NMON_Analyser User Guide for V6.0

Preface

NMON_Analyser is designed to complement NMON (**N**igel's **Mon**itor) in analysing and reporting performance problems; it produces graphs for virtually all sections of output created using the "spreadsheet output" mode of NMON as well as doing some additional analyses for ESS, EMC and FAStT subsystems. It will also work with files produced by topasout and with other tools that produce data in "NMON" format.

For Windows users: It is written using Excel 2016 and shoud work with Excel 2007 or later.   It may work also on Excel 2003 with the required Microsoft updates to support .xlsx files
https://support.microsoft.com/kb/924074?wa=wsignin1.0

For Mac users: It is tested using Excel for Mac 2016. It should work with earlier versions such as Excel 2011 but has not been recently tested.

Note #1: Recent Microsoft Excel for Mac updates (released in Jan 2019) have broken nmon Analyzer (prior versions) and users will receive overflow errors. This version is intended to work around these issues. As a result, some graphs may appear at the top of the sheet instead of the bottom.

Note #2: If you wish to collect very large nmon files (10+ MB) and attempt to use the NMON Analyser against them, then you should install 64-bit Microsoft Excel 2013 or 2016. Installing 64-bit Excel is critical or you will likely run out of memory and it will fail. You also need to have 8 GB RAM and preferably more.

NMON was originally written by Nigel Griffiths (nag@uk.ibm.com) and is now (since AIX 5.3 TL09 and AIX 6.1 TL02) part of topas. NMON_Analyser was originally written by Stephen Atkins with contributions from many people – including Ralf Schmidt-Dannert and Markus Fehling, both of IBM. Currently nmon Analyzer is maintained by Ron McCargar (mccargar@us.ibm.com.)

Support for NMON_Analyser is provided on a best efforts basis. Please direct questions to the User Forum (see below) rather than contact the author direct.

Links:
    NMON_Analyser home page
    https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Power%20Systems/page/nmon_analyser
    User Forum
    https://www.ibm.com/developerworks/community/forums/html/forum?id=11111111-0000-0000-0000-000000000749&ps=25

New in V6.0

- Work around added for overflow error on Mac OS when trying to chart graphs on several sheet like DISK_SUMM. Appears to be a bug with Excel on Mac OS only.

- On Mac OS the graphs were drawn too wide, made a change so by default they are drawn to consume a more reasonable amount of space similar to how it is done on Windows.

- Fixed graph titles/descriptions for DISKAVGRIO, DISKAVGWIO: they were wrong per AIX documentation here: https://www.ibm.com/support/knowledgecenter/en/ssw_aix_72/com.ibm.aix.prftools.io_stat.htm

New in V5.5

- Fixed an issue where the DISK_SUMM metrics were wrong if more than 150 disks existed in the data.

New in V5.4

- When running on MAC, showing progress during the run is much faster now. The SHOWMACPROGRESS option now defaults to YES. Note: Using SHOWMACPROGRESS = YES will slow execution somewhat (maybe 30-50% slower) but the benefit is you will clearly know that the process is not "hung".

- On MACs, fixed batch processing of multiple files related to use of : instead / for subdirectories.

- On Macs, fixed issues relating to merging of multiple files because of usage of : instead of /

New in V5.2.2

- When using the merge function its possible the filename generated may incorrectly contain a : colon causing the analyzer to fail.

New in V5.2.1

- Tested on Excel 2016:   Please note you may experience the following issue "texttocolumns method of range class failed" a possible work around:

"I installed Excel 2016 on my mac. I open excel , enable macros run the latest analyzer. It fails with "texttocolumns method of range class failed". Then i see in the background that the Sheet2 has no data in column A, i.e. data didn't get copied into Excel from the .nmon file. I put in a 1 into first cell of column A (could add any data really to avoid the error) and hit the button to continue. Then it fails like it should saying "no valid nmon data".

I shut down Excel and I came back in again later rerun the analyzer and I did not get the problem it seemed like a 1 time thing, it ran through it successfully. Tested it a few more times, it doesn't happen anymore. Uncertain what caused this to fail the 1st time but FYI"

New in V5.2

- Add Merge support when running on a Mac.

New in V5.12

- Fixed problems if running on Mac Excel 2016

New in V5.1

- If Merge = Yes then the 1st file's data would appear twice incorrectly.
- The Avg and Max CPU% on the SYS_SUMM column was being calculated incorrectly.
- The default save format used on a Mac may not have been set correctly in some cases.
- If importing multiple .nmon files that contain multiple "AAA,x86," records popup errors would be shown to the user.

New in V5.0.2

- Added support to run nmon Analyzer on a Mac with Excel 2011 or higher. Please note the following limitations: Pivot charts don't work, Merge option doesn't work.
- Fixed a problem when trying to import very large .nmon files it may fail.
- On the AAA sheet renamed the "elapsed" value given to "Analysis time" (the elapsed time it took Excel to analyze the nmon file)
- Enabling PIVOT charts would crash the analyzer on Excel 2010 and higher and has been fixed. (Analyser - PIVOT = YES)
- On the Analyser's Settings - DELIM value the user can now specify any value desired.
- In ProcessFile() added a check for EOF before doing the import to possibly avoid an error.
- Graphs that show Avg/Wavg/Max will now show Avg/Wavg on the Y1 axis and Min/Max as lines on the Y2 axis. Also if the Analyzer Output option SyncAxis is YES then the Y1 and Y2
- Several graphs were not sorted previously based on the Analyser's Settings - SORTDEFAULT value but are now.
- Added On Error Resume Next in a few spots in order to be better able to analyze huge data sets without failing prematurely.
- Added new options SHOWMACPROGRESS and SHOWLINUXCPUUTIL

New in V4.7

- Added a check for Power 8 and SMT mode 8. In previous versions no check exists which defaults the number of CPUs shown to 2 on the CPU_SUMM sheet.

New in V4.6

- If an INTERVAL range was specified the CPU_SUMM sheet would contain no data.
- If the SORTDEFAULT setting was set to YES then sheets like NETPACKET would incorrectly display a SortKey value on the graph.
- If the LPAR sheet didn't exist, then SYS_SUMM graph would not display CPU% as expected.
- On the SYS_SUMM sheet if the LPAR sheet exists, the CPU% column will be now what it was back with analyzer version 3.4a or Avg CPU% "=M4/LPAR!C2*100." and Max CPU% = "=M4/LPAR!C2*100.".
- If the nmon data contained ERROR tags, the resulting ERROR sheet would incorrectly list thousands of columns at the end containing 'N/A'

New in V4.3

- Added a new setting called SORTDEFAULT. This setting indicates if the 1st graph on 'default' sheets (ones not handled else where by the Analyzer) are sorted. **Note:** If Yes, this will also reorder the columns.
- Added support for the CPU_ALL, CPUnnn Steal% column that is being added to nmon for Linux.

New in V4.2

- On the LPAR sheet on the 1st graph the Unfolded VPs value was incorrectly removed.
- Support for checking for EMC/PowerPath disks was broken. The EMC graphs are available again.

- The EMC check problem caused the disk graphs showing Avg, Wavg, Max to not be sorted with the highest Avg + Wavg first.
- Many additional sheets now have their Avg, Wavg, Max graphs sorted with the highest Avg + Wavg first (FC*, etc)

New in V4.0

- The elapsed time value produced on AAA was incorrect previously. It did not include the time to open the file(s) and sort the data.
- Added a new option on the Analyser sheet called BIGDATA. Use this if you have large lines (up to 32K) or > 1048576 rows.
- The analyzer would fail on Excel 2007 if > 1048576 rows were encountered. Use BIGDATA=YES.
- The maximum allowed rows per .nmon file is now unlimited (based on available Excel/PC memory)
- The SORTINP setting has been removed. Data is now always sorted.
- On the SYS_SUMM sheet the CPU% values were wrong and have been changed to the sum of users% + Sys% for each line.
- On the SYS_SUMM sheet added the following additional fields from the LPAR tab (it it exists): Virtual CPUs, Other LPARs, Pool CPUs, Entitled, Weight
- On the LPAR tab the 1st and 3rd graphs showed incorrect values.
- Updated nmon analyzer to check for either hdiskpower, emcpower or just power when determining if EMC data is present.
- Analysis supports any number of columns up to the Excel maximum of 16684. On large data sets (> 10 MB) you most likely will require 64-bit Windows however and a minimum of 8 GB of RAM.
- NPIV header was incorrectly shown at the bottom of the sheet previously.
- NPIV and SEALICTRAFFIC sheets will no longer show any graphs (due to potential huge size.)

