



nmon & topas for Performance Monitoring

Presentation Version 21



Nigel Griffiths
IBM Power Systems
Advanced Technology Support, Europe

© 2012 IBM Corporation

Let talk about

nmon

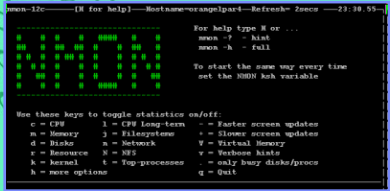
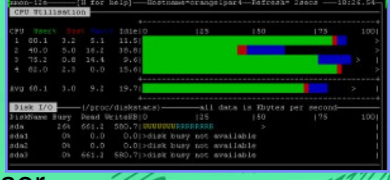
Year 16

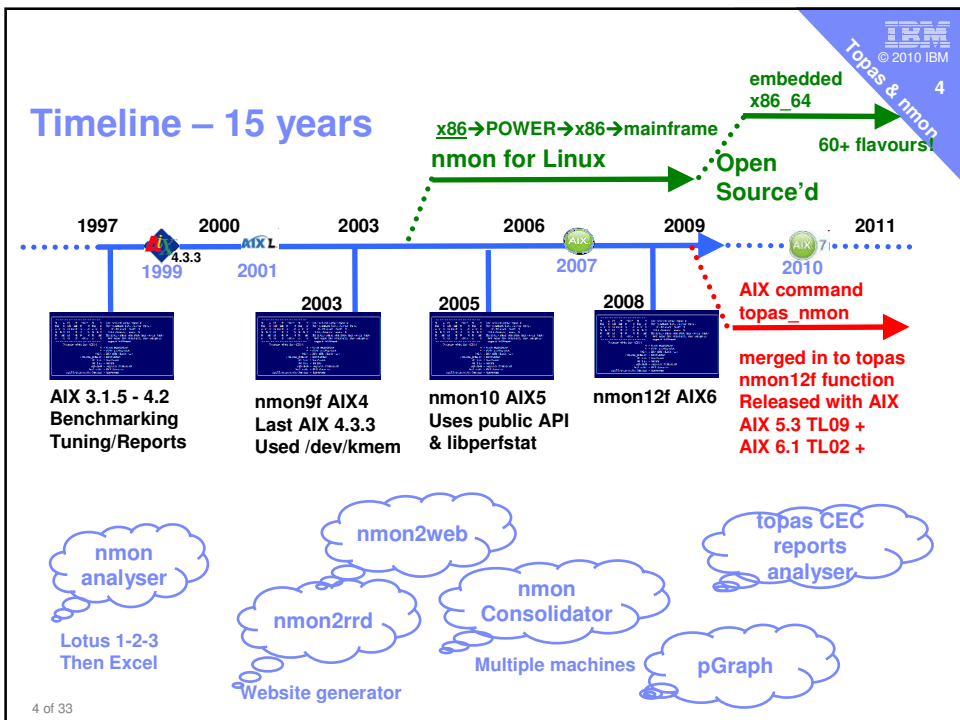
2 of 33

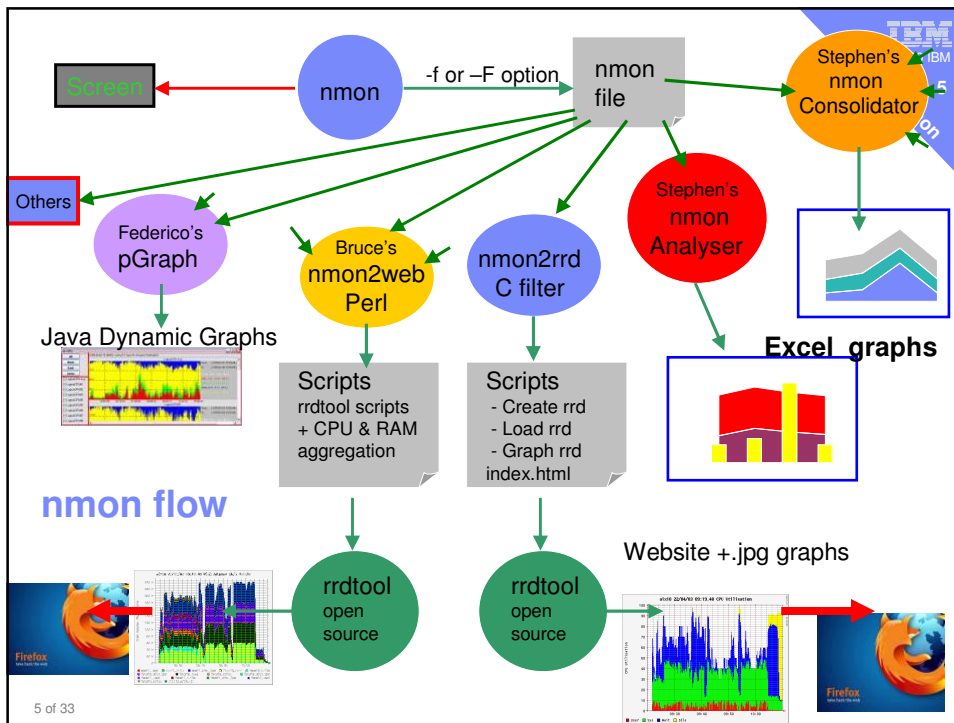


nmon for Linux v14g now Open Source

- nmon for Linux on Power then added x86, Z, x86_64
- Released by IBM - July 2009
 - <http://nmon.sourceforge.net>
 - Binaries for 66 Linux flavours Largely thanks for MPG
- Compile for your Linux flavour
 - Cell blades, disk engines,
 - Diskless, embedded ++
 - 6000 lines of C code → 200 KB
- Fully supported by the nmon analyser
- On a Linux Distro's?
 - Ubuntu/Debian (get-apt), IBM on PowerLinux, x86 & secret, ...





Topas & nmon
© 2010 IBM
6

nmon is half the story

**The other half is due to the nmon Analyser
Excel spreadsheets from Stephen Atkins, UK**

Another personal time project

The sole designer/developer:

- nmon analyser
- nmon Consolidation
- Topas CEC analyser

6 of 33

Email 6: What 1st things do I look for in an nmon file?

- Firmware level (BBBP lscnf)
- AIX level - release, TL and SP (BBBP oslevel -s)
- Not using classic nmon (AAA)

- If any over 2 years old the nmon user are either a newbie or a time waster.
- From this I know how professional they are!!

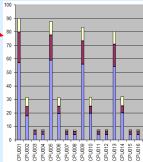
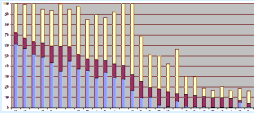
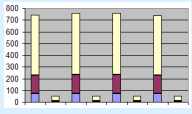
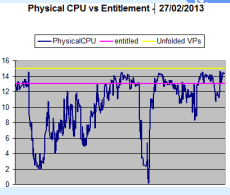
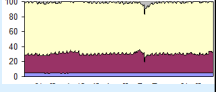
- If you don't service your car for three+ years:
Why should the car maker spend time investigating your problem or take the blame for the top speed !

- Then period of time covered and the interval
 - Long intervals hide peaks

- Then what version of the Analyser was used
 - Only use the latest 3.4a

1 of 3





Then What 2nd things do I look for in an nmon file?

- LPAR tab physical CPU use
 - Shared Uncapped – how often over Entitlement
 - Most utilisation numbers are pointless now IMHO
- CPU_SUMM tab 
 - Shows use of SMT
- DISKBUSY tab 
 - No disks over 50%
 - Most disks active – i.e. not old school disk layout
- IOADAPT tab 
 - Disk I/O even across FC adapters
- MEM tab 
 - Check free memory is NOT large. If it is – tune to use it
- MEMNEW tab 
 - Is filesystem cache use as expected. RDBMS DIO/CIO should be small

2 of 3

Then What 3rd things do I look for in an nmon file?

