt	Logl IO Rt Cum	File IO Rt
/	/dev/vg00/lvol3	8kb	2.8	43.5	0.0
/stand	/dev/vg00/lvol1	8kb	6.4	0.3	1.4
/var	/dev/vg00/lvol9	8kb	3.5	4.3	0.0
/var/mail	/dev/vg02/lvol1	8kb	0.0	13.2	0.0
/usr1	/dev/vg02/lvol2	8kb	0.0	0.0	0.0
/usr	/dev/vg00/lvol8	8kb	4.2	43.7	0.0
/tmp	/dev/vg00/lvol7	8kb	0.0	3.1	0.0
/opt	/dev/vg00/lvol6	8kb	0.0	38.8	0.0
/opt/vsifax3	/dev/vg00/lvol10	8kb	0.0	0.0	0.0
/home	/dev/vg01/lvol1	8kb	0.0	15.4	0.0
/Backup	/dev/vg02/lvol4	8kb	0.0	0.0	0.0
/patches	patch:*patches	8kb	0.0	0.0	0.0
/patches/OLD	patch:*/OLD	8kb	0.0	0.0	0.0
/docs	/dev/vg03/lvol1	8kb	0.0	0.0	0.0
lvm swap device	/dev/vg00/lvol2	na	0.0	0.0	0.0
lvm swap device	/dev/vg02/lvol3	na	0.0	0.0	0.0

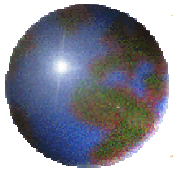

Performance and Kernel Tuning

Performance tuning

- What can you change:
  - Move busy filesystems to other disks
  - Move disks to additional channels
  - Add more processors (and/or change processor speed)
  - Add more RAM (application dependent)
  - Change large data area programs to use large memory pages (chatr) to reduce TLB misses
  - Use AutoPort Aggregation (APA) for parallel data transfers (and fallback reliability)

Memory Management

- 32/64 bit
- 32 bit Data area
- 32 bit Shared memory
- Memory windows

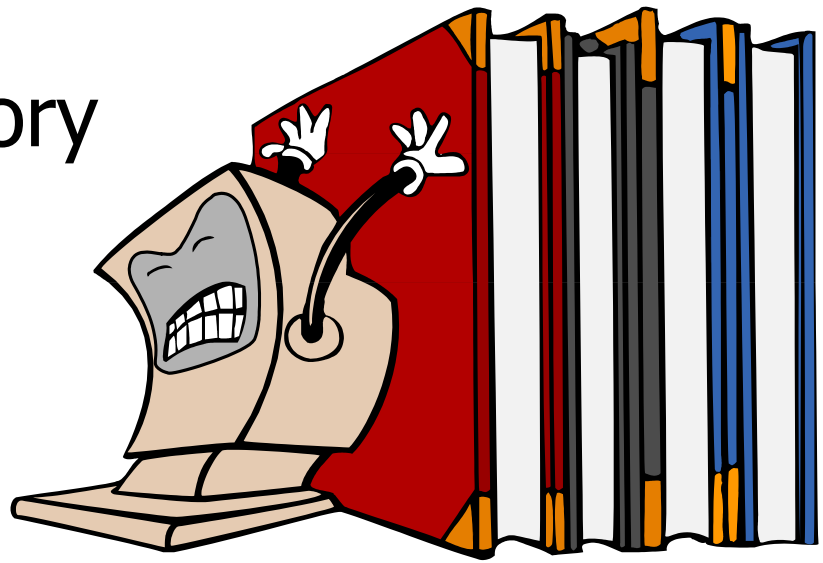

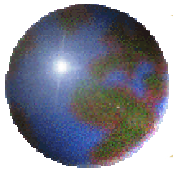

Memory Management

- 32/64 bit
  - HP-UX operating system
  - Hardware dependent
    - 32bit only
    - 32/64bit (either)
    - 64bit only (new machines)
  - 32bit programs run in either
  - 64bit programs only for 64bit HP-UX
  - 64bit programs remove memory map limitations, not a performance feature

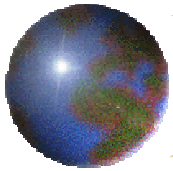

Memory Management

- 32 bit Data area
  - 32bit programs have four 1000meg quadrants
  - Default data area = 960megs apx.
  - Compiler or chatr options for EXEC_MAGIC allow quadrant 1 and 2 to be combined for about 1750 megs total data area
  - Adjust maxdsiz if necessary
  - Documentation in /usr/share/doc for mem_mgt and proc_mgt (missing in 11i)

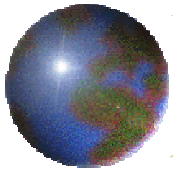

Memory Management

- 32 bit Shared memory
  - Most common limitation for 32bit database programs
    - Oracle: SGA is shared memory
  - Similar limits (960megs and 1750 megs)
  - Use SHMEM_MAGIC option for 1750 meg access
    - (all sharing processes must match)
  - **ALL** 32bit shared memory processes have one map (fragmentation, memory mapped files, shared libraries)
  - Use **ipcs –bmop** for a snapshot (does not show fragmentation)
  - Get **shminfo** from the hprc.external.hp.com ftp site: ftp://contrib:9unsupp8@hprc.external.hp.com/sysadmin

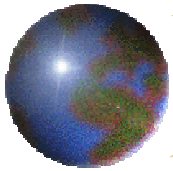

Memory Management

- Memory windows
  - Reserves a separate memory map to remove fragmentation issues
  - Memory-mapped files, shared libraries are in memory window 0 (plus ‘normal’ shared memory processes)
  - Need patches for 11.0 to enable
  - Set kernel param for quantity of windows
  - Start ALL sharing programs with memory window startup command.

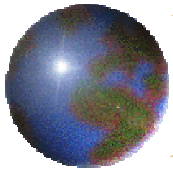

Disk space management

- Root filesystem
  - VG00 size and management
    - HP-UX only, possibly user $HOME
    - 3-6Gb typical
    - New systems have 36/72Gb internal (too big)
    - Or use external disks
      - SCSI vs. fibre channel
      - Boot issues for external arrays
      - NAS versus SAN
    - Mirroring
    - Ignite/UX and VG00
    - Striping VG00

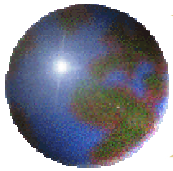

Disk space management

- Small versus large disks
  - JBODs
  - Arrays
- Large lvols
  - Large files versus thousands of small files
  - Flat versus hierarchical directory structure
    - Performance
    - System impact (from users)
    - Managing files and backups

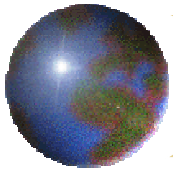

Disk space management

- Filesystem types:
  - **HFS** (Fast FileSystem - BSD, circa 1984)
  - **VxFS** (Veritas FileSystem, aka Journaled FileSystem) or **JFS**
  - **CDFS** (CDROM FileSystem - ISO 9660 only)
  - **NFS** (Network FileSystem)

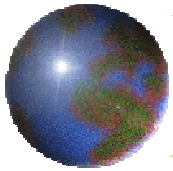

Disk space management

- **HFS** (Fast FileSystem - BSD, circa 1984)
  - McKusick HPF and fragmentation
  - fsck proportional to directory size
  - Inode count fixed at creation
  - resizing
    - up = umount, lvextend, extendfs
    - down = umount, backup, lvreduce, newfs, reload from backup, mount

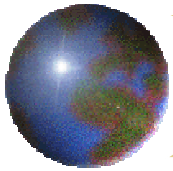

Disk space management

- **JFS** (Veritas filesystem (not Volume Manager)
  - Aka: Journaled FileSystem or JFS
  - Inodes created as needed
  - Very busy filesystem may need defrag
  - Online (Advanced) JFS option needed for online resizing and advanced options
  - Better overall performance than HFS

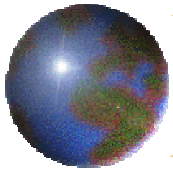

Disk space management

- **CDFS** (CDROM filesystem)
  - ISO 9660 only (see below)
  - 8.3 UPPERCASE filenames plus ;1 version
  - 11.0 patch for –ocdcase to lower UPPERCASE and remove ;1 version (but no long filenanes
  - HP makes special (non-standard) CD's for Core, Application, SupportPlus, etc that have long filenames
  - PFS (Portable FileSystem) to perform RockRidge or long filename translation
  - Recent patches add –orr to natively handle long filenames
  - No direct support for Joliet, audio, multimedia, DVD
  - No direct support for CD writers

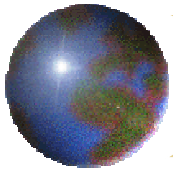

Disk space management

- **NFS** (Network FileSystem)
  - Network-based filesystem (server and clients)
  - Stateless so performance issues exist
  - Does require tuning for heavy usage
  - Version 2 = no largefiles
  - Version 3 = largefile-capable
  - nfsstat very useful
    - **biod** and **nfsd** quantity
    - Network quality
  - NFS Performance book by Dave Olker

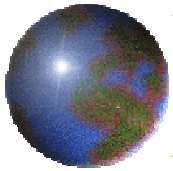

Disk space management

- VG00 filesystem design
  - / (root) lvol
    - HP-UX only, rest = mountpoints
    - Static: size to about 256 megs
    - **/etc /dev /sbin** (that's all!)
  - **/stand** lvol
    - Only HFS filesystem
    - Holds bootable kernels (current and previous)
    - 60-200 megs

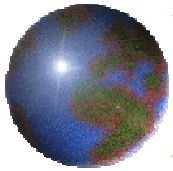

Disk space management

- VG00 filesystem design - other HP-UX mountpoints
  - **/usr** = HP-UX and local executables, libraries
    - **/bin and /lib**
    - **/usr/local** versus **/usr/contrib**
    - Perhaps 400-1000 megs
  - **/opt** = application installation directories
    - /opt/app-name
      - bin doc lib lbin etc man…
    - Perhaps 500-1500 megs
  - **/tmp** = system temporary directory (note: system)
    - Not for users, scripts, etc
    - Perhaps 200-400 megs

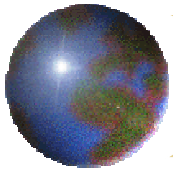

Disk space management

- VG00 filesystem design - other HP-UX mountpoints
  - **/var** = variable directory
    - Most critical directory in HP-UX
    - When full, daemons and processes abort
    - Used by many unrelated subsystems
    - Ideally, separate mountpoint lvols will be created for the most intrusive directories:
      - **/var/mail**
      - **/var/spool**
      - **/var/tmp**
      - **/var/adm**
      - **/var/adm/crash**
      - **/var/adm/sw**

Disk space management

- VG00 filesystem design - other HP-UX mountpoints
  - **/home** = user home directories
    - Highly variable
    - When full, cause little impact to production processes
    - May require active space management or quotas

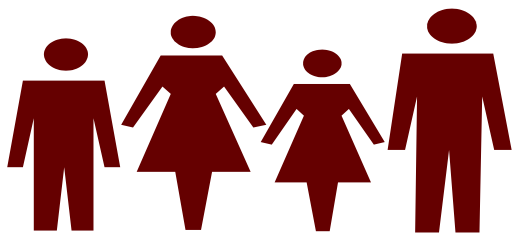

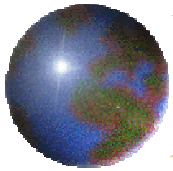

Disk space management

- XVG
  - Contributed Xwindow program
  - Visualize disk layout
  - Highlight swap, free space, filesystem types
  - Highlight non-contigous areas
  - Download from:

    ftp://contrib:9unsupp8@hprc.external.hp.com/sysadmin

Xvg Display Program

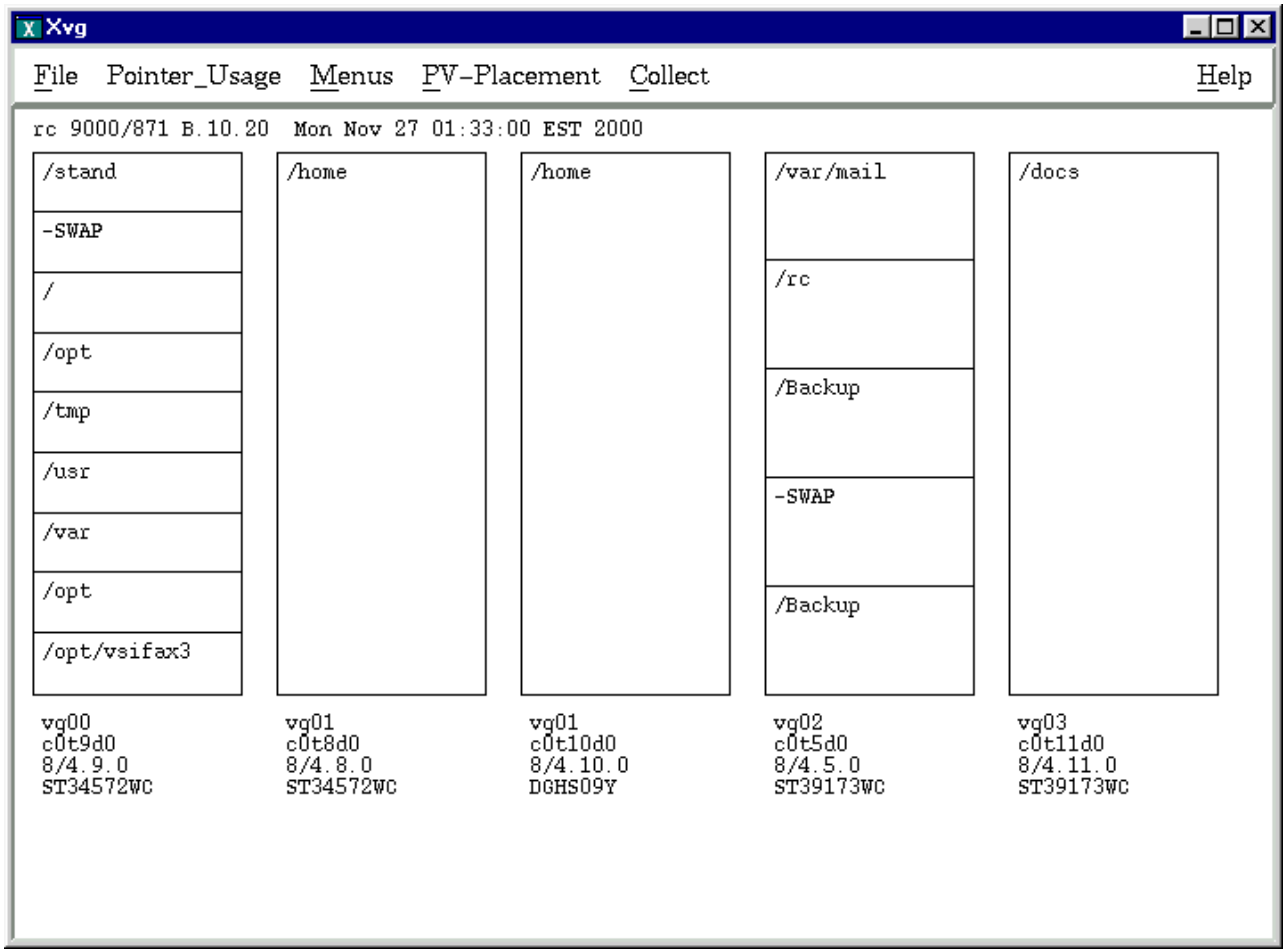

Xvg Display Program

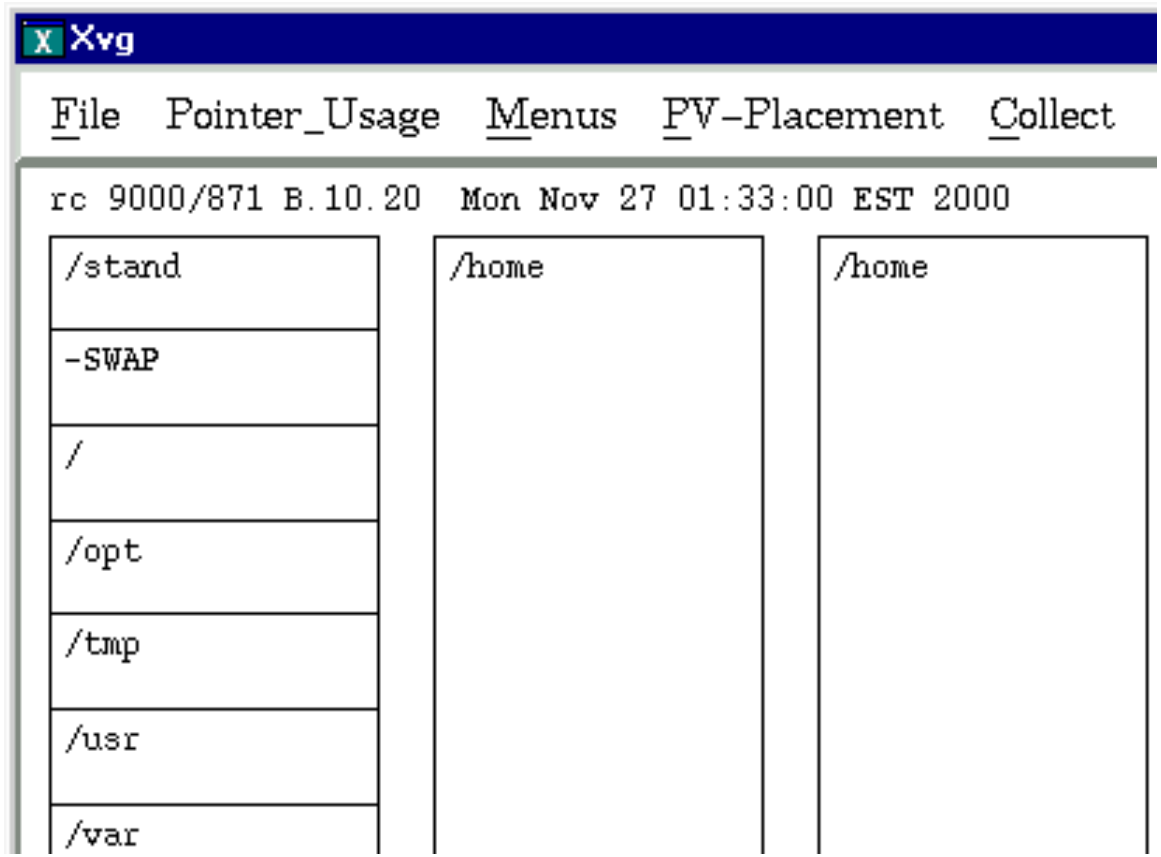

Xvg Display Program

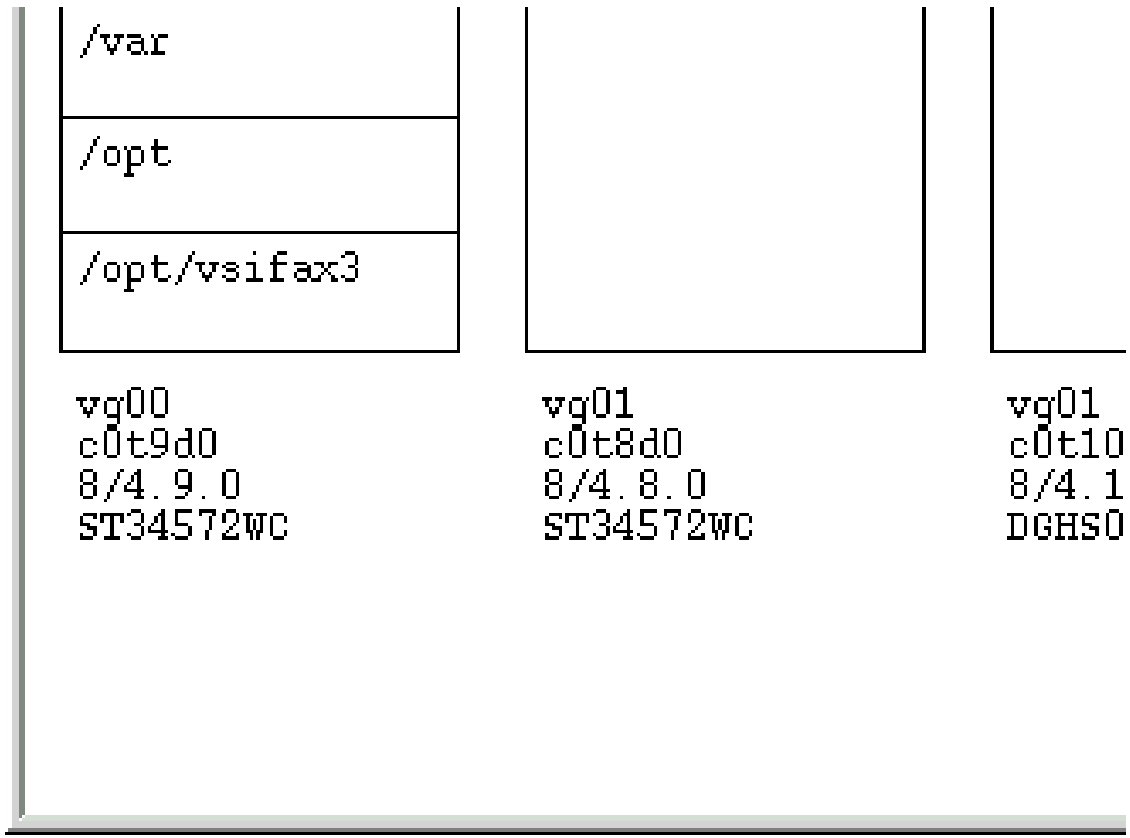