New in V3.5

- Added SCATTER option to include/exclude TOP CPU% by PID graph to avoid crash with Excel 2013

New in V3.4

- Improved support for Excel 2007 and 2010
- Support for 64-bit Windows
- Ability to specify time values instead of interval numbers
- Wildcard support in LIST
- Limit the overall number of CPU, PCPU and SCPU sheets generated
- Add graphs for the PCPU and SCPU sheets
- Correct handling of the DISKRXFER sheet
- Fix a bug with the MERGE option
- Automatically include CPU_SUMM, DISK_SUMM and SYS_SUMM in LIST

Table of Contents

Collecting data using NMON

Please make sure you have the latest versions of both NMON and NMON_Analyser before starting a new engagement. If you want automatic notification of a new release of the Analyser send a note to steve_atkins@uk.ibm.com and I'll add you to my distribution list. Nigel maintains a similar list for NMON.

You will need to have root privileges in order to get a complete set of data on the BBBP sheet. In order to collect data for the DISKBUSY sheets you need to make sure that iostat data collection is enabled:

> chdev -l sys0 -a iostat=true

For spreadsheet output mode (comma separated values) use the following flags when invoking nmon:

-f	spreadsheet output format [note: default -s300 -c288]
	Output file is <hostname>_**YYYYMMDD_HHMM.nmon**
-F <filename>	same as -f but user supplied filename
-c <number>	number of snapshots
-d	requests disk service and wait times (DISKSERV and DISKWAIT)
-i <percent>	Ignore processes using less than this amount of CPU when generating TOP section – useful for reducing data volumes
-g <filename>	file containing disk group definitions
-l <dpl>	number of hdisks per sheet - defaults to 150, maximum 250. See notes
-m <dir>	NMON changes to this directory before saving the file
-r <runname>	goes into spreadsheet file [default hostname]
-s <seconds>	interval between snap shots
-x	capacity planning (15 mins for 1 day = -fdt -s900 -c96)
-t	include top processes in the output
-T	as **–t** plus saves command line arguments in UARG section
-A	include data for async I/O (PROCAIO) sections
-D	prevents DISK sections being produced (useful when Disk Groups are being used because there are too many hdisks to process)
-E	stops ESS sections being produced (necessary when Disk Groups are being used because there are too many vpaths to process)
-J	prevents JFS sections being produced (prevents Excel errors when you have more than 255 filesystems)
-L	includes LARGEPAGE section
-N	include NFS sections
-S	include WLM sections with subclasses
-W	include WLM sections without subclasses
-Y	include SUMMARY section (very efficient alternative to **–t** if PID level data is not required)

> example: nmon_aix51 -F asterix.nmon -r Test1 -s6 -c12

Notes:
1. The **–f** (or **–F**) flag must appear first.

2. The value of the **-l** flag controls the number of hdisks per sheet on the DISK sheets and per line on the BBBD sheet.   There are two factors to consider when choosing this value. Excel has a limit of 256 columns per sheet; however, both NMON and NMON_Analyser use some columns, so the upper limit is really 250.   The second factor is that Excel VBA has an upper limit of 2048 bytes for input line length.   This particularly affects users of EMC systems that use long hdisk names (e.g. *hdiskpower123*).   The default of 150 is safe for such systems. Other users may set the value to 250 in order to reduce the number of output sheets.
3. Consider the value of the **-s** flag very carefully.   The shorter the interval between snapshots, the more variable the values for each resource will be. If you use an interval of 1 second, don't be surprised to see many of your disks hitting 100% busy for short periods. For normal monitoring, 10-minute intervals (-c 600) provide a good balance.
4. The graphs produced by NMON_Analyser look best when the number of snapshots (specified by the **-c** flag) is 300 or less.
5. The TOP section (produced by specifying the **-t** flag) can generate large amounts of output and the size of the output can grow exponentially if a large value is specified for the **-c** flag. If you want the TOP section then specify no more than 250 snapshots - ideally less.

Collecting data using topas (xmwlm)

You need AIX V5.3 TL5 Service pack 4 with APAR IY87993 or later.   Note that only the output using the −a flag can be analysed. In particular, cross-partition statistics cannot be analysed; if you wish to get a report for the entire machine, collect data from each LPAR separately and then use NMON_Consolidator to merge the data. Commands like the following are required to collect the data.

```
topasout -a /etc/perf/daily/xmwlm.yymmdd
:
cp /etc/perf/daily/xmwlm.yymmdd_01 hostname.ddmmyy.topasout.csv
```

Using NMON_Analyser

- FTP the input file to your PC – ideally using the ASCI or TEXT options to make sure that lines are terminated with the CRLF characters required by Windows applications.
- Open the NMON_Analyser spreadsheet and specify the options you want on the "Analyser" and "Settings" sheets (see below). Save the spreadsheet if you want to make these options your personal defaults.
- Click on the "Analyse nmon data" button and find/select the .nmon file(s) to be processed. You may select several files in the same directory. If you wish to process several files in different directories you may wish to consider using the "FILELIST" option described below.
- You may see the message SORT command failed for "filename" if the file has >65K lines and the filename (or directory name) contains blanks or special characters. Either rename the file/directory or just pre-sort the file before using the Analyser.

Analyser options

GRAPHS   The first option is either **ALL** or **LIST**.   If the value is **LIST** then only those sheets which appear in the LIST on the Settings sheet will have graphs drawn for them. This option is particularly useful if the graphs are to be printed/published or to reduce the amount of memory/fonts/disk space required when analysing files from large systems.

The second is either **CHARTS**, **PICTURES**, **PRINT** or **WEB**.   The meaning of these are as follows:

- **CHARTS**   produce Excel charts in-place on the selected sheets
- **PICTURES**   graphs will be produced on a separate "Charts" sheets as pictures. Selecting this option can reduce the size of the output file by up to 90%.
- **PRINT**   implies PICTURES. Pictures will be printed to the designated printer (see "Printing Options" below)
- **WEB**   implies PICTURES. Automated web publishing (see "Web Publishing" below)

INTERVALS   specifies the first and last time interval to be processed.   Intervals outside this range will be discarded after parsing.   Note that these are numbers between 1-999999 and are not time values.   Setting a value of 2 for the first interval is useful in discarding the very large numbers that often appear at the start of an NMON collection run with AIX.   If you have used a splitter program on the input file, or if you are analysing data from a LINUX system, then you should leave this as 1.

TIMES   specifies the first and last time/date to be processed.   Samples outside this range will be discarded after parsing.   They can be specified in any form recognised by Excel as time/date values; e.g.

14:00:10	**16:15:30**
4-Aug-12	**6-Aug-12**
18:00 28/6/12	**04:00**

Notes:
- Specifying a date without a time is the same as specifying a time of 00:00.
- If the second time is less than the first then Analyser will assume the second time is in the next calendar day
- If there are no qualifying intervals in the file a message will be issued: "Invalid values for FIRST/LAST - values reset to 1/999999" and the entire file will be processed.