- NET tab
 - Is it pretty!
 - Does it peak at a well know network speed?
- Page tab
 - Should be low but do see high filesystem paging these days = memory mapped files.
- POOL tab
 - See if Pool is exhausted
- Proc
 - Compare RunQ with SMT thread available
- Top - if present
 - Check is the top CPU processes expected



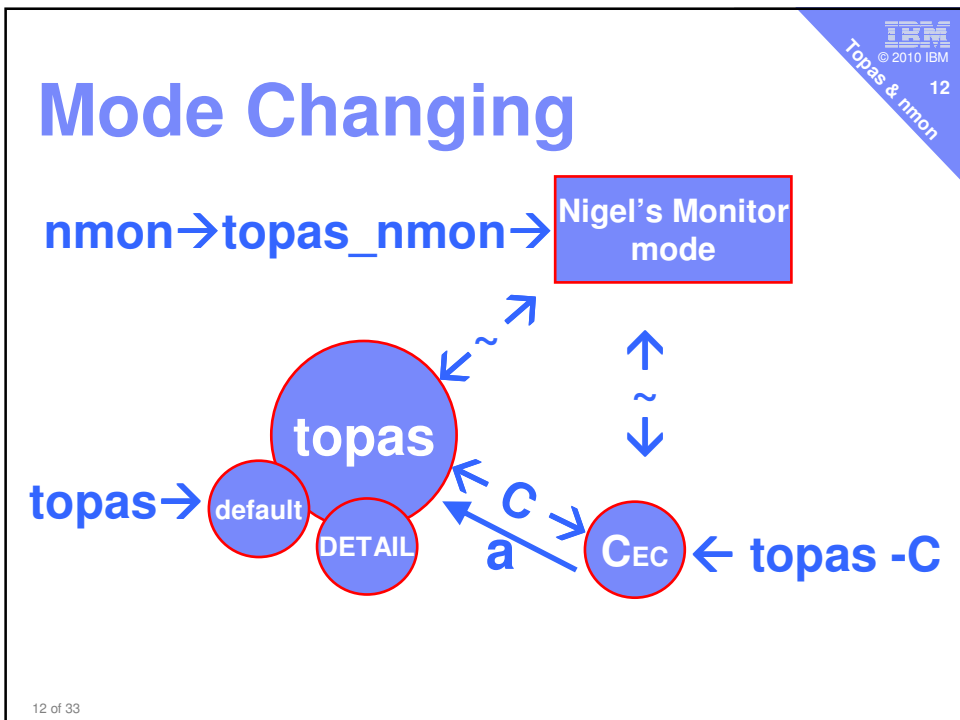

3 of 3

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

10 of 33



Let talk about

topas

IBM
© 2010 IBM

13

Topas & rmon

13 of 33

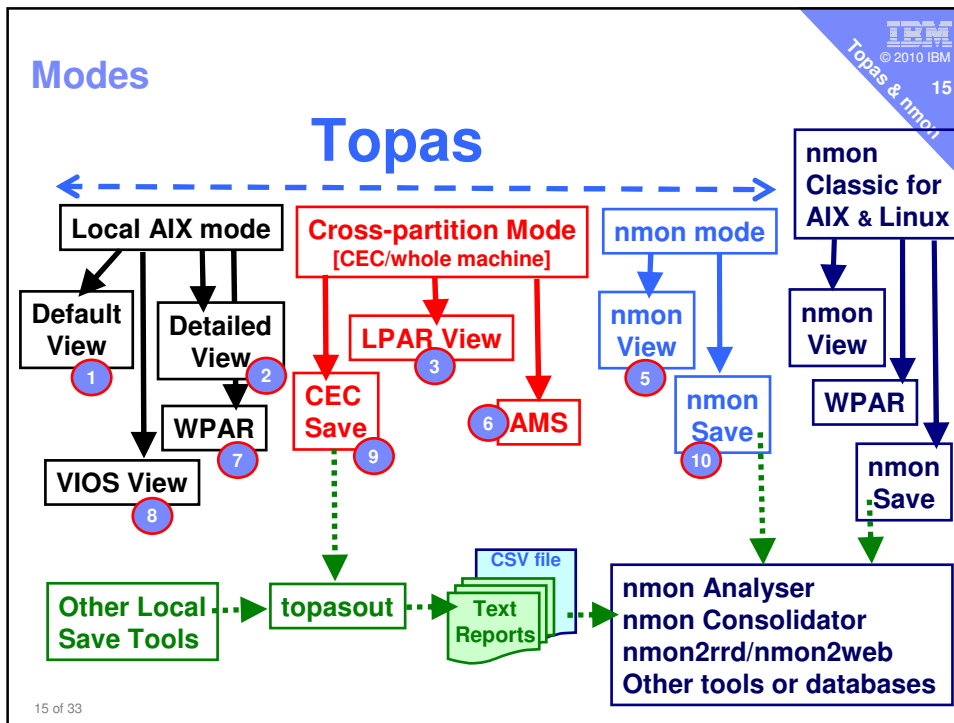
topas problems

No marketing
No Redbook
Manual is 30+ pages

10 different online modes
7 file formats (with friends)

KB-W
AQW
MWT
PSz
Vcsw
%bdon
MPuse
hpit
physb
pmem

RTFM



Topas Local AIX

16 of 33



© 2010 IBM
Topas & nmon
17

```

Topas Monitor for host: diamond6          EVENTS/QUEUES  FILE/TTY
Sat Jan 30 16:44:11 2010 Interval: 2      Cswitch 744  Readch 19.0K
                                           Syscall 23472 Writech 12391
CPU  User% Kern% Wait% Idle% Physc  Entc  Reads 5148 Rawin 0
ALL  51.0  11.6  0.0  37.5  0.86  107.7  Writes 52 Ttyout 98
                                           Forks 16 Igets 0
Network KBPS I-Pack O-Pack KB-In KB-Out Execs 16 NameI 107K
Total 5.8 16.9 14.0 2.7 3.1 Runqueue 2.0 Dirblk 0
                                           Waitqueue 0.0

Disk  Busy% KBPS  TPS KB-Read KB-Write  MEMORY
Total 0.0 2.0 0.0 0.0 2.0  PAGING  Real,MB 4096
                                           Faults 15850 % Comp 47
Filesystem KBPS  TPS KB-Read KB-Write  Steals 0 % Noncomp 15
Total 19.0K 5.0K 19.0K 12.0  PgsPIn 0 % Client 15
                                           PgsPOut 0
Name  PID  CPU% PgSp Owner  PageIn 0 PAGING SPACE
java  426214 1.0 373.1 root  PageOut 0 Size,MB 512
inetd  299154 0.8 3.5 root
topas  442408 0.6 4.1 root
xmtopas 393364 0.3 2.0 root
syncd  143522 0.1 0.6 root
xmgc  45078 0.0 0.4 root
getty  278682 0.0 0.6 root
rll  114744 0.0 0.9 root
  
```

```

Topas Monitor for host: diamond6          EVENTS/QUEUES  FILE/TTY
Sat Jan 30 17:12:04 2010 Interval: 4      Cswitch 644  Readch 16.2M
                                           Syscall 20883 Writech 10919
CPU  User% Kern% Wait% Idle% Physc  Entc  Reads 4564 Rawin 0
ALL  51.0  11.6  0.0  37.5  0.86  107.7  Writes 47 Ttyout 31
                                           Forks 14 Igets 0
Network KBPS I-Pack O-Pack KB-In KB-Out Execs 14 NameI 94K
Total 5.8 16.9 14.0 2.7 3.1 Runqueue 1.2 Dirblk 0
                                           Waitqueue 0.0

Disk  Busy% KBPS  TPS KB-Read KB-Write  MEMORY
Total 0.0 2.0 0.0 0.0 2.0  PAGING  Real,MB 4096
                                           Faults 14282 % Comp 47
Filesystem KBPS  TPS KB-Read KB-Write  Steals 0 % Noncomp 15
Total 19.0K 5.0K 19.0K 12.0  PgsPIn 0 % Client 15
                                           PgsPOut 0
Name  PID  CPU% PgSp Owner  PageIn 0 PAGING SPACE
java  426214 1.0 373.1 root  PageOut 0 Size,MB 512
inetd  299154 0.8 3.5 root
topas  442408 0.6 4.1 root
xmtopas 393364 0.3 2.0 root
syncd  143522 0.1 0.6 root
xmgc  45078 0.0 0.4 root
getty  278682 0.0 0.6 root
rll  114744 0.0 0.9 root
  
```

cndfp

17 of 33

```

Topas Monitor for host: diamond6          Interval: 4 Sat Jan 30 18:04:00 2010
USER  PID  PPID PRI NI RES  RES SPACE  TIME  PRI  I/O  VTH  COMMAND
root  426214 1 80 20 95525 22 95525 19:15 4.2 0 0 java
root  303296 421990 84 0 17 3 30 0:01 4.7 0 0 ncdu
root  311540 303296 82 0 16 3 29 0:01 3.7 0 0 ncdu
root  229154 217234 85 20 891 11 891 27:49 0.4 0 2212 inetd
root  315936 299154 64 20 455 25 455 0:00 0.2 0 547 xmtopas
root  442422 413838 58 41 1173 92 1173 0:00 0.0 0 0 topas
root  278682 1 60 20 147 21 147 0:15 0.0 0 0 getty
root  119744 0 37 41 240 0 140 0:15 0.0 0 0 gll
root  229412 1 60 20 112 0 112 0:19 0.0 0 0 random
root  143822 1 60 20 148 2 148 0:18 0.0 0 19 syncd
root  1 0 60 40 180 11 180 0:00 0.0 0 0 init
ocsonsole 372944 344306 82 20 9770 18 9770 0:10 0.0 0 0 java
root  295142 1 60 20 304 0 304 0:05 0.0 0 0 rpc.Lock
root  250014 217234 60 20 1031 176 1031 0:00 0.0 0 0 sendmail
root  274592 217234 60 20 861 10 861 0:00 0.0 0 0 syslogd
root  110646 0 36 41 112 0 112 0:00 0.0 0 0 netcm
root  286866 217234 60 20 1373 53 1373 0:00 0.0 0 0 aixnsmbd
root  45078 0 60 41 112 0 112 0:20 0.0 0 0 xmgc
root  262276 1 60 20 783 6 783 0:05 0.0 0 0 slp_srvr
  
