



# 62

# AIX

*December 2000*

---

## **In this issue**

- 3 Command-line processing
- 9 Performance monitoring with GlancePlus
- 23 Using Samba to integrate AIX and Windows
- 27 SP/2 node migration
- 32 A utility to terminate and clean up user processes
- 34 Maintain network details – part 2
- 52 AIX news

update

# AIX Update

---

## Published by

Xephon  
27-35 London Road  
Newbury  
Berkshire RG14 1JL  
England  
Telephone: 01635 550955  
From USA: 01144 1635 33823  
E-mail: harryl@xephon.com

## North American office

Xephon/QNA  
Post Office Box 350100, Westminster CO  
80035-0100, USA  
Telephone: (303) 410 9344  
Fax: (303) 438 0290

## Contributions

If you have anything original to say about AIX, or any interesting experience to recount, why not spend an hour or two putting it on paper? The article need not be very long – two or three paragraphs could be sufficient. Not only will you actively be helping the free exchange of information, which benefits all AIX users, but you will also gain professional recognition for your expertise and that of your colleagues, as well as being paid a publication fee – Xephon pays at the rate of £170 (\$250) per 1000 words for original material published in AIX Update.

To find out more about contributing an article, see *Notes for contributors* on Xephon's Web site, where you can download *Notes for contributors* in either text form or as an Adobe Acrobat file.

## Editor

Harold Lewis

## Disclaimer

Readers are cautioned that, although the information in this journal is presented in good faith, neither Xephon nor the organizations or individuals that supplied information in this journal give any warranty or make any representations as to the accuracy of the material it contains. Neither Xephon nor the contributing organizations or individuals accept any liability of any kind howsoever arising out of the use of such material. Readers should satisfy themselves as to the correctness and relevance to their circumstances of all advice, information, code, JCL, scripts, and other contents of this journal before making any use of it.

## Subscriptions and back-issues

A year's subscription to *AIX Update*, comprising twelve monthly issues, costs £180.00 in the UK; \$275.00 in the USA and Canada; £186.00 in Europe; £192.00 in Australasia and Japan; and £190.50 elsewhere. In all cases the price includes postage. Individual issues, starting with the November 1995 issue, are available separately to subscribers for £16.00 (\$23.00) each including postage.

## AIX Update on-line

Code from *AIX Update* is available from Xephon's Web page at [www.xephon.com/aixupdate.html](http://www.xephon.com/aixupdate.html) (you'll need the user-id shown on your address label to access it).

---

© Xephon plc 2000. All rights reserved. None of the text in this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior permission of the copyright owner. Subscribers are free to copy any code reproduced in this publication for use in their own installations, but may not sell such code or incorporate it in any commercial product. No part of this publication may be used for any form of advertising, sales promotion, or publicity without the written permission of the publisher. Copying permits are available from Xephon in the form of pressure-sensitive labels, for application to individual copies. A pack of 240 labels costs \$36 (£24), giving a cost per copy of 15 cents (10 pence). To order, contact Xephon at any of the addresses above.

*Printed in England.*

## Command-line processing

I have often found myself pondering a complex command line and wondering: “do I need a dot between two backslashes, or should it be a back quote followed by a hyphen?” Well, perhaps not that exact question, though the arcane punctuation of the Unix command line can lead to some very unusual-looking commands.

Something that will help sort out what is happening in a command line, and improve your ability to construct a command line, is an understanding of how the command line is interpreted.

The command line is the input to the shell. The shell (**sh**, **ksh**, **cs****h**, or any other variant) reads the input line, untangles it, and then attempts to execute portions of it that are recognized as commands. The sequence the shell goes through to untangle a command provides an interesting insight into shell programming and will help you learn some new tricks.

Each of these steps is covered in more detail in a moment, but first look at the sequence of evaluation of a command line.

- 1 History substitution (except for the Bourne Shell).
- 2 Splitting words, including special characters.
- 3 Updating the history list (except for the Bourne Shell).
- 4 Interpreting single and double quotes.
- 5 Alias substitution (except for the Bourne shell)
- 6 Redirection of input and output (`<`, `>`, `|`, etc).
- 7 Variable substitution (variables starting with `$`).
- 8 Command substitution (commands inside back quotes).
- 9 Filename expansion (file name wildcards).

Note that the Bourne shell is missing the steps to do with history and alias substitution. These are not part of the Bourne shell.

## HISTORY SUBSTITUTION

If you have history set up in the Korn shell (**ksh**), C shell (**cs**h), or any other shell, then each command line is saved in a history file before being executed. You can review your previous commands by typing:

```
$ history
```

The resulting list of commands is preceded by a number, as in:

```
13 ls *.txt
14 cd $HOME
15 ls *.log
```

In the Korn Shell, you can usually recover a history command by typing **r** followed by the number of the command in the history file. Typing **r 13** in the example above would repeat the command **ls \*.txt**.

In the C shell, use an exclamation mark instead of an **r**, with no intervening space, as in **!13**.

The first thing the shell does when processing a command line is check for command substitutions from the history file. If a request is found, the command is looked up in the history file and recovered. The recovered command is then used to create a new command line. There is much more to command history than this simple example suggests, but that's a subject for a separate article!

## SPLITTING WORDS

A 'word' on a Unix command line is a token that is recognized by the shell as an element of a command. Consider, for example, the following command, which does a 'long listing' of the current directory and searches for 'mjb' in any line of the returned directory information:

```
ls -l|grep mjb
```

The 'words' in this command are: **ls**, **-l**, **|**, **grep**, and **mjb**.

A word can also be a quoted string. In the command below, a long directory listing is scanned for files created on 'Sep 07'.

```
ls -l|grep "Sep 07"
```

In this case the words are **ls**, **-l**, **|**, **grep**, and **"Sep 07"** (the last item is

treated as one word as it is enclosed in quotes). So, in this step, commands are separated into ‘words’ or tokens, as I did in the two examples above.

## UPDATE THE HISTORY LIST

Once the words in a command line are identified, the command is written to the end of the history file, if you’re using a shell that includes a command history facility and this facility is enabled.

## SINGLE AND DOUBLE QUOTES

If a word is enclosed in double or single quotes, the word is tagged so that variables in it are either expanded or left as-is, depending on which type of quotation mark is used (variables surrounded by single quotes are left as-is, while ones in double quotes are expanded). To test these two different types of behaviour for yourself, enter the following commands with no quotes, single quotes, and double quotes:

```
echo $PATH
echo '$PATH'
echo "$PATH"
```

The first displays the value of the *\$PATH* variable; the second displays the word ‘\$PATH’, and the third again displays the value of the *\$PATH* variable, as in the following examples:

```
echo $PATH
/bin:/usr/bin:/my/bin:.
```

```
echo '$PATH'
$PATH
```

```
echo "$PATH"
/bin:/usr/bin:/my/bin:.
```

As an exercise, try the following commands using double quotes around single quotes and single quotes around double quotes. When in doubt about the effect of something in a command line, experiment.

```
echo $PATH
/bin:/usr/bin:/my/bin:.
```

```
echo "'$PATH'"           # single quotes around double quotes
"$PATH"
```

```
echo "$PATH"           # double quotes around single quotes
'/bin:/usr/bin:/my/bin:.'
```

Bear in mind that, even though we've discussed variable expansion here, the actual expansion of variables in words takes place at a later stage in the processing of a command line. By the end of this stage all that's happened is that words have been marked according to whether they are in single or double quotes. Variable expansion has to wait a couple more steps.

## ALIAS SUBSTITUTION

Alias substitution is the next stage of command line processing.

The first word of each command is checked against the alias list. In the example above (**ls -l|grep mjb**), **ls** and **grep** are checked in the alias list and any alias substitution is performed. An alias is a method of substituting one command for another. For example the following command:

```
alias ll 'ls -l '
```

creates a 'new' command **ll**, so that the command:

```
ll *.txt
```

it is equivalent to:

```
ls -l *.txt.
```

## PIPES AND REDIRECTION

At this point the shell looks through the words for pipe and redirection commands (**|**, **>**, **<**, **>>**, and others). If it finds one, it creates the pipe or establishes the redirection.

## VARIABLE EXPANSION

Now, at last, variables in the command line are expanded, based on the type of quotes (if any) that enclose them. Thus *\$PATH* and any other variable in the command is replaced by its value (if the variable is not enclosed by single quotes).

## COMMAND SUBSTITUTION

Command substitution involves looking for ‘back quotes’ (``). If back quotes are found, then everything between them, commands and arguments, is executed as a command, the results of which are substituted for the back-quoted string in the outer command.

Consider, for example, this command:

```
ls -l `ls -p|grep /`|more
```

**ls -p** produces a directory listing in which any directory is marked with a trailing slash. The sample listing below includes two directories, *somestuff* and *xdir*, as indicated by the trailing slash.

```
ls -p
STARTUP
file.txt
file2.txt
somestuff/
xdir/
```

Adding **grep /** to the command line selects only those lines that contain a trailing slash.

```
ls -p|grep /
somestuff/
xdir/
```

By enclosing the whole of **`ls -p|grep /`** in back quotes, the command is executed and the results are handed to the enclosing command as arguments. Thus, the following example:

```
ls -l `ls -p|grep /`|more
```

is equivalent to:

```
ls -l somestuff/ xdir/|more
```

This causes a page-by-page listing of all subdirectories within the requested directory, and this is how the command substitution phase of command processing works.

## WILDCARDS

Wildcards used in filenames are located and expanded. These are the standard wildcard characters ‘\*’ and ‘?’, as well as ‘bracketed wildcards’, as in:

```
ls -l [abc]*
```

which provides a listing of any files or directories that start with ‘a’, ‘b’, or ‘c’.

## EXECUTION

Finally, with everything expanded, substituted, shampooed, and clean-shaven, the command is executed. Not surprisingly, this is the last step in the shell processing sequence.

One additional thing is worth mentioning: when the command processor encounters a command between back quotes (in step 8), it isolates the quoted command and applies steps 1 through 9 to it.