MERGE   **Note: This option doesn't work on a Mac.**

specifying **YES** here results in NMON_Analyser merging all of the input files to form a single file.   The input files must be unsorted. By default the Analyser will delete the TOP and UARG sections during the merge process; specify **TOP** to prevent this from happening but be aware that processing time will be increased and that if the TOP section exceeds the maximum number of lines per sheet (depending on the version of Excel) then data will be lost. Specify **KEEP** to stop the merged file from being deleted at the end of the run. Specify **ONLY** if you simply want to create a merged file for future analysis.

PIVOT   **Note: This option doesn't work on a Mac.**
specifying **YES** here results in NMON_Analyser creating a Pivot Chart from the specified sheet after all other processing (including printing/publishing the other charts) has completed. See additional parameters on the "Settings" sheet.

SCATTER   Indicates if Scatter charts shall be included. Applies to TOP CPU% by PID chart only. This option is to avoid a possible Excel 2013 crash.

BIGDATA   Set to YES if your data has > 1048576 rows or large lines (up to 32K.) Set this to NO to improve performance but the Analyzer will fail if a large line is encountered and data will also be truncated at 1048576 rows.

ESS   specifying **NO** here results in NMON_Analyser bypassing the additional analysis performed for ESS subsystems. This will result in faster analysis and can allow larger files to be analysed successfully when "out of memory" errors occur.

SYNCAXIS   On graphs that show Avg/Wavg on Y1 and Min/Max on Y2 the Maximum scales will be set to the same value.

SHOWMACPROGRESS   **Note: This option has no effect on Windows.** On a MAC refreshing the screen to show progress in status bar will make things 3 to 5 times slower!

SHOWLINUXCPUUTIL   When analyzing Linux data a value of YES indicates that all of the CPUUTILnnn sheets will be included. CPUUTIL_ALL will always be included (if it exists.)

FILELIST   the name of a control file containing a list of nmon output files to be processed by the Analyser. Leave this field blank for normal operation. The name must be fully qualified (e.g. `c:\nmon\testcases\filenames.txt`). The names specified in the text file must contain full path information. Wildcard characters may be included in the filenames so long as they conform to Windows standards. For example:

```
c:\nmon\testcases\*.nmon
```

NB: if you save the spreadsheet with a value in this field, the Analyser will automatically begin execution the next time you open it. This is defined as "batch mode" (see Appendix). You can stop the execution by pressing Ctrl+Break or by deleting/renaming the Control File.

Batch Processing Options

The following fields can be found on the "Settings" sheet.

REPROC   Change this to **NO** if you want to bypass processing of input files which may have been processed in a previous run. This is useful if you make use of wildcards in the batch control file. Note that the REPROC option only takes effect when you have specified more than one input file.

OUTDIR   the name of an existing directory in which output files will be saved by default. This is primarily intended for batch operation (see **FILELIST** above) but also works to set the default directory for interactive sessions. If the directory does not exist (or OUTDIR is blank) then output files are saved back to the same directory as the corresponding input files.
Example: `C:\NMON\Analyser\Output\`

Formatting Options

The following fields can be found on the "Settings" sheet.

BBBFont   Enter the name of a fixed pitch font to be used for formatting the BBBC and
                 BBBP sheets. `Courier` provides acceptable results.

GWIDTH   Change the values in this row to make the generated graphs bigger or smaller.
                 The default value of 0 means that the Analyser will dynamically size the graphs
                 according to your screen size, font settings or page size.   Be careful not to set a
                 value larger than your page width when printing.

GHEIGHT   Change the values in this row to make the generated graphs bigger or smaller. If
                 you specify a value here you must also specify a value for GWIDTH.

LIST   A comma-separated list a sheets for which the Analyser is to draw graphs. Only
                 used if the GRAPHS option is set to **LIST**.   The list can contain any valid wildcard
                 characters recognised by Excel, e.g. "EMC*"
                 Note: graphs are always drawn for SYS_SUMM, CPU_SUMM and
                 DISK_SUMM

NOLIST   The default is **KEEP**. If you change this to **DELETE** then all sheets which do not
                 appear in LIST will be deleted after analysis. This can dramatically reduce the
                 size of files that are to be kept for long periods.

CPUmax   Specifies the maximum number of CPUnnn, PCPUnnn and SCPUnnn sheets that
                 will be generated.   The default value of 0 will choose all sheets for an LPAR
                 using dedicated processors and number equal to the SMT mode for an LPAR
                 using shared processors.

REORDER   specifying **YES** here results in NMON_Analyser reordering the sheets to improve
                 navigation to more relevant information.

TOPDISKS   the maximum number of hdisks/vpaths to include on disk graphs. A value of 0
                 produces graphs containing all the hdisks on a sheet (up to 250).   Graphs
                 containing more than 50 hdisks will be automatically scaled to fit and may
                 therefore exceed the size of the screen.

xToD   Format to be used for timestamps on Time of Day graphs. Anything acceptable to
                 Excel as a Number Format Code may be entered.   Default is **hh:mm**. Note that
                 the date is also available within the timestamp and you may therefore use
                 something like **dd-mmm-yy hh:mm** if, for example, you have merged multiple
                 NMON files together. If you use something other than the default string you may
                 need to increase the value of GHEIGHT – experiment with different values if you
                 don't see what you expect.

SORTDEFAULT This setting indicates if the 1st graph on 'default' sheets (ones not handled else
                 where by the Analyzer) are sorted. **Note:** If Yes, this will also reorder the
                 columns.

Pivot Chart

These parameters are used to construct a pivot chart.   The required parameters are: Sheetname,
PageField, RowField, ColumnField, DataField and xlFunction (can be **COUNT, SUM, MIN, AVG,
MAX**). This is primarily useful for the TOP and SUMMARY sheets but might prove useful for
other, possibly user-supplied, data sheets.

Printing Options

The following fields can be found on the "Settings" sheet. Note - these only take effect if you
select **PRINT** for the OUTPUT option on the Analyser sheet.

LSCAPE   Change to **YES** if you want the Analyser to set the page orientation to Landscape. By default the Analyser will fit one chart per page when printing landscape.

COPIES   Set to the number of copies to be printed.

PRINTER   The name of the printer. Specifying a value of **PREVIEW** will cause the Analyser to invoke the Excel print preview function - useful for testing.   You may also specify **DEFAULT** to print to the system default printer or the name of network printer.

The Analyser adds page headers and footers.

Web Publishing Options

The following fields can be found on the "Settings" sheet.

PNG   Change to **NO** if your browser can't handle the PNG graphics format. Graphics will be generated as GIF files.

SUBDIR   If this is **YES** then all supporting files, such as background textures and graphics, are organized in a separate folder. If this is **NO** then supporting files are saved in the same folder as the Web page.

WEBDIR   the name of an existing directory in which HTML files will be saved by default. If the directory does not exist (or OUTDIR is blank) then output files are saved back to the same directory as the corresponding input file.
                  Example: `C:\NMON\Analyser\HTML\`

Interpreting the output sections

Notes on the "Weighted Average" as used in the Analyser

Several graphs produced by the Analyser show average, weighted average and maximum values. Although everyone understands averages and maximums, the concept of a weighted average is a little more difficult to grasp.

One of the problems we are faced with in analysing sample data is that resources on the target system may be idle for long periods during the collection. For example, the NMON data collection may be started some time before the system reaches peak utilisation and may not be stopped until the workload being monitored has long since finished.   Although this does not affect the maximums it can severely affect the accuracy of the averages.

The idea of a weighted average is to apply a weighting factor to each snapshot to indicate how relevant that snapshot is to the average.   In NMON_Analyser, we use the value of the measurement itself as the weighting factor.   In effect, this produces a figure that shows how busy a resource is *when it is active*. For example: a database log disk is only active during the middle part of a benchmark.   We record the following figures for %tm_act (DISKBUSY):

Snapshot	%tm_act
1	0
2	0
3	0
4	0
5	0

6	6.1
7	6.3
8	6.5
9	9.1
10	5.9
11	0
12	0
13	0
14	0
15	0

The average for this set of data is 2.3 and the weighted average is 7.0.   The weighted average gives a better picture of how busy the disk is while logging is taking place.   NMON_Analyser uses the weighted average as a sort key when sorting the contents of disk sheets.