```

```

Topas Monitor for host: diamond6          Interval: 4 Sat Jan 30 19:03:04 2010
VolumeGroup  TPS  Kb_read  Kb_wrth  Kbps
rootvg  228  8 28804 14.1K
LogicalVolume  TPS  Kb_read  Kb_wrth  Kbps
hd1  0  0  0  0.0
hd10opt  0  0  0  0.0
hd2  0  0  0  0.0
hd3  2  4  0  2.0
hd4  0  0  4  28800 14.1K
hd5  0  0  0  0.0
hd6  0  0  0  0.0
hd8  1  0  4  2.0
hd9var  0  0  0  0.0
lvg_dumpv  0  0  0  0.0
livedump  0  0  0  0.0
  
```

```

Topas Monitor for host: diamond6          Interval: 4 Sat Jan 30 17:59:48 2010
Network  KBPS  I-Pack  O-Pack  KB-In  KB-Out
eth0  0.2  2.6  0.0  0.2  0.1
eth1  0.0  0.0  0.0  0.0  0.0
  
```

```

Topas Monitor for host: silverveer01 Interval: 2 Sun Jan 31 01:31:39 2010
Link  Busy% KBPS  TPS KB-R  KB-W  APT  BPT  RD-W  RD-W  A2W  A2D
eth0  67.0  23.0K 184.0 0.0 0.0 7.2 23.0K 3.0 26.5 0.0 0.0
eth1  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  24.4  0.0  0.0
eth2  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
eth3  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
eth4  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
eth5  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
eth6  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
eth7  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
eth8  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
eth9  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
eth10 0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
eth11 0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
eth12 0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
eth13 0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
eth14 0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
eth15 0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
  
```

```

Topas Monitor for host: diamond6          Interval: 4 Sat Jan 30 17:59:20 2010
Filesystem  KBPS  TPS  KB-R  KB-W  Open  Create  Lock
usr  14.4K  3.7K 14.4K  0.0  1K  0  0
tmp  3.1K  675.0  2.1K  9.0  689  0  285
var  0.0  0.0  0.0  0.0  0  0  0
/home  0.0  0.0  0.0  0.0  0  0  0
/var/adb/cas/livedump  0.0  0.0  0.0  0.0  0  0  0
admin  0.0  0.0  0.0  0.0  0  0  0
proc  0.0  0.0  0.0  0.0  0  0  0
opt  0.0  0.0  0.0  0.0  0  0  0
  
```

E
D
V
F

18 of 33

Topas VIOS view

topas VIOS 2.1+ SEA Network Monitoring


Hit "E" for shared Ethernet adapter(s) → SEA

Direction of traffic in this case from Virtual Ethernet out of the VIOS

- ent2 virtual 14KB/s incoming
- ent0 = physical = 14 KB/s outgoing
- ent10 = SEA = 14KB/s in and 14 KB/s out

```

Topas Monitor for host:  bronze_ivm Interval:  2  Fri Dec  5 08:08:19 2008
=====
Network
ent10 (SEA)              KBPS    I-Pack  O-Pack  KB-In  KB-Out
| \--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
  
```



Another Example two SEA + HEA

```


Topas Monitor for host:  bluevio44  Interval:  2  Mon Nov  9 18:28:12 2009
=====
Network
ent5 (SEA)                KBPS    I-Pack  O-Pack  KB-In   KB-Out
| \--ent1 (PHYS)         30.9    35.5    35.0    15.5    15.5
| \--ent4 (VETH)         15.5    25.0    10.5    14.7    0.8
| \--ent4 (VETH)         15.5    10.5    24.5    0.8     14.6
ent6 (SEA)                1.9     15.5    16.0    0.9     1.0
| \--ent0 (HEA PHYS)     1.0     15.5    0.5     0.9     0.1
| \--ent3 (VETH)         0.9     0.0     15.5    0.0     0.9
lo0                       0.0     0.0     0.0     0.0     0.0
  
```

topas VIOS 2.1+ Storage Monitoring

Hit "D" and then "d" → You can see VIOS virtual/physical Disks
And which hosts adapters is generating the I/O

```

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  AOD
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
  
```



CEC Whole machine Cross Partition All LPAR View

```

Topas CEC Monitor          Interval: 10          Sun Jan 31 02:12:49 2010
Partitions Memory (GB)      Processors
Shr: 5   Mon:11.9   InUse: 9.5   Shr:2.5   PSz: 4   Don: 0.0   Shr_PhysB 0.02
Ded: 0   Avl: -    Ded: 0   APP: 4.0   St1: 0.0   Ded_PhysB 0.00

Host      OS   M   Mem InU Lp   Us Sy Wa Id   PhysB  Vcsw Ent  %EntC Phi  pmem
-----
-----shared-----
silvrvios1  A61 U 2.0 0.8 2   0 0 0 99   0.01 284 0.50 1.4 0 -
silvrvios2  A61 U 1.9 1.4 4   0 0 0 99   0.01 427 0.50 1.3 2 -
silver5     A61 UM 3.0 3.0 2   0 0 0 99   0.01 217 0.50 1.2 0 1.59
silver8     A61 U 2.0 1.2 4   0 0 0 99   0.01 0 0.50 1.0 0 -
silver4     A61 UM 3.0 3.0 8   0 0 0 100  0.00 0 0.50 0.0 0 1.12

Host      OS   M   Mem InU Lp   Us Sy Wa Id   PhysB  Vcsw %st1 %bst1
-----
-----dedicated-----

```

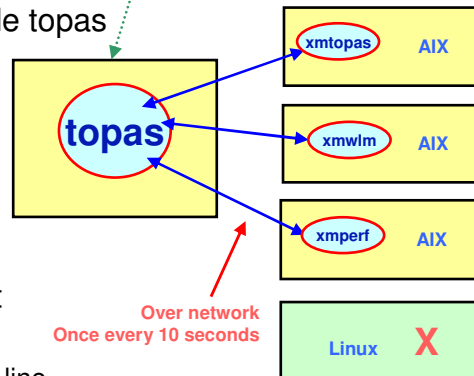
23 of 33

Cross Partition View

- LPAR: topas -C
- VIOS
 - topas then hit "C"
 - topas -ceview (ugly)
- Totally different program inside topas



Once only
- Can take 20 seconds
- Newer version polls the local network for AIX LPARs



Over network
Once every 10 seconds

- Only LPARs on same sub-net
 - Edit \$HOME/Rsi.hosts
 - Fully hostname.domain one per line

24 of 33

Active Memory Sharing view

AMS has a renewed focus due to new De-Duplication features

LPAR: vmstat -h

```
# vmstat -h 10
```

System configuration: lcpu=2 mem=2048MB ent=0.50 mmode=shared apsz=4.00GB

memory		page		faults		cpu		hypv-page			
r b	ava fra	re pi po fr sr cy in sy cs us sy id wa pc ec	hpi hpl	pmem loan							
0 0	190419 173073	0 0 0 0 0 0	2 149 159	0 1 99 0 0.01 1.6	0 0	1.20 0.80					
0 0	190419 173073	0 0 0 0 0 0	2 24 152	0 0 99 0 0.00 0.8	0 0	1.20 0.80					
0 0	190419 173073	0 0 0 0 0 0	1 19 166	0 0 99 0 0.00 0.8	0 0	1.20 0.80					
0 0	207225 189696	0 0 0 0 0 0	6 334 196	35 2 64 0 0.18 36.9	25 53	1.33 0.67					
0 0	207227 189694	0 0 0 0 0 0	2 39 170	50 1 50 0 0.25 50.8	7 15	1.33 0.67					
0 0	207227 189694	0 0 0 0 0 0	5 20 164	50 1 50 0 0.25 50.8	0 0	1.33 0.67					

Logical Memory Statistics

Monitoring 1 LPAR for AMS is "insane"

Hypervisor Page-Ins Faults/s

Time waiting for hypervisor page-ins (in milliseconds)

Physical Memory → pmem

Loaned Memory

If ams_loan_policy=0 (off) this will be zero

topas -C AMS → Have to monitor all LPARS

No details for Linux on POWER or IBM i = no daemon

Regular extra "g" details

```

Topas CEC Monitor          Interval: 10          Wed Dec 3 10:15:06 2008
Partition Info  Memory (GB)  Processor  Virtual Pools : 0
Monitored : 4  Monitored : 8.0  Monitored :2.0  Avail Pool Proc: 3.7
UnMonitored: -  UnMonitored: -  UnMonitored: -  Shr Physical Busy: 0.28
Shared : 4  Available : -  Available : -  Ded Physical Busy: 0.00
Uncapped : 4  UnAllocated: -  UnAllocated: -  Donated Phys. CPUs 0.00
Capped : 0  Consumed : 6.5  Shared : 2  Stolen Phys. CPUs : 0.00
Dedicated : 0  Dedicated : 0  Hypervisor
Donating : 0  Donated : 0  Virt. Context Switch: 976
Pool Size : 4  Phantom Interrupts : 1
    
```