This process also happens on commands separated by semicolons (;). Steps 1 through 9 are performed on each separate command. You can test this yourself by selecting a directory with only a few files, and then issuing the command to display a list of files in another directory:

```
echo files in /chosen/dir are ; echo `cd /chosen/dir; ls *`
```

What is echoed is the list of files in */chosen/dir*. The asterisk argument to **ls** is obviously not expanded until after the **cd** command. If this were not the case, the asterisk would be expanded using the list of files in the current directory, instead of the target directory.

Commands within back quotes and commands separated by semicolons are both processed with all steps from 1 through 9 as though they were separate commands.



# Performance monitoring with GlancePlus

GlancePlus is a tool from Hewlett-Packard for carrying out performance monitoring and producing performance alerts. It's available for AIX, and also for the HP-UX and Solaris operating systems. The program is supplied with two kinds of user-interface: one that runs in character-mode (**glance**) and a graphical one based on Motif (**gpm**). GlancePlus enables a system administrator to:

- Interpret performance data and identify bottlenecks using a rules-based 'advisor'.
- View metrics at the global-, application-, and process-level.
- Use alarms to monitor the system.
- Define filters to see only information that's of interest.
- Define the order of columns of data displayed and their sort order.

GlancePlus can be integrated with various network management systems, such as H-P's OpenView, IBM's NetView, and AT&T's OneVision – the registration and script files necessary for integration are included with the product. Integration enables GlancePlus to be invoked from a network management console on any suitable system.

## INSTALLATION

According to its release notes, GlancePlus runs on all versions of AIX from 4.1.5 to 4.3.2. However, I was able to execute it on AIX 4.3.3 (with Maintenance Level 04) without any version-specific issues. You need about 10 MB of free space in directory */usr/lpp/* to install the program successfully.

Below is the installation procedure for the program.

- 1 Log in as root.
- 2 Check whether a previous version of the tool has been run – this can be done with the command:

```
/usr/lpp/perf/bin/perfstat
```

Ask all users of either **glance** or **gpm** who are using the program to quit it.

- 3 Check whether H-P's MeasureWare software agent is running using the following command:

```
ps alxww|fgrep -v fgrep|fgrep mwa
```

If **mwa** is running, terminate it using the command:

```
/usr/lpp/perf/bin/mwa stop
```

- 4 Check that at least 10 MB of free space is available in the */usr/lpp* directory. If you cannot allocate enough disk space, create a symbolic link */usr/lpp/perf* in the directory that points to a filesystem that has sufficient free space. For instance, if you have free space in */opt*, then use the following commands to create the link:

```
mkdir /opt/perf;  
ln -s /usr/lpp/perf /opt/perf
```

- 5 Mount the program's CD-ROM using the following command:

```
mount -rv cdrfs /dev/cd0 /cdrom
```

Change the current directory to */cdrom* and ensure that you can access a file called *install.gp*.

- 6 Execute the installation script:

```
./install.gp
```

- 7 Change directory from the CD-ROM, then unmount the CD-ROM using the command:

```
umount /dev/cd0
```

- 8 If **mwa** was stopped in Step 3, restart it with the following command:

```
/usr/lpp/perf/bin/mwa start
```

- 9 Execute the command **perfstat -v** to verify that the installation completed successfully. The sample output below is from my system after GlancePlus was installed successfully.

```
*****
```

```

** perfstat for host1 on Fri Sep 15 10:22:44 WET 2000
*****
Listing the installed RPM products:
  HP OpenView GlancePlus A.01.02.03

listing version strings for performance tool files:
-----

NOTE: The following software version information can be
compared with the version information shown in the
/usr/lpp/perf/ReleaseNotes file(s).

GlancePlus executables in directory /usr/lpp/perf/bin
      glance  A.01.02.03      06/07/99 AIX 4.1+
      gpm     A.01.02       08/13/97 AIX 4.1+

GlancePlus library in directory /usr/lpp/perf/lib
      libnums.a B.02.19.03    07/05/99 AIX 4.1+

***** (end of perfstat -v output) *****

```

If you encounter problems while installing or running GlancePlus, refer to the document `/usr/lpp/perf/ReleaseNotes/gpm.msg`, which contains the product's *Installation Troubleshooting Guide*.

GlancePlus installs files using the following directory structure on AIX:

- `~username/`  
Configuration start-up files.
- `/usr/lpp/perf/bin`  
Binaries and shell scripts.
- `/usr/lpp/perf/help/gp/C`  
On-line help files.
- `/usr/lpp/perf/help/gp/C/gpmgraphics`  
On-line help graphics files.
- `/usr/lpp/perf/lib`  
Library files.
- `/usr/lpp/perf/lib/nls/msg/`

Message catalogs.

- */usr/lpp/perf/man/man1*

**man** pages.

- */usr/lpp/perf/newconfig*

Temporary copies of configuration and template files.

- */usr/lpp/perf/newonfig/app-defaults/C*

A temporary copy of **gpm**'s default settings file.

- */usr/lpp/perf/paperdocs/gp/C*

Compressed printable documents (see the references at the end of this article for details).

- */usr/lpp/perf/ReleaseNotes*

Release notes and installation documentation.

- */var/opt/perf*

Product configuration and status files.

- */var/opt/perf/app-defaults/C*

The production copy of **gpm**'s default settings file.

- */var/opt/perf/datafiles*

**gpm**'s default configuration file.

## GLANCEPLUS PERFORMANCE DISPLAYS

GlancePlus lets you view many aspects of your system's performance, including:

- Running processes
- CPU
- Memory
- Swap space

```

B3696A GlancePlus A.01.02.03 12:45:59 rsg301 7012/G30 Current Avg High
-----
CPU Util  S  SU  U  | 46% 41% 48%
Disk Util  0  0  | 10% 16% 63%
Mem Util  S  |240% 240% 240%
Swap Util  U  | 2% 2% 2%
-----
PROCESS LIST
Process Name PID PPID Pri Name CPU Util Cum CPU Block IO Rate RSS On
-----
find 17278 1 0 root 3.2/14.1 165.2 0.0/ 0.0 44kb SLEEP
glance 23108 20224 0 root 0.4/ 0.4 2.9 0.0/ 0.0 1.8mb SLEEP
glbd 14458 4440 0 root 0.4/ 0.0 164.3 0.0/ 0.0 1.1mb SLEEP
ialld 18060 4440 0 root 0.2/ 0.2 914.7 0.0/ 0.0 564kb SLEEP
ialnd 18576 4440 0 root 0.2/ 0.1 288.4 0.0/ 0.0 2.5mb SLEEP
kproc 3948 1 0 root 0.0/ 0.0 0.0 0.0/ 0.0 7.1mb SLEEP
kproc 4152 0 0 root 0.0/ 0.0 3.7 0.0/ 0.0 7.1mb SLEEP
kproc 4644 1 0 root 0.0/ 0.0 0.0 0.0/ 0.0 7.1mb SLEEP
kproc 1032 0 0 root 0.4/ 0.1 672.6 0.0/ 0.0 7.0mb SLEEP
kproc 4934 0 0 root 0.0/ 0.0 1.9 0.0/ 0.0 7.1mb SLEEP
kproc 1290 0 0 root 0.0/ 0.0 37.4 0.0/ 0.0 7.1mb SLEEP
? - Commands Menu Page 1 of 2
Proclist CPU Rpt Mem Rpt Disk Rpt NextKeys AppList Help Exit

```

Figure 1: GlancePlus's process list report

- Disk
- LAN
- NFS
- I/O (by filesystem, physical disk, logical volume)
- System table utilization.

Figure 1 shows GlancePlus's process list report (this and other screenshots in this article are taken from the character-mode interface). This is the default screen shown when **glance** is invoked.

The processes list provides an overview of system resources and active processes. The top section of the screen (the title and histogram) is common to many GlancePlus screens. The bottom section displays a summary of active processes.

The first line of the display shows GlancePlus's product name and version number, the time, the computer's host name, and the system type. The example shown is running GlancePlus version A.01.02.03 on server 'rsg301', which is an RS/6000 7012/G30.

Next come four bar charts that are displayed using reverse video. The first bar shows overall CPU utilization – the bar is divided into

fragments that are indicated by the following letters:

- 1 *S* indicates the amount of time spent on system activity, such as context switching and system calls.
- 2 *U* indicates the amount of time spent running user processes.

At the far right of each bar is a row of figures showing the current, average (calculated from the time GlancePlus was started), and highest percentages.

The second bar shows disk utilization for the busiest mounted disk. The bar shows the percentage of file system and virtual memory disk I/O during the update interval.

The fourth bar shows system memory utilization divided into three parts:

- 1 *S* indicates memory used by the system.
- 2 *U* indicates memory used by user programs and data.
- 3 *B* indicates memory used for file mapping.

Note that memory data in the *Current*, *Avg*, and *High* columns is wrong – this is a known bug in GlancePlus running on AIX 4.3.

The fourth bar shows swap space information, showing the percentage of total swap in use, and the *PROCESS LIST* section shows the most active processes on the system. The headings are as follows:

- *Process Name*  
The name or abbreviation used to load the executable program.
- *PID*  
The process identification number.
- *PPID*  
The PID of the parent process.
- *Pri*  
The priority of the process – the lower the number, the higher the priority.

- *User Name*  
The name of the user who started the process.
- *CPU Util*  
The first number is the percentage of CPU consumed by the process over the current interval. The second number is the percentage of CPU utilization accounted for by the process since GlancePlus started. Note that each processor contributes 100% to CPU capacity, so my two-processor system has a maximum CPU capacity of 200%!
- *Cum CPU*  
The CPU time (in seconds) used by the process since it started.
- *Block IO Rate*  
The first number is the average block I/O rate per second over the last update interval, and the second number is the average block I/O rate since GlancePlus was started.
- *RSS*  
The RAM consumed by the process in KB. All memory used by the process is counted, including data, stack, text, and shared memory segments.
- *Block On*  
The reason the process was last blocked (unable to run). The table below lists the most common reasons for process blocking.

Block conditions and their reasons:

<i>DIED</i>	The process terminated during the sampling interval.
<i>LOCK</i>	The process is waiting for serialization of a physical lock.
<i>NEW</i>	The process was created during the sampling interval.
<i>NONE</i>	The process is not blocked and is ready to run.
<i>OTHER</i>	The process is waiting for some reason that's not recognized by GlancePlus.