Note that, occasionally, NMON generates very large numbers for the first interval and this heavily skews the weighted average values.   If you see this problem occurring then change the value for the FIRST parameter to 2 in order to exclude the first interval completely.   This is only a problem for AIX.

SYS_SUMM

This section is entirely generated by the Analyser and contains a useful summary of data taken from other sheets.   Note the the avg/max values for User%, Sys%, Wait% and Idle% are independent and will not add up to 100%.   The CPU% column shows avg/max values for the sum of usr%+sys% during each interval.   For micro-partitions, the values are shown as a percentage of the Virtual Processors – they do not relate to the CPU_ALL values.

For non-partitioned or dedicated CPU partitions the graph shows the total CPU Utilisation (%usr + %sys) together with the Disk I/O rate (taken from the DISKXFER sheet) by time of day.   For micro-partitions the graph shows the number of physical CPUs being used instead of CPU%.

The value "Max:Avg" is simply the maximum value divided by the average. If monitored over a long period of time the value for CPU% can be useful in spotting a system reaching saturation level (the ratio will steadily decrease). If you have historical data then the value can also be useful in determining how much latent demand is present in a system running at saturation level.

AAA

The AAA section is generated by NMON at the start of the data collection and contains information about the system and NMON itself – contents vary by release, following is for 12e

AIX   this is the release / maintenance level of AIX being used on the target system as reported by the lslpp command.
build   the particular build of nmon used to collect this data
command   the command line used to invoke NMON and the date when it finished.

cpus   the number of CPUs in the system and the number active at the start of data collection.

date   date at the start of the collection.

disks_per_line the setting of the **-l** flag on the NMON command line or the default value.

hardware   the processor technology used in the target system.

host   the hostname of the target system.

Interval   the time (in seconds) between snapshots**.**

kernel   information from the kernel - useful in identifying the type of kernel (32-bit or 64-bit) and whether this is an LPAR.

LPARNumberName   the LPAR number followed by the name

MachineType machine type and model of the system

progname   the name of the NMON executable.

runname   taken from the NMON command line if specified using the -r parameter, else defaults to hostname.

SerialNumber the machine serial number

snapshots   the number of snapshots - this is used by the Analyser.   The Analyser will modify this value to match the number of snapshots actually found in the input file.

steal   this value will be a 1 if running on Linux and Steal% CPU is provided in the CPUnnn and CPU_ALL data.

subversion   detailed information about the nmon version used to collect this data – including the date and time it was created.

time   time as shown by the system clock at the start of the collection. Also see **date** value**.**

timestampsize the number of characters used for timestamps in each record. The default is 5.

TL   the Technology Level of the AIX release

user   the name of the user executing the NMON command.

version   the version of NMON used to collect this data.

VIOS   the release/maintenance level if this is a VIOS LPAR.

analyser   the version of NMON_Analyser used to generate the output file together with the elapsed time (in seconds) for processing this file.

environment the version of Excel you are using.

parms   the values of most of the user options specified on the "Analyser" sheet.

settings   the values of most of the user options specified on the "Settings" sheet.

Analysis time the execution time of the Analyser. I use this for tuning.

NodeName   the name of the LPAR's node name

NMON_Analyser deletes the **NOTES** lines generated by NMON.

BBBB

The BBBB sheet lists all of the disks listed in the ODM together with the capacity (in Gbytes) and the adapter type (SCSI/SSA/Fibre) as reported by lsdev.   Note that some fibre-attached devices do not report their capacity to AIX.

NMON_Analyser deletes the column containing the sort key generated by NMON on all BBB sheets.

BBBC

The BBBC sheet shows the output from the lspv command for all local disks at the start of the data collection. The Analyser highlights the hdisk name using a bold font and sets the sheet to use the fixed-pitch font specified on the NLS sheet (default `Courier`) in order to improve readability.

BBBD

The BBBD sheet shows a list of all I/O adapters listed in the ODM together with the hdisks addressed through that adapter.

BBBE

The BBBE sheet contains data extracted from the lsdev command and shows the mapping between vpaths and hdisks. NMON_Analyser uses this information to construct the ESSBUSY, ESSRSIZE and ESSWSIZE sheets.

BBBG

The BBBG sheet contains details of the NMON disk group mappings.

BBBL

The BBBL sheet is only produced if the operating system is running in a partitions and contains details of the configuration of the LPAR at the start of the collection run.

BBBN

The BBBN describes each network adapter in the system and shows the name, speed and MTU size.

BBBP

The BBBP sheet contains the un-interpreted output from the emstat and lsattr commands. Note that to get output from these commands requires NMON to be running with root privileges. The Analyser sets the sheet to use the fixed-pitch font specified on the NLS sheet (default `Courier`) in order to improve readability.

BBBR

This sheet records dynamic LPAR reconfiguration events during the collection run.

BBBV

This sheet lists all of the volume groups present at the start of the collection run.

CPUnnn

These sheets show %usr, %sys, %wait and %idle by time of day for each logical processor. Note that for micropartitions the Idle% and Wait% figures will include times when the physical processor was ceded to the shared pool.
NMON_Analyser generates a graph and a column headed "CPU%" containing the sum of %usr and %sys for use on the CPU_SUMM sheet.   The Analyser also adds blank intervals for CPUs

that are varied online during the collection interval so that the graphs cover the entire collection period.

If **REORDER** is set to **YES** on the Control sheet the Analyser will move all CPUnn sheets to the end of the file.

CPU_ALL

This sheet shows the average utilisation for all physical threads by time of day. **Note that for micro-partitions CPU% is a measure of utilisation vs the entitlement but is adjusted by libperfstat so that it never exceeds 100%; this makes the numbers virtually useless for analysis of uncapped partitions and you may choose to look at the charts on the LPAR sheet instead.** Note also that micro-partitions generally record very little Idle% or Wait% because they will normally cede their timeslice to other LPARs rather than waste CPU time waiting for work. If the **CPUmax** value is lower than the number of physical threads in use then the graph title will show the number of threads excluded.

The second graph shows the number of active CPUs by time of day and is useful in determining whether CPUs have been varied on/off during the collection period. If the system has SMT enabled then the Analyser shows logical CPUs rather than physical CPUs and the legend reports "SMT=ON".

CPU_SUMM

The Analyser generates this sheet from data on the CPUnnn sheets. It gives a breakdown of CPU Utilisation by thread (logical processor) and by core over the collection period. The chart can be very useful in identifying situations in which the system is thread-starved (i.e. too few threads to fully utilise the logical processors) or where the workload is dominated by a small number of single-threaded processes.

Note that if CPUs have been dynamically reconfigured during the collection period, these figures reflect only those intervals when the CPU was varied on.

General notes for DISK, ESS, EMC, FASt and DG sheets

The DISK sheets record device statistics for each hdisk in the system. If there are more than <dpl> hdisks in the system (see "Collecting data using NMON" above) NMON generates multiple output sections. DISKBUSY will therefore contain device busy statistics for the first group of hdisks, DISKBUSY1 for the next, DISKBUSY2 for the next and so on.

The Analyser will normally sort the contents of the sheet in ascending sequence using the weighted average values as a sort key. However, if a storage subsystem is detected then sorting of the DISK sheets is disabled and the ESS/EMC/FASt/DG sheet contents are sorted instead.

If the number of hdisks (or vpaths) on the sheet exceeds the value specified for **TOPDISKS** (see "Analyser options" above), the graph will only show information for the specified number of disks and a warning will appear in the graph title.

