



# nmon (Nigel's Monitor)

- the basics
- new features
- now part of AIX itself

Nigel Griffiths  
IBM Europe



Presentation Version 14



# Trademarks

**The following are trademarks of the International Business Machines Corporation in the United States, other countries, or both.**

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

▪ **Some, all or none of this presentation might, may or will be true or not, as applicable.**

For a complete list of IBM Trademarks, see [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml):

\*, AS/400®, e business(logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

**The following are trademarks or registered trademarks of other companies.**

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

\* All other products may be trademarks or registered trademarks of their respective companies.

## Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

## nmon Why and Principles?

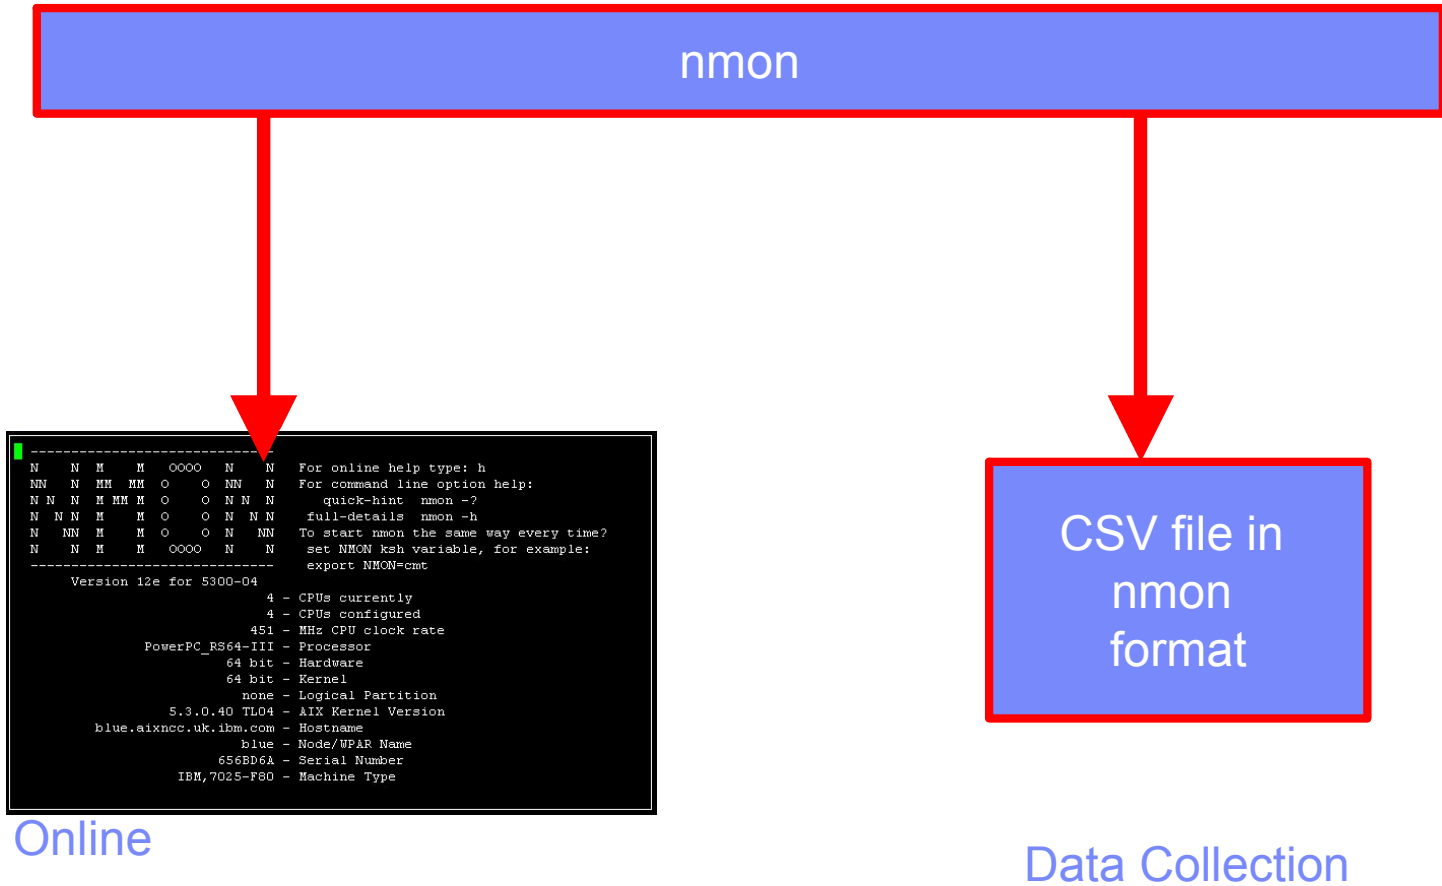
- Why nmon?
  - AIX Monitoring & Tuning → Simple, small & safe
  - Produce benchmark reports with graphs (I am lazy)
  - For myself (personal pet project) but everyone wanted a copy!!
  