- PRI** The process is waiting in a run queue.
- SLEEP** The process is waiting for an event to complete.
- TIMER** The process is waiting for a timer to time out.
- TRACE** The process received a signal to stop because its parent is tracing it.
- VM** The process is waiting for a virtual memory operation to complete.
- ZOMB** The process has terminated; however, its parent is not waiting for its termination.

GlancePlus's default refresh interval is five seconds – this can be changed when the program is invoked using the **-j** command-line switch with a numerical argument. The operator can force an immediate refresh with the key sequence **CTRL+L**.

Note that there may be more than one page of processes to display. The message at the bottom right-hand corner of the screen indicates which page you're viewing. You can scroll forward to the next page with **f** and back with **b** – navigation between the tool's various screens is done using one-letter commands. Use the command **?** to display a screen with a summary of commands (Figure 2).

```

B3696A GlancePlus A.01.02.03 12:47:24 rsg301 7012/630 Current Avg High
-----
CPU Util  S ██████████ SU  U  | 38% 41% 48%
Disk Util  D ██████████  | 27% 20% 63%
Mem Util   S ██████████ S  | 240% 240% 240%
Swap Util  U ██████████  | 2% 2% 2%
-----
GlancePlus Commands Menu

? - Commands Menu  b - Page Backward (or -)  < - Display Previous Screen
! - Invoke Shell  F - Page Forward (or +, space)  > - Display Next Logical Scr
h - Online Help  q - exit (or e)  z - Reset Statistics to Zero
p - Print Toggle  r - Refresh Screen (or ^L)  <cr>- Update Current Screen

g - Process List  d - Disk Report  A - Application List
a - CPU By Processor  u - IO By Disk  T - Transaction Tracker
c - CPU Report  i - File System Capacity  S - Select Appl / Trans
m - Memory Report  v - IO By Logical Volume  s - Select Process
t - System Tables  N - NFS Global Activity  F - Process Open Files
w - Swap Space  l - Network By Interface  R - Process Resources
H - Alarm History
j - Adjust Refresh Interval  o - Process Threshold Options  y - Renice Process

```

*Figure 2: GlancePlus commands*



```

B3096@ GlancePlus 8.01.02.03 14:15:00 rsg001 7012/630 Current Avg High
-----
CPU Util 50 | 11% 15% 51%
Disk Util 0 | 67% 8% 100%
Mem Util 5 | 240% 240% 240%
Swap Util 0 | 2% 2% 2%
-----
State          Current  Average  High      Time      Users=  2
-----
User           1.8     4.4     21.3     0.09     263.03
System         9.2     10.8    48.0     0.46     650.19
Idle          89.0     84.8    100.0    4.47     5092.28
Syscall Rate 1206.6  1024.3  4711.4
Interrupt Rate 121.8  190.2  909.2
CSwitch Rate  104.2  125.7  624.4
CPU Queue     1.0     0.9     5.0
Load Average  0.1     0.8     2.6
Top CPU user: PID 12098, Find      28.4% cpu util

```

Figure 3: The process list

The tool's on-line help system can be invoked with **h** command – information about navigation between different screens and about the statistics displayed appears in a text window.

If the process list shows that your CPU is overworked, you should check the CPU report screen (Figure 3), which can be displayed using the **c** command.

This screen reports CPU-related statistics for each type of usage (*user*, *system*, and *idle*). The following columns are displayed:

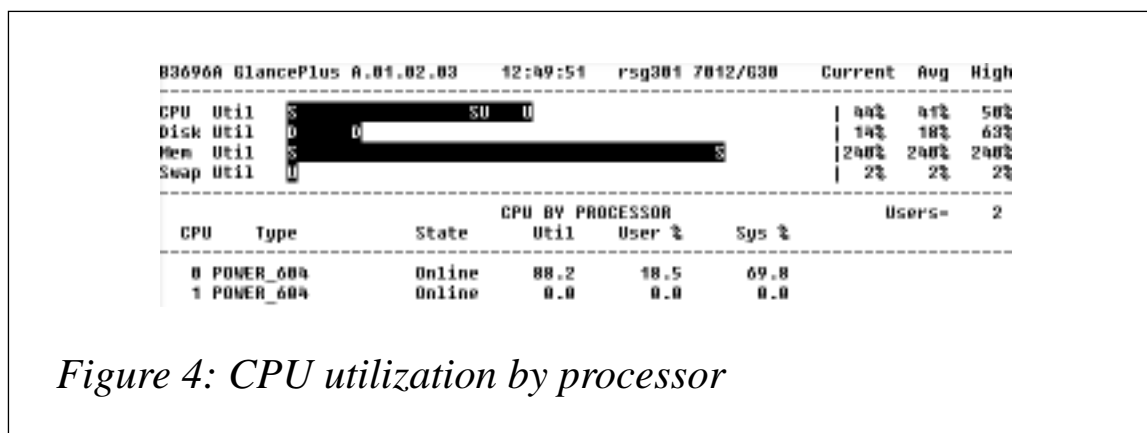
- *Current*  
Displays the percentage of CPU accounted for by this state over the last interval.
- *Average*  
Displays the percentage of CPU time spent in this state since GlancePlus started.
- *High*  
Displays the highest percentage of CPU time spent in this state since GlancePlus started.

- *Time*  
Displays the CPU time spent in this state over the last interval.
- *Cum Time*  
Displays the total CPU time spent in this state since GlancePlus started.

Additional statistics displayed in the report are the system call rate, interrupt rate, context switching rate, CPU (run) queue, and load average.

The last line of the report shows information about the process that used most CPU resources during the last interval.

Command **a** brings up a report on CPU utilization distribution between CPUs installed on your server (Figure 4).



*Figure 4: CPU utilization by processor*

Command **m** brings up a memory utilization report. The statistics are shown in terms of counts, rather than percentages (Figure 5).

The following five counters are shown for each memory management event:

- *Current*  
The number of times this event occurred during the last interval.
- *Cumulative*  
The sum of all counts for this event since GlancePlus started.

```

B2696A GlancePlus 8.01.02.03 14:12:24 rsg301 7012/630 Current Avg High
CPU Util | 0% 10% 51%
Disk Util | 1% 8% 100%
Mem Util | 240% 240% 240%
Swap Util | 2% 2% 2%
-----
MEMORY REPORT Users= 2
Event Current Cumulative Current Rate Cum Rate High Rate
-----
Page Faults 4 2568746 0.4 437.7 2233.0
Paging Requests 6 213233 0.7 36.4 874.4
Page Ins 2 197763 0.2 33.8 873.4
Page Outs 4 15470 0.4 2.6 191.6
Paged In (kb) 8kb 772.5mb 0.9 135.2 3493.6
Paged Out (kb) 16kb 60.4mb 1.9 10.5 766.4
Swap Ins 0 67 0.0 0.0 4.4
Swap Outs 0 1478 0.0 0.2 51.6
Swapped In (kb) 0kb 268kb 0.0 0.0 17.6
Swapped Out (kb) 0kb 5.8mb 0.0 1.0 206.4

```

Figure 5: The memory report screen

- *Current Rate*  
The number of events per second.
- *Cum Rate*  
The total number of events.
- *High Rate*  
The highest rate recorded.

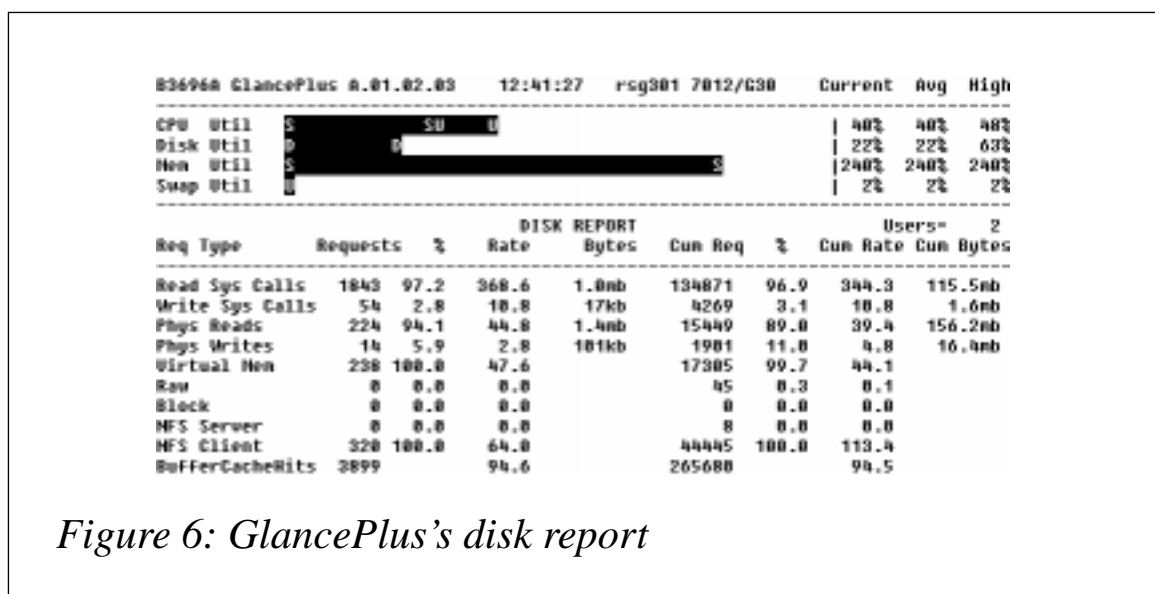
Below are brief descriptions of the memory management events for which statistics are provided.

- *Page Faults*  
A fault occurs when a process tries to access a page that's not in RAM.
- *Paging Requests*  
The number of paging requests.
- *Page Ins/Page Outs*  
The number of pages of data moved from virtual memory (disk) to RAM and vice-versa.

- *Paged In (kb)/Paged Out (kb)*  
The amount of data paged in and out.
- *Swap Ins/Swap Outs*  
The number of processes swapped into and out of memory.
- *Swapped In (kb)/Swapped out (kb)*  
The amount of data swapped into and out of memory when entire processes are deactivated.

The following additional values are also displayed in the memory report screen:

- *Phys Mem*  
The total RAM installed in your system.
- *Avail Mem*  
The amount of physical memory available.
- *Free mem*  
The amount of unallocated physical memory.
- *Buf cache*  
The amount of physical memory used by buffer cache. This



should be an absolute minimum on AIX thanks to its extensive use of memory-mapped files.