Xvg Display Program

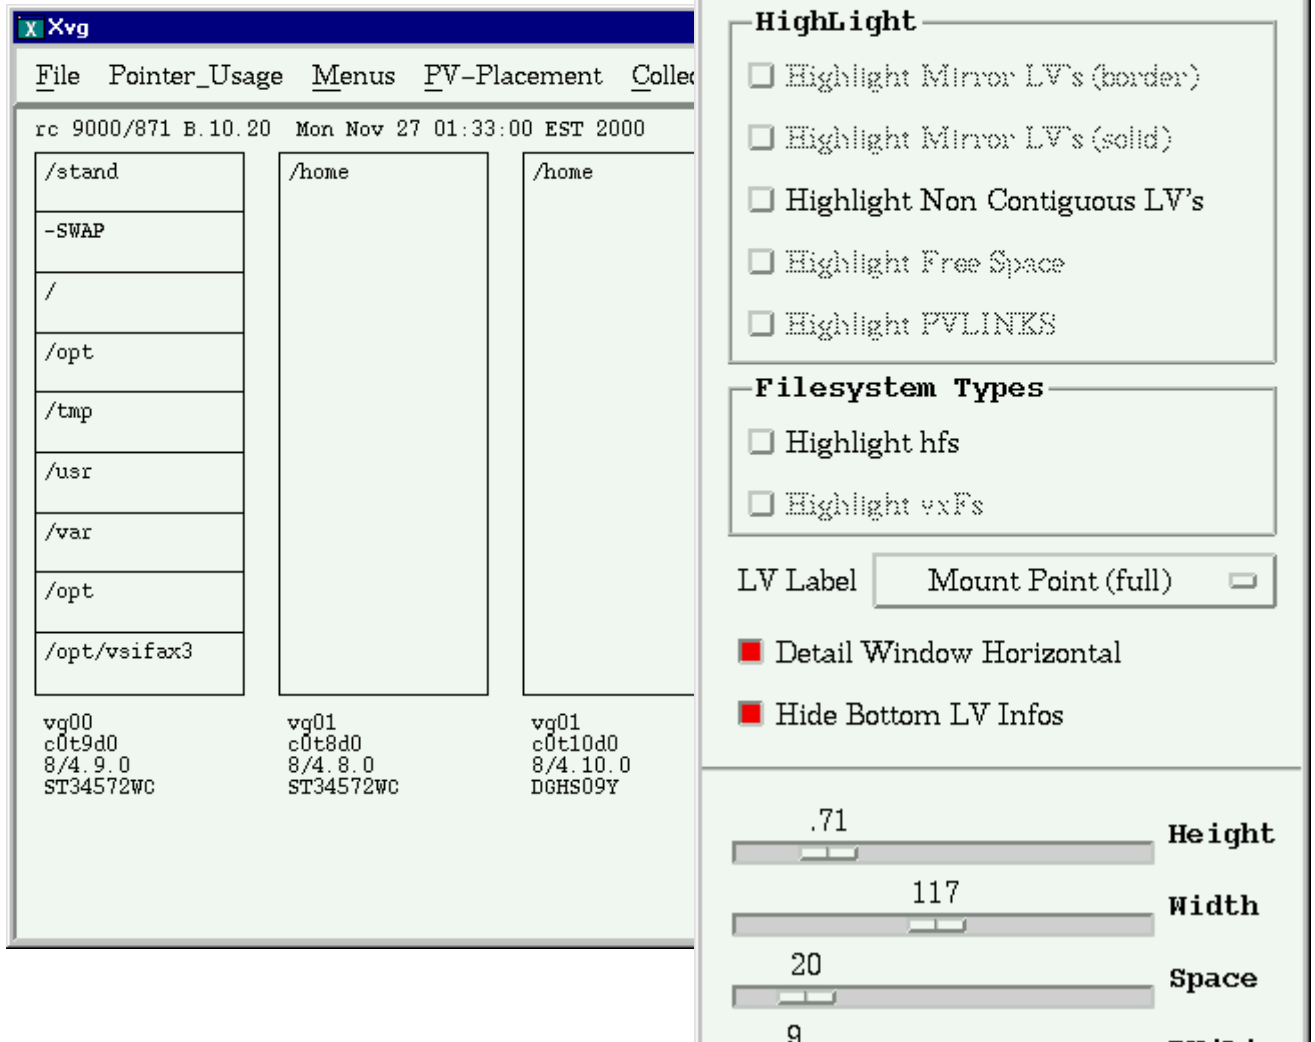

Xvg Display Program

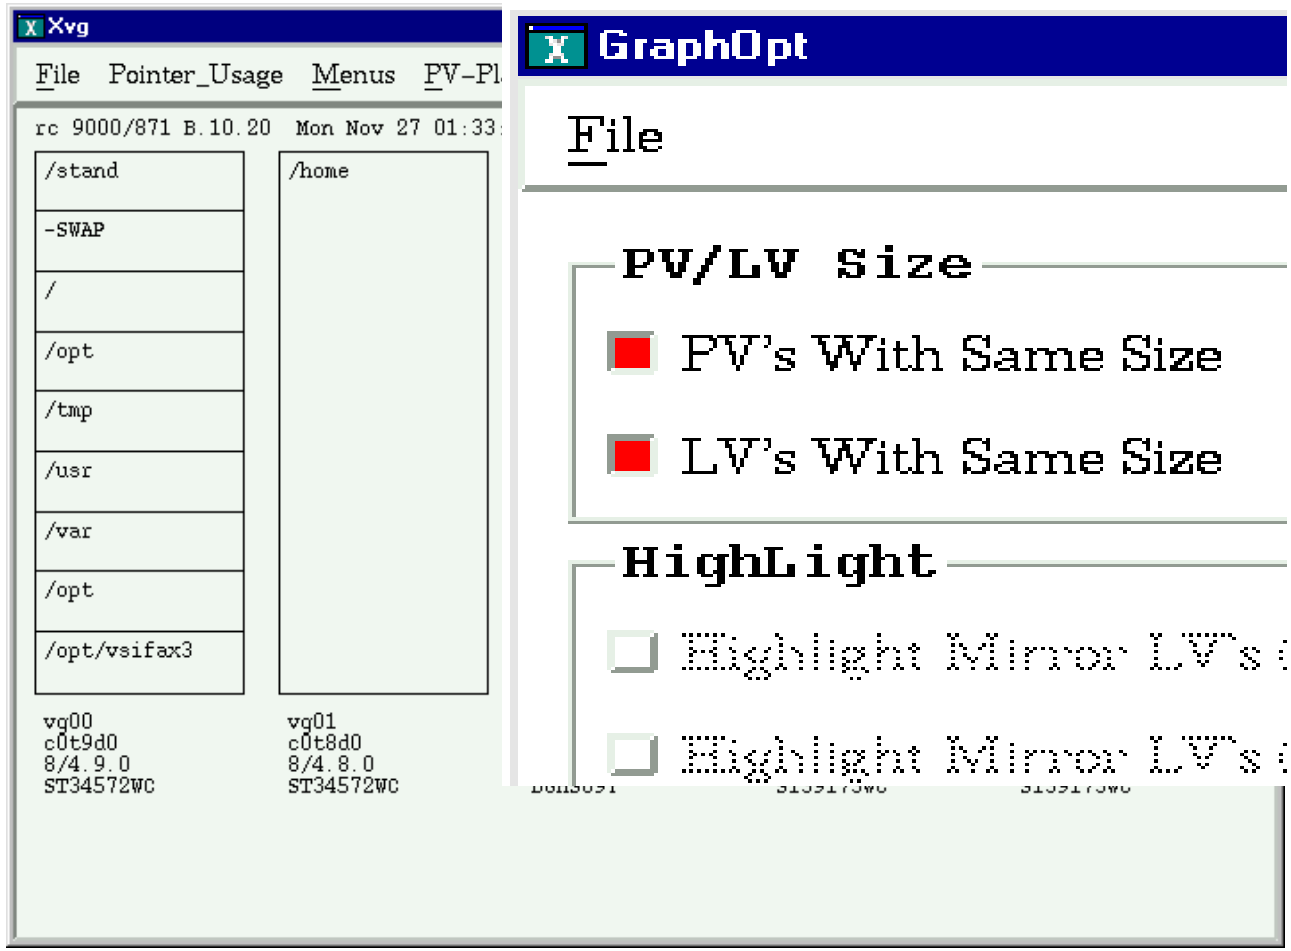

Xvg Display Program

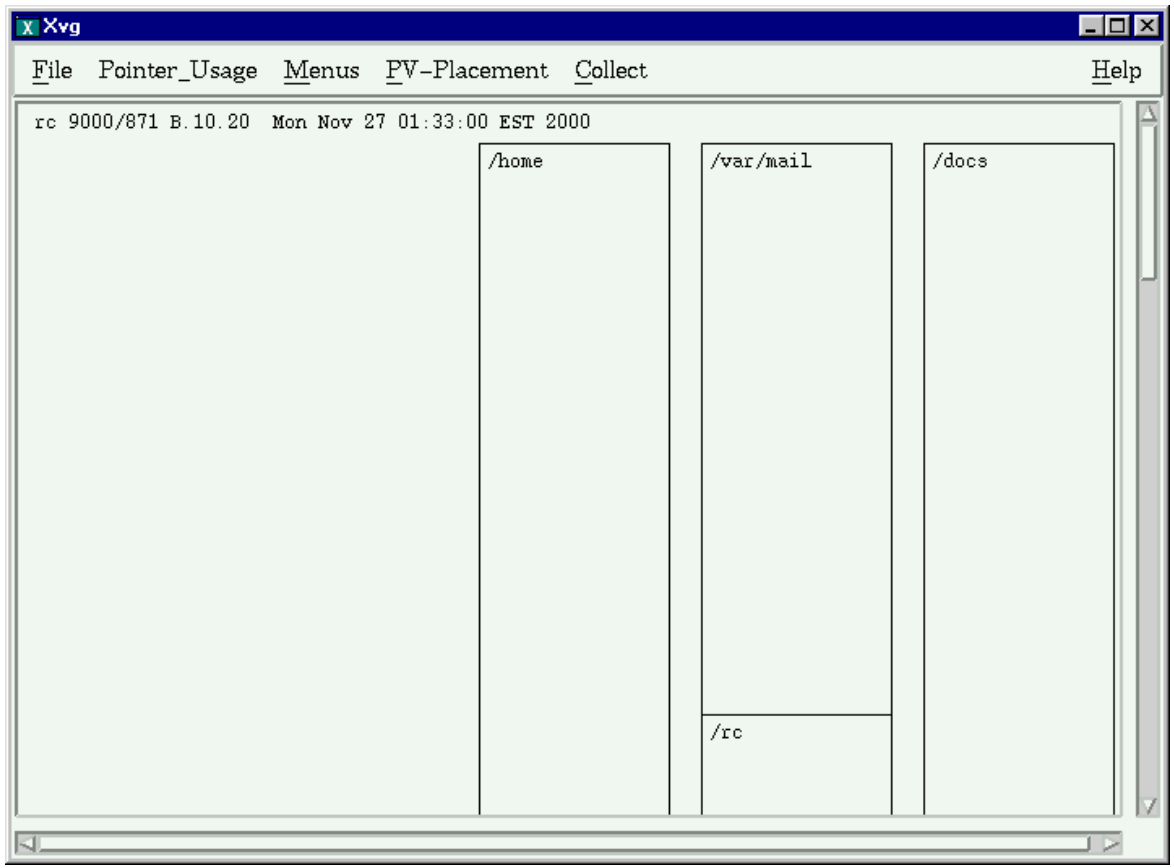

Xvg Display Program