- Design principles
  - Zero installation time & simple to use
  - For the performance expert – max info on the screen
  - Less than 1% CPU
  - For large machines
    - 64 CPUs, 4000+ disks, 35,000 processes

# nmon basics

- Freely available
- Performance monitor for
  - AIX 5 and 6
  - AIX 4 (via older nmon version)
  - Linux
  - POWER, x86 & mainframe



# nmon – Online or Data Collector

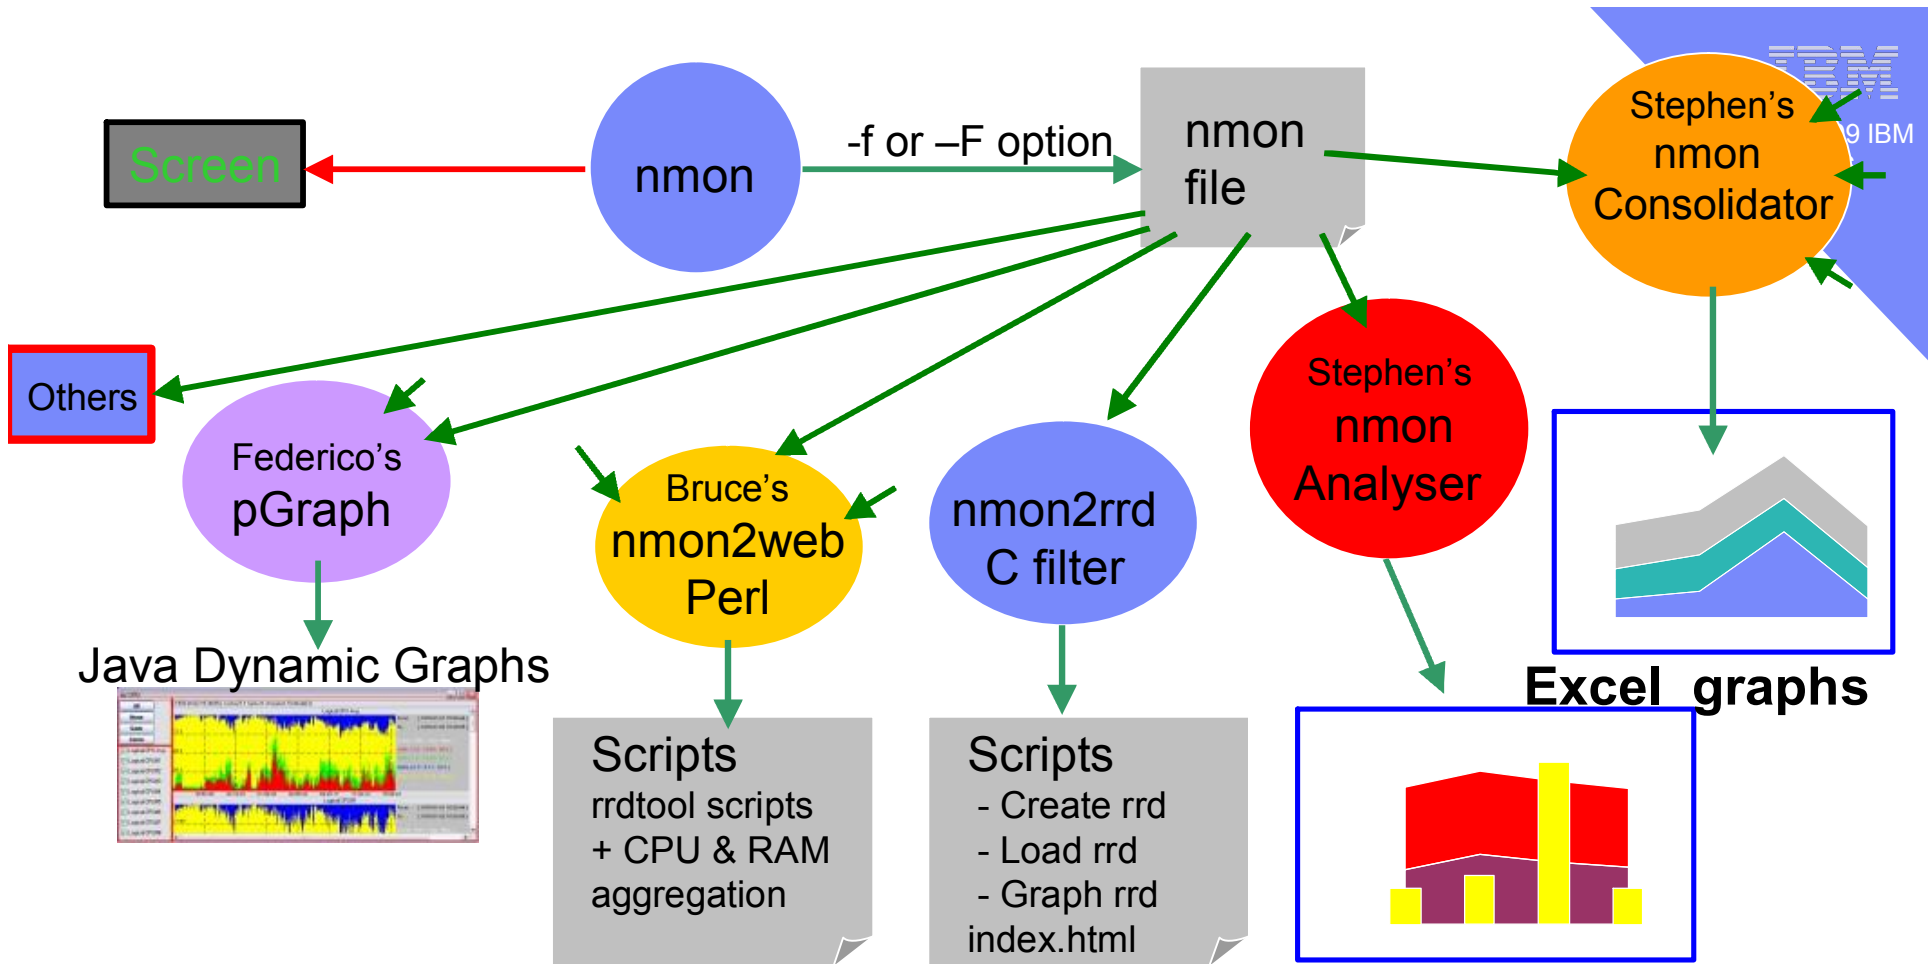


```
N N M M OOOO N N For online help type: h
NN N MM MM O O NN N For command line option help:
N N N M MM M O O N N N quick-hint nmon -?
N N N M M O O N N N full-details nmon -h
N NN M M O O N NN To start nmon the same way every time?
N N M M OOOO N N set NMON ksh variable, for example:
export NMON=cmt
-----
Version 12e for 5300-04
          4 - CPUs currently
          4 - CPUs configured
          451 - MHz CPU clock rate
PowerPC_RS64-III - Processor
          64 bit - Hardware
          64 bit - Kernel
          none - Logical Partition
          5.3.0.40 TLQ4 - AIX Kernel Version
blue.aixcc.uk.ibm.com - Hostname
          blue - Node/WPAR Name
          656ED6A - Serial Number
          IBH,7025-F80 - Machine Type
```