Host	OS	M	Mem	InU	Lp	Us	Sy	Wa	Id	PhysB	Vcsw	Ent	%EntC	Phi	pmem
-----shared-----															
silver_vios1	A61	U	2.0	1.9	8	08	1	0	90	0.01	873	0.40	2.6	0	-
silver_lpar2	A61	UM	2.0	1.4	2	49	1	0	49	0.26	244	0.50	51.3	1	1.30
silver_lpar3	A61	UM	2.0	1.6	2	0	0	0	99	0.01	294	0.50	1.4	0	0.78
silver_lpar4	A61	UM	2.0	1.7	2	0	0	0	99	0.01	220	0.50	1.3	0	0.87
silver_lpar5	A61	UM	2.0	1.7	2	0	0	0	99	0.01	218	0.50	1.1	0	1.01

27 of 33

Logical Memory → Mem

Inu→InUse LPAR Working Set

Shared memory LPAR
Physical Memory → pmem

Dedicated
Memory
LPAR

AMS Pool Level: topas -C then hit "m"

```

Topas CEC Monitor          Interval: 10          Wed Dec 3 10:56:12 2008
Partitions  Memory (GB)  Memory Pool (GB)  I/O Memory (GB)
Mshr: 4  Mon: 8.0  InUse: 6.8  Mpsz: 4.0  MPuse: 4.0  Entl: 308.0Use: 47.9
Mded: 0  Avl: 1.2  Pools: 1
    
```

mpid	mps	mpus	mem	memu	iome	iomu	hpi	hpit
0	4.00	3.99	8.00	6.82	308.0	47.9	32	31

Pool Size

Pool Used

Machine memory

Machine memory Used

In use I/O Memory

I/O Memory → assign for device driver DMA I/O

For more on I/O Memory see:
• lparstat -m
• lparstat -me

Hypervisor Page-in/second

Hypervisor Page-in Time (msecs)

28 of 33

topasout

Takes binary topas recording files /etc/perf/...

Reports for

- LPAR,
- All Partitions or
- Whole Machine

In formats

- text report
- CSV data files

Manual page worse than topas ☹

topasout

Machine or LPAR - Data reports

- topasout -c|-s|-a /etc/perf/daily/hostname_YYMMDD.topas
- Example: topasout -a /etc/perf/daily/silver8_090304.topas
- The output overwrites the file /etc/perf/daily/hostname_YYMMDD.topas_01
 - c = comma separated →but one timestamp and value per line
 - s = spreadsheet →space separated 100's of stats per line
 - a = analyser →.csv file for nmon analyser – or – use topas_nmon directly

Machine or LPAR - Text reports

- topasout -R OPTION [-i MM -b HHMM -e HHMM] /etc/perf/daily/hostname_YYMMDD.topas
- OPTION = detailed|summary|poolinfo|disk |lan|mempool|adapter|vadapter|vios|vios_adapter
- Example: topasout -R summary /etc/perf/daily/silver8_090304.topas >/tmp/LPARsummary.txt
- Text reports via standard out, so you need to redirect the output to a file

Cross-Partition -Text reports

- topasout -R detailed | -R summary [-i MM -b HHMM -e HHMM] /etc/perf/topas_cec.YYMMDD
- Example: topasout -R detailed topas_cec.090305 >/tmp/CECdetailed.txt
- Text reports via standard out, so you need to redirect the output to a file

For system use only, do NOT run these manually.

- topasout -R daily | -R weekly

Notes: YYMMDD = year, month, day MM is minutes HHMM is hours & minutes

Whole machine performance graphs topasrec → topasout → CEC analyser (Excel)

```

Topas CEC Monitor          Interval: 10          Tue Dec 21 10:46:34 2010
Partitions Memory (GB)    Processors
Pbr: 9  Mon: 119 InUse:27.0 Shr: 8  PSz: 26  Don: 0.0 Shr PhysB 0.17
Ded: 1  Avi:  -          Ded: 6  APP: 25.8 St1: 0.0 Ded PhysB 0.00

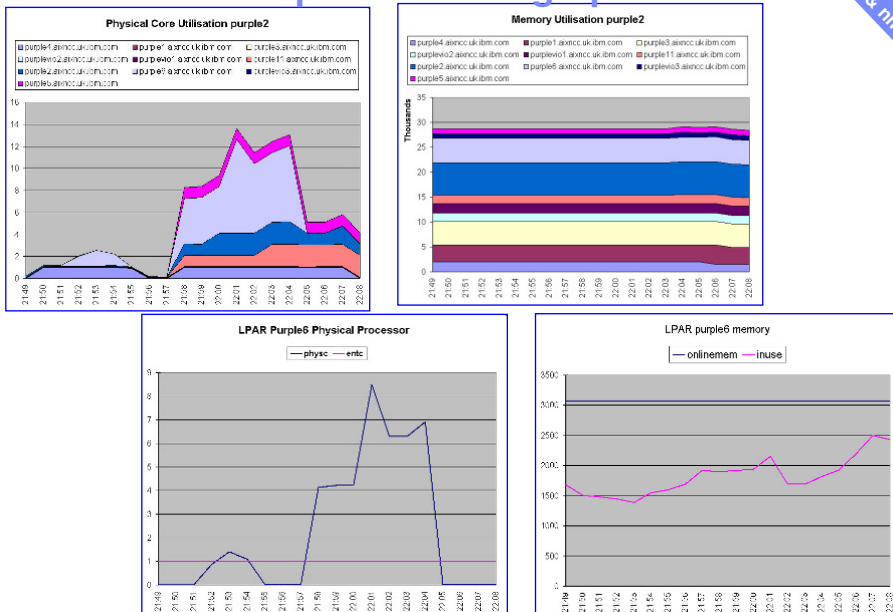
-----
Host      OS  Mod Mem InU Lp  Us Sy Wa Id  PhysB Vcsw Ent  EntC Phi  pmem
-----
purple2   A61 Ued 6.5 6.4 12  0  0  0 98  0.04 969  2.00  2.0  0  -
purplevio1 A61 Ued 2.0 1.8  8  0  3  0 95  0.04  0  0.50  7.6  0  -
purplevio2 A61 Ued 2.0 1.7  8  0  2  0 97  0.02  0  0.50  4.5  0  -
purple1   A61 Ued 4.0 3.3 32  0  0  0 99  0.02  0  2.00  1.0  0  -
purple6   A61 Ued 64 4.8 40  0  0  0 99  0.01  0  1.00  1.5  0  -
purple11  A71 Ued 3.0 1.6  8  0  1  0 98  0.01 552  0.50  2.3  0  -
purplevio3 A61 Ued 2.0 0.9  8  0  1  0 98  0.01 425  0.50  2.3  0  -
purple4   A61 Ued 2.0 2.0  2  0  0  0 99  0.00 645  0.50  0.8  0  -
purple5   A53 U-d 2.0 1.0  2  0  0  0 99  0.00  0  0.50  0.6  0  -
  
```

topas -C working OK but want to record it

1. <http://tinyurl.com/topas-cec>
 - Download CEC Analyser from Stephen Atkins
2. `topasrec -C -o sample -s 120 -c 720` (every 2 minutes all day)
 - Creates sample_cec_101222_2149.topas
3. `topasout -a sample_cec_101222_2149.topas`
 - Creates sample_cec_101222_21492.topas.csv
4. Start CEC Analyser & load the .csv output file & you get →

31 of 33

Whole machine performance graphs



32 of 33

New Stuff



IBM
© 2010 IBM

Topas & nmon
33

1. nmon analyser improvements from Stephen Atkins
 - Make sure you use the latest version
 - <http://tinyurl.com/nmon-analyser>
2. Topas Active Memory Expansion
3. nmon Active Memory Expansion file output
4. Topas Memory Affinity
5. nmon gets WPAR
6. Too Many CPUs problem
7. nmon online disk columns
8. nmon file output more stats

33 of 33

purple7.aixncc.uk.ibm.com

Active Memory Expansion - AME

- Rapidly growing adoption
- Nigel's opinion
 - “It's a no brainer & here to stay”
 - “Future POWER processors will ... err!
Blow your socks off!”


34 of 33

IBM
© 2010 IBM

Topas & nmon
34

Active Memory Expansion Conceptual Model

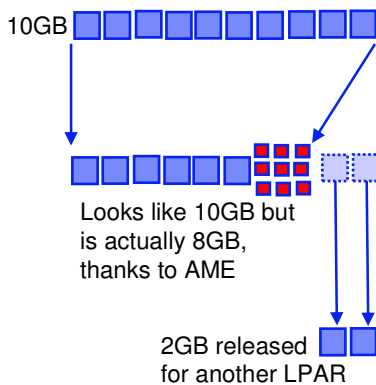
Memory Pages




Dynamically adjusted depending
on compression ratio & target

What is your Plan?