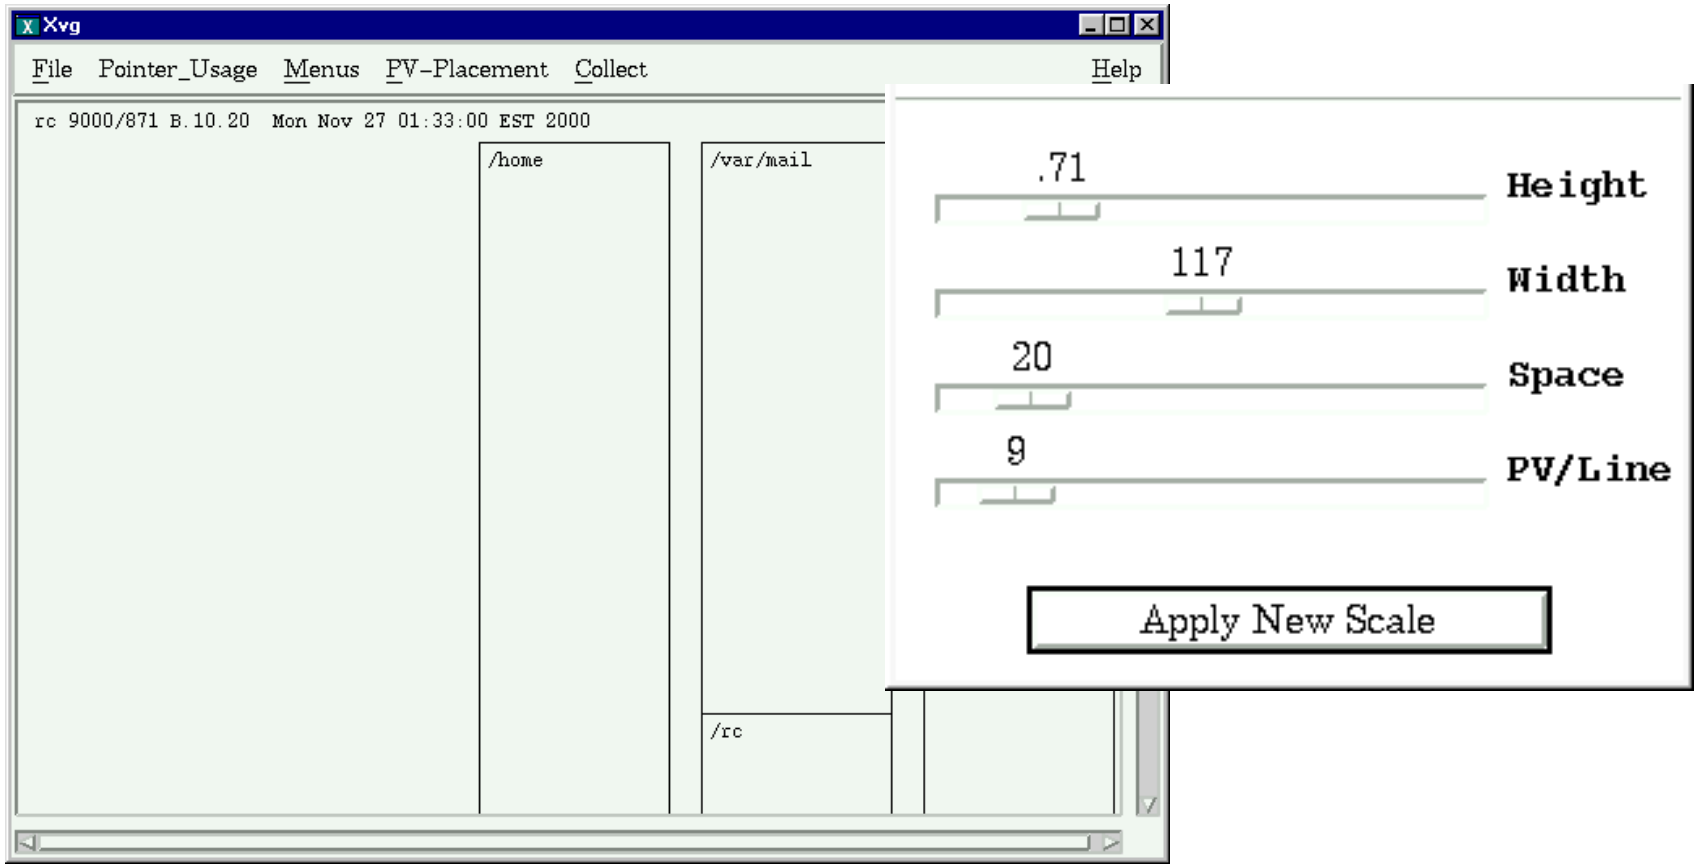

Xvg Display Program

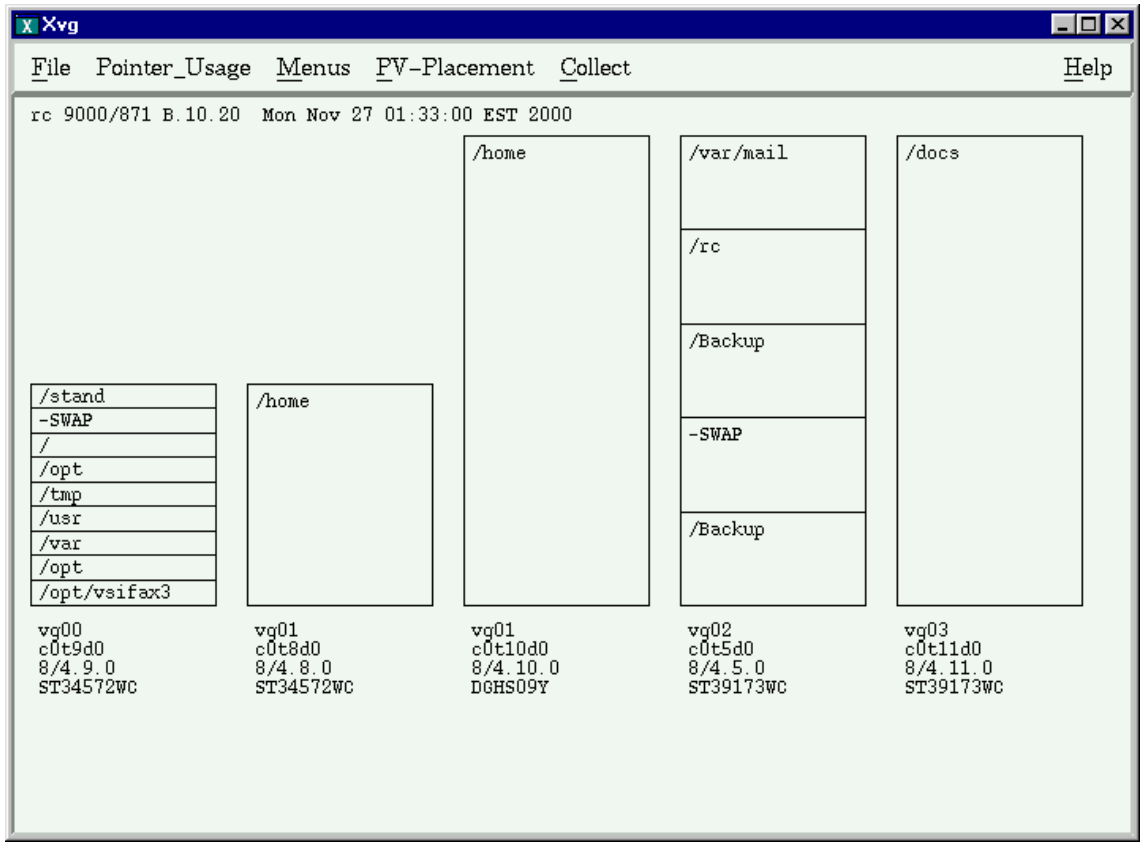

Xvg Display Program

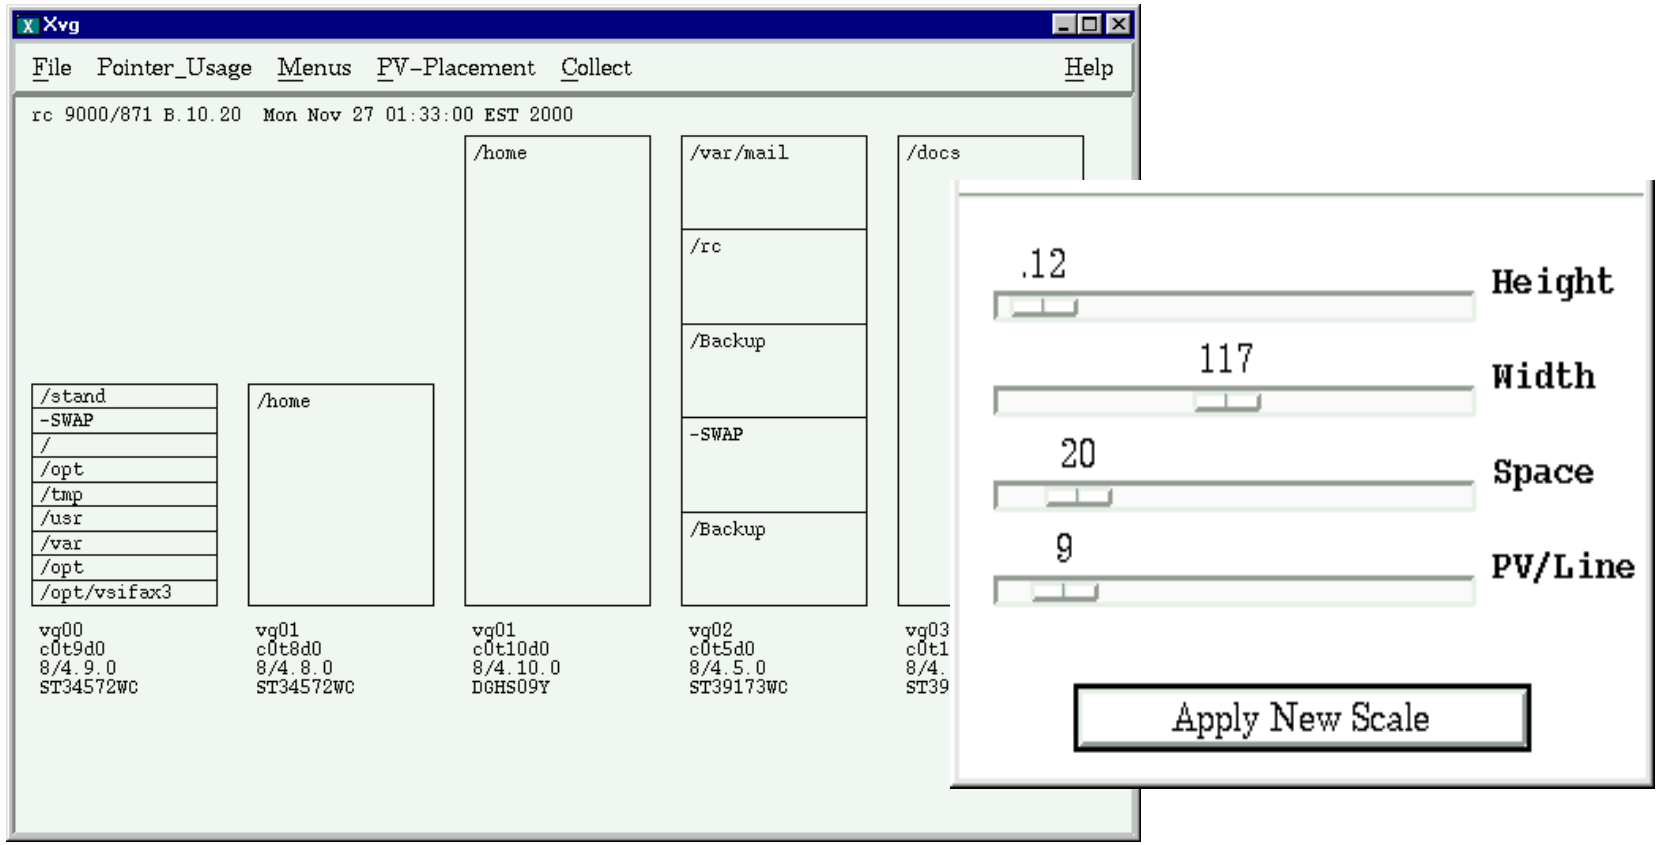

Xvg Display Program

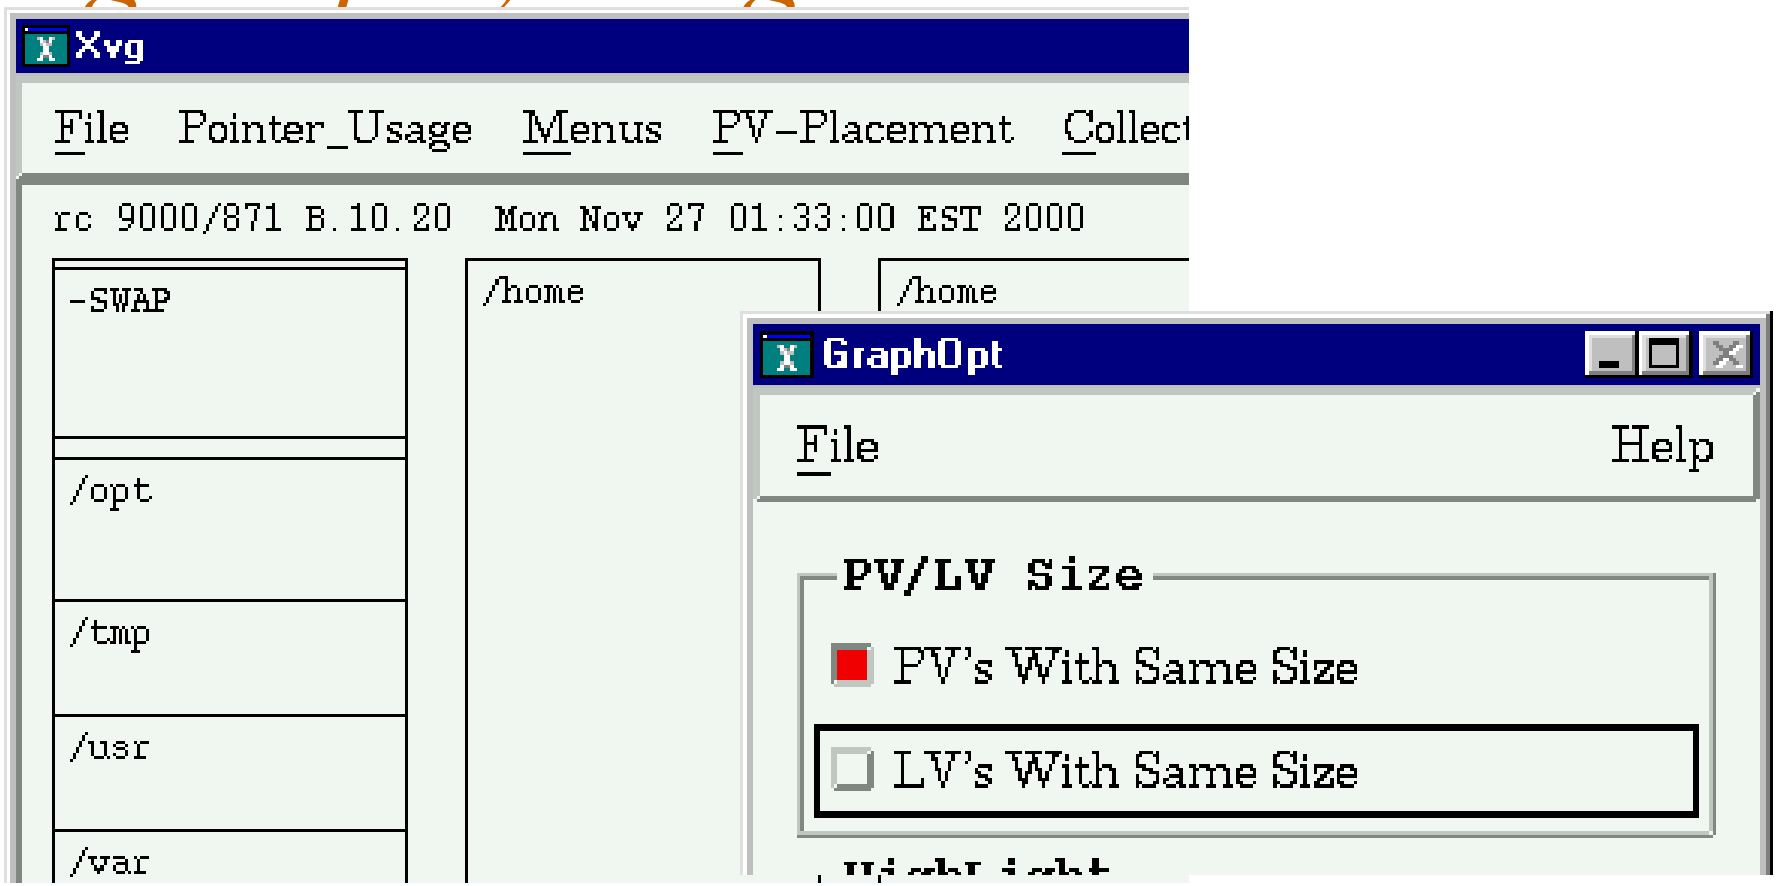

Xvg Display Program