Online

CSV file in  
nmon  
format

Data Collection



# nmon flow



nmon is half the story



**At least, 50% of the success of nmon  
is down to the nmon Analyser from Stephen Atkins**

## nmon Consolidator v1.4.1



- Produces overview charts for CPU, Memory, Network, Disk
  - Fast & simple code, pre-generated charts
- Multiple nmon/topasout files for:
  - ALL LPARs of one machine
  - Clustered system – Oracle RAC / HPC
  - Modelling changes from Dedicated to Shared CPU LPARs
  - Server consolidation Modelling
  - Reporting actual or potential savings from virtualisation
- Trend charts - processing multiple files from 1 LPAR

## Quick Demonstration Here

- **nmon online**
- **nmon analyser**

## nmon Offered to AIX Labs in Austin, Texas

nmon started in 1996

Many customers have wanted IBM/AIX Support

- Lab attitude was classic “Not invented here” syndrome & not a problem determination tool – fair enough

I have provided Support

- I hope quite well or brilliantly

“Perhaps we don’t want AIX support”

- I can produce a new version over a weekend
- I can add features “as and when” I want them
- Labs don’t get the “back porting” concept → only latest TL

Decided happy to continue for years to come but then ...



AIX Labs:  
“nmon will be  
assimilated!”

# Code stuffing



Lines of code from 6500 → 9000

## nmon on POWER6 & AIX6 + New Features for V12

### User Requested

1. Disk Service Times ★
2. Selecting Particular Disks ★
3. Time Drift ★
4. Multiple Page Sizes ★ 4KB, 64KB, ++
- ★ 5. Timestamps in UTC & no. of digits
- ★ 6. More Kernel & Hypervisor Stats \*
7. High Priority nmon

### Advanced, POWER6 and AIX6 items

8. Virtual I/O Server SEA via entstat
9. Partition Mobility (POWER6)
10. WPAR & Application Mobility (AIX6)
11. Dedicated Donating (POWER6)
12. Folded CPU count (SPLPAR)
13. Multiple Shared Pools (POWER6)
14. Fibre Channel stats via fcstat