Memory Shrinking
to release RAM for other uses



Memory Growing
for RAM optimisation & performance



topas

```

Topas Monitor for host:  diamond3
Fri Apr 23 16:50:11 2010 Interval: 2

CPU  User%  Kern%  Wait%  Idle%  Physc  Entc
ALL  45.9   21.3   1.3    31.5   0.63   78.5

Network  KBPS  I-Pack  O-Pack  KB-In  KB-Out
Total    7.4   26.4    27.4    2.4    4.9

Disk  Busy%  KBPS  TPS  KB-Read  KB-Writ
Total  28.0  1287.2  321.0  1283.2    4.0

FileSystem  KBPS  TPS  KB-Read  KB-Writ
Total        2.0  19.4    2.0    0.0

Name  PID  CPU%  PgSp  Owner
nmem64  442598  19.6  100.3  root
cmemd   36882  12.5   0.2  root
nmem64  360456  10.7  100.3  root
lrud    16392   6.0   0.1  root
nmem64  450560  4.3  100.3  root
nmem64  446464  3.5  100.3  root
nmem64  376908  3.0  100.3  root
java    237688  0.1  58.9  root
rmd     381132  0.1   9.2  root
xclock  401516  0.1   3.6  root
Xvnc    295052  0.0   7.8  root
.java   352468  0.0  39.1  pconsole

EVENTS/QUEUES  FILE/TTY
Cswitch 1048  Readch 1137
Syscall 881  Writech 3686
Reads   4    Rawin   0
Writes  14   Ttyout  309
Forks   0    Igets   0
Execs   0    Namei   3
Runqueue 0.5  Dirbik  0
Waitqueue 0.0

MEMORY
Real.MB 1280
% Faults 7924
% Comp 98
Steals 7904
% Noncomp 0
PgspIn 316
% Client 0
PgspOut 0
PageIn 320
PAGING SPACE
PageOut 0
Size.MB 512
Size 321
% Used 3
% Free 97

AME
TMEM.MB 1024  WPAR Activ 0
CMEM.MB 181  WPAR Total 0
EF(T/A) 1,2/1,2  Press: "h"-help
CI:7.2K CO:7.4K  "q"-quit
  
```

TMEM = True Memory
CMEM = Compressed Memory
CI = Compressed Page In
CO = Compressed Page Out

EF = Expansion Factor
T = Target
A = Actual

37 of 33

nmon capture to file then Analyser

- MEM tab
 - Size of the Compressed pool (MB)
 - Size of true memory (MB)
 - Expanded memory size (MB),
 - Size of the Uncompressed pool (MB)
- MEMNEW tab
 - Compressed Pool%
- PAGE tab
 - Paging rates (pages per second) but this time very quickly in & out of the compressed memory area
 - Compressed pool pgins - other tools like topas call this CI
 - Compressed pool pgouts - other tools like topas call this CO

38 of 33

Topas AME - Very small change

```

Topas Monitor for host:green3
Thu Oct 6 17:48:40 2011 Interval:2
EVENTS/QUEUES FILE/TTY
Cvrtch 290 Readch 3703
Syscall 219 Wrtch 287
CPU User% Kern% Wait% Idle% Physc Entc% Reads 29 Ravin 0
Total 99.6 0.2 0.0 0.0 2.00 199.91 Writes 0 Tryout 287
Network BPS I-Pkts O-Pkts B-In B-Out Excs 0 Name1 7
Total 1.20K 10.00 1.50 836.0 392.0 Runqueue 9.00 Dirblk 0
Waitqueue 0.0
Disk Busy% BPS TPS B-Read B-Writ MEMORY
Total 0.0 0 0 0 0 PAGING Real, MB 2048
Faults 0 % Comp 53
FileSystem BPS TPS B-Read B-Writ Steals 0 % Noncomp 40
Total 3.62K 28.50 3.62K 0 PgsIn 0 % Client 40
PgsOut 0
WLM-Class (Active) CPU% Mem% Blk-I/O% PageIn 0 PAGING SPACE
System 0 65 0 PageOut 0 Size, MB 512
Shared 0 5 0 Sio 0 % Used 1
% Free 99
Name PID CPU% PgsP Class AME.....
ncpu 7536784 25.0 148K wp06 TMEM 2.00GPPAR Activ 2
ncpu 9044056 0.0 144K wp06 CMEM 512.00GPPAR Total 2
ncpu 8388632 0.0 108K wp06 EF [1/A] 1.0/1.0Press: "h"-help
ncpu 9306134 0.0 136K wp06 CI: OCO: 0 "q"-quit
  
```

- Units have moved
 - TMEM, MB 2000 → TMEM 2.00G

Memory Affinity & SRAD

- See my session
 - PE07
 - POWER7, AIX, Affinity and Performance

topas -M or
topas and the hit M or
nmon then ~ then M

Power 770 again ...

```
# lssrad -av
REF1  SRAD      MEM      CPU
0
    0  25054.75   0-11 28-31 40-43 56-59 72-75
    3   6705.50   52-55 68-71
1
    1  17679.00   12-15 20-23 32-35 44-47 60-63 76-79 104-107
    2  14193.00   24-27 36-39 48-51 64-67
```

Yes it is a bit of a mess to highlight some things

```
Topas Monitor for host: purple1 Interval: 2 Mon Aug 22 04:24:30 2011
-----
REF1  SRAD  TOTALMEM  INUSE  FREE  FILECACHE  HOMETHRDS  CPUS
-----
0     0    24.5G    6919.9  17.7G  90.7        224        0-11 28-31 ...
    3    6705.5   2051.4  4654.1  17.9    359        52-55 68-71
1     2    13.9G    4118.7   9.8G  49.5        379        24-27 36-39 ...
    1    17.3G    4784.2  12.6G   61.8    217        12-15 20-23 ...
-----
CPU   SRAD  TOTALDISP  LOCALDISP%  NEARDISP%  FARDISP%
-----
0     0     70         85.7        11.4        2.9
1     0     62         100.0       0.0         0.0
2     0     12         100.0       0.0         0.0
3     0     32         100.0       0.0         0.0
4     0     23         95.7        0.0         4.3
5     0     52         100.0       0.0         0.0
6     0     36         100.0       0.0         0.0
7     0     20         100.0       0.0         0.0
8     1     67         79.1        20.9        0.0
9     1     31         93.5        6.5         0.0
10    1     18         77.8        22.2        0.0
11    1     11         100.0       0.0         0.0
12    1     62         38.7        61.3        0.0
13    1     65         52.3        47.7        0.0
14    1     14         100.0       0.0         0.0
15    1     15         60.0        40.0        0.0
16    2     85         76.5        2.4         21.2
17    2     56         82.1        17.9        0.0
18    2     23         100.0       0.0         0.0
19    2     20         95.0        5.0         0.0
```

You know this bit now
Note: different order to lssrad

Topas Monitor for host: purple1 Interval: 2 Mon Aug 22 04:24:30 2011

IBM © 2010 IBM Topas & nmon 43

REF1	SRAD	TOTALMEM	INUSE	FREE	FILECACHE	HOMETHRDS	CPUS
0	0	24.5G	6919.9	17.7G	90.7	224	0-11 28-31 ...
	3	6705.5	2051.4	4654.1	17.9	359	52-55 68-71
1	2	13.9G	4118.7	9.8G	49.5	379	24-27 36-39 ...
	1	17.3G	4784.2	12.6G	61.8	217	12-15 20-23 ...

CPU	SRAD	TOTALDISP%	LOCALDISP%	NEARDISP%	FARDISP%
0	0	197	26.9	69.0	4.1
1	0	30	73.3	26.7	0.0
2	0	7	100.0	0.0	0.0
3	0	33	100.0	0.0	0.0
4	0	70	85.7	11.4	2.9
5	0	62	100.0	0.0	0.0
6	0	12	100.0	0.0	0.0
7	0	32	100.0	0.0	0.0
8	0	23	95.7	0.0	4.3
9	0	52	100.0	0.0	0.0
10	0	36	100.0	0.0	0.0
11	0	20	100.0	0.0	0.0
12	1	67	79.1	20.9	0.0
13	1	31	93.5	6.5	0.0
14	1	18	77.8	22.2	0.0
15	1	11	100.0	0.0	0.0
20	1	62	38.7	61.3	0.0
21	1	65	52.3	47.7	0.0
22	1	14	100.0	0.0	0.0
23	1	15	60.0	40.0	0.0
24	2	85	76.5	2.4	21.2
25	2	56	82.1	17.9	0.0
26	2	23	100.0	0.0	0.0
27	2	20	95.0	5.0	0.0

SRAD Memory View

CPUs but limited room

Home processes/threads

File system cache RAM (numperm)

Topas Monitor for host: purple1 Interval: 2 Mon Aug 22 04:24:30 2011

IBM © 2010 IBM Topas & nmon 44

REF1	SRAD	TOTALMEM	INUSE	FREE	FILECACHE	HOMETHRDS	CPUS
0	0	24.5G	6919.9	17.7G	90.7	224	0-11 28-31 ...
	3	6705.5	2051.4	4654.1	17.9	359	52-55 68-71
1	2	13.9G	4118.7	9.8G	49.5	379	24-27 36-39 ...
	1	17.3G	4784.2	12.6G	61.8	217	12-15 20-23 ...