If **REORDER** is set to **YES** on the Control sheet the Analyser will move all but the DISKBUSY, DISK_SUMM and DISKSERV sheets to the end of the file when a storage subsystem or Disk Group is detected.   The rationale behind this being that the ESS/EMC/FASt/DG sheets contain the most useful data and that the DISK sheets merely replicate it.   However, the DISKBUSY sheet is useful for checking the activity on system disks (normally local) and the DISK_SUMM sheet gives total data rates for the system (local + subsystem disks).

Note that NMON and NMON_Analyser can only handle a maximum of 250 vpaths in a system. If you have more than this then you need to use the Disk Groups feature of NMON to select the vpaths that are of primary interest.

EMC/PowerPath subsystems

NMON_Analyser detects the presence of an EMC/PowerPath subsystem by scanning the input file for the string "hdiskpower" before starting the analysis.   EMC PowerPath creates devices called "hdiskpower*n*" which each map to multiple hdisks.   They are therefore comparable to the vpaths generated by ESS/SDD.   However, unlike vpaths, hdiskpower devices appear to AIX (and, therefore NMON) as real disks and NMON therefore records their activity on the DISK sheets.

NMON_Analyser removes all hdiskpower entries from DISK sheets and moves them to new sheets beginning with the letters EMC.   For example, hdiskpower entries found on DISKBUSY1 are simply moved to a new sheet called EMCBUSY1. If a DISK sheet contains only hdiskpower devices, that sheet is simply renamed.

All device statistics reported by the Analyser (e.g. System I/O rates) are correct and as NMON excludes hdiskpower activity from the IOADAPT statistics these figures are also correct.

FAStT subsystems

NMON_Analyser detects the presence of a FAStT subsystem by scanning the input file for the string "dac" before starting the analysis. They are handled in the same way as EMC/PowerPath subsystems except that the dac devices are moved to sheets with names beginning "FASt". Note that the Analyser is not able to correctly handle systems having both EMC and FASt subsystems.

DGBUSY

This sheet records the average value of device busy for each hdisk in the NMON Disk Group.

DGREAD

This sheet records the average data rate (Kbytes/sec) for read operations to each NMON Disk Group.

DGSIZE

This sheet records the average data transfer size (block size), in Kbytes, for read/write operations to each NMON Disk Group.

DGWRITE

This sheet records the average data rate (Kbytes/sec) for write operations to each NMON Disk Group.

DGXFER

This sheet records the total I/O operations per second to each NMON Disk Group.

DISKBSIZE

These sheets record the number of disk blocks, for read/write operations on each hdisk in the system. If this number is not very close to the stripe size for the device there may be a problem that could be solved by increasing the value of numclust.

DISKBUSY

These sheets record device busy for each hdisk in the system. This is the same as the %tm_act value recorded by iostat. Note that if this sheet contains all zero values then it means you forgot to enable iostat collection before starting nmon:

        chdev -l sys0 -a iostat=true

DISKREAD

These sheets record the data rate (Kbytes/sec) for read operations on each hdisk in the system.

DISKSERV

These sheets record the service times (in milliseconds) for read/write transfers to each hdisk in the system.

DISKWAIT

These sheets record the queue times (in milliseconds) for read/write transfers to each hdisk in the system.

DISKWRITE

These sheets record the data rate (Kbytes/sec) for write operations on each hdisk in the system.

DISKXFER

These sheets record the I/O operations per second for each hdisk in the system.   This is the same as the tps value recorded by iostat.

DISK_SUMM

The Analyser creates this sheet.   It shows the total data rates (reads and writes) in Kbytes/sec plus total I/O rates for all hdisks in the system.   The figures on this sheet are accurate for all systems including ESS, EMC, FASt and HDS configurations.

These data are displayed on the chart; the IO/sec data are also graphically displayed on the AAA sheet.

DONATE

This sheet records physical processor usage and donation to the shared pool. Only present for dedicated LPARs running on POWER6 systems.

EMCBSIZE/FAStBSIZE

This sheet records the average data transfer size (blocksize), in Kbytes, for read/write operations to each esoteric device in a system using EMC/PowerPath or FAStT.

EMCBUSY/FAStBUSY

These sheets record device busy for each esoteric device in a system using EMC/PowerPath or FAStT.

EMCREAD/FAStREAD

These sheets record the data rate (Kbytes/sec) for read operations to each esoteric device in a system using EMC/PowerPath or FAStT.

EMCWRITE/FAStWRITE

This sheet records the data rate (Kbytes/sec) for write operations to each esoteric device in a system using EMC/PowerPath or FAStT.

EMCXFER/FAStXFER

These sheets record the I/O operations per second to each esoteric device in a system using EMC/PowerPath or FAStT.

EMCSERV/FAStSERV

The Analyser creates this sheet.   It shows estimated services times (not response times) for each esoteric device over the collection interval.   The service time is derived from the device busy and the transfer rate taken from the corresponding BUSY and XFER sheets.   Intervals where the transfer rate is below **SVCXLIM** are ignored in order to improve the accuracy of the estimate.

ERROR

This sheet shows all errors reported by nmon during the collection run.   FIRST/LAST intervals are ignored for this sheet.   Timestamps are not converted to time values.

ESSBSIZE

This sheet is only present if **ESS** is set to **YES** and records the average data transfer size (blocksize), in Kbytes, for read/write operations on each vpath in the system. The data on this sheet are calculated by NMON_Analyser as the average of the DISKBSIZE values for component hdisks as recorded on the BBBE sheet.

ESSBUSY

This sheet is only present if **ESS** is set to **YES** and records device busy for each vpath in a system using ESS. The data on this sheet are calculated by NMON_Analyser as the average of the DISKBUSY values for component hdisks as recorded on the BBBE sheet.

ESSREAD

This sheet records the data rate (Kbytes/sec) for read operations on each vpath in the system. This information is provided by NMON.

ESSWRITE

This sheet records the data rate (Kbytes/sec) for write operations on each vpath in the system. This information is provided by NMON.

ESSXFER

These sheets record the I/O operations per second for each vpath in the system.   This information is provided by NMON.

FCREAD

This sheet records the data rate (MBytes/sec) for read operations on each Fibre Channel adapter in the system.

FCWRITE

This sheet records the data rate (MBytes/sec) for write operations on each Fibre Channel adapter in the system.

FCXFERIN

This sheet records the read operations per second for each Fibre Channel adapter in the system.

FCXFEROUT

This sheet records the write operations per second for each Fibre Channel adapter in the system.

FILE

This sheet contains a subset of the fields reported by NMON on the Kernel Internal Statistics panel.   These are the same values as reported by the sar command.

All fields are rates/sec.

iget   translations of i-node numbers to pointers to the i-node structure of a file or device. This is reported as **iget/s** by the sar -a command. Calls to iget occur when a call to to namei has failed to find a pointer in the i-node cache. This figure should therefore be as close to 0 as possible.

namei   calls to the directory search routine that finds the address of a v-node given a path name. This is reported as **lookuppn/s** by the sar -a command.

dirblk   number of 512-byte blocks read by the directory search routine to locate a directory entry for a specific file. This is reported as **dirblk/s** by the sar -a command.

readch   characters transferred by read system call. This is reported as **rchar/s** by the sar -c command

writech   characters transferred by write system call. This is reported as **wchar/s** by the sar -c command.

ttyrawch   tty input queue characters. This is reported as rawch/s by the sar -y command.

ttycanch   tty canonical input queue characters. This field is always 0 (zero) for AIX Version 4 and later versions.

ttyoutch   tty output queue characters. This is reported as outch/s by the sar -y command

NMON_Analyser produces two graphs - one showing rates/sec for readch and writech by time of day and one showing rates/sec for iget, namei and dirblk.

FRCA

This sheet is only generated if FRCA is loaded on the target system. NMON_Analyser produces a graph showing the cache hit ratio (as a percentage). If FRCA is not loaded, NMON generates no data and the Analyser consigns the redundant header record to the StrayLines sheet.