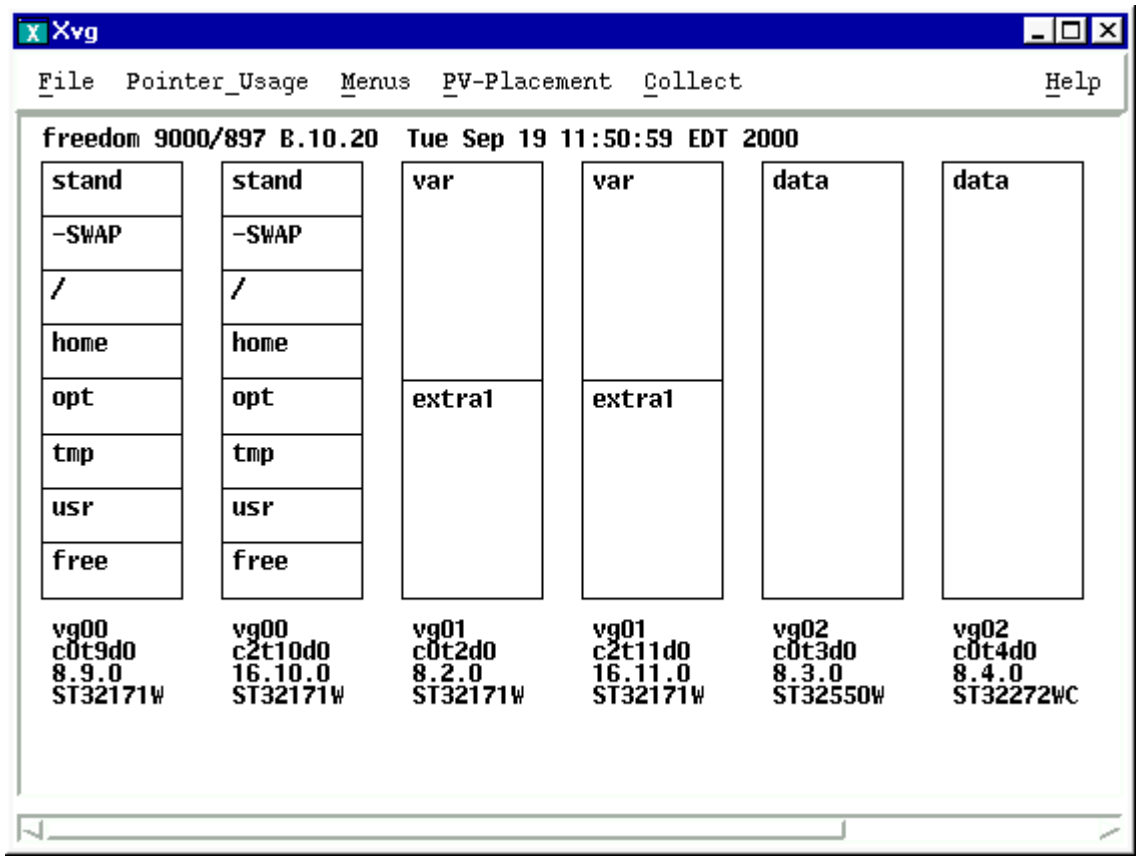

Xvg - continued

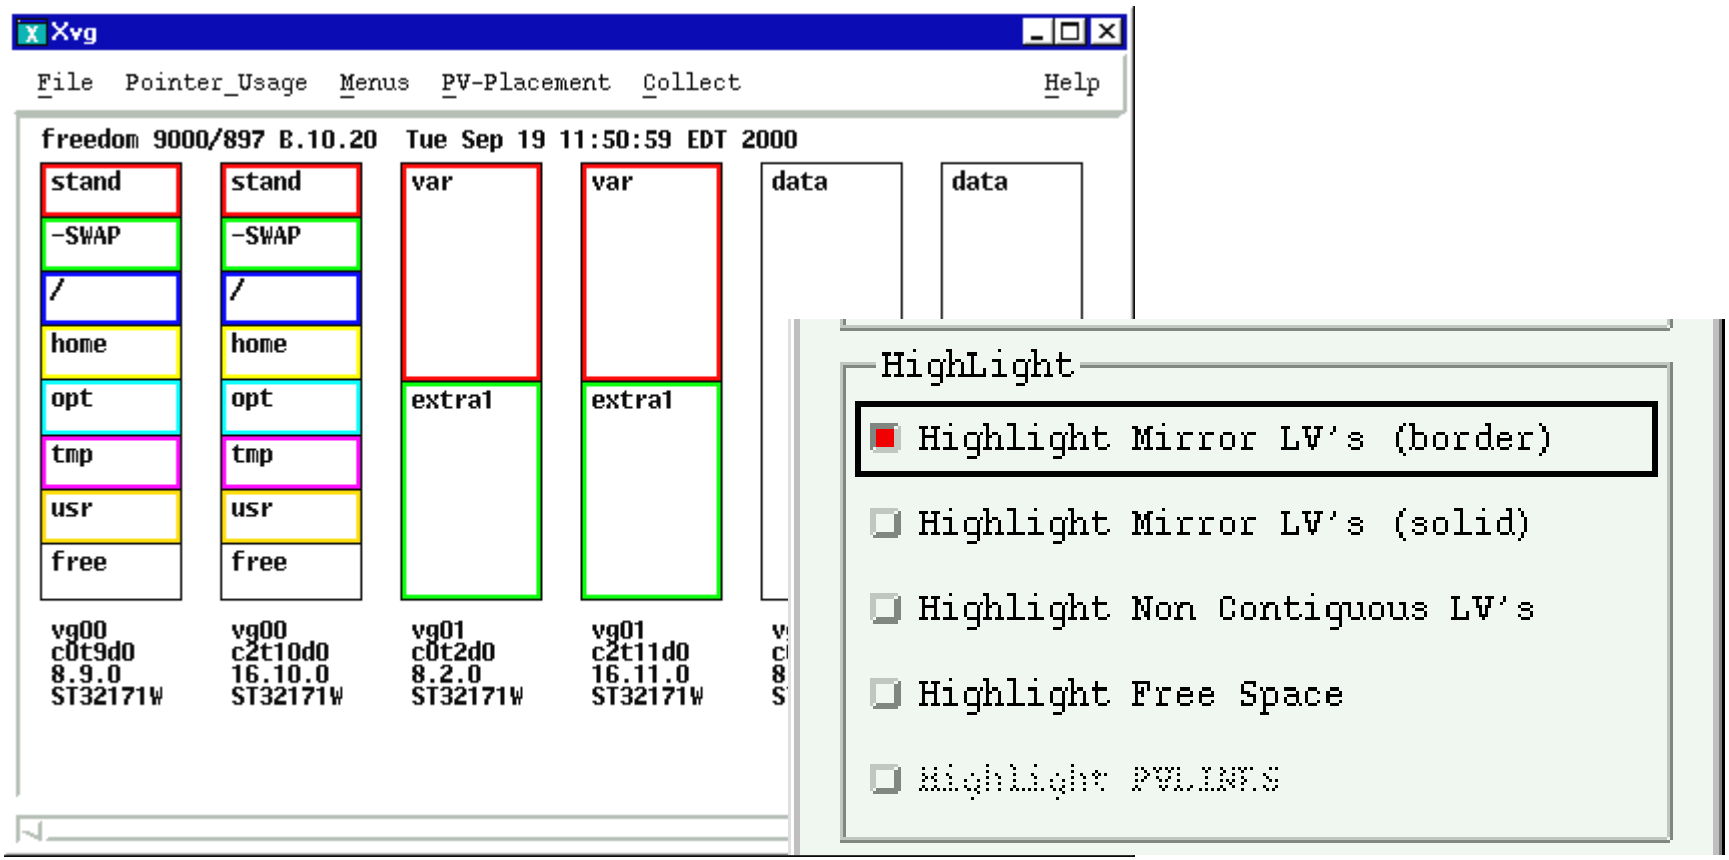

Xvg - continued

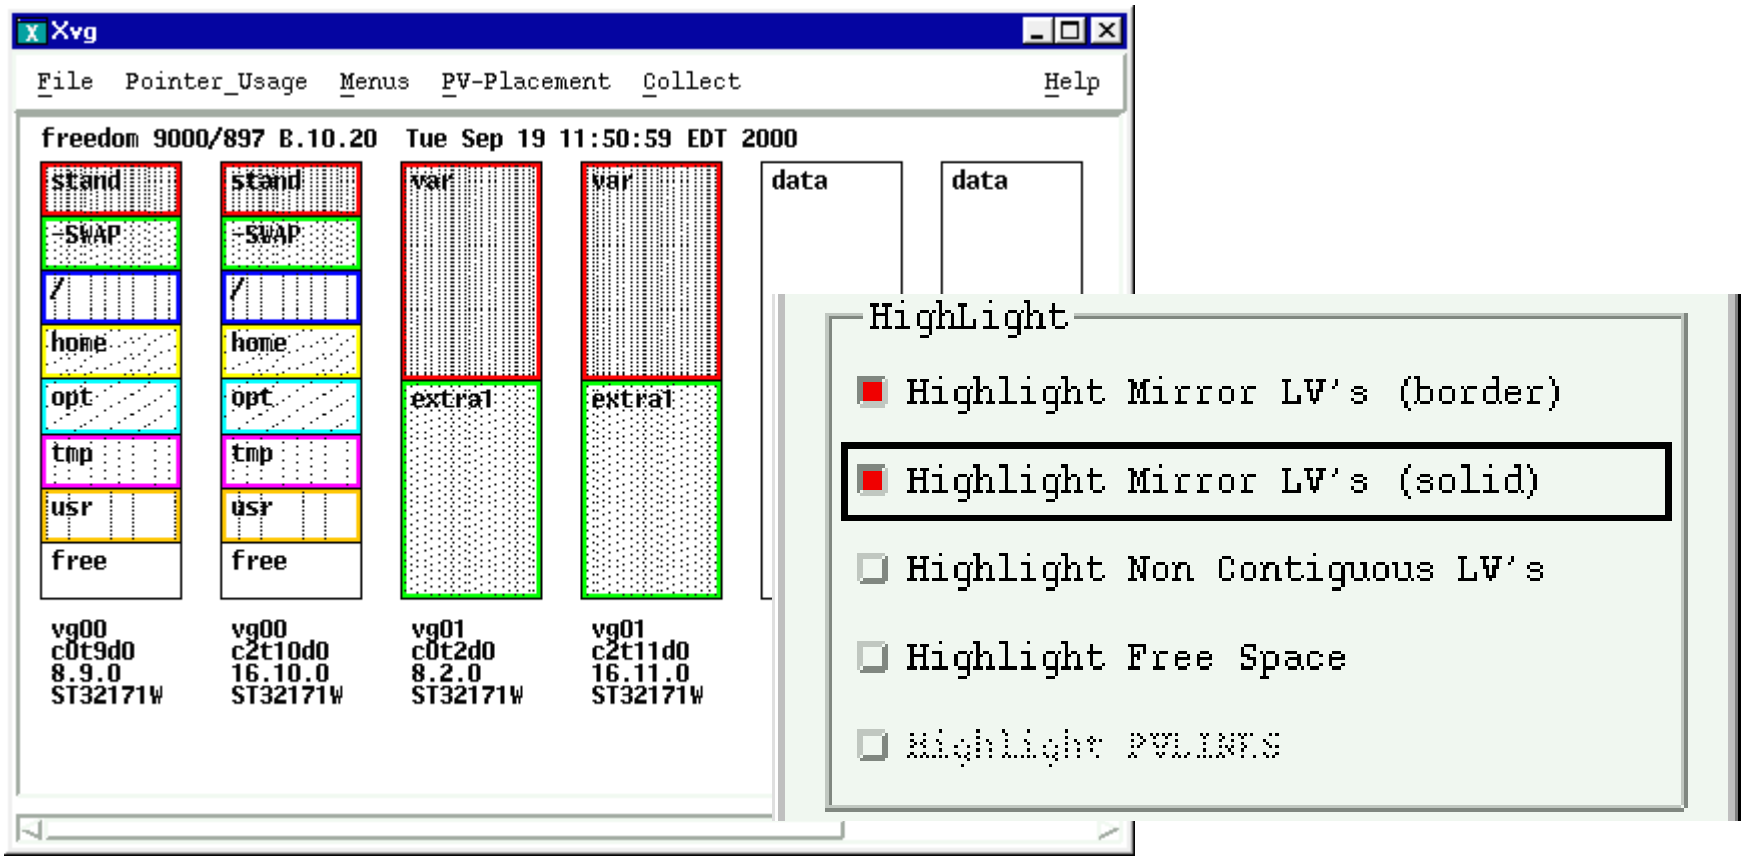

Xvg - continued

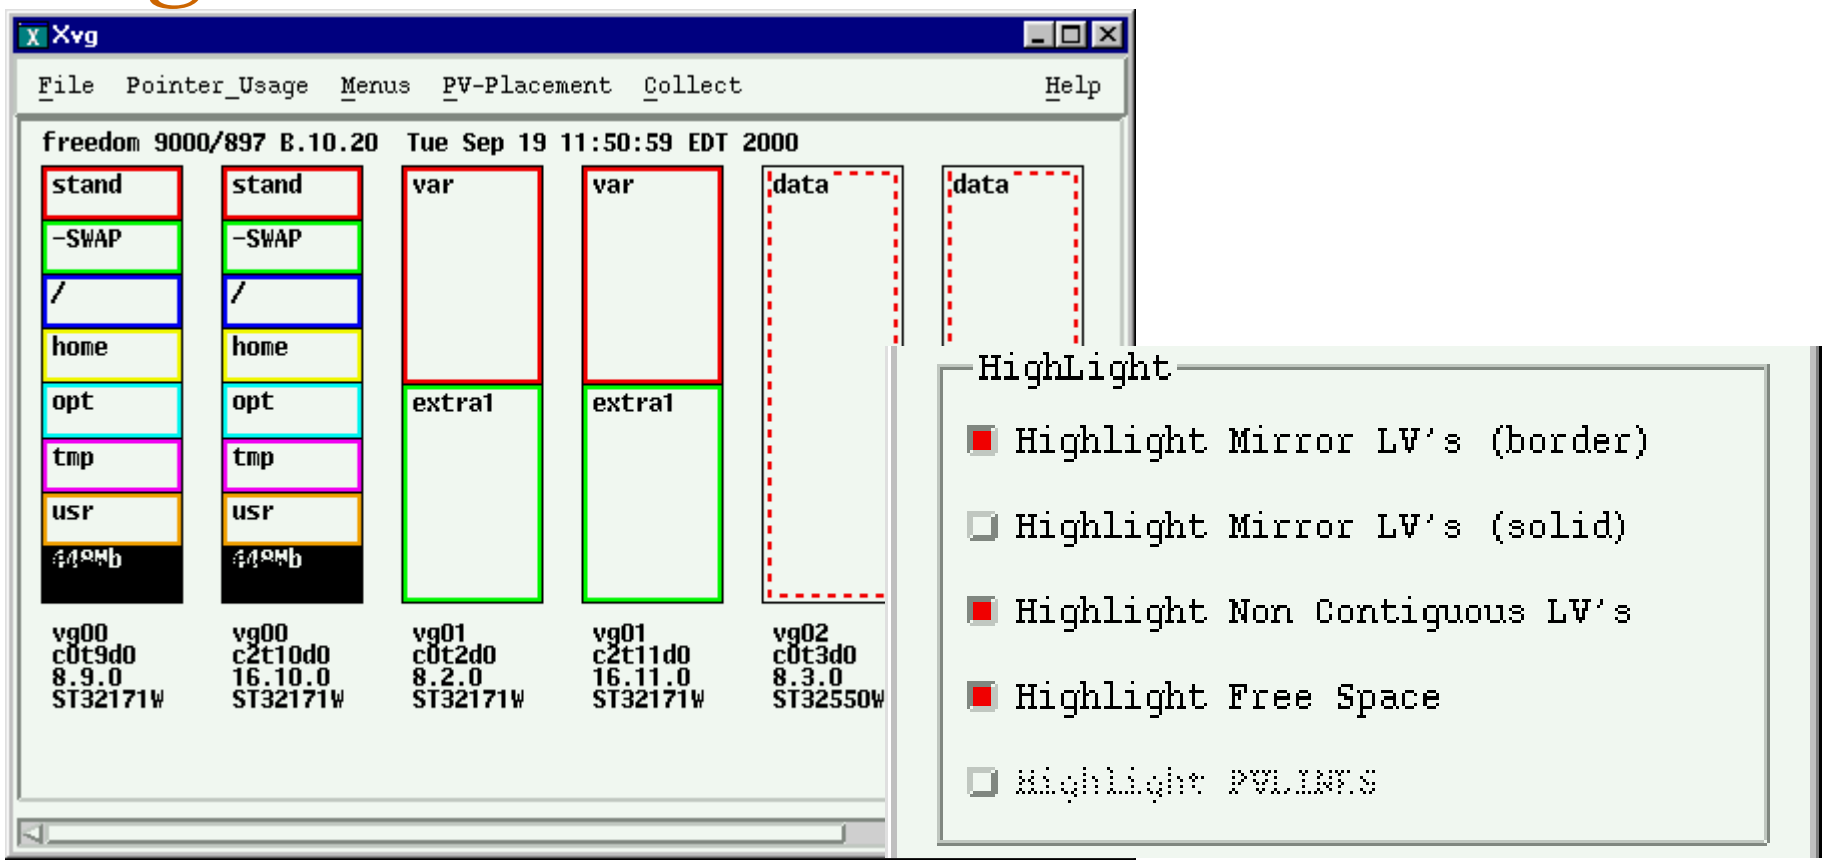

Xvg - continued