CPU	SRAD	TOTALDISP%	LOCALDISP%	NEARDISP%	FARDISP%
0	0	197	26.9	69.0	4.1
1	0	30	73.3	26.7	0.0
2	0	7	100.0	0.0	0.0
3	0	33	100.0	0.0	0.0
4	0	70	85.7	11.4	2.9
5	0	62	100.0	0.0	0.0
6	0	12	100.0	0.0	0.0
7	0	32	100.0	0.0	0.0
8	0	23	95.7	0.0	4.3
9	0	52	100.0	0.0	0.0
10	0	36	100.0	0.0	0.0
11	0	20	100.0	0.0	0.0
12	1	67	79.1	20.9	0.0
13	1	31	93.5	6.5	0.0
14	1	18	77.8	22.2	0.0
15	1	11	100.0	0.0	0.0
20	1	62	38.7	61.3	0.0
21	1	65	52.3	47.7	0.0
22	1	14	100.0	0.0	0.0
23	1	15	60.0	40.0	0.0
24	2	85	76.5	2.4	21.2
25	2	56	82.1	17.9	0.0
26	2	23	100.0	0.0	0.0
27	2	20	95.0	5.0	0.0

**SMT=4 → 1 physical core
Are all logical CPUs in use?**

Mostly Near memory access 😊

Most work on 1st SMT

Move the cursor here to order the Logical CPUs

Topas New: Affinity stats → M (memory!)

Power 770

LPAR = 64 GB
Across 3 SRAD's

CPU's are
0- 7 → 8 = 2
8-11 → 4 = 1
16-23 → 4 = 1
VP = 4

VP=10 across 16 CPUs
-We don't have a SRAD
for 10 CPUs
VP=10 SMT=4 40 VCPUs

```

Topas Monitor for host: purple6 Interval: 2 Thu May 27 07:35:32 2010
-----
EF1 SRAD TOTALMEM INUSE FREE FILECACHE HOMETHRDS CPUS
-----
0 0 30.6G 3431842.9 28.8G 295 87.7 305 0-7
1 1 18.8G 1921206.2 17.7G 180 54.0 200 8-11
2 2 12.6G 129841.1 11.8G 120 36.1 192 20-23

CPU SRAD TOTALDISP LOCALDISP% NBRDISP% FARDISP%
-----
12 0 106 2.8 97.2 0.0 Excellent
0 0 102 37.3 23.5 37.2 Good
24 0 91 4.4 0.0 95.6 Blisteringly
16 1 90 42.2 0.0 57.6 Fast
4 0 54 64.8 35.2 0.0
28 1 33 39.4 0.0 60.6
8 1 27 56.3 0.0 43.7
32 2 27 96.3 3.7 0.0
20 2 15 93.3 6.7 0.0
17 1 9 88.9 0.0 11.1
33 2 9 88.9 11.1 0.0
9 1 9 88.9 0.0 11.1
29 1 8 87.5 0.0 12.5
21 2 8 87.5 12.5 0.0
26 0 1 100.0 0.0 0.0
27 0 1 100.0 0.0 0.0
36 0 1 100.0 0.0 0.0
37 0 1 100.0 0.0 0.0
38 0 1 100.0 0.0 0.0
39 0 1 100.0 0.0 0.0
3 0 1 100.0 0.0 0.0
10 1 1 0.0 0.0 100.0
11 1 1 0.0 0.0 100.0
7 0 1 100.0 0.0 0.0
2 0 1 100.0 0.0 0.0
35 2 1 0.0 100.0 0.0
6 0 1 100.0 0.0 0.0
18 1 1 0.0 0.0 100.0
19 1 1 0.0 0.0 100.0
    
```

45 of 33

Workload Partition - WPAR

Was in nmon classic but not topas_nmon

- Now in topas_nmon 4Q 2011

- Below is topas

Below is topas -@ (WPAR view)

```

Topas Monitor for host:green3
Thu Oct 6 19:51:31 2011 Interval:2
-----
CPU 0% 0% 0% 0% 0% 0% 0% 0% 0% 0% 0% 0% 0% 0% 0% 0% 0% 0% 0% 0%
Total 99.7 0.3 0.0 0.0 2.00 199.90
Network 1.59K 17.50 1.00 1.20K 390.8
Total 1.59K 17.50 1.00 1.20K 390.8
Disk Busy 0 TPS B-Read B-Writ
Total 0.0 0 0 0
Filesystem 2.4K 28.0K 2.4K 0
Total 2.4K 28.0K 2.4K 0
WPAR 0 Mem Bk-I/O
wp0 0 2 0
wp06 100 4 0
-----
Name PID CPU PPRU WPAR
ncpu 10551514 25.0 152K wp06
ncpu 4063342 25.0 128K wp06
ncpu 7530784 25.0 148K wp06
ncpu 7602390 25.0 128K wp06
ncpu 7667810 0.0 116K wp06
ncpu 8388932 0.0 108K wp06
ncpu 9306134 0.0 136K wp06
ncpu 9044056 0.0 144K wp06
ncpu 2145338 0.0 100K Global
topas 9371880 0.0 2.24K Global
finched 1572912 0.0 76.0K Global
cpu-load 3094118 0.0 244K Global
cload 7202000 0.0 1.52M wp02
cload 7884546 0.0 1.52M wp06
    
```

```

Topas Monitor for host:green3 Interval: 2 Thu Oct 6 19:54:49 2011
-----
WPAR 0 Mem Bk-I/O
wp06 99.78 4.01 0.00
-----
USER PID PPRU PRI NI RES RES SPACE TIME CPU I/O OTH COMMAND
root 4003342 10551514 162 0 52.0K 12.0K 128K 40:38 25.0 0 0 ncpu
root 10551514 10551514 161 0 52.0K 12.0K 128K 41:28 25.0 0 0 ncpu
root 9044056 10551514 162 0 52.0K 12.0K 144K 41:27 25.0 0 0 ncpu
root 9306134 10551514 162 0 53.0K 12.0K 136K 41:32 25.0 0 0 ncpu
root 7530784 10551514 162 0 48.0K 12.0K 148K 41:27 0.0 0 0 ncpu
root 8388932 10551514 162 0 56.0K 12.0K 108K 41:13 0.0 0 0 ncpu
root 7667810 10551514 162 0 56.0K 12.0K 116K 41:06 0.0 0 0 ncpu
root 7602390 10551514 162 0 60.0K 12.0K 168K 41:04 0.0 0 0 ncpu
root 9371880 10092648 60 20 2.30M 716K 2.30M 0:00 0.0 0 0 topas
root 7884546 5636174 156 20 1.52M 240K 1.52M 0:00 0.0 0 0 cload
root 7202000 7077948 60 20 1.52M 240K 1.52M 0:00 0.0 0 0 cload
    
```

Move to WPAR name & hit f
to see its processes

nmon WPAR stats are back in AIX 7 TL1/AIX6 TL7

- WPAR hit @ & WLM hit W

Memory stats
numperm=
filesystem cache

Kernel stats
Process switches
Fork new process

Type =
System or App
State =
Active or Defined

Includes the resource
constraints & current
values

```

-topas nmon -k=Kernel-stats -Host=green3 Refresh=2 secs 00:07.04
CPU-Utilisation-Small-View EntitledCPU= 1.00 UsedCPU= 1.555
Logical CPUs 0-----25-----50-----75-----100
CPU User% Sys% Wait% Idle%|
0 33.3 1.3 0.0 65.4|
1 42.3 0.0 0.0 57.7|
2 0.0 0.0 0.0 100.0|>
3 0.0 0.0 0.0 100.0|>
4 61.5 1.3 0.0 37.2|
5 12.8 0.0 0.0 87.2|
6 0.0 0.0 0.0 100.0|>
7 0.0 0.0 0.0 100.0|>
EntitleCapacity/VirtualCPU +-----+-----+-----+-----+
EC+ 62.9 0.9 0.0 36.2|
VP 48.9 0.7 0.0 28.1|
EC= 15.5% VP= 77.8% +-No Cap---|-SMT=4-----100% VP=2 CPU+

Workload Partition (WPAR)
Number of WPARs 2
WPAR-Name Type State Physical-CPU Memory-in-MB
wp06 Sys Active 0.00 0.01 64.3 35.9 0.0 14 0.0
wp02 Sys Active 0.96 0.00 32.7 8.5 2.6 54 0.0

Work-Load-Manager (WLM)
Classes=2 Mode=Active
Class Name CPU MEM Rsets ProcTotal ClassTotal
|---Used-Percent---|---Desired---|---Shares---|Proc's T I L
wp06 0.5 4.5 0.0 100 98 100 -1 -1 -1 20 0 0 0
wp02 52.4 2.2 0.0 100 98 100 -1 -1 -1 21 0 0 0
Total percentage 52.9 6.7 0.0 Tier, Inheritance & LocalShm=T I L
    
```