### Housekeeping items

15. Bug fixes – small or fine tuning, see nmon wiki for a list
16. Network packet sizes now saved to file
17. Warnings from nmon like: network overflow
18. TOP Online - sizes in KB, MB, GB to keep columns aligned
19. Fast abort in capture mode – SIGUSR1 during config collection stops nmon
20. Return an error & on invalid options – return codes to check it finished normally
21. Adapter "not available" work around – see nmon FAQ wiki about these AIX bugs
22. EMC hdiskpower renamed to "power" to save screen space
23. NFS v4 for AIX 5.3 ML5+

```

nmon12  —b=Black&White— Host=ssc10— Refresh=8 secs— 14:01.31
Workload Partition (WPAR)
Number of WPARs= 10
WPAR-Name      Type   State | Physical-CPU|Memory-in-MB |
User System|InUse numperm| RunQ pSwitch Fork
wp06           Sys   Defined |              |              |
wp07           Sys   Defined |              |              |
wp08           Sys   Defined |              |              |
wp09           Sys   Defined |              |              |
wp05           Sys   Active | 0.00  0.00| 20.2   3.7| 0.0    3   0.0
wp04           Sys   Active | 0.00  0.00| 19.7   3.7| 0.0    3   0.0
wp02           Sys   Active | 0.00  0.00| 22.9   4.7| 0.0    3   0.0
wp10           Sys   Defined |              |              |
wp03           Sys   Active | 0.27  0.00| 21.2   3.9| 1.0    24  0.0
wp0001        Sys   Active | 0.00  0.00| 18.8   0.3| 0.0    3   0.0

Work-Load-Manager (WLM)
Classes=5   Mode=Active
Class Name  |---Used-Percent---| |---Desired---| |---Shares---| |Proc's T I L
wp04        | 0.2  0.7  0.0 100 98 100 | -1  -1  -1 | 0 0 0 0
wp05        | 0.0  0.8  0.0 100 98 100 | -1  -1  -1 | 0 0 0 0
wp02        | 0.0  0.9  0.0 100 98 100 | -1  -1  -1 | 0 0 0 0
wp03        | 34.0 0.8  0.0 100 98 100 | -1  -1  -1 | 0 0 0 0
wp0001      | 0.0  0.7  0.0 100 98 100 | -1  -1  -1 | 0 0 0 0
Total percentage 34.2 3.8 0.0
Tier, Inheritance & LocalShm=T I L

Top-Processes-(171)  Mode=3 [1=Basic 2=CPU 3=Perf 4=Size 5=I/O 6=Cmnds]
PID      %CPU ResSize Char      WLM-Class Command
Used      KB      I/O
565290   27.4   168    0          wp03 ncpu
569394   0.4    9700   0          System nmon
90156    0.1    36096  0          System wlmshed
270484   0.1    20088  2          System /usr/sbin/xmwm -L
389266   0.1    3924   0          wp02 /opt/freeware/sbin/gmond -p /va
585808   0.1    3836   0          wp03 /opt/freeware/sbin/gmond -p /va
499750   0.1    3836   0          wp04 /opt/freeware/sbin/gmond -p /va

```

App/Sys

Defined/Active

Usr/Sys% CPU

RAM use & FScache

RunQ/pSwich and Fork

Dynamic WPAR class'

Default No WLM control

Top Procs WPAR name

System/Default Etc =Global AIX

## Direct Fibre-Channel and VIOS SEA

Shared Ethernet with 'O'

Fibre Channel with '^' → also regular LPAR & FC tape drives

- Actually using more expensive entstat and fcstat commands

```
nmon12-----h=Help-----Host=silver_vios1---Refresh=2 secs---10:03.36
```

---

```
Shared-Ethernet-Adapter
```

Number	Adapter Name	Receive KB/s	Transmit KB/s	-Requests-		--Size KB--	
				In	Out	In	Out
1	ent3	141.4	1.4	211.6	22.5	67.0	1024.0
	Totals	0.0	0.0 MB/s	21.6	22.5		

---

```
Fibre-Channel-Adapter
```

Number	Adapter Name	Receive KB/s	Transmit KB/s	-Requests-		--Size KB--	
				In	Out	In	Out
1	fcs0	9945.8	0.0	420.2	0.0	17.9	
	Totals	0.0	0.0 MB/s	0.0	0.0		

---

## I am internal beta tester for AIX 6.1 TL02 (Nov 2008)

- I do product introduction need access to “new stuff”
- Installed AIX 6.1 TL02 & was using nmon for a couple of days before I realised it was topas\_nmon!!

- Syntax, toggle and layout perfect
- “Gob-smacked”



# nmon → topas\_nmon

```
-----  
N   N   M   M   OOOO   N   N   For online help type: h  
NN  N  MM  MM  O   O  NN  N   For command line option help:  
N N  N  M  MM  M  O   O  N N  N   quick-hint  nmon -?  
N  N  N  M   M  O   O  N  N  N   full-details nmon -h  
N   NN  M   M  O   O  N   NN   To start nmon the same way every time?  
N   N  M   M  OOOO   N   N   set NMON ksh variable, for example:  
-----                               export NMON=cmt  
  
TOPAS-NMON  
  
4 - CPUs currently  
4 - CPUs configured  
4204 - MHz CPU clock rate  
PowerPC_POWER6 - Processor  
64 bit - Hardware  
64 bit - Kernel  
1,silver_lpar3 - Logical Partition  
6.1.2.0 TL02 - AIX Kernel Version  
silver_lpar3 - Hostname  
silver_lpar3 - Node/WPAR Name  
Serial Number  
IBM,8203-E4A - Machine Type
```