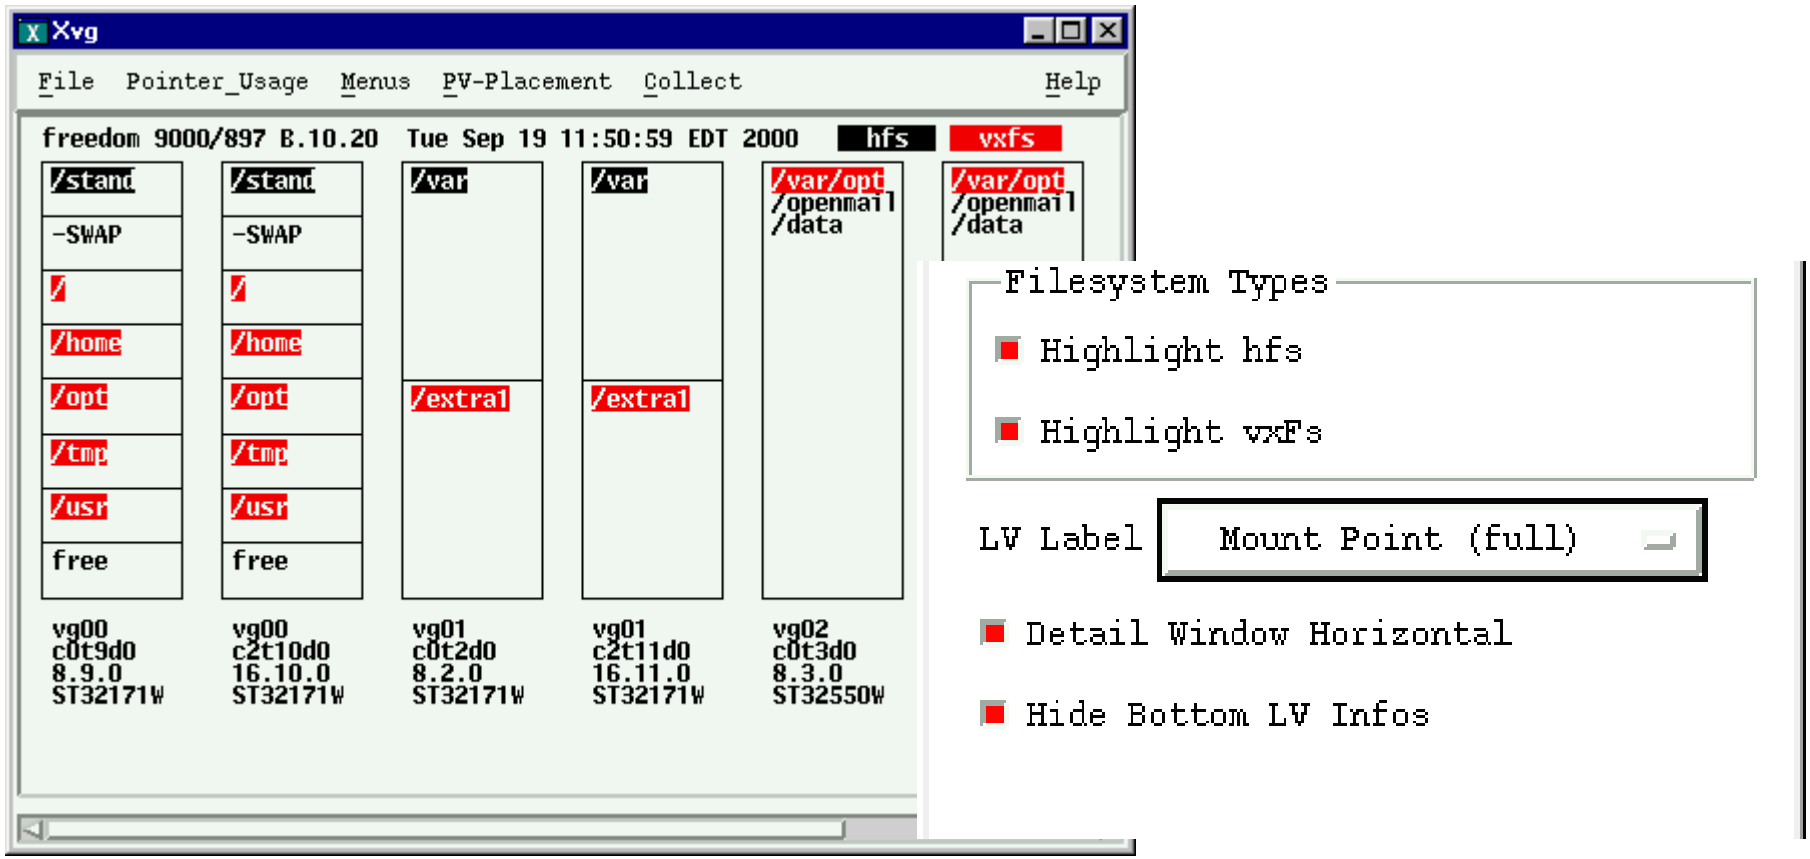

Xvg - continued

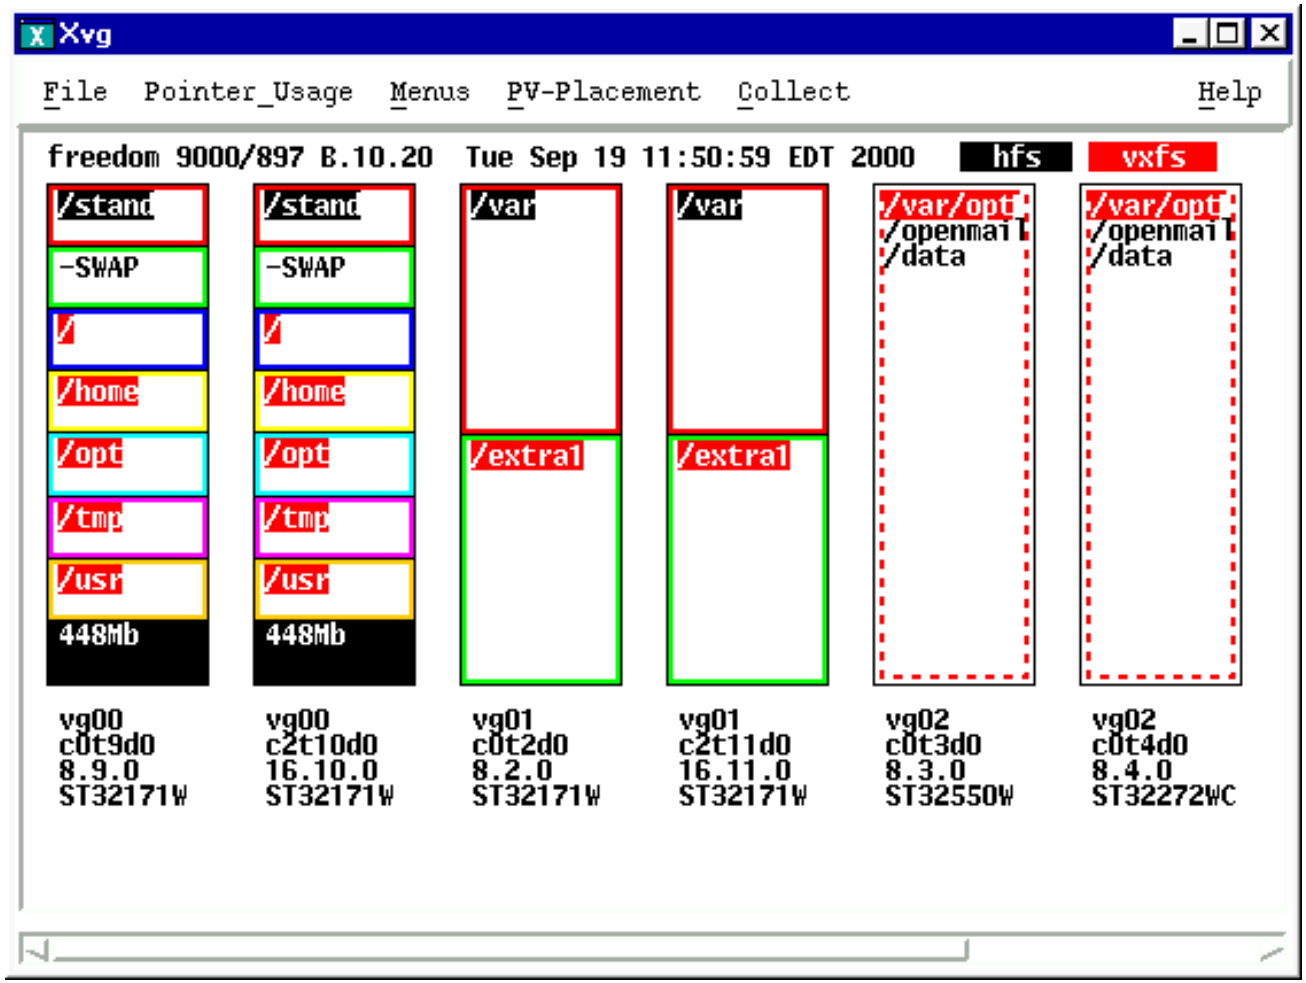

Xvg - continued

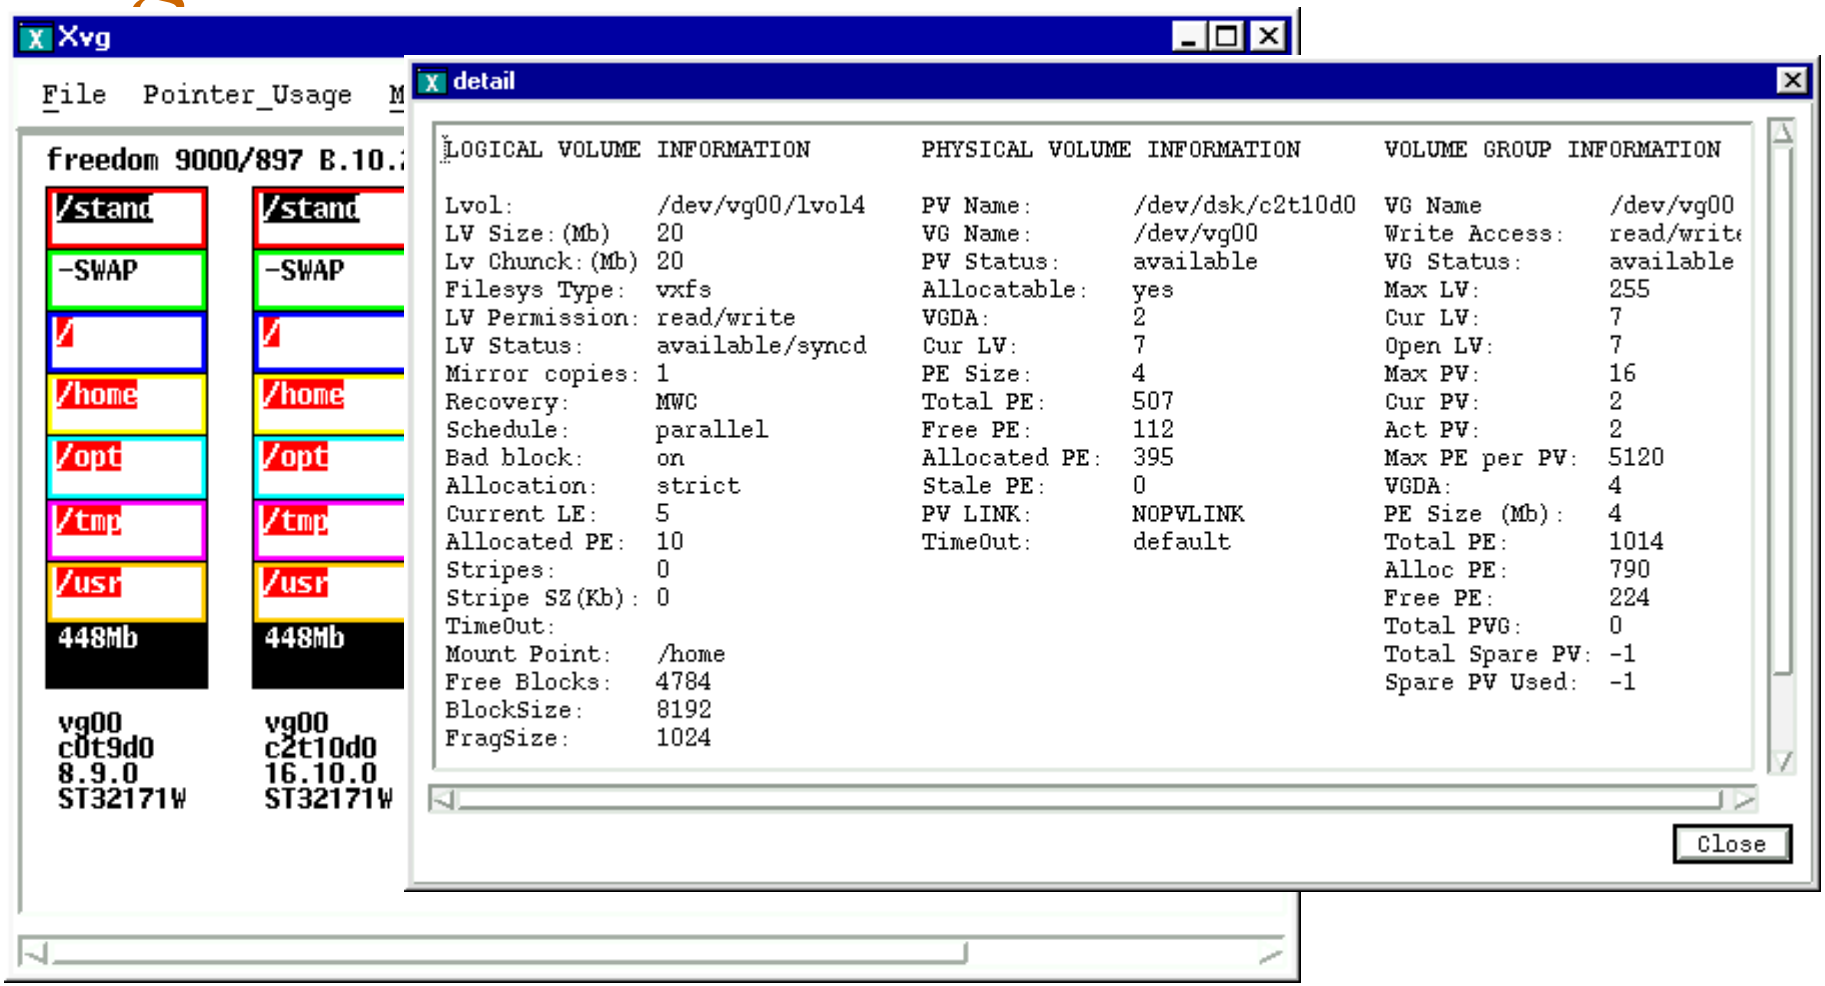

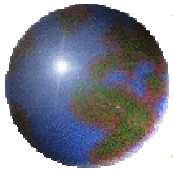

Logfiles: Where did the space go?