Use the **d** command to view the disk report screen (Figure 6).

There are eight statistics shown for various events related to logical and physical accesses to all disks that are accessible to the system. The statistics are:

- *Requests*  
The total number of requests of a particular type over the last interval.
- *Requests %*  
The percentage of this type of disk event relative to other types.
- *Rate*  
The average number of requests of this type per second.
- *Bytes*  
The total number of bytes transferred as a result of this type of event over the last interval.
- *Cum Req*  
The cumulative number of requests since GlancePlus started.
- *Cum Req %*  
The relative number of this type of disk event since GlancePlus started.
- *Cum Rate*  
The total number of requests.
- *Cum Bytes*  
The total number of bytes transferred as a result of this type of event since GlancePlus started.

Below are descriptions of the disk events for which these statistics are provided.

- *Read Sys Calls*  
The number of read system calls.
- *Write Sys Calls*  
The number of write system calls.
- *Phys Reads*  
The number of physical reads to disk.
- *Phys Writes*  
The number of physical writes to disk.
- *Virtual Mem*  
The amount of physical I/O as a result of virtual memory management activity.
- *Raw*  
The volume of 'raw mode' disk I/O.
- *Block*  
The total number of block I/O operations in the interval.
- *NFS Server*  
The level of I/O activity associated with the NFS service.
- *NFS Client*  
The level of I/O activity associated with the NFS client.
- *BufferCacheHits*  
The number of file I/O operations satisfied by buffer cache.

## REFERENCES

- 1 There are two **man** pages associated with GlancePlus:

`gpm(10)`

`glance(1)`

- 2 Two on-line help guides are also provided:
  - *The gpm On-line Help User's Guide*
  - *The GlancePlus On-line Help User's Guide.*
- 4 *The GlancePlus Adviser's Manual:*  
/usr/lpp/perf/paperdocs/gp/C/adviser.ps.Z
- 5 *The GlancePlus Concepts Manual:*  
/usr/lpp/perf/paperdocs/gp/C/concepts.ps.Z
- 6 *The Performance Metrics Definition Manual:*  
/usr/lpp/perf/paperdocs/gp/C/metrics.ps.Z

---

*Alex Polak*  
*System Engineer*  
*APS (Israel)*

© Xephon 2000

---

## **Using Samba to integrate AIX and Windows**

I was recently asked how three PC users in an AIX environment could share Excel files. Apparently, it was getting tiresome to transfer the files by diskette. My customer had expected an answer along the lines of electronic file transfer, as the PCs are connected to an AIX system that runs their accounting and order processing software. These programs run on Windows 95/98 PCs with Hummingbird's X server. The customer even mentioned that they would accept a solution in which user *A* transfers a file to the AIX system and then yells across the room to user *B* that the file is on the AIX system ready for download. They began drooling when I suggested that the files could reside on the AIX system but appear to be on a common virtual PC drive which they could read from and write to. I would name the drive *K:* in honour of the manager, whose surname begins with 'K'. And,

when the company's bookkeeper heard that no costly software or hardware would be needed, he asked for a common drive *M:*. You can guess with which letter his last name begins!

#### WHAT IS SAMBA?

The software I used is Samba. Samba is a suite of programs that allows an AIX system – or, for that matter, just about any Unix or Unix-like system – to provide resources, such as files and printers, to DOS- or Microsoft Windows-based systems. This software was mentioned in the August 1999 issue of *AIX Update* as a means of connecting AIX and NT systems.

Files and directories belonging to the AIX filesystem and AIX printers can be accessed directly from the Windows system. Samba allows users to open a file directly inside a Windows program, as though it were stored locally, and to use an AIX-controlled printer.

#### THE GNU PUBLIC LICENSE

Samba is supplied under the terms of the GNU Public License. This means that Samba is 'free' software. 'Free' in this sense means free from restriction, not cost. In other words, Samba's development team, headed by its founder Andrew Tridgell, who now lives in Australia, welcome contributions that expand Samba's capabilities. More information concerning GNU software can be found at the Free Software Foundation's Web site at <http://www.fsf.org>.

#### SMB VERSUS NFS

Samba uses the SMB (Server Message Block) protocol that Microsoft adopted for its native file sharing. In a Unix context, SMB is NetBIOS running over TCP/UDP. While NFS is oriented more towards machine-based file sharing, Samba shares are authenticated on a per-user basis, like **ftp** or **telnet**. Another major difference between the two is that Samba's primary transport is TCP instead of UDP. The most noticeable consequence of this is that Samba clients will not keep a connection open to a server if the server goes down.



## GETTING SAMBA

You can get a copy of Samba from <http://samba.ana.edu.au> or one of the many mirror sites whose addresses are on Samba's home site (<http://www.samba.org>). There are a few books around that contain not only detailed information on Samba but also a CD-ROM with a complete copy of the Samba **ftp** distribution, including source code, documentation, and binaries for most Unix operating systems, which (of course) includes IBM's AIX. Notable among them are *Teach yourself Samba in 24 Hours* by Samba Team members Jerry Carter and Richard Sharpe, *Samba: Integrating UNIX and Windows* by John Blair, also a Samba Team member (both books are published by SAMS), and *Using Samba* by Robert Eckstein, David Collier-Brown, and Peter Kelly (published by O'Reilly).

## CONFIGURATION OPTIONS

Samba can be administered from the command line with the help of a simple text editor. Fortunately no complex files are involved in its configuration. Only one file, *smb.conf*, which is usually located in */usr/local/samba/lib/*, needs to be edited.

## A NETWORKED DRIVE EXAMPLE

The following example *smb.conf* configuration file implements the virtual PC drive *K:*. The drive is physically on the AIX system in directory */home/samba\_K*.

### SMB.CONF (1)

```
[global]
  security = user
  log file = /usr/local/samba/log.%m
  smb passwd file = /etc/passwd

[K]
  comment = Pete's group drive
  path = /home/samba_K
  writeable = yes
  valid users = @staff
  locking = yes
  create mode = 0660
  directory mode = 0770
```

Our example requires the user to use his AIX user name and password to gain access to drive *K:*. Unlike its Windows 95 predecessor, for security reasons, Windows 98 will not allow you to send plain-text passwords – the password is encrypted by default. However, Samba servers can be configured to require plain-text passwords. You will not be able to connect to Samba servers that are configured in this way from Windows 98 clients unless you set a Registry entry to enable plain-text passwords. To do this, add the entry *EnablePlainTextPassword* (type DWORD) in the following location and set its value to '1':

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\VxD\VnetSUP
```

## PRINTER EXAMPLE

Our next example *smb.conf* configuration file lets your PC access printers configured on and connected to your AIX system.

## SMB.CONF (2)

```
[global]
    local printers = yes

[ps]
    printable = yes
    comment = AIX connected printer
    public = yes
    browseable = yes
    printer driver = IBM 4019 LaserPrinter PS39
```

Note that the printer driver parameter should specify the exact name of the Windows printer driver.

## OTHER SAMBA POSSIBILITIES

While allowing your PCs to access AIX files and printers is a wonder in itself, there are more sophisticated possibilities:

- Using a Windows tape drive to back up your AIX system.
- Synchronizing PC clocks with your AIX Samba server.
- Printing from an AIX system to a Windows print server.

- Having the AIX Samba server send WinPopup messages.

## SAMBA AND WINDOWS 2000

Samba 2.0.7 is the latest stable release of Samba. This version has been tested on Windows 2000 and has no *known* issues with that release of Windows.

---

*Werner Klauser*  
*Klauser Informatik (Switzerland)*

© Xephon 2000

---

## SP/2 node migration

Moving an operating system and applications off one SP node and onto another requires a fair amount of planning. In many cases, the process cannot be fully tested prior to actual implementation, which makes the task all the more daunting. Nevertheless, some planning can help make things go more smoothly. While the approach outlined here is the one our company used a number of times to migrate Model 604 and 604E (in other words ‘old’) high nodes to Silver wide nodes, the general steps apply to most other types of node.

Given that the AIX level of our high nodes before the migration was 4.2.1 (with latest maintenance applied), and that the minimum level of AIX on a Silver node is also 4.2.1, no upgrade was required. However, other migrations may require an AIX upgrade first – for example, the new POWER3 nodes require AIX 4.3.3 and PSSP 3.1.x.

Another point to note is that parts of the approach laid out in this article are based on ideas from an article in IBM’s AIX support Web site – to view the article, first point your browser to the following URL (note the use of the continuation character, ‘>’, below to indicate a formatting line break):

```
http://techsupport.services.ibm.com/support/rs6000.support/  
> databases
```

then select *RS/6000 SP Resources*, followed by *332 MHz SMP Node Migration Tips and Hints*. There are also scripts available at this Web site to help save and import volume group information.

## DOCUMENT YOUR SYSTEM

There is a lot of information that you need to gather about your system. This step is very important – for instance, some of the information is needed when you import non-root volume groups, while other information is simply not going to be available from the donor node once the transition is made. To this end:

- Document all of your physical volumes, noting their current location, volume groups, and PVIDs.
- Document all volume groups, including quorum settings and major numbers.
- Document */etc/filesystems*.
- Collect LV information using the command:

```
lsvg | lsvg -i -l
```

- Document any customized device settings – for example, if you have a switch, then issue the following command to gather information about it:

```
lsattr -E -l css0
```

- Document your network options settings using the command:

```
no -a
```

- Document your MTU settings using the command:

```
netstat -I
```

- Document your network adapter settings.
- Document your name service settings (the files you need include */etc/hosts*, */etc/resolv.conf*, and */etc/netsvc.conf*).
- Document your user information – the files needed for this include */etc/passwd*, */etc/group*, */etc/security/passwd*, */etc/security/group*, and */etc/security/user*.

- Document the attributes of any devices you've customized. This can be done, for instance, with the command:

```
lsattr -El rmt2
```

- Gather SSA loop information (I recommend you carry out a physical inspection and label cables at both ends).
- Document your external SCSI disk configuration (again, I recommend you carry out a physical inspection and label the cables at both ends).
- Collect your preferred performance monitoring information (using **sar**, **vmstat**, **iostat**, etc.) for at least two weeks before you migrate to the new node. This data can be useful if you encounter performance problems on the new node.
- Document your TTY settings. If you use 128-port adapters, you'll find some helpful information on reconfiguring ports after a **mksysb** restore at the following URL:

```
http://techsupport.services.ibm.com/support/rs6000.support/  
➤ databases
```

(Once at this URL, select *Hints and Tips*, followed by *AIX TechDocs*, then use the filters to select *TTY* and the article *Restoring 128-port devices after mksysb system restore*.)

- Plan your back-ups for *rootvg* and all other non-*rootvg*-based applications.
- Find and document all programs that use *uname* settings for licenses and the like.
- If you choose to change the hostname and/or the alias of your enterprise network adapter (if one exists), you need to track down and change all references to them. Examples of where such references may exist include Oracle's *listener.ora* and *tnsnames.ora* files, Sybase's */usr/sybxxx/interfaces* file, and so on.
- Make sure you have updated device drivers for your OS level available for installation on the receiving node immediately following the install. This is especially important when going

from MCA to PCI nodes. Drivers are available for download from <http://www.rs6000.ibm.com/support/micro/>.

## MIGRATION STEPS

- 1 At the donor node, stop and/or kill all non-root processes and all non-essential root processes. This is especially important for applications that happen to run from */usr* or other *rootvg*-based filesystems.
- 2 On the donor node:
  - Back up your non-*rootvg* data.
  - Unmount all non-*rootvg* filesystems.
  - ‘Vary off’ all non-*rootvg* volume groups.
- 3 Create a clean **mksysb** back-up of the donor node using the following steps:
  - i Run the following commands (in sequence):

```
rm /etc/niminfo*  
odmget -q name=sp CuAt > /tmp/odmsave  
odmdelete -q name=sp -o CuAt  
mkdir /tmp/ssp; mv /etc/ssp/*name* /tmp/ssp
```
  - ii Make sure that the */etc/hosts* file has all nodes’ reliable hostnames.
  - iii Establish whether there are any */etc/inittab* entries that you would like to comment out, thus preventing them from running on the receiving node.
  - iv Create the **mksysb** image using the following two commands (again, issue them in sequence):

```
cp /tmp/ssp/*name* /etc/ssp  
odmadd /tmp/odmsave
```
- 4 Power off the donor node.

- 5 Run **install** (**spbootins -r install -l node#, nodecond**, etc) on the receiving node. Note that it's a good idea to try out this step beforehand to iron out any stumbling blocks and to get the timing down.
- 6 If necessary, change any enterprise network IP addresses to the old node's settings and then, if your network folks deem it necessary to have the addresses, flush the **arp** tables in the appropriate routers.
- 7 Shut down the receiving node and recable any disk subsystems and other devices, as appropriate. Power on the node.
- 8 Perform all device customizations documented previously.
- 9 Import and 'vary on' your volume groups.
- 10 Reconcile */etc/inittab* using the data gathered in step 3iii.
- 11 You may want to reboot at this point, depending on the device attributes you've changed (this will also demonstrate whether there are any issues to do with rebooting).

#### FOLLOW-UP ITEMS

There are a number of miscellaneous follow-up items that I have on my list – your list will probably be a variation of this one.

- Check **sendmail** configuration.
- Re-register the node with the Service Director.
- Update */etc/motd* (ours contain node-specific information).
- If you changed the IP addresses of the node's adapters, update them in the SDR. This can be done using **smitty node\_data**, then *Additional Adapter Information*.
- Re-register the client with ADSM.

---

*David Miller*  
*Database Architect*  
*Baystate Health Systems (USA)*

© Xephon 2000

---

## A utility to terminate and clean up user processes

When an application hangs, the normal solution is to kill all its processes. With some applications (Lotus Notes, for example), this isn't sufficient – it's necessary to remove from memory all the IPC resources (semaphores, shared memory, and queues) allocated by the application. I have coded a Perl utility (**kill\_user**) that kills all user processes and removes all IPC resources allocated to the user whose name is passed as an argument to the program when it's invoked. **kill\_user** starts by sending a *SIGHUP* signal to each user process; after five seconds, a *SIGKILL* is issued to any remaining processes. Thereafter, a *SIGKILL* signal is sent every second to each process that's still alive until all the processes terminate (a maximum of eight further signals being sent). The IPC resources allocated to the user are then removed using AIX's **ipcrm** command.

Note the use of the continuation character, '›', in the code below to indicate a formatting line break that's not present in either the original source code or the code that can be downloaded from Xephon's Web site (<http://www.xephon.com/aixupdate.html>).

### KILL\_USER

```
#!/usr/local/bin/perl
use strict;

main:
{
    exit unless defined($ARGV[0]);
    exit unless $ARGV[0] ne "root";

    kill_process($ARGV[0]);
    kill_ipc($ARGV[0]);
}

sub kill_process
{
    my ($user) = @_ ;
    my $loop=0;
    my $kill=1;
    my $signal=-1;
    my $sleep=5;
```



```

while(($loop < 10) and ($kill eq 1))
{
    $loop++;
    $kill=0;
    open(USERPS,"/usr/bin/ps -o pid,args -u $user | /usr/bin/sort
    > -n +1 |");

    while(<USERPS>)
    {
        if (/PID COMMAND/) {next;};
        if (/^\s*$/) {next;};

        if (m%^\s*(\d+)\s+(.+)%)
        {
            {
                system("/usr/bin/kill $signal $1");
                $kill=1;
            }
        }

        close(USERPS);
        if ($kill eq 1) { sleep($sleep);}
        $sleep=1;
        $signal=-9;
    }
}

sub kill_ipc
{
    my ($user) = @_ ;

    open(IPCS,"/usr/bin/ipcs |");

    while(<IPCS>)
    {
        {
            if (m%(\\S)\s*(\d+)\s+\\S+\\s+\\S+\\s+(\\S+)\s+\\S+\\s*%)
            {
                next if $3 ne $user;
                if ($1 eq 's') {system("/usr/bin/ipcrm -s $2");}
                if ($1 eq 'm') {system("/usr/bin/ipcrm -m $2");}
                if ($1 eq 'q') {system("/usr/bin/ipcrm -q $2");}
            }
        }
    }

    close(IPCS);
}

```

---

*Fernando Carvalho*  
*System Engineer (Portugal)*

© Xephon 2000

---

## Maintain network details – part 2

This month's instalment concludes this article on a TCP/IP management utility for AIX. The first part appeared in last month's issue of *AIX Update*. Note the use of the continuation character, '›', in the code below to indicate a formatting line break that's not present in the original source code.

### MND.SH (CONTINUED)

```
#####
# Name      : DisplaySpecificSubsystemStatus
#
# Overview  : Displays full details of a specific subsystem.
#
# Notes     : 1. A validation is performed on client address, so '*.*'
#             is displayed as 'UNKNOWN'
#####
DisplaySpecificSubsystemStatus ()
{
  DisplayMessage I "${WORKING}" "N"
  while true
  do
    clear
    echo "Enter subsystem name (l = list of values):\c"
    read SUBSYSTEM
    case ${SUBSYSTEM} in
      l|L ) DisplayListOfValues "SUBSYSTEM" ;
            if [ "${SELECTED_VALUE}" = "" ]
            then
              : ;
            else
              SUBSYSTEM="${SELECTED_VALUE}" ;
              break ;
            fi ;;
      "" ) DisplayMessage E "${INVALID_ENTRY}" ;;
      * ) break ;;
    esac
  done
  DATETIME=`date "+%d/%m/%Y at %H:%M:%S"`
  HEADER="Subsystem details ${DATETIME}"
  FormatUnderscores "${HEADER}"
  echo "          $HEADER" > ${REPORT_FILE}
  echo "          $UNDERScore" >> ${REPORT_FILE}
  lssrc -l -s ${SUBSYSTEM} >> ${REPORT_FILE}
  view ${REPORT_FILE}
}
```

```

PrintFile "${REPORT_FILE}"
}
#####
# Name      : ProcessOption
#
# Overview  : Processes a menu option.
#
# Input     : Menu name
#####
ProcessOption()
{
MENU_NAME="$1"
if [ "${MENU_NAME}" = "MND" ]
then
    case $OPTION in
        5) DisplayOSI7LayerModel ;;
        10) DisplaySpecificNetworkOption ;;
        15) DisplayAllNetworkOptions ;;
        20) SetSpecificNetworkOption ;;
        25) DisplayMNSMenu ;
            break ;;
        30) DisplayFileConfigurationMenu ;;
        99) clear; ProcessExit $SEC ;;
        *) DisplayMessage E "${INVALID_ENTRY}"
    esac
elif [ "${MENU_NAME}" = "TCPIP" ]
then
    case $OPTION in
        5) DisplayAllTCPIPConnections ;;
        10) DisplayRoutingTableInformation ;;
        15) ShowRouteForRemoteHost ;;
        20) ListPortsInUse ;;
        25) DisplayNotesOnIPAddresses ;;
        30) DisplaySpecificSubsystemStatus ;;
        90) break;
            DisplayMNDMenu ;;
        99) clear; ProcessExit $SEC ;;
        *) DisplayMessage E "${INVALID_ENTRY}"
    esac
elif [ "${MENU_NAME}" = "FCONFIG" ]
then
    case $OPTION in
        5) ViewEditConfigurationFile "/etc/hosts" ;;
        10) ViewEditConfigurationFile "/etc/services" ;;
        15) ViewEditConfigurationFile "/etc/networks" ;;
        20) ViewEditConfigurationFile "/etc/gateways" ;;
        25) ViewEditConfigurationFile "/etc/inetd.conf" ;;
        30) ViewEditConfigurationFile "/etc/gated.conf" ;;
        35) ViewEditConfigurationFile "/etc/resolv.conf" ;;
        40) ViewEditConfigurationFile "/etc/protocols" ;;
    esac

```