IOADAPT

For each I/O adapter listed on the BBBC sheet, contains the data rates for both read and write operations (Kbytes/sec) and total number of I/O operations performed. On AIX 5.1 and later, this information is reported by the iostat -A command. NMON_Analyser reorders the columns on the sheet for easier graphing.

The Analyser generates three graphs. Note the charts can be easily converted to line charts if required. Simply right click on the white space within the chart area, then select Chart Type>Line>OK.

IP

This sheet only appears for topasout.

JFSFILE

For each file system, this sheet shows what percentage of the space allocation is being used during each interval. These figures are the same as the **%Used** value reported by the df command. The column headings show the mount point; sheet BBBC can be used to cross-reference to the file system/LV.

JFSINODE

For each file system, this sheet shows what percentage of the Inode allocation is being used during each interval. These figures are the same as the **%Iused** value reported by the df command. The column headings show the mount point; sheet BBBC can be used to cross-reference to the file system/LV.

LAN

This sheet only appears for topasout.

LARGEPAGE

The graph shows Usedpages and Freepages over time.

Columns on the sheet are as follows:

Freepages   the number of large pages on the free list.

Usedpages   the number of large pages currently in use.
Pages   the number of large pages in the pool.
HighWater   the maximum number of pages used since the last reboot.
SizeMB   the size of a large page in Mbytes.

LPAR

The first graph shows the number of physical processors used by this partition vs the entitlement. For an uncapped partition the number of physical processors may exceed the entitlement but can never exceed the number of virtual processors allocated to the partition. For AIX the graph also shows the number of unfolded virtual processors (AIX will "fold" - stop dispatching work to - excess processors in order to minimise scheduling costs).

Note that the ratio of physical processor to entitlement (shown as %entc in the output of the lparstat command) will generally be higher than CPU% on the CPU_ALL sheet. The reason for this is that a partition that is within its entitlement may wait for a short period of time before ceding a processor that enters an I/O wait or becomes idle. This can eliminate unnecessary context switches.

The second graph shows CPU utilisation as a percentage of virtual processors – for AIX this is broken down in to usr%, sys% and wait%. This level of detail is not available for Linux or releases of NMON prior to version 12.

The third graph is only present for AIX systems and shows cpu utilisation of the shared pool by this and other partitions. The area marked "UsedPoolCPU%" represents the percentage of the shared pool that has been used by this partition, while the area marked "other% " represents the percentage used by all other partitions. Note that if the partition is not authorised to see utilisation of the shared pool then the pool will appear to be 100% utilised.

Columns on the sheet are as follows:
PhysicalCPU   physical cores consumed by AIX
VirtualCPUs   number of Virtual CPUs allocated to the LPAR
logicalCPUs   number of threads (i.e. Virtual CPUs multiplied by SMT mode)
poolCPUs   number of cores in the pool that this LPAR occupies
entitled   the number of cores guaranteed to be available to this LPAR
weight   the priority of this LPAR when competing for unused CPU cycles
PoolIdle   the number of unused cores in the Pool
usedAllCPU%   percentage of active cores in the machine that this LPAR is using
usedPoolCPU   percentage of cores in the pool that this LPAR is using
SharedCPU   1 if the LPAR is not a dedicated CPU
Capped   1 if the LPAR is capped
EC_User%   percentage of Entitlement used in User mode
EC_Sys%   percentage of Entitlement used in System mode
EC_Wait%   percentage of Entitlement waiting for I/O
EC_Idle%   percentage of Entitlement used in Idle mode
VP_User%   percentage of Virtual CPU used in User mode

VP_Sys%	percentage of Virtual CPU used in System mode
VP_Wait%	percentage of Virtual CPU waiting for I/O
VP_Idle%	percentage of Virtual CPU spent in Idle mode
Folded	no. of Virtual CPUs unused for efficiency reasons
CPU_Pool_id	the Id of the CPU pool, useful if there are multiple pools

MEM

The main graph on this sheet shows the amount of Real Free memory in Mbytes by time of day. This would be the same as dividing the **fre** values reported by vmstat over the same interval by 256. The small graph shows the amount of real memory. This is useful in determining if dynamic reconfiguration has been used during the collection period.

For AIX, other columns on the sheet are as follows:

Real Free	the percentage of real pages on the free list.
Virtual Free	the percentage of unallocated virtual slots on the paging spaces.
Real Free (MB)	the amount of memory on the free list in Mbytes.
Virtual Free (MB)	the amount of unallocated space on the paging spaces.
Real Total (MB)	the total amount of memory available to AIX.
Virtual Total (MB)	the total amount of space allocated for paging spaces.

Note: you can calculate the amount of memory used during an interval simply by subtracting the **Real Free (MB)** value from the **Real Total (MB)** value. This will, however, include file pages. The graph on the MEMUSE sheet gives a more accurate assessment of memory used by programs (computational pages).

MEMUSE

Except for **%comp**, the values on this sheet are the same as would be reported by the vmtune command.

%numperm	the percentage of real memory allocated to file pages.
%minperm	value specified on the vmtune command or system default of 20%. This will normally be constant for the run unless the vmtune or rmss commands are used during collection.
%maxperm	value specified on the vmtune command or system default of 80%. This will normally be constant for the run unless the vmtune or rmss commands are used during collection.
minfree	the minimum number of pages AIX is to keep on the free list. Specified on the vmtune command or system default of **maxfree - 8**.
maxfree	the maximum number of pages AIX is steal in order to replenish the free list. Specified on the vmtune command or system default.
%comp	The percentage of real memory allocated to computational pages. NMON_Analyser calculates this value. Computational pages are those backed by page space and include working storage and program text segments. They exclude data, executable and shared library files.

The Analyser generates two graphs. The first shows the split between computational and file pages by time of day. The second plots the values of **%numperm, %minperm, %maxperm** and **%comp**.

If **%numperm** falls below **%minperm** then computational pages will be stolen.   If **%numperm** rises above **%maxperm** then computational pages cannot be stolen. Low values for both **%minperm** and **%maxperm** indicate that the system has been tuned for a database server.   You may also want to check the setting of STRICT_MAXPERM on the BBBP sheet (if present).

MEMNEW

The graph shows the allocation of memory split into the three major categories: pages used by user processes, file system cache, and pages used by the system (kernel).

Process%   the percentage of real memory allocated to user processes
FSCache%   the percentage of real memory allocated to file system cache
System%   the percentage of real memory used by system segments
Free%   the percentage of unallocated real memory
User%   the percentage of real memory used by non-system segments

MEMPAGES4K/64K/16MB/16GB

These sheets are only present for AIX. They show various statistics for different page sizes in use within the system.   The Analyser will delete the sheets for pages sizes not currently in use.

For the MEMPAGES64KB sheet a graph is drawn showing the use of both 4KB and 64KB pages within the system – starting with AIX V6.1 (and supporting hardware) these page sizes are selected dynamically based on memory access patterns.

MEMREAL

This sheet only appears for topasout. The Analyser adds a column showing Real Free memory in Mbytes.

MEMVIRT

This sheet only appears for topasout.

NET

This sheet shows the data rates, in Kbytes/sec, for each network adapter in the system (including SP switch if present). This is the same as produced by the netpmon –O dd command. NMON_Analyser adds one column for each adapter showing the total data rate (read + write) and two columns showing Total Read and Total Write. Note that the Total Write is calculated as a negative number for graphing.

The Analyser generates three graphs.   The first graph shows total network traffic broken down as Total-Read and Total-Write. The writes are shown below the X-Axis.

Note the area chart can be easily converted to a line chart if required.   Simply right click on the white space within the chart area, then select Chart Type>Line>OK.

NETPACKET

This sheet shows the number of read/write network packets for each adapter. This is the same as produced by the netpmon –O dd command.

NETSIZE

This sheet shows the average packet size in bytes for each network adapter in the system.