- Don't look for big files!
- Look for big directories (10,000 small files)
- Use du as in:

```
du -kx /
9         /lost+found
1682      /etc/lvmconf
464       /etc/hpC2400
2         /etc/switch
8         /etc/eisa
18        /etc/skel/.elm
1         /etc/skel/Mail
```

(oops…not sorted)

Where did the space go?

- Use du as in (sorted):

```
du -kx / | sort -rn | more
36048   /
15215   /sbin
14196   /etc
8032    /etc/lp
7927    /etc/lp/interface
6864    /etc/lp/interface/model.orig
6509    /root
3868    /sbin/fs
```

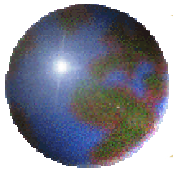

Database directories

- Databases
  - files vs raw access
  - raw: document! (`/etc/fstab`)

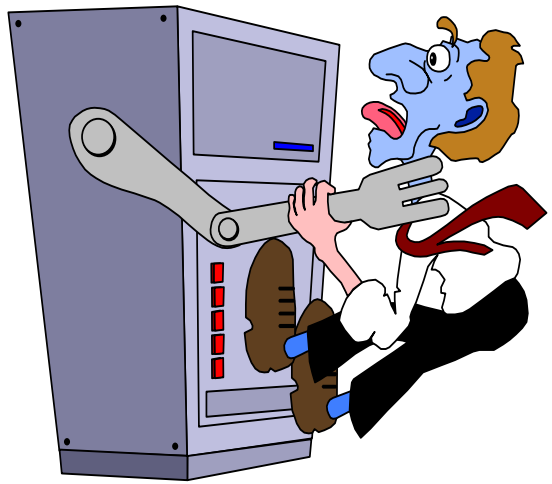

Database directories

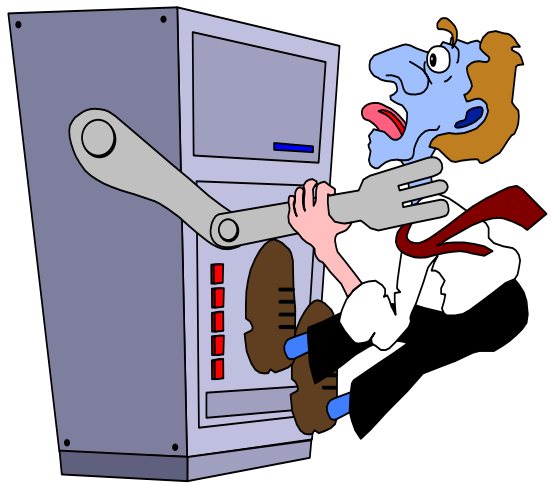

- Databases
  - files vs raw access
  - raw: document! (`/etc/fstab`)
  - fixed file count and size
    - minfree (10% and 0%)
    - inodes (HFS only)

Database directories

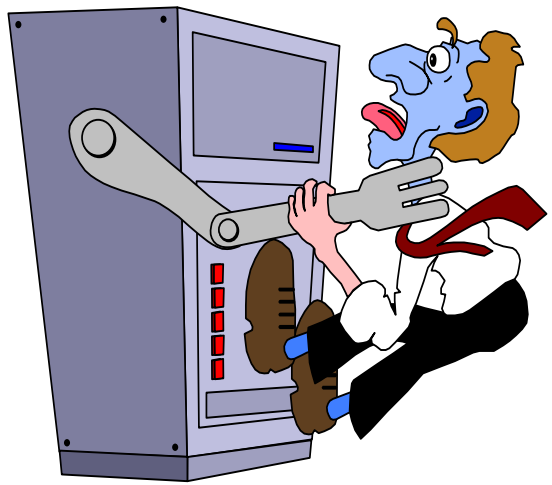

- Databases
  - files vs raw access
  - raw: document! (`/etc/fstab`)
  - fixed file count and size
    - minfree (10% and 0%)
    - inodes (HFS only)
  - Mount options (advanced VxFS)
    - Table space (/u01 /u02 …)
      - **minfree=direct,convosync=direct,nodatainlog**

Database directories

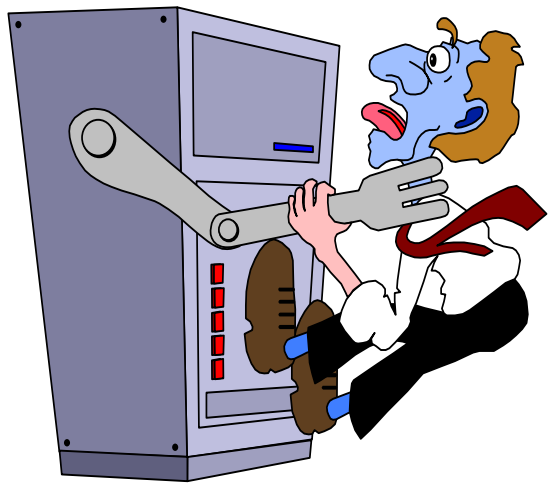

- Databases
  - files vs raw access
  - raw: document! (`/etc/fstab`)
  - fixed file count and size
    - minfree (10% and 0%)
    - inodes (HFS only)
  - Mount options (advanced VxFS)
    - Table space (/u01 /u02 …)
      - **minfree=direct,convosync=direct,nodatainlog**
    - Redo logs, archived redo logs, indexes
    - executables

Growing directories

- New application installs
  - Ask vendor for inventory and size
  - Require all apps to install in /opt
  - Or create a symlink (ln -s)
- logfiles
  - Always grow
  - Require regular trim
- core (**ulimit –Sc 0**) and other files

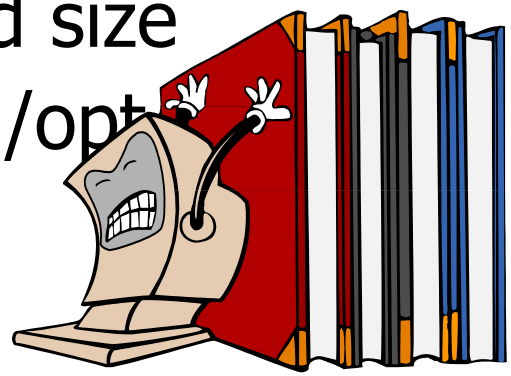

Log Files

- /var/adm:
  - acct/
  - automount.log
  - btmp
  - crash/
  - cron/
  - diag/
  - dmesg.log
  - eisa/
  - inetd.sec
  - lp/
  - netstat_data
  - nettl.LOG00
  - ps_data
  - ptydaemonlog
- /var/adm:
  - rbootd.log
  - rpc.lockd.log
  - rpc.statd.log
  - sbtab
  - shutdownlog@ -> /etc/shutdownlog
  - streams/
  - sulog
  - sw/
  - syslog/
  - vtdaemonlog
  - wtmp
  - wtmpold

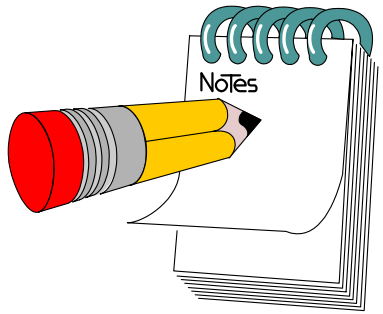

Core Files

- Identify with file command:

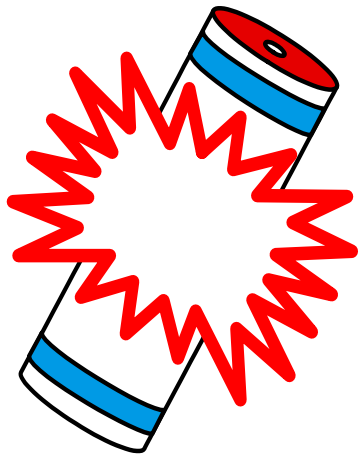

```
# sleep 999 &
# ps
11244 ttyp3  0:00 sleep
# kill -SIGQUIT 11244
# ll core
-rw-------1 root sys 243140 Nov 27 16:16 core
# file core
core: core file from 'sleep'-received SIGQUIT
```

Core Files - cleanup

- Find and remove:
  - `find / -fsonly vxfs -type f -name core -exec rm {} \;`
- add `-mtime` for development (save for a while)

  `-mtime +3` (older than 3 days)

  `-mtime -3` (newer than 3 days)

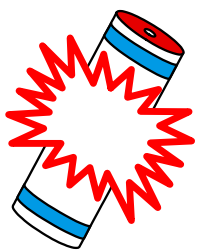

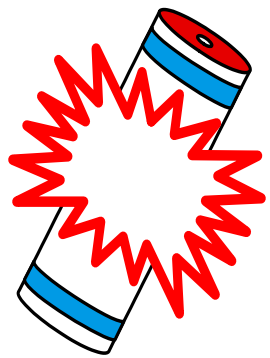

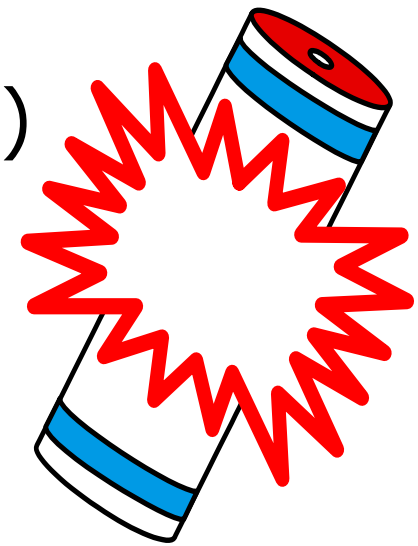

Cron Monitoring

- diskspace.sh:
  - monitors all mountpoints
  - selectable limits
  - automated notification
  - run from cron
  - Download from ftp site

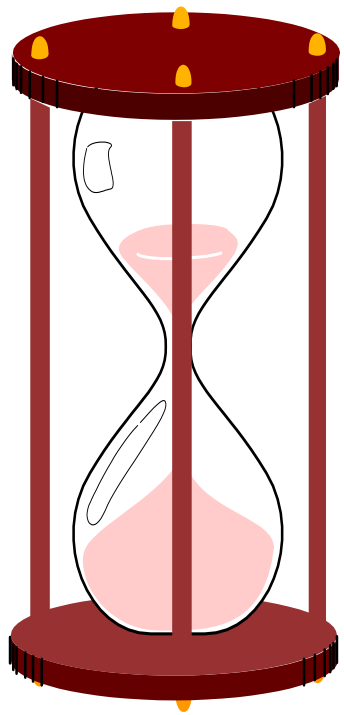

Logfiles

- Characteristics and locations
- syslog, the big one
- User Activity logs (*tmp)
- Miscellaneous logs

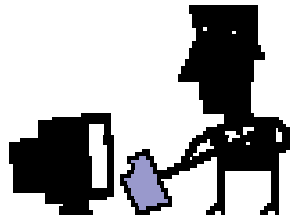

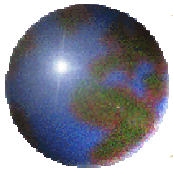

Logfiles - what to do?

- Can’t live with ‘em, can’t live without ‘em
- Why?
- Where?
- How?
- Huh?

Filesystem Full!

- White paper (9.x and 10.x/11.x)

  ftp://contrib:9unsupp8@hprc.external.hp.com

- find versus du
  - find = files only
  - du = directory sums:

du -x /home | sort -rn > /usr/tmp/du.home

Where?

- /var/adm
  - Most system logfiles
  - diag logs
- /var/adm/lp
  - log, lpd.log, lpana.log
- (and others in /var/adm)
- applications!