nmon WPAR stats are back in AIX 7 TL1/AIX6 TL7

- Not just WLM hit W

Busy with
WPAR work

These are WLM classes
and WPAR stats
appear to be missing

```

-topas nmon -W=WLM -Host=green3 Refresh=2 secs 17:08.11
CPU-Utilisation-Small-View EntitledCPU= 1.00 UsedCPU= 1.999
Logical CPUs 0-----25-----50-----75-----100
CPU User% Sys% Wait% Idle%|
0 20.0 1.0 0.0 79.0|
1 79.5 0.0 0.0 20.5|
2 0.0 0.0 0.0 100.0|>
3 0.0 0.0 0.0 100.0|>
4 100.0 0.0 0.0 0.0|
5 100.0 0.0 0.0 0.0|
6 0.0 0.0 0.0 100.0|>
7 0.0 0.0 0.0 100.0|>
EntitleCapacity/VirtualCPU +-----+-----+-----+-----+
EC+ 79.5 0.3 0.0 20.2|
VP 79.5 0.3 0.0 20.2|
EC= 199.9% VP= 100.0% +-No Cap---|-SMT=4-----100% VP=2 CPU+

Work-Load-Manager (WLM)
Classes=5 Mode=Active
Class Name CPU MEM Rsets ProcTotal ClassTotal
|---Used-Percent---|---Desired---|---Shares---|Proc's T I L
Unclassified 0.0 8.5 0.0 100 98 100 -1 -1 -1 0 0 0 0
Unmanaged 0.0 30.6 0.0 100 99 100 -1 -1 -1 0 0 0 0
Default 0.0 8.1 0.0 100 98 100 -1 -1 -1 3 0 0 0
Shared 0.0 5.2 0.0 100 98 100 -1 -1 -1 0 0 0 0
System 0.2 64.7 0.0 100 99 100 -1 -1 -1 79 0 0 0
Total percentage 0.2 86.5 0.0 Tier, Inheritance & LocalShm=T I L
    
```


Inside a WPAR

- Bonkers – like driving a car with your eyes closed

```
Topas Monitor for host:green3  
Thu Oct 6 04:52:15 2011 Interval:2  
Kernel 88.95 ##### |  
User 11.05 ##### |  
Wait 0.00 |  
Idle 0.00 |  
Phyc = 0.00 |  
%Resc= 0.28  
Network BPS I-Pkts O-Pkts B-In B-Out |  
Total 1.00K 1.00 1.00 40.03 984.8 |  
FileSystem BPS TPS B-Read B-Writ |  
Total 7.89K 5.00 7.89 0 |  
Name PID CPU% PgSp |  
topas 6488072 95.2 1012K |  
ksh 10027132 3.8 572K |  
clcomd 7864546 0.7 1.52M |  
sendmail 3997892 0.3 1.08M |  
srcmstr 5636174 0.0 760K |  
init 1 0.0 732K |  
rpc.stat 6684796 0.0 1.43M |  
errdaemon 6881304 0.0 664K |  
EVENTS/QUEUES FILE/TTY  
Cswitch 7 Readch 9293  
Syscall 453 Writech 930  
Reads 6 Rawin 0  
Writes 1 Ttyout 927  
Forks 0 Igets 0  
Execs 0 Namei 21  
unqueue 0 Dirblk 0  
Waitqueue 0.0  
PAGING MEMORY  
Faults 568 % Comp 53  
Steals 0 % Noncomp 40  
PgspIn 0 % Client 40  
PgspOut 0  
PageIn 0 PAGING SPACE  
PageOut 0 Size,MB 512  
Sios 0 % Used 1  
% Free 99  
AME  
THEM 2.00Press: "h"-help  
CMEM 512.00 "q"-quit  
EF(T/A) 1.0/1.0  
CI: OCO: 0
```

Relative to WPAR CPU limit (like %ent)
- Not documented !

49 of 33

Too Many CPUs

Power 795 = 256 POWER7


SMT=4 → 4 Logical CPUs
make it even worse

```
Logical CPU  
0 0.0 51.5 0.0 48.0#####>  
1 0.0 51.9 0.0 73.9#####>  
2 0.0 56.1 0.0 73.9#####>  
3 0.0 56.3 0.0 73.9#####>  
4 0.0 54.8 0.0 45.2#####>  
5 0.0 58.0 0.0 43.7#####>  
6 0.0 58.1 0.0 41.3#####>  
7 0.0 55.1 0.0 41.3#####>  
8 48.0 46.4 0.0 5.4#####>  
9 0.0 51.0 0.0 96.0#####>  
10 0.0 54.4 0.0 96.0#####>  
11 0.0 51.5 0.0 96.0#####>  
12 0.0 50.7 0.0 49.1#####>  
13 0.0 59.1 0.0 49.1#####>  
14 0.0 59.4 0.0 40.8#####>  
15 0.0 48.9 0.0 3.1#####>  
16 0.0 50.0 0.0 45.9#####>  
17 0.0 56.6 0.0 43.4#####>  
18 0.0 56.3 0.0 43.7#####>  
19 0.0 51.7 0.0 40.1#####>  
20 0.0 50.8 0.0 49.2#####>  
21 0.0 59.7 0.0 40.1#####>  
22 0.0 59.2 0.0 40.1#####>  
23 0.0 49.7 0.0 90.7#####>  
24 0.0 59.2 0.0 40.1#####>  
25 0.0 49.7 0.0 90.7#####>  
26 0.0 57.5 0.0 42.1#####>  
27 0.0 57.6 0.0 42.1#####>  
28 0.0 51.3 0.0 49.7#####>  
29 0.0 50.8 0.0 45.2#####>  
30 0.0 59.5 0.0 40.8#####>  
31 0.0 59.4 0.0 40.8#####>  
32 0.0 54.9 0.0 45.7#####>  
33 0.0 54.6 0.0 40.1#####>  
34 0.0 58.8 0.0 43.2#####>  
35 0.0 58.0 0.0 43.2#####>  
36 0.0 56.4 0.0 43.2#####>  
37 0.0 58.0 0.0 42.1#####>  
38 0.0 56.2 0.0 41.1#####>  
39 0.0 59.0 0.0 42.0#####>  
40 0.0 59.5 0.0 49.7#####>  
41 0.0 51.1 0.0 41.1#####>  
42 0.0 59.3 0.0 49.7#####>  
43 0.0 59.5 0.0 40.1#####>  
44 0.0 58.0 0.0 40.1#####>  
45 0.0 58.2 0.0 44.1#####>  
46 0.0 58.0 0.0 42.1#####>  
47 0.0 56.7 0.0 43.1#####>  
48 0.0 56.2 0.0 43.1#####>  
49 0.0 50.5 0.0 49.1#####>  
50 0.0 57.8 0.0 42.1#####>  
51 0.0 59.9 0.0 42.1#####>  
52 0.0 51.0 0.0 49.1#####>  
53 0.0 59.5 0.0 40.1#####>  
54 0.0 60.0 0.0 49.1#####>  
55 0.0 50.3 0.0 45.1#####>  
56 0.0 59.3 0.0 40.1#####>  
57 0.0 59.6 0.0 49.1#####>  
58 0.0 51.3 0.0 48.1#####>  
59 0.0 41.9 0.0 53.1#####>  
60 0.0 51.3 0.0 47.1#####>  
61 0.0 56.2 0.0 45.1#####>  
62 0.0 56.2 0.0 41.1#####>  
63 0.0 58.1 0.0 47.1#####>  
64 0.0 47.5 0.0 52.1#####>  
65 0.0 56.7 0.0 41.1#####>  
66 0.0 59.9 0.0 40.1#####>  
67 0.0 51.2 0.0 48.1#####>  
68 0.0 59.7 0.0 49.1#####>  
69 0.0 59.6 0.0 40.1#####>  
70 0.0 49.1 0.0 53.1#####>  
71 0.0 51.3 0.0 45.7#####>  
72 0.0 59.9 0.0 49.1#####>  
73 0.0 59.9 0.0 40.1#####>  
74 0.0 51.6 0.0 59.1#####>  
75 0.0 49.1 0.0 50.1#####>  
76 0.0 51.6 0.0 59.1#####>  
77 0.0 49.1 0.0 50.1#####>
```