# topas nmon Install

Clarification:  
Do nothing as it is  
**Installed by Default**  
with AIX



## topas nmon Command

Typical place  
users but the  
nmon shell script

\$ nmon

→ script /usr/bin/nmon

→ binary /usr/bin/topas\_nmon

→ hard link to topas & topasrec

\$ topas

→ and hit ~

\$ nmon -f -s 300 -c 288

Just works as normal

## Menu based access to nmon functions

smitty topas, websm or pConsole

All run off the same template

- Performance & Resource → Configure Topas

# Using: smitty topas

```
# smitty topas
```

```
-----
Add Host to topas external subnet search file (Rsi.hosts)
List hosts in topas external subnet search file (Rsi.hosts)
List active Recordings
Start New Recording
Stop Recording
List completed recordings
Generate Report
```

```
Start Persistent local recording
Start Persistent CEC recording
Start local recording
Start CEC recording
```

```
Period
Day
Hour
Custom
```

```
Binary
nmon
```

**Start now  
& once only**

**Via Cron  
every day**

```
Start nmon recording
[Entry Fields]
Type of Recording
nmon
* Length of Recording
custom
* Recording Interval in seconds [300] #
* Number of Samples [288] #
* Overwrite existing recording no +
Output Path []
Running Priority [] #
* Recording time zone Local time+
Include top processes Include top process+
```

```
Start Persistent XXXXX recording


Type or select values in entry fields.
Press Enter AFTER making all desired
changes.

[Entry Fields]
Type of Recording CEC
Length of Recording
persistent
* Recording Interval in seconds [60] #
* Number of Days to store per file [1] #
* Number of Days to retain [31] #
Output Path []
* Overwrite existing recording file no +
```

## File Post Processing

So do the following work?

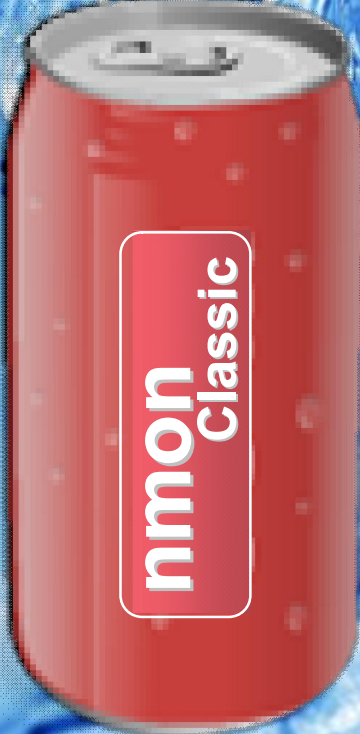
- nmon analyser
- nmon consolidator
- nmon2rrd → C program to generate graphs + html
- nmon2web → Perl program - ditto



**I assume so,  
users have the  
code to fix it!**

## Smitty and pConsole

- smitty topas or pConsole
- Run off the same template
  - Performance & Resource → Configure Topas



## nmon Support

### Use new Official nmon

- AIX 5.3 TL09+ & AIX 6.1 TL02+
- You have **full** AIX Support with SWMA
- Raise a PMR

### Older AIX releases use nmon Classic

- Please ONLY version 12e+
- Google: AIX “Performance Tool Forum”