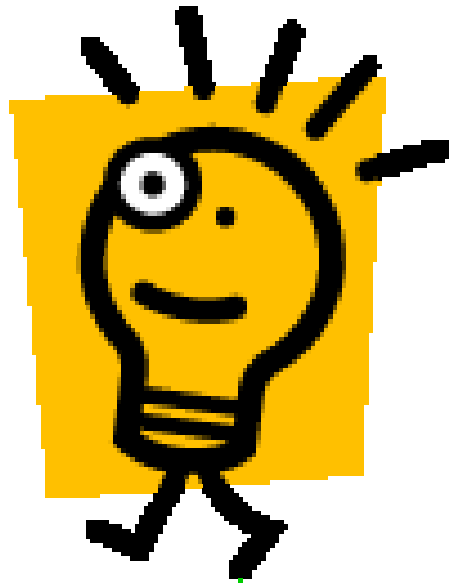

Syslog - everything's here

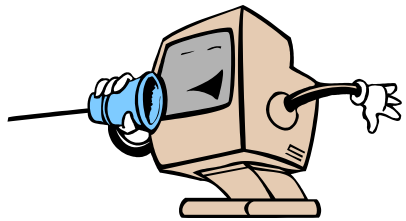

- /var/adm/syslog/syslog.log
  - Bootup info
  - daemons, kernel, errors, info
  - applications
  - user-defined messages

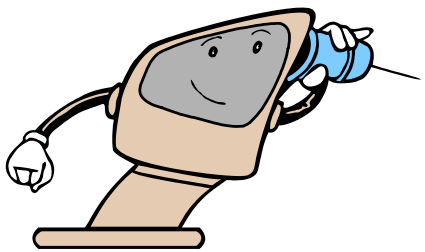

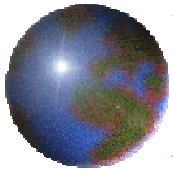

Logfiles and crash dumps

- Example entries:

```
Jan 19 02:00:01 bambam named[3790]: reloading nameserver
       \               \        \              \___ message
        \               \        \_process[PID]
         \               \_machine
          \_Date/time
```

syslog boot messages

- `syslogd: restart`
- `vmunix: 8/16 bus_adapter`
- `vmunix: Logical volume 64, 0x3 configured as ROOT`
- `vmunix: Memory Information:`
- `vmunix: physical page size = 4096 bytes, logical page size = 4096`
- `vmunix: Physical: 131072 Kbytes, lock[illegible]0036 Kbytes, available`

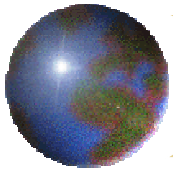

More syslog oneliners

- `nettl[593]: nettl starting up.`
- `inetd[749]: Reading configuration`
- `inetd[749]: ftp/tcp: Added service, server /usr/lbin/ftpd`
- `inetd[19760]: auth/tcp: Connection from i3125ljm`
- `xntpd[864]: offset -0.001348 freq -26.23103 comp 6`
- `ftpd[8204]: User hassell: Login incorrect`
- `syslogd: going down on signal 1`

Still more syslog

- `snmpd[5706]: EXCEPTIONS: bad version: 1`
- `last message repeated 3 times`
- `do_timeouts() - cannot unmount XXX:/`
- `do_timeouts() check symlink on XXX`
- `fs is busy - keep the link` (automount -v msg)
- `bootpd[266]: found 15.17.186.105 e3109569`
- `bootpd[266]: request from hardware address 52415320D044` (bootp patch or /etc/syslog.conf)

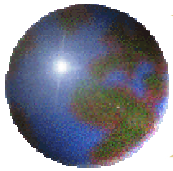

syslog.conf

- facility:
  - `kern, mail, lpr, daemon, auth, local0...local7, mark`
- level:
  - `debug, info, notice, warning, err, alert, emerg, crit, none`
- destination/action:
  - pathname/file: `/var/adm/mylogfile`
  - device: `/dev/console`
  - usernames (logged in)
  - * (every user)

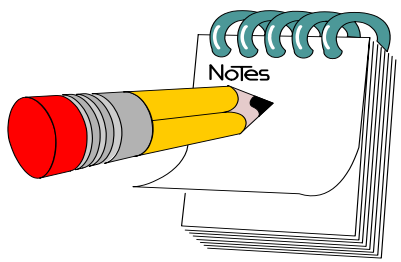

syslog examples

```
kern,mark.debug  /dev/console
mail.debug       /var/adm/syslog/mail.log
*.info;mail.none /var/adm/syslog/syslog.log
*.alert                  /dev/console
*.alert                  root,eric, hassell
*.emerg                  *
*.emerg                  @ admin.hp.com
```

logger

- shell script log entries:

```
logger  [-t tag] [-p pri] [-f file] msg
```

  - `logger -t MyScript -p lpr.warn Testing`
  - `logger -t blh-cron -p kern.alert "step 3 ok"`
  - `logger -p local3.info "News throttled"`
- Test syslog.conf entries
- Debug scripts with no tty output

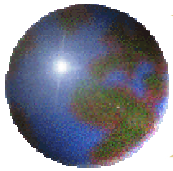

Facility Logging

- Idea: separate syslog files
  - `syslog.kern`
  - `syslog.mail`
  - `syslog.lpr`
  - `syslog.daemon`
  - `syslog.auth`
  - etc...

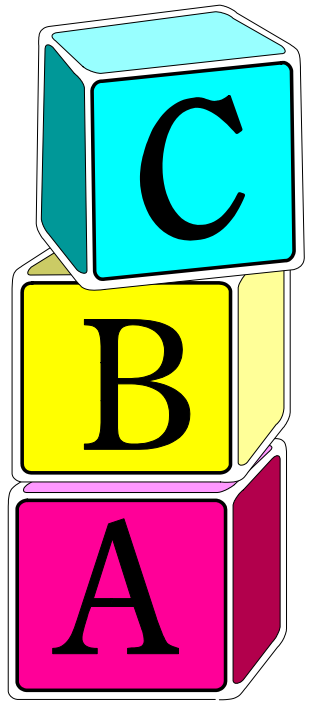

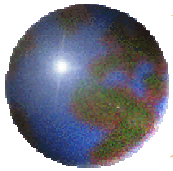

User Activity Logs

- `utmp` (zeroed at reboot)
- `wtmp` (cumulative!)
- `btmp` (create!)
- `last, who, write, login`
- `fwtmp` to decode (`/usr/sbin/acct/fwtmp`)

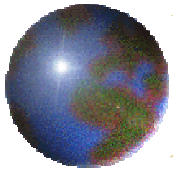

*The *tmp files*

- *tmp files are all in binary format
- `/etc/utmp`
  - active login/logout
  - rebuilt at bootup
  - needs valid logouts
  - user programs or actions can corrupt

    (`must login at lowest level shell`…)

*More *tmp files*

- `/var/adm/wtmp`
  - comprehensive and cumulative:
    - logins/logouts
    - init changes
    - boot time
    - accounting
  - Always needs trimming!
  - View: `last -R -<qty> login-name`

*Still more *tmp files*

- `/var/adm/btmp`:
  - Bad logins by UID and tty port
  - Never allow group/user readable
  - `lastb` (section 1, not 1m, hence warning)
  - watch for sudden jumps in size and repeated failures for root or other users

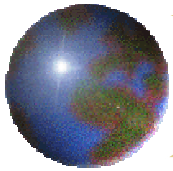

fwtmp

- Program is hiding:
  - `/usr/sbin/acct/fwtmp`
- decodes `utmp, wtmp, btmp` into ASCII
- no errors!
- Converts either direction

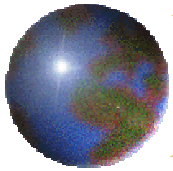

lp logfiles

- `/var/adm/lp`:
  - `log` (`lpsched -v` for verbose entries)
  - `lpd.log` (`rlpdaemon -l` to enable logging)
  - `lpana.log` (`lpsched -a` to add `lpana` logging)

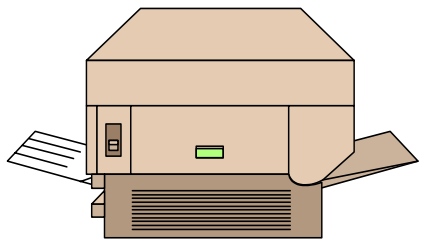

cron

- `/var/adm/cron/log`
- Format:

```
> CMD: command to run
> owner PID c/a start time
< owner PID c/a start time
```

- Sample:

```
>  CMD: /usr/sbin/dmesg - >> /var/adm/dmesg.log
>  root 10034 c Sun Mar  4 01:11:00 EST 2001
<  root 10034 c Sun Mar  4 01:11:00 EST 2001
>  CMD: $HOME/cronfiles/bin/badboy
>  root 10206 c Sun Mar  4 01:15:00 EST 2001
<  root 10206 c Sun Mar  4 01:15:01 EST 2001
```

sulog

- `/var/adm/sulog`
  - logs every su change (good or bad)
  - `/etc/securetty` to force su - root
    - Use: `console` (not `/dev/console`)
    - man `login`

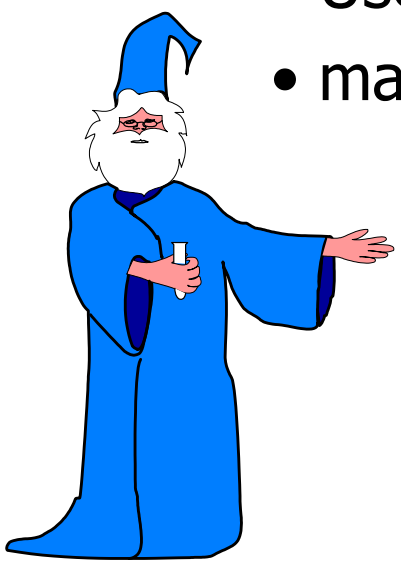

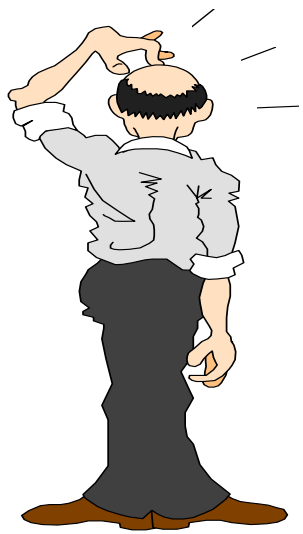

/etc/shutdownlog

- Make sure it exists
- logs every normal shutdown
- logs abnormal restarts and can decode a panic (crash) if `/var/adm/crash` is OK

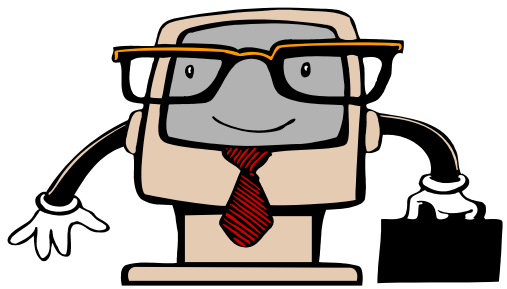

mail logs

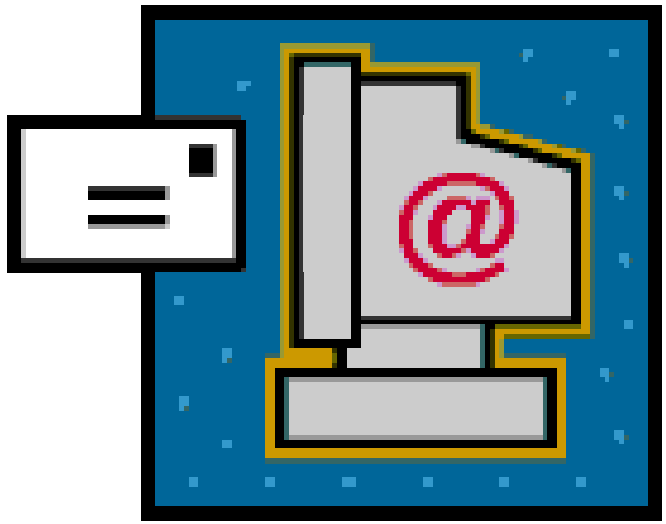

- `/var/adm/syslog/mail.log`
  - email transactions
  - Sample:

```
Jan 19 22:58:28 hpuerca sendmail[2524]:
    AA025201: msgid=<19980..
```

more mail logs

- `from=<baae26@mpset.com.mx`
- `to=<rex7@bambam.atl.hp.com>, delay=00:00:05, stat=Sent, mailer=local`
- `to=jo_amstrong@hp.com, delay=10:00:04, stat=Deferred: <this_is_not@my.email.address.com>... unresolvable host name my.email.address.com,check your configuration., mailer=tcp, MX host=palsmtpx.hp.com., address=[15.81.184.10]`

Install/SD logs

- Software Distributor
  - `/var/adm/sw`
    `swagent.log  swconfig.log swinstall.log`
    `swpackage.log swremove.log swagentd.log`
    `swcopy.log swmodify.log swreg.log`
    `swverify.log`
- `cleanup` program
- `Swmodify:`

```
swmodify -x patch_commit=true <patch_name>
```

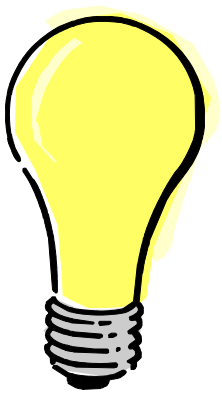

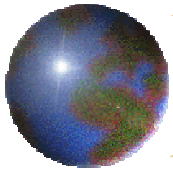

logtrim

- Trim common logfiles
  - archive older information
    - log -> log.1
    - log.1 -> log.2 etc
  - compress the archives
  - ensure no data loss
  - correct permissions/ownerships
  - available from:

`ftp://contrib:9unsupp8@hprc.external.hp.com`

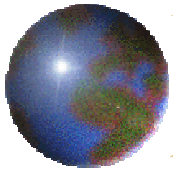

Useful tools (ftp site)

- autosum
- bdfmegs
- cstm-cpu
- cstm-mem
- cstm-periph
- cstm-all
- errno
- jetinfo
- laninfo
- ljdisplay
- lls
- llt
- lpinfo
- loadmedia
- mx
- psbyname
- psgrep, pslgrep
- psram
- pwgrep
- remshall
- restart
- syslogdecode
- viman

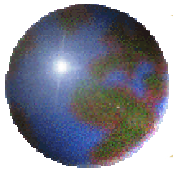

Additional resources

- Web help:

  At: `docs.hp.com /hpux/onlinedocs/939/KCParms/ KCparams.OverviewAll.html`
- `/usr/share/doc` (11.0, not 11i)
- `itrc.hp.com`
- `software.hp.com`
- `eproducts.hp.com`
- `education.hp.com`
- `docs.hp.com`
- `licensing.hp.com`
- `partsurfer.hp.com`
- Interex-Netherlands: sysadmin mail list

```
echo subscribe hpux-admin |
mailx -sSubscribe  majordomo@dutchworks.nl
```