NFS sheets

There are separate sheets for NFS2, NFS3 and NF4 client/server.   The Analyser will delete empty sheets.

PAGE

This sheet has the paging statistics as recorded by NMON.

faults	the number of page faults per second. This is not a count of page faults that generate I/O, because some page faults can be resolved without I/O.
pgin	the total rate/sec of page-in operations to both paging space and file systems during the interval.
pgout	the total rate/sec of page-out operations to both paging space and file systems during the interval.
pgsin	the rate/sec of page-in operations from paging space during the interval. This is the same as the **pi** value reported by vmstat. If **pgsin** is consistently higher than **pgsout** this may indicate thrashing.
pgsout	the rate/sec of page-out operations to paging space during the interval. This is the same as the **po** value reported by vmstat.
reclaims	from NMON 10 onwards this field is the same as the **fr** value reported by vmstat and represents the number of pages/sec freed by the page-replacement routine.
scans	the number of pages/sec examined by the page replacement routine. This is the same as the **sr** value reported by vmstat. Page replacement is initiated when the number of free pages falls below **minfree** and stops when the number of free pages exceeds **maxfree**.
cycles	the number of times/sec the page replacement routine had to scan the entire Page Frame Table in order to replenish the free list. This is the same as the **cy** value reported by vmstat but note that vmstat reports this number as an integer whereas nmon reports it as a real number.
fsin	calculated by the Analyser as **pgin-pgsin** for graphing
fsout	calculated by the Analyser as **pgout-pgsout** for graphing
sr/fr	calculated by the Analyser as **scans/reclaims** for graphing

NMON_Analyser produces two graphs.   The first shows paging operations to/from paging space. The ideal here would be no more than 5 operations/sec per page space (see the BBBC sheet for details). The second graph shows the scan:free rate.   Memory may be over-committed when this figure is >4 although you also need to examine the MEM and PAGE sheets as well.

POOLS

This sheet contains information about the shared pool in which the LPAR is running.   Most of the data will only be present if "Allow performance information collection." is set in the LPAR properties.

shcpus_in_sys	the number of cores allocated to the global shared pool
max_pool_capacity	the maximum number of VPs defined for this pool
entitled_pool_capacity	the entitlement for this pool (includes reserve entitlement)
pool_max_time	same as **max_pool_capacity** but may vary in value if the pool definition is changed during the collection run.
pool_busy_time	the average number of cores in use by this shared pool during the interval
shcpu_tot_time	the average number of cores available to the global shared pool (including shared dedicated resources?) during the interval
shcpu_busy_time	the average number of core in use within the global shared pool
pool_id	this identifier of this pool
entitled	the entitlement of this LPAR

PROC

This sheet contains a subset of the fields reported by NMON on the Kernel Internal Statistics panel. The **RunQueue** and **Swap-in** fields are average values for the interval. All other fields are rates/sec:

RunQueue	the average number of kernel threads in the run queue. This is reported as **runq-sz** by the sar -q command and is reported as **RunQueue** on the nmon Kernel Internal Statistics panel. A value that exceeds 3x the number of CPUs may indicate CPU constraint.
Swap-in	the average number of kernel threads waiting to be paged in. This is reported as **swpq-sz** by the sar -q command.
pswitch	the number of context switches. This is reported as **pswch/s** by the sar -w command.
syscall	the total number of system calls. This is reported as **scall/s** by the sar -c command.
read	the number of read system calls. This is reported as **sread/s** reported by the sar -c command.
write	the number of write system calls. This is reported as **swrit/s** by the sar -c command.
fork	the number of fork system calls. This is reported as **fork/s** by the sar -c command.
exec	the number of exec system calls. This is reported as **exec/s** by the sar -c command.
rcvint	the number of tty receive interrupts. This is reported as **revin/s** by the sar -y command.
xmtint	the number of tty transmit interrupts. This is reported as **xmtin/s** by the sar -y command.
sem	the number of IPC semaphore primitives (creating, using and destroying). This is reported as **sema/s** by the sar -m command.
msg	the number of IPC message primitives (sending and receiving). This is reported as **msg/s** by the sar -m command.

NMON_Analyser produces three graphs - one showing the average length of the **RunQueue** and the number of **swap-ins** by time of day, another showing rates/sec for **pswitch** and **syscalls** by time of day and a third showing rates/sec for **forks** and **execs**.

The graph for forks/execs can be useful when monitoring web server systems.

PROCAIO

This sheet contains information about the number of asynchronous I/O processes available and active (i.e. using more than 0.1% of the CPU) .   It also shows the amount of CPU being used by the AIO processes during the collection interval.

Two graphs are produced.   The second uses two y-axes.   The number of running aio processes is shown against the first axis and the amount of cpu used is shown against the second.

RAWCPUTOTAL, RAWLPAR

These sheets contain a dump of various counters such as context switches and phantom interrupts.

TCPUDP

This sheet only appears for topasout.

TOP

This sheet is only generated if you specify the -t flag on the NMON command line. The output is similar to that produced using the ps v command.   Note that, because of the limitation of having only 65,000 lines on a single sheet, some data may be omitted for very large files and this may mean that entire PIDs or even commands may be missing from the analysis.

Note that data are only present for processes that consumed a significant amount of CPU during an interval. The TOP sheet does not represent a complete view of the system.

NMON_Analyser does the following:
* reorders the columns for easier processing.
* Sorts the data on the sheet into **COMMAND** name order - using **TIME** as a minor sort key.
* Creates a table at the end of the sheet summarising the data by command name and used for graphing.

You can see the detail section by scrolling to the top of the sheet.   The summary table is largely obscured by the graphs and so you will need to move (or delete) them for easier viewing.

PID   in the detail section this is the process ID of a specific invocation of a command. In the summary table this is the command name.

%CPU   in the detail section this is the utilisation of a single processor (rather than of the system) by that PID during the interval. In the summary table this is the average amount of CPU used by all invocations of the command during the collection period.

%Usr   in the detail section this is the average amount of User-mode CPU used by that PID during the interval.

%Sys   in the detail section this is the average amount of Kernel-mode CPU used by that PID during the interval.

Threads   the number of (software) threads being used by this command.

Size   the average amount of paging space (in Kbytes) allocated for the data section (private segment + shared library data pages) for one invocation of this command. This is the same as the **SIZE** figure on the ps v command. Note that if **Size** is greater than **ResData** it means some working segment pages are currently paged out.

ResText   the average amount of real memory (in Kbytes) used for the code segments of one invocation of this command. Note that multiple concurrent invocations will normally share these pages.

ResData   the average amount of real memory (in Kbytes) used for the data segments of one invocation of this command. A method of calculating real memory usage for a command is **ResText** + (**ResData** * **N**).

CharIO   this is the count of bytes/sec being passed via the read and write system calls. The bulk of this is reading and writing to disks but also includes data to/from terminals, sockets and pipes. Use this to work out which processes are doing the I/O.

%RAM   this is an indication of what percentage of real memory this command is using. This is **(ResText + ResData) / Real Mem**; it is the same as the **%MEM** value on the ps v command.   Due to rounding/truncation, and the large amounts of memory in modern systems, this is usually 0.

Paging   sum of all page faults for this process. Use this to identify which process is causing paging but note that the figure includes asynchronous I/O and can be misleading.

Command   name of the command

WLMClass   name of the Workload Partition or Workload Manager superclass to which this command has been allocated (64-bit kernel only).

IntervalCPU   generated by the Analyser.   In the detail section this shows the total amount of CPU used by all invocations of a command in the time interval. It is calculated as the sum of CPU used by all PIDS running the same command divided by the number of active processors (physical cores) available during the interval. In the summary section this is broken down as Average, Weighted Average and Maximum and is used to generate the graph.

WSet   generated by the Analyser.   In the detail section this shows the total amount of memory used by all invocations of a command recorded in the time interval. It is calculated as **ResText** + (**ResData** * **N)** (where "N" is the number of copies of this command running concurrently during the interval). In the summary section this is broken down as Minimum, Average and Maximum and is used to generate the graph.