50 of 33

Too many CPUs

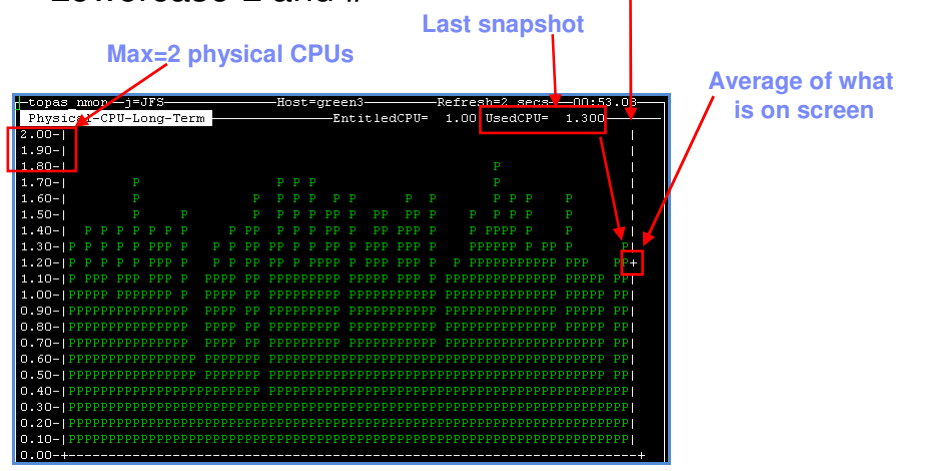
- Power 795 = 256 cores = 1024 logical CPUs
- Can't watch this on screen 1024 pixel screen!
- Actually logical CPUs not so interesting
 - Once we know they are all in use
- Want to see whole VM physical CPU use hit **l** (L)
- CPU Long term view
 - 70 snapshots
 - But tops out at 100%
 - Even if say 8 times over Entitlement



51 of 33

Too many CPUs


- Long term Physical CPU view
- Lowercase L and #




Stopped all work

- Long term Physical CPU view
 - Autoscaling


Max=2




Max=0.5



Max=0.031



Max=0.06



nmon file output

Started with: nmon -f

nmon More data saved for AIX6 TL7/AIX7 TL1

In the BBBP section = AIX Parameters

- **Uptime** → long AIX has been running
 - BBBP,000,uptime
 - BBBP,001,uptime," 01:13AM up19:18,2 users, load average: 1.91, 1.54, 1.18"
- **Issrad -av** → SRAD and CPU-Memory Affinity
 - BBBP,174,Issrad
 - BBBP,175,Issrad,"REF1 SRAD MEM CPU"
 - BBBP,176,Issrad,"0"
 - BBBP,177,Issrad," 0 1699.37 0-7"
- **Next service pack: emgr -l**
 - For determining if or when eFixes are installed
- **I forget when added but it includes**
 - **mpstat -b** output = Memory Affinity stats Local, Near, Far references
 - **ifconfig & netstat** – may have been there a while!!

nmon More data saved for AIX6 TL7/AIX7 TL1

Even more Disk Statistics

- Below are splitting out the old DISKXFERS to read and write
 - DISKRIO - Disk IO Read per second
 - DISKWIO - Disk IO Writes per second
- **Experimental**
 - DISKAVGRIO - Disk IO Average Reads per second
 - DISKAVGWIO - Disk IO Average Writes per second

nmon More data saved for AIX6 TL7/AIX7 TL1


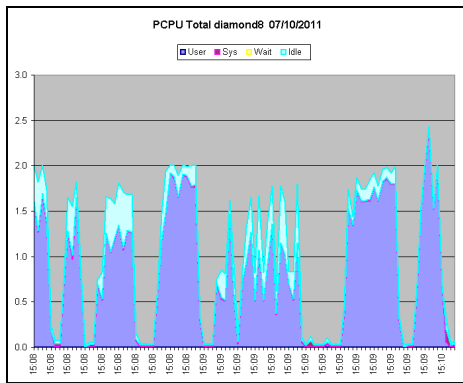
- CPU PURR and SPURR based stats
 - PCPU_ALL and PCPU<number 1 to logical processors>
 - SCPU_ALL and SCPU<number 1 to logical processors>
- Power Saving Mode via Systems Director AEM

On my POWER7 Power 750

```
# pmcycles -M                               ← Undocumented
This machine runs at 3550 MHz                 ← No power saving
--- Switch on Power saving here ---
# pmcycles -M
This machine runs at 2126 MHz                 ← Power saving mode
# pmcycles -M
This machine runs at 3862 MHz                 ← Over-clocking
```

57 of 33

topas_nmon PURR/SPURR for AIX7 TL1 on POWER7 E=0.8 SMT=4 VP=2 with Power Saving On via Systems Director AEM



Note: Goes over the VP. It has more clock ticks due to over clocking by ~9% - I think!!

58 of 33

Online Disks

nmon then
DDDD

- New columns



```

-topas nmon -p=Partitions -Host=green3 -Refresh=2 secs -21:05.44
Disk-KBytes/second-(K=1024,M=1024*1024)
Disk  Busy  Read  Write Transfers  Size Peak%  Peak KB/s qDepth
Name   KB/s   KB/s   /sec  KB      Read+Write or N/A
hdisk0 38%    0.0 12032.0  94.0 128.0  38%  12610.9  --
cd0    0%     0.0   0.0    0.0   0.0   0%   0.0      --
Totals (MB/s) Read=0.0 Write=11.8 Size(GB)=64 Free(GB)=56

-topas nmon -W=WLM -Host=green3 -Refresh=2 secs -21:05.52
Disk-KBytes/second-(K=1024,M=1024*1024)
Disk  Size Free Disk Disk Volume Disk
Name  GB  GB Paths Adapter Group Description
hdisk0 64 56 1 vscsi0 rootvg Virtual SCSI Disk Drive
cd0    0 0 0 vscsi0 None Virtual SCSI Optical Served
Totals (MB/s) Read=0.0 Write=0.0 Size(GB)=64 Free(GB)=56

-topas nmon -W=WLM -Host=green3 -Refresh=2 secs -21:07.32
Disk - Service times and Queues
Disk  Read Service Write Service Wait ServQ WaitQ ServQ
Name  milli-seconds milli-seconds Size Size Full
hdisk0 0.0 0.0 4.1 0.0 0.1 0.0 0.0
cd0    0.0 0.0 0.0 0.0 0.0 0.0 0.0
Totals (MB/s) Read=0.0 Write=14.1 Size(GB)=64 Free(GB)=56

-topas nmon -3=Top-by-CPU-use -Host=green3 -Refresh=2 secs -21:07.50
Disk-KBytes/second-(K=1024,M=1024*1024)
Disk  Busy  Read  Write 0-----25-----75-----100
Name   KB/s   KB/s |
hdisk0 43%    0 14464|#####>
cd0    0%     0 0|
Totals 0 14464+-----|-----|-----+
    
```

59 of 33

nmon

- While online hit [to start a capture and then] to stop

```

-topas nmon -A=&sync-I/O -Host=purple7 -Refresh=2 secs -04:29.58
CPU-Utilisation-Small-View EntitledCPU= 4.00 UsedCPU= 0.027
Logical CPUs 0-----25-----50-----75-----100
CPU User% Sys% Wait% Idle%|
0 0.0 2.0 0.0 98.0|s >
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| >
4 0.0 5.4 0.0 94.6|ss >
5 0.0 0.1 0.0 99.9| >
6 0.0 0.1 0.0 99.9| >
7 0.0 0.1 0.0 99.9| >
8 0.0 9.1 0.0 90.9|ssss >
9 0.0 1.0 0.0 99.0| >
10 0.0 1.0 0.0 99.0| >
11 0.0 1.0 0.0 99.0| >
12 0.0 0.0 0.0 100.0|>
13 0.0 0.1 0.0 99.9|>
14 0.0 0.1 0.0 99.9|>
15 0.0 0.1 0.0 99.9|>
16 0.0 2.2 0.0 97.8|s >
17 0.0 1.0 0.0 99.0| >
18 0.0 1.0 0.0 99.0| >
-On demand Recording started. Press '^]' to stop this recording \
    
```

60 of 33

topas

- Hit Space to FREEZE the output
 - if you want to study the numbers in detail

```
Topas Monitor for host: purple7-AIX71D-1136Bval: 2 Thu Oct 13 04:36:11 2011
-----
#PF1  SRAD  TOTALMEM  INUSE  FREE  FILECACHE  HOMETHRDS  CPUS
-----
0      0    7.46G    1.00G    6.46G    305M    293.0    0-7 12-15
1      1    5.59G    759M    4.85G    231M    299.0    8-11
2      2    1.95G    291M    1.66G    80.6M    133.0    24-27
-----
CPU    SRAD  TOTALDISP  LOCALDISP%  NEARDISP%  FARDISP%
-----
0      0    228.0    22.4    0.0    77.6
1      0    139.0    19.4    0.0    80.6
3      0    102.0    11.8    0.0    88.2
2      0    85.00    2.4    0.0    97.6
24     2    61.00    100.0    0.0    0.0
13     0    59.00    88.1    0.0    11.9
19     1    44.00    100.0    0.0    0.0
17     1    41.00    100.0    0.0    0.0
8      1    40.00    100.0    0.0    0.0
18     1    40.00    100.0    0.0    0.0
```