```

45) ViewEditConfigurationFile "/etc/rc.tcpip" ;;
50) ViewEditConfigurationFile "/etc/rc.net" ;;
55) ViewEditConfigurationFile "/etc/hosts.equiv" ;;
60) ViewEditRemoteHostFile ;;
90) break;
    DisplayMNDMenu ;;
99) clear; ProcessExit $SEC;;
    *) DisplayMessage E "${INVALID_ENTRY}"
esac
elif [ "${MENU_NAME}" = "MNS" ]
then
    case $OPTION in
        5) DisplayNetworkSubsystemGroups ;;
        10) DisplayTCPIPMenu ;
            break ;;
        90) break ;
            DisplayMNDMenu ;;
        99) clear; ProcessExit $SEC ;;
        * ) DisplayMessage E "${INVALID_ENTRY}"
    esac
fi
}
#####
# Name      : DisplayOSI7LayerModel
#
# Overview : Displays the ISO seven-layer model
#####
DisplayOSI7LayerModel ()
{
    clear
    echo "      ${RVON} OSI seven-layer model ${RVOFF}"
    echo ""
    echo "      ${RVON}7 Application Layer${RVOFF}"
    echo " The Application layer consists of application programs"
    echo " and the user interface. Features in this layer exchange"
    echo " meaningful information with the user."
    echo ""
    echo "      ${RVON}6 Presentation Layer${RVOFF}"
    echo " The Presentation layer defines the representation of data,"
    echo " so that data is exchanged in a consistent format."
    echo ""
    echo "      ${RVON}5 Session Layer${RVOFF}"
    echo " The Session layer sets up and terminates communications"
    echo " on the network. It also manages the dialogue between"
    echo " users and systems. It is at this level that user and"
    echo " machine names are interpreted."
    echo ""
    echo " Press any key to continue..."
    read DUMMY
    echo "      ${RVON}4 Transport Layer${RVOFF} "

```

```

echo " The Transport layer controls the quality and reliability"
echo " of the data transmission. Packets are sequenced and"
echo " acknowledged in this layer. An example of a layer 4"
echo " protocol is TCP."
echo ""
echo "      ${RVON}3 Network Layer${RVOFF}"
echo " The Network layer routes data through the network. It"
echo " allows any properly configured pair of nodes on an"
echo " internetwork to communicate. The network layer calculates"
echo " routes and controls congestion. An example of a layer 3"
echo " protocol is NetWare's IPX."
echo ""
echo "      ${RVON}2 Data Link Layer${RVOFF}"
echo " The Data Link layer packages and unpackages data for"
echo " transmission across a single link. It deals with data"
echo " corruption (through checksums) and coordination of"
echo " shared media. An example of layer 2 protocol is Ethernet."
echo ""
echo " Press any key to continue..."
read DUMMY
echo "      ${RVON}1 Physical Layer${RVOFF}"
echo " The Physical layer establishes a physical connection"
echo " between a computer and the network. It also controls the"
echo " transmission of information and specifies the mechanical"
echo " and electrical characteristics of the protocol in terms of"
echo " connector size, pin assignments, and voltage levels. An"
echo " example of a layer 1 protocol is RS-232."
echo ""
echo " Press any key to continue..."
read DUMMY
}
#####
# Name      : DisplaySpecificNetworkOption
#
# Overview : Displays value of specific network option.
#####
DisplaySpecificNetworkOption ( )
{
DATETIME=`date "+%d/%m/%Y at %H:%M:%S"`
HEADER="Report on specificnetwork option value on ${DATETIME}"
FormatUnderscores "${HEADER}"
echo "      $HEADER" > ${REPORT_FILE}
echo "      $UNDERSCORE" >> ${REPORT_FILE}
while true
do
clear
echo "Enter network option (1 = list of values):\c"
read NO
case $NO in
1|L ) DisplayListOfValues "NO" ;

```

```

        if [ "${SELECTED_VALUE}" = "" ]
        then
            : ;
        else
            NO="${SELECTED_VALUE}" ;
            break ;
        fi ;;
    "" ) DisplayMessage E "${INVALID_ENTRY}" ;;
    * ) break ;;
esac
done
COMMAND="no -o ${NO}"
$COMMAND > ${ERROR_FILE} 2>&1
if [ $? -ne 0 ]
then
    DisplayMessage E "${COMMAND_FAILURE}"
    OS_ERM=`cat ${ERROR_FILE}`
    DisplayMessage E "${OS_ERROR}"
    return $FALSE
else
    cat ${ERROR_FILE} >> ${REPORT_FILE}
fi
view ${REPORT_FILE}
}
#####
# Name      : ValidNetworkOption
#
# Overview  : Validates a network option.
#
# Input     : String (containing network option)
#
# Returns   : $TRUE or $FALSE
#####
ValidNetworkOption ()
{
    PARAM="$1"
    no -a | cut -d'=' -f1 > ${TEMP_FILE}
    if grep "${PARAM}" ${TEMP_FILE} > /dev/null 2>&1
    then
        return $TRUE
    else
        return $FALSE
    fi
}
#####
# Name      : GetCurrentNetworkOptionValue
#
# Overview  : Gets the current value of a specific network option in
#            the global variable $CNOV.
#

```

```

# Input      : String containing network option.
#####
GetCurrentNetworkOptionValue ()
{
PARAM="$1"
no -a > ${TEMP_FILE}
CNOV=`grep "${PARAM}" ${TEMP_FILE} | cut -d'=' -f2`
}
#####
# Name      : SetSpecificNetworkOption
#
# Overview : Sets the value of a specific network option.
#####
SetSpecificNetworkOption ()
{
while true
do
clear
echo "Enter network option (l=list of values):\c"
read NO
case $NO in
l|L ) DisplayListOfValues "NO" ;
if [ "${SELECTED_VALUE}" = "" ]
then
: ;
else
NO="${SELECTED_VALUE}" ;
break ;
fi ;;
"" ) DisplayMessage E "${INVALID_ENTRY}" ;;
* ) if ! ValidNetworkOption "${NO}"
then
DisplayMessage E "${INVALID_NO}" ;
else
break ;
fi ;;
esac
done
GetCurrentNetworkOptionValue "${NO}"
while true
do
clear
echo "Enter new value of ${NO} (current value is ${CNOV}):\c"
read NNOV
case $NNOV in
"" ) DisplayMessage E "${INVALID_ENTRY}" ;;
* ) break ;;
esac
done
COMMAND="no -o ${NO}=${NNOV}"

```

```

$COMMAND > ${ERROR_FILE} 2>&1
if [ $? -ne 0 ]
then
    DisplayMessage E "${COMMAND_FAILURE}"
    OS_ERM=`cat ${ERROR_FILE}`
    DisplayMessage E "${OS_ERROR}"
    return $FALSE
else
    DisplayMessage I "${COMMAND_SUCCEEDED}"
fi
}
#####
# Name      : DisplayAllNetworkOptions
#
# Overview : Displays values of all network options
#####
DisplayAllNetworkOptions ()
{
    DATETIME=`date "+%d/%m/%Y at %H:%M:%S"`
    HEADER="Report on network options values on ${DATETIME}"
    FormatUnderscores "${HEADER}"
    echo "      $HEADER" > ${TEMP_FILE}
    echo "      $UNDERSCORE" >> ${TEMP_FILE}
    no -a >> ${TEMP_FILE}
    echo "" >> ${TEMP_FILE}
    echo "Option Glossary " >> ${TEMP_FILE}
    echo "===== " >> ${TEMP_FILE}
    echo "" >> ${TEMP_FILE}
    cat <<! >> ${TEMP_FILE}
    arpqsiz
    =====
    Specifies the maximum number of packets to queue while waiting for
    ARP responses. The default value is one. This attribute is supported
    by Ethernet 802.3, Token Ring, and FDDI. The arpqsiz value is
    increased to a minimum value of five when path MTU discovery is
    enabled. Note that the value does not decrease automatically if path
    MTU discovery is disabled. This attribute applies to AIX Versions
    4.1.5, 4.2.1, and later. arpqsiz is a runtime attribute.

    arptab_bsiz
    =====
    Specifies Address Resolution Protocol (ARP) table bucket size. The
    default value is seven. arptab_bsiz is a load-time attribute.

    arptab_nb
    =====
    Specifies the number of ARP table buckets. The default value is 25.
    arptab_nb is a load-time attribute.