## nmon Development



nmon classic – may be a 12f then stop  
ALL new development in topas\_nmon



nmon is now  
**GREEN**  
fully recycled  
source code



No animals  
harmd in the  
testing of nmon



# Blatant Advert for Hands-On Movies



## AIX6 & Power6 - the movie!

**10 on WPAR**

**8 on AIX**

**7 on POWER6 - SPLPAR/SMT**

**10 on IVM**

**5 on Cool stuff**

**15 on PowerVM**

Each movie 5 to 40 minutes


Downloadable files in .wmv format

Watch actual use of these new features

In total ~11 hours

Google™ → AIX movie

## Search Google for: AIX Movie

- Ten on Workload Partitions (WPAR)
- Ten on Integrated Virtualisation Manager (IVM)
- AIX6 Role based Access Control (RBAC)
- AIX6 JFS2 with the no log option
- AIX6 JFS2 with in-line snapshot option
- AIX6 Probevue new wave monitoring
- AIX6 Encrypted File System
- AIX6 System Director pConsole
- AIX6 Security with aixpert
- AIX6 Security with fpm
- POWER6 Decimal Floating Point (DFP)
- HMC 7 Hardware Management Console
- POWER6 Partition Mobility
- POWER6 Memory Protection Keys
- POWER6 Partition Priority
- POWER6 Multiple Shared CPU Pools
- POWER6 Shared CPU & Pool Monitoring
- Ganglia Performance Monitoring
- Linux for Power PowerVM Lx86 Translator
- nmon - Get You Started 
- nmon - version 12
- nmon now in AIX
- POWER Introduction
- HMC Introduction
- Create the first Logical Partition (LPAR)
- Installing AIX from CDROM
- Installing AIX from NIM
- Dynamically changing LPARs
- Ethernet Options with Virtual Ethernet
- Virtual Ethernet Setup
- Disk Options with Virtual SCSI
- Virtual Disk Setup
- How many CPUs in a Power Chip?
- Are Logical CPUs Real? What is SMT?
- What is a Virtual Processor? And controlling Shared Processor LPARs?
- VIOS 2.1 Principles
- VIOS 2.1 Features

## **Topas 2 secrets:**

- 1) nmon interworking**
- 2) Virtual I/O Server  
specific statistics**

topas -C

```
Topas CEC Monitor                      Interval: 10                      Thu Nov 13 04:06:54 2008
Partitions Memory (GB)                 Processors
Shr: 4   Mon: 8.0  InUse: 6.1  Shr: 2  PSz: 4   Don: 0.0  Shr PhysB 0.02
Ded: 0   Avl:   -           Ded: 0  APP: 4.0  Stl: 0.0  Ded PhysB 0.00

Host      OS  M  Mem InU Lp  Us  Sy  Wa  Id  PhysB  Vcsw Ent  %EntC PhI  pmem
-----
silver_lpar3 A61 UM 2.0 1.6 2 45 0 0 55 0.01 210 0.50 1.1 0 -
silver_lpar2 A61 UM 2.0 1.3 2 72 0 0 27 0.00 219 0.50 0.9 0 1.03
silver_lpar4 A61 UM 2.0 1.6 2 0 0 0 99 0.00 218 0.50 0.8 0 0.97
silver_lpar5 A61 UM 2.0 1.6 2 10 0 0 89 0.00 230 0.50 0.7 0 1.15
silver8      A61 U 2.0 1.3 4 0 1 98 0.01 927 0.50 2.2 4 -
```