User   generated by the Analyser if a UARG sheet is present. This contains the name of the user running the process.

Arg   generated by the Analyser if a UARG sheet is present. This contains the complete argument string entered for the command.

The Analyser generates four graphs using data in the generated table:
• A graph showing Average, Weighted Average and Maximum CPU Utilisation by command
• A graph showing Minimum, Average and Maximum Memory Utilisation by command
• A graph showing Average, Weighted Average and Maximum CHARIO by command
• A graph showing the CPU utilisation for each PID for each interval as a scatter chart. Note that this chart is only produced if there are fewer than 32,000 lines on the TOP sheet. See below for notes on interpreting this chart.

Interpreting the %Processor by PID chart

The purpose of the chart is to provide a link to the UARG sheet so that you can discover precisely which invocation of a command was responsible for using the CPU. It shows the processor utilisation (utilisation of a single CPU) by each PID captured on the TOP sheet. Note that a process can use more than 100% of a single CPU if it is multi-threaded.

Active PIDs will create a cluster of points on the chart.   The highest point will show the maximum amount of CPU used during any one snapshot.   To find out which PID a point refers to, move the mouse to position the cursor above it and Excel will display a coordinate pair.   The first coordinate is the PID – use this to refer to the UARG sheet to find precisely which command was being executed.

UARG

This sheet has the first 1,000 commands executed during the collection period.   The commands are listed in time order.   Note that commands appearing in the first interval may have been executing prior to the start of the collection.

PID	the process ID of a specific invocation of a command
PPID	the parent process ID
COMM	the command being executed
THCOUNT	the number of threads started by this process
USER	the name of the user running this process
GROUP	the name of the group to which the user belongs
FullCommand	the full command string entered by the user

VM

This sheet is only present for Linux systems and contains a dump of the /proc/vmstat file values. The two graphs shows file-backed paging (pgpgin/pgpgout) and swap space activity.

WPAR sheets

These sheets are only present for AIX V6 and record data for each Workload Partition in the system.   They are not currently generated by the topas version of nmon.

WLM sheets

Contain details of CPU, Memory and I/O bandwidth used by each Superclass/Subclass defined to WLM during the collection run.   The Analyser will extract subclasss data and create a new set

of sheets for each class with more than one subclass. These sheets will be named "WLMCPU.class" etc.

If this is a Micro-partition then the Analyser will also create a set of WLMPCPU sheets which will show the physical processor utilisation rather than %CPU utilisation.

ZZZZ

The Analyser uses the information on this sheet to automatically convert all NMON time stamps to actual time of day for easier analysis. For NMON10 or later a column is added which contains the date and time as a single value and this is used for the conversions. The number of rows on this sheet is used by the Analyser to reset the "snapshots" value on the AAA sheet in case the nmon run was terminated with SIGUSR2.

Error Handling

Error handling in NMON_Analyser is rudimentary.   The Analyser can handle many input file errors, but occasionally the analysis will halt leaving you staring at a dialog box. Should this happen, please accept my apologies. However, before sending me a copy of the input file, please read the following:

Common problems

The most commonly reported problems arise from invalid input files.   We also get problems reported where, for whatever reason, lines have been truncated, split or even duplicated. NMON_Analyser attempts to trap these errors and will report them on the "StrayLines" sheet. Check this sheet if the run stops unexpectedly.

- `'No valid input! NMON run may have failed.'`

The most common cause of this message is that the NMON run failed and there really is no valid input.   NMON initialises the output file by writing all of the section headings.   If it subsequently fails, you will get an output file that consists purely of headings - with no data. Check the file by loading it into a word processor or, indeed, a spreadsheet (as a .csv file) before you send it to me.

- `'Unexpected end of file.'`

This is only reported when processing files containing more lines than can be stored on one sheet and when SORTINP is set to NO. It is usually caused by the fact that lines are being terminated with a CR character instead of CRLF. Change your FTP option to ASCII or TEXT when sending the file to your PC. This problem only shows up with large files because the Analyser uses a different technique to read them than that used to read smaller files.

Strangely, one of the most common problems I get is caused by the fact that the Analyser can't handle files with a single data interval very well.   If you want to test the package just let NMON run for a few minutes to get a reasonable data sample!

Known bugs/problems (V3.4, topas_nmon)

- When analysing systems with a very large number of disks, Excel can stop with error "No more fonts can be added" or "Insufficient Resources".   Set **GRAPHS** to **LIST** and select only those sheets you are interested in using the **LIST** option on the Settings sheet.
- The data can be sorted incorrectly with some versions of nmon (notably 14g) which generate variable length timestamps.   Process each file separately and specify **MERGE=YES** to get the data sorted correctly.
- The PIVOT option does not seem to work with Excel 2007 or later
- The Analyser will crash if you edit the .csv file using Excel prior to the run and the file contains a TOP section.   If you need to edit the input file, use a word processor.
- There are some issues with processing files from systems having both ESS and EMC subsystems attached.
- When analysing ESS subsystems with more than 253 vpaths some vpath data will be missing from the output. Use NMON Disk Groups to combine several vpaths into a single unit for reporting and use the **–E** flag to prevent the ESS sections from being produced. Alternatively, set **ESS** to **NO** on the Control sheet and ignore warning messages about data truncation – only the first 253 vpaths will appear in the output.

How to report a problem

Post the relevant information on the nmon forum (see the link in the introduction). It can help to include a copy of the *original, unmodified* .nmon input file, plus the incomplete output spreadsheet, as a compressed (zipped) file but please consider your system security before doing so.

If you have the ability to capture a screenshot then a copy of any dialog boxes also proves useful on occasion.

Note: Development is currently done on Microsoft Excel 2003 (11.8320.8221) SP3. It may not be possible to fix problems arising from the use of different releases.

Excel/VBA Resources/Links

This is a good source for Excel tips and it has some VBA examples as well:
http://exceltips.vitalnews.com/

This is the home of an excellent reference book:
http://www.exceltip.com/

And of course
http://www.microsoft.com

Appendix: Notes on Batch Operation

If you regularly process large numbers of files, the operation of NMON_Analyser can be completely automated.   Simply create a text file containing a list of nmon file names (using wild card characters as appropriate) and enter the name of this file into the FILELIST field of the Analyser control sheet.   Specify the name of an existing directory in the OUTDIR field if you want all of the output files to end up in one place.   Save the NMON_Analyser spreadsheet under a new name (this is recommended so that you can still use NMON_Analyser for interactive sessions).
Now create Windows .bat files to invoke Excel (see the samples below).

After processing the last input file, the Analyser will automatically close down Excel.   Note, however, that this only happens if you load a copy of the Analyser that has a saved FILELIST name **and** if there are no other open spreadsheets.   This allows you to use the FILELIST option safely during an interactive session.

Sample .bat files

These sample batch control files are designed to use the pscp file from the PuTTY suite written and maintained by Simon Tatham at http://www.chiark.greenend.org.uk/~sgtatham/putty/
My thanks to Jamie Dennis for providing them.

getcsv.bat

```
cd \NMON\RawData
del *.csv
d:\progra~1\putty\pscp -p -r -l userid host://Performance/NMON/Rawdata/*.csv
.
```

analyser.bat

```
D:
cd \NMON\FinishedData
del d:\NMON\FinishedData\*.xls
"D:\NMON\nmon analyser batch.xls"
```

putxls.bat

```
D:
cd \NMON\FinishedData
d:\progra~1\putty\pscp -p -r -l userid *.xls host:/Performance/NMON/
```

Control.txt

```
d:\NMON\RawData\*.csv
```

NMON Analyser Batch.xls

OUTDIR   d:\NMON\FinishedData\
FILELIST   d:\NMON\FinishedData\Control.txt