    arpt_killc
    =====

```



Specifies the number of minutes before an ARP entry is completely deleted. The default is 20 minutes. `arpt_killc` is a run-time attribute.

#### bcastping

Allows responses to ICMP echo packets to the broadcast address. '0' turns this feature off and '1' turns it on. The default is 0. `bcastping` is a run-time attribute.

#### directed\_broadcast

Specifies whether to allow a directed broadcast to a gateway. A value of '1' allows packets to be directed to a gateway and broadcast on a network on the other side of the gateway. `directed_broadcast` is a run-time attribute.

#### icmpaddressmask

Specifies whether the system responds to an ICMP address mask request. If the default value '0' is set, the network silently ignores any ICMP address mask request that it receives. `icmpaddressmask` is a run-time attribute.

#### ie5\_old\_multicast\_mapping

A value of '1' specifies that IP multicasts on token ring should be mapped to the broadcast address rather than a functional address. The default value is 0. `ie5_old_multicast_mapping` is a run-time attribute.

#### ifsize

Specifies the maximum number of network interface structures per interface. The default number is 8. `ifsize` is a load-time attribute.

#### ipforwarding

Specifies whether the kernel should forward packets. The default value '0' prevents forwarding of IP packets when they are not for the local system. The value '1' enables IP forwarding. `ipforwarding` is a run-time attribute.

#### ipfragttl

Specifies the time-to-live of IP fragments. The default value is 60 half-seconds. `ipfragttl` is a run-time attribute.

#### ipqmaxlen

Specifies the number of packets received that can be queued on the IP protocol's input queue. `ipqmaxlen` is a load-time attribute.

### ipignoreredirects

Specifies whether to process redirects that are received. The default value of '0' indicates that redirects are to be processed, while the value '1' indicates they are to be ignored. This option applies only to AIX Version 4.2.1 or later. `ipignoreredirects` is a run-time attribute.

### ipsendredirects

Specifies whether the kernel should send redirect signals. '1' indicates they are to be sent, and '0' indicates they are not to be sent. `ipsendredirects` is a run-time attribute.

### ipsrcrouteforward

Specifies whether the system forwards source-routed packets. The default value '1' allows the forwarding of source-routed packets, while the value '0' causes them to be discarded. This attribute applies only to AIX Version 4.2.1 or later.

### ipsrcrouterrecv

Specifies whether the system accepts source-routed packets. The default value '0' causes all source-routed packets destined for this system to be discarded. A value '1' allows source-routed packets to be received. This attribute applies only to AIX Version 4.2.1 or later.

### ipsrcroutesend

Specifies whether applications can send source-routed packets. The default value '1' allows source-routed packets to be sent. A value '0' causes `setsockopt()` to return an error if an application attempts to set the source routing option and removes any source routing options from outgoing packets. This attribute applies only to AIX Version 4.2.1 or later.

### maxttl

Specifies the time-to-live of RIP packets. The default is 255 seconds. `ipsendredirects` is a run-time attribute.

### net\_malloc\_police

Specifies the size of the `net_malloc/net_free` trace buffer. If the value of this variable is non-zero, all `net_malloc` and `net_frees` will be traced in a kernel buffer by system trace hook `HKWD_NET_MALLOC`. Additional error checking is also be enabled. This includes checks for freeing a free buffer, alignment, and buffer overwrite. The default value is '0' (policing off). Values of `net_malloc_police` larger than 1024 result in the allocation of that many items in the kernel buffer

for tracing. net\_malloc\_police is a run-time attribute.

#### nonlocsrcroute

=====

Tells the Internet Protocol that strictly source-routed packets may be addressed to hosts outside the local network. A default value of '0' disallows addressing to outside hosts. The value '1' allows packets to be addressed to outside hosts. Loosely source-routed packets are not affected by this attribute. nonlocsrcroute is a run-time attribute.

#### pmtu\_default\_age

=====

Specifies the default delay (in minutes) before the path MTU for UDP paths is checked for a lower value. A value of zero allows no ageing. The default value is 10 minutes. The pmtu\_default\_age value can be overridden by UDP applications. This attribute applies only to AIX Version 4.2.1 or later. pmtu\_default\_age is a run-time attribute.

#### pmtu\_rediscover\_interval

=====

Specifies the default delay (in minutes) before the path MTU values of UDP and TCP paths are checked for a higher value. A value of '0' allows no path MTU rediscovery. The default value is 30 minutes. This attribute applies only to AIX Version 4.2.1 or later. pmtu\_rediscover\_interval is a run-time attribute.

#### rfc1122addrchk

=====

Performs address validation as specified by RFC 1122 (Requirements for Internet Hosts Communication Layers). The default value of '0' indicates no address validation is performed, and '1' indicates that address validation is performed. rfc1122addrchk is a run-time attribute.

#### rfc1323

=====

Enables TCP enhancements as specified by RFC 1323 (TCP Extensions for High Performance). The default value '0' disables the RFC enhancements on a system-wide basis, while a '1' specifies that all TCP connections shall attempt to negotiate RFC enhancements. The SOCKETS application can override the default behaviour on individual TCP connections, using the setsockopt subroutine. rfc1323 is a run-time attribute.

#### route\_expire

=====

Specifies whether the route expires. The value '0' allows no route expiration, which is the default. Negative values are not allowed. This attribute applies only to AIX Version 4.2.1 or later. route\_expire is a run-time attribute.

#### rto\_length

Specifies the TCP 'Retransmit Time Out' length used in calculating both the factors and the maximum number of allowable re-transmits used in TCP data segment re-transmits. rto\_length is the total number of time segments; the default is 13. rto\_length is a load-time attribute.

#### rto\_limit

Specifies the TCP 'Retransmit Time Out' limit used in calculating both factors and the maximum allowable re-transmits used in TCP data segment re-transmits. rto\_limit is the number of time segments from rto\_low to rto\_high (the default is 7). rto\_limit is a load-time attribute.

#### rto\_low

Specifies the TCP 'Retransmit Time Out' low value used in calculating both factors and the maximum allowable re-transmits used in TCP data segment re-transmits. rto\_low is the low factor, the default being 1. rto\_low is a load-time attribute.

#### rto\_high

Specifies the TCP 'Retransmit Time Out' high value used in calculating factors and the maximum allowable re-transmits used in TCP data segment re-transmits. rto\_high is the high factor, the default being 64. rto\_high is a load-time attribute.

#### sb\_max

Specifies the maximum buffer size allowed for a socket. The default is 65,536 bytes. sb\_max is a run-time attribute.

#### somaxconn

Specifies the maximum listen backlog. The default is 1,024 bytes. somaxconn is a run-time attribute. This attribute applies only to AIX Versions 4.1.5, 4.2 or later.

#### subnetsarelocal

Determines whether a packet address is on the local network. This attribute is used by the subroutine in\_localaddress. The default value '1' specifies that addresses that match the local network mask are local. If the value is '0', only addresses matching the local subnetwork are local. subnetsarelocal is a run-time attribute.

#### tcp\_keepidle

Specifies the length of time to keep the connection active, measured in

half-seconds. The default is 14,400 half seconds (7,200 seconds or two hours). `tcp_keepidle` is a run-time attribute.

#### `tcp_keepinit`

---

Sets the initial time-out value of a TCP connection. This value is defined in half-second units, the default being 150 or 75 seconds. This attribute can be changed to any value using the `-o` flag. `tcp_keepinit` is a run-time attribute.

#### `tcp_keepintvl`

---

Specifies the interval (in half seconds) between packets being sent to validate the connection. The default is 150 half-seconds (75 seconds). `tcp_keepintvl` is a run-time attribute.

#### `tcp_mssdflt`

---

The default maximum segment size used in communicating with remote networks. For AIX Version 4.2.1 or later, `tcp_mssdflt` is used only if path MTU discovery is not enabled or path MTU discovery fails to discover a path MTU. `tcp_mssdflt` is a run-time attribute. The default value is 512.

#### `tcp_ndebug`

---

Specifies the number of `tcp_debug` structures. The default is 100. `tcp_ndebug` is a run-time attribute.

#### `tcp_pmtu_discover`

---

Enables or disables path MTU discovery for TCP applications. A value of '0' disables path MTU discovery and a value of '1' enables it. The default value is '0'. This attribute applies only to AIX Version 4.2.1 or later. `tcp_pmtu_discover` is a run-time attribute.

#### `tcp_recvspace`

---

Specifies the system default socket buffer size for receiving data. This affects the window size used by TCP. Setting the socket buffer size to 16 KB (16,384 bytes) improves performance over standard Ethernet and Token Ring networks. The default value is 4,096. However, a value of 16,384 is set automatically by either `rc.net` or `rc.bsdnet`, if Berkeley-style configuration is used. Low bandwidth networks, such as Serial Line Internet Protocol (SLIP), and high bandwidth networks, such as Serial Optical Link, have different optimum buffer sizes. The optimum size is the product of the media bandwidth and the average packet round-trip time. The `tcp_recvspace` attribute must specify a socket buffer size less than or equal to the setting of the `sb_max` attribute. `tcp_recvspace` is a run-time attribute.

## tcp\_sendspace

---

Specifies the system default socket buffer size for sending data. This affects the window size used by TCP. Setting the socket buffer size to 16 KB (16,384 bytes) improves performance over standard Ethernet and Token Ring networks. The default value is 4,096. However, a value of 16,384 is set automatically by either rc.net or rc.bsdnet, if Berkeley-style configuration is used. Low bandwidth networks, such as Serial Line Internet Protocol (SLIP), and high bandwidth networks, such as Serial Optical Link, have different optimum buffer sizes. The optimum buffer size is the product of the medium bandwidth and the average packet round-trip time. The tcp\_sendspace attribute must specify a socket buffer size less than or equal to the setting of sb\_max. tcp\_sendspace is a run-time attribute.

## tcp\_ttl

---

Specifies the time-to-live for TCP packets. The default is 60 'ticks' (a tick is one hundredth of a minute). tcp\_ttl is a run-time attribute.

## thewall

---

Specifies the maximum memory, in kilobytes, that is allocated to the memory pool. The default value is 1/8 of real memory or 65,536 kilobytes (64 megabytes), whichever is smaller. thewall is a run-time attribute.

## udp\_pmtu\_discover

---

Enables or disables path MTU discovery for UDP applications. UDP applications must be specifically written to utilize path MTU discovery. A value of '0' disables the feature, while a value of '1' enables it. The default value is '0'. This attribute applies only to AIX Version 4.2.1 or later. udp\_pmtu\_discover is a run-time attribute.

## udp\_recvspace

---

Specifies the system default socket buffer size for receiving UDP data. The default is 41,600 bytes. The udp\_recvspace attribute must specify a socket buffer size less than or equal to the setting of sb\_max. udp\_recvspace is a run-time attribute.

## udp\_sendspace

---

Specifies the system default socket buffer size for sending UDP data. The default is 9,216 bytes. The udp\_sendspace attribute must specify a socket buffer size less than or equal to the setting of sb\_max. udp\_sendspace is a run-time attribute.

## udp\_tta

---

Specifies the time-to-live for UDP packets. The default is 30 seconds. `udp_ttl` is a run-time attribute.

#### `udpcksum`

=====

Allows UDP checksum to be turned on or off - '0' turns it off and '1' turns it on (the default is '1'). `udpcksum` is a run-time attribute. Note that, if you use the `tcp_recvspac`, `tcp_sendspace`, `udp_recvspac`, or `udp_sendspace` to specify a socket buffer size larger than the default value of `sb_max`, then you must set the `sb_max` to an equal or greater value. Otherwise, the socket system call returns an `ENOBUFS` error message when an application tries to create a socket.

### Streams Tuneable Attributes

The following Streams Tuneable attributes apply only to AIX Version 4.2 or later.

#### `lowthresh`

=====

Specifies the maximum number of bytes that can be allocated using the `alloca()` call with `BPRI_LO` priority. When the total amount of memory allocated by `net_malloc()` reaches this threshold, the `alloca()` request with `BPRI_LO` priority returns '0'. `lowthresh` represents a percentage of `thewall`, and you can set its value from '0' to '100'. This is a run-time attribute and the default value is set to 90 (90% of `thewall`).

#### `medthresh`

=====

Specifies the maximum number of bytes that can be allocated using the `alloca()` call with `BPRI_MED` priority. When the total amount of memory allocated by `net_malloc()` reaches this threshold, the `alloca()` request with `BPRI_MED` priority returns '0'. `medthresh` represents a percentage of `thewall`, and you can set its value from '0' to '100'. This is a run-time attribute and the default value is set to 95 (95% of `thewall`).

#### `nstrpush`

=====

Specifies the maximum number of modules that you can push onto a single stream (the minimum value is '8'). This is a load-time attribute and the default value is set to '8'.

#### `psebufcalls`

=====

Specifies the maximum number of `bufcalls` to be allocated by the Streams subsystem. The Stream subsystem allocates a certain number of `bufcall` structures at initialization. This allows users to register requests to `bufcall()` even when `alloca()` fails. You are not allowed to lower this value until the system reboots, at which time it returns to its default value. This is a run-time attribute, and its default value is '20'.

## pseintrstack

Specifies the maximum size of the interrupt stack allowed by Streams when running in OFFLEVEL. Sometimes, when a process running in a level other than INTBASE enters a Stream, it encounters a stack overflow problem if the interrupt stack is too small. Setting this attribute to an appropriate value reduces the chances of a stack overflow. This is a load-time attribute and its default value is 0x3000.

## psetimers

Specifies the maximum number of timers that Streams can allocate. The Streams subsystem allocates a certain number of timer structures at initialization so that the Streams driver or module can register timeout() calls. You are not allowed to lower this value until the system reboots, at which time it returns to its default value. This is a run-time attribute and the default value is '20'.

## strctlsz

Specifies the maximum number of bytes that a single system call can pass to a Stream to place in the control part of a message (in an M\_PROTO or M\_PCPROTO block). A putmsg() call with a control part larger than this will fail with ERANGE. This is a run-time attribute and the default value is '1024'.

## strmsgsz

Specifies the maximum number of bytes that a single system call can pass to a Stream to place in the data part of a message (in M\_DATA blocks). Any write() call that exceeds this is broken into multiple messages. A putmsg() call with a data part larger than this fails with ERANGE. This is a run-time attribute and its default value is '1024'.

## strthresh

Specifies the maximum number of bytes Streams are normally allowed to allocate. When this threshold is breached, only users with appropriate privileges may open Streams, push modules, or write to Streams devices, otherwise calls return an ENOSR. The threshold applies only to the output side and does not affect data coming into the system (eg the console continues to work properly). A value of zero means that there is no threshold.

The strthresh attribute represents a percentage of the thewall attribute and its value ranges from '0' to '100'. The thewall attribute indicates the maximum number of bytes that can be allocated by Streams and Sockets using the net\_malloc() call. When you change thewall attribute, the threshold gets updated accordingly.



## strturncnt

=====

This attribute specifies the maximum number of requests for module-level or elsewhere-level Streams synchronization that are to be handled by the thread that's currently running. Module-level synchronization works in such a way that only one thread can run in the module at any time and all other threads trying to acquire the same module will enqueue their requests and leave. After the thread that's currently running completes its work, it dequeues previously enqueued requests, one by one, and invokes them. If many requests are enqueued on the list, then the thread that's currently running, which has to serve everyone, spends most of the time serving others and not itself. To avoid this, the thread serves only strturncnt number of threads; once this threshold is breached a separate kernel thread wakes up and invokes all pending requests. This is a run-time attribute and its default value is '15'.