topas

```
Topas Monitor for host: orange
Mon Mar 16 08:04:05 2009 Interval: 2

CPU  User%  Kern%  Wait%  Idle%  Phyc
1    0.0    33.1    0.0    66.9    0.00
0    19.4    62.2    0.0    18.5    0.01
2    0.0    27.5    0.0    72.5    0.00
3    0.0    29.3    0.0    70.7    0.00

Network  KBPS    I-Pack  O-Pack  KB-In  KB-Out
Total    0.7     2.5     2.5     0.2    0.5

Disk     Busy%    KBPS    TPS  KB-Read  KB-Writ
Total    0.0     0.0     0.0  0.0     0.0

FileSystem  KBPS    TPS  KB-Read  KB-Writ
Total       0.0     0.0  0.0     0.0

EVENTS/QUEUES  FILE/TTY
Cswitch        184  Readch        0
Syscall         71  Writech       322
Reads           0  Rawin         0
Writes          1  Ttyout       322
Forks           0  Igets         0
Execs           0  Namei         2
Runqueue        1.0  Dirblk        0
Waitqueue        0.0

MEMORY
Real,MB        2048
% Comp         64.4
% Noncomp      21.1
% Client       21.1

PAGING
Faults         0
Steals         0
PgspIn         0
PgspOut        0
PageIn         0
PageOut        0
Sios           0

PAGING SPACE
Size,MB        512
% Used         1.1
```

```
topas_nmon Host=orange Refresh=2 secs 08:01.24
CPU-Utilisation-Small-View EntitledCPU= 0.40 UsedCPU= 0.008
Logical CPUs 0-----25-----50-----75-----100
CPU User% Sys% Wait% Idle%|
0 0.0 0.0 0.0 100.0| >
1 0.0 0.0 0.0 100.0| >
2 0.0 0.0 0.0 100.0|>
3 0.0 0.0 0.0 100.0|>
EntitleCapacity/VirtualCPU +-----+-----+-----+-----+
EC 0.4 1.0 0.0 0.7|-----|-----|-----|-----|
VP 0.1 0.2 0.0 0.1|-----|-----|-----|-----|
EC= 2.1% VP= 0.4% +--No Cap--|--Folded=1--|-----100% VP=2 CPU+
```

nmon  
→ topas\_nmon



# VIOS monitoring via topas

Hit "E" Virtual Ethernet including SEA

```
Topas Monitor for host:   bronze_ivm  Interval:    2   Fri Dec  5 08:08:19 2008
=====
Network                  KBPS      I-Pack    O-Pack    KB-In    KB-Out
ent10 (SEA)              31.3      43.0      43.0      15.9     15.4
  | \--ent2 (VETH)       15.9      15.5      27.0      14.3     1.6
  | \--ent0 (PHYS)       15.5      27.5      16.0      1.6     13.9
lo0                       0.0       0.0       0.0       0.0     0.0
```

Hit "D" and then "d" Virtual SCSI over the VIOS

```
Topas Adapter View      :   bronze_ivm  Interval:    2   Fri Dec  5 08:29:17 2008
=====
Adapter                 KBPS      TPS      KB-R      KB-W
sissas0                 9.9K      92.0     0.0       9.9K
vhost0                   0.0       0.0     0.0       0.0
vhost1                   9.9K     184.0    92.0      92.0
=====
Vtargets/Disks         Busy%     KBPS     TPS      KB-R     ART     MRT     KB-W     AWT     MWT     AQW     AQD
hdisk0                  0.0       0.0     0.0     0.0     0.0     3.6     0.0     0.0     62.6    0.0     0.0
hdisk1                  25.0     9.9K    92.0     0.0     0.0     3.5     9.9K    4.2     28.9    0.0     0.0
hdisk2                  0.0       0.0     0.0     0.0     0.0     0.0     0.0     0.0     0.0     0.0     0.0
hdisk3                  0.0       0.0     0.0     0.0     0.0     0.0     0.0     0.0     0.0     0.0     0.0
hdisk4                  0.0       0.0     0.0     0.0     0.0     0.0     0.0     0.0     0.0     0.0     0.0
```

## And Finally the End

### Google: nmon wiki

- Download
- FAQ & Manual
- Analyser / Consolidator
- Performance Tools Forum
  - Questions mostly about nmon 😊

### Google: AIX Movies

- nmon and loads more

Run the latest version of nmon

- Changing one file is hardly an upgrade!

Uncapped Shared Processor LPARs (SPLPAR) monitor  
Physical CPU use

→ **Not Utilisation**