```
!  
chmod 444 ${TEMP_FILE}  
view ${TEMP_FILE}  
PrintFile ${TEMP_FILE}  
}  
#####  
# Name      : DisplayNETFMenu  
#  
# Overview : Displays the menu for network file management  
#####  
DisplayNETFMenu ( )  
{  
    clear  
    echo ""  
    echo "      #####"  
    echo "      #                                     #"  
    echo "      #      ${NETF}      #"   
    echo "      #                                     #"  
    echo "      #          1. View  /etc/.rhosts  file  #"   
    echo "      #          5. Edit  /etc/.rhosts  file  #"   
    echo "      #         10. View  /etc/hosts    file  #"   
    echo "      #         15. Modify /etc/hosts    file  #"   
    echo "      #         20. View  /etc/services file  #"   
    echo "      #         25. Modify /etc/services file  #"   
    echo "      #         30. View  /etc/networks file  #"   
    echo "      #         35. Modify /etc/networks file  #"   
    echo "      #         40. View  /etc/sockcf   file  #"   
    echo "      #         45. Modify /etc/sockcf   file  #"   
    echo "      #          98. Previous menu          #"   
    echo "      #          99. Exit                   #"   
    echo "      #                                     #"  
    echo "      #####"  
    echo "      Please enter option (h for help)--->\c"
```

```

GetOption NETF
ProcessOption NETF
}
#####
# Name      : RootUser
#
# Overview  : The function checks whether the user is root.
#
# Returns   : TRUE if user is root, FALSE otherwise
#####
RootUser ()
{
USER=`id | cut -d'(' -f2 | cut -d')' -f1`
if [ "${USER}" = "root" ]
then
    return $TRUE
else
    return $FALSE
fi
}
#####
# Name      : ProcessExit
#
# Overview  : Removes any temporary files and makes a graceful exit.
#
# Input     : Exit Code
#####
ProcessExit ()
{
EXIT_CODE="$1"
rm -f /tmp/mnd*
clear
exit ${EXIT_CODE}
}
#####
# Name      : main
#
# Overview  : Implements the processing structure.
#####
main ()
{
InitializeVariables
trap "HandleInterrupt " $SIGINT $SIGTERM $SIGHUP $SIGTSTP
if ! RootUser
then
    DisplayMessage E "${ROOT_USER}"
    ProcessExit $FEC
fi
DisplayMNDMenu
}
main

```

## SAMPLE OUTPUT

Report on all active TCP/IP connections  
to server on 29/02/2000 at 22:34:20

```
=====
Local_Port  Service                Foreign_Address        State
1526        listener              137.213.40.55.35160   ESTABLISHED
35160       listener              137.213.40.55.1526   ESTABLISHED
23          telnet                137.213.43.121.3672  ESTABLISHED
1526        listener              137.213.40.7.2278    ESTABLISHED
1526        listener              137.213.40.7.2277    ESTABLISHED
1526        listener              137.213.40.47.2719   ESTABLISHED
1526        listener              137.213.10.103.1927  ESTABLISHED
32786       listener              137.213.40.55.199    CLOSE_WAIT
199         smux                  127.0.0.1.32772     ESTABLISHED
32772       smux                  127.0.0.1.199       ESTABLISHED
199         smux                  127.0.0.1.32770     ESTABLISHED
32770       smux                  127.0.0.1.199       ESTABLISHED
```

Total No of Connections = 12

Report on ports in use on 29/02/2000 at 22:34:55

```
=====
23
199
1526
32770
32772
32786
35160
```

Route details for orafins on 29/02/2000 at 22:37:06

```
=====
trying to get source for orafins
source should be 137.213.40.55
1 refmche2.uk34.bull.co.uk (137.213.40.42) 9 ms 4 ms 4 ms
2 refmhem2-s0.uk03.bull.co.uk (137.213.248.5) 16 ms 14 ms 14 ms
3 refmhem5.uk03.bull.co.uk (137.213.193.150) 16 ms 15 ms 17 ms
4 129.184.251.25 (129.184.251.25) 111 ms 107 ms 109 ms
5 129.184.251.18 (129.184.251.18) 182 ms 181 ms 182 ms
6 129.181.205.219 (129.181.205.219) 185 ms 189 ms 191 ms
7 orafins (129.181.200.250) 188 ms 202 ms 188 ms
```

---

*Arif Zaman*  
*DBA/System Administrator*  
*High-Tech Software Ltd (UK)*

© Xephon 2000

---

# AIX news

---

Singlepoint Systems has announced AlarmPoint Version 3.2 an event notification system with versions that send alerts via pagers, phone messages, and e-mail. The system now supports two-way pagers, allowing it to receive instructions via a menu selection on the pager. It also supports Symon electronic message display boards.

The new release includes clients for AIX, HP-UX, IRIX, Linux, Solaris, Tru64 Unix, and Windows 95, 98, NT 4, and 2000. Out now, the Phone Edition costs US\$10,000 while the Paging Edition costs US\$3,000.

*For further information contact:*  
Singlepoint Systems, 7094 Almaden Rd,  
San Jose, CA 95120, USA  
Tel: +1 408 262 0205  
Fax: +1 408 268 0531  
Web: <http://www.singlepointsys.com>

\* \* \*

Zero G has announced InstallAnywhere Version 3.5, a tool for creating multi-platform deployment software that now includes Java 2 V1.3 compatibility and Linux support. Available in Enterprise, Standard, and 'Now!' editions, the software builds multi-platform installation tools that are capable of installing and configuring software on virtually any client or server platform.

The installer is able to recognize the platform on which it's running and tailor its installation to the target system. This allows the creation of universal installers that can deploy software to multiple platforms via the

Internet or CD-ROM. The platforms supported include AIX, HP-UX, Solaris, Linux, Windows (all versions), Mac OS, and virtually any other Java-enabled client.

The software can also install a Java virtual machine as part of an installation and automatically build Web Installer HTML pages for deploying via intranets and the Internet.

Priced on a per-developer, per-system basis, Enterprise Edition costs US\$2,000, Standard Edition US\$1000, while the 'Now!' version is free.

*For further information contact:*  
Zero G Software Inc, 514 Bryant Street, San  
Francisco, California 94107, USA  
Tel: + 1 415 512 -7771  
Fax: +1 415 723 7244  
Web: <http://www.zerog.com>

\* \* \*

IBM, i2 Technologies, and Ariba have announced that Ariba Buyer 7.0 now interoperates with i2 Discovery Catalog and Content Exchange 1.5 (formerly of Aspect) using the cXML (commerce XML) specification. Ariba Dynamic Trade 2.0 interoperates with i2 Discovery Catalog and Content Exchange 1.5 to develop and deploy trading functions in e-markets. Ariba Buyer 7.0 now supports AIX, DB2 Universal Database, and MQSeries integration, while Ariba Dynamic Trade 2.0 now runs on AIX and DB2. Ariba Marketplace now supports Tivoli eMarketplace Manager and Marketplace 7.0 supports Websphere B2B Integrator.



**xephon**