STARTING TO LOOK AT A PERFORMANCE PROBLEM

This presentation at:
http://www.circle4.com/papers/common-perfprobs.pdf

Jaqui Lynch
lynchj@forsythe.com

FORSYTHE

AGENDA

How to avoid performance crit sits!
Where to start when there is a problem
Performance Tools

FORSYTHE

AVOIDING PROBLEMS

FORSYTHE

UNDERSTAND YOUR WORKLOAD

Are you about speed?

- Speed is distance over time or performance
- Affected by clock speed, memory and I/O bandwidth, etc
- Basically how much can I push through one core
- Higher frequency cores
- May run better with SMT2 or SMT or dedicated cores

Or throughput?

- Volume over time or capacity
- How many concurrent things can I push through
- Affected by pipelining and SMT

Architect accordingly

Check for gating factors that could impact use of SMT

- i.e. is there one thread that controls all work?

FORSYTHE

APPLICATIONS AND SPLPARS

Applications do not need to be aware of Micro-Partitioning

Not all applications benefit from SPLPARs

Applications that may not benefit from Micro-Partitioning:

- Applications with a strong response time requirements for transactions may find Micro-Partitioning detrimental:
  - Because virtual processors can be dispatched at various times during a timeslice
  - May result in longer response time with too many virtual processors:
    - Each virtual processor with a small entitled capacity is in effect a slower CPU
  - Compensate with more entitled capacity (2-5% PUs over plan)
- Applications with polling behavior
- CPU intensive application examples: DSS, HPC, SAS

Applications that are good candidates for Micro-Partitioning:

- Ones with low average CPU utilization, with high peaks:
  - Examples: OLTP, web applications, mail server, directory servers

In general Oracle databases are fine in the shared processor pool

For licensing reasons you may want to use a separate pool for databases

5

FORSYTHE

UNDERSTAND SMT

SMT

- Threads dispatch via a Virtual Processor (VP)
- SMT1: Largest unit of execution work
- SMT2: Smaller unit of work, but provides greater amount of execution work per cycle
- SMT4: Smallest unit of work, but provides the maximum amount of execution work per cycle
- On POWER7, a single thread cannot exceed 65% utilization
- On POWER6 or POWER5, a single thread can consume 100%
- Understand thread dispatch order

- VPs are unfolded when threshold is reached
  - P5 and P6 primary and secondary threads are loaded to 80% before another VP unfolded
  - In P7 primary threads are loaded to 50%, then VPs are unfolded. Secondary threads are used when VPs are all in use for primary threads. When secondaries are loaded to 50% tertiary threads are dispatched

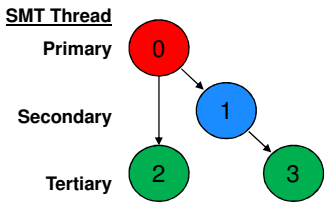

Diagram courtesy of IBM

6

FORSYTHE

ENTITLEMENT AND VPS

Utilization calculation for CPU is different between POWER5, 6 and POWER7
VPs are also unfolded sooner (at lower utilization levels than on P6 and P5)

This means that in POWER7 you need to pay more attention to VPs
- You may see more cores activated a lower utilization levels
- But you will see higher idle
- If only primary SMT threads in use then you have excess VPs

Try to avoid this issue by:
- Reducing VP counts
- Use realistic entitlement to VP ratios
  - 10x or 20x is not a good idea
  - Try setting entitlement to .6 or .7 of VPs
- Ensure workloads never run consistently above 100% entitlement
- Too little entitlement means too many VPs will be contending for the cores

- **Performance may (in most cases, will) degrade when the number of Virtual Processors in an LPAR exceeds the number of physical processors**

7

FORSYTHE

UNCAPPED VS CAPPED

Capped LPARs can cede unused cycles back but can never exceed entitlement
Uncapped LPARs can exceed entitlement up to the size of the pool or the total virtual processors, whichever is smaller
Unused capacity is ceded back
User defined weighting (0 to 255) is used to resolve competing requests
Weights are share based
- 2 LPARs need 3 cores each
- Only 3 cores available
- If A is 100 and B is 200 then A gets 1 core and B gets 2 cores

Use common sense when planning your use of weights and remember the default is128

- Prod VIO 192
- Prod 160
- Test/Dev 128

- Have a plan, not necessarily this one – document it well

8

FORSYTHE

General Server Sizing thoughts

Correct amount of processor power
Balanced memory, processor and I/O
Min, desired and max settings and their effect on system overhead
Memory overhead for page tables, TCE, etc that are used by virtualization
Shared or dedicated processors
Capped or uncapped
If uncapped – number of virtual processors
Do not starve your VIO servers!
Set entitlement and VPs correctly
Be cautious of sizing studies – they tend to undersize memory and sometimes cores and usually do not include the VIO server needs
Consider whether the workload will play well with shared processors
Never underestimate the power of common sense

Scale by rPerf (or other benchmark data) NOT by ghz when comparing boxes

9

FORSYTHE

Adapter Priorities affect Performance

Power 770 Layout		9117-MMC												
CEC		Top	123456 has GX cables			Bottom	2468ab			5877 pcie only I/O Drawer 123487				
	Slot	Desc	Pri	Alloc	Slot	Desc	Pri	Alloc	Slot	Desc	Pri	Alloc	IOC	
	C1	8GB DP fibre	1	lpar1	C1	8GB DP fibre	1	lpar1	C1	8GB DP fibre	1	vio1	1	
	C2	4PT 10/100/1000	3	lpar1	C2	4PT 10/100/1000	3	lpar1	C2	4PT 10/100/1000	3		1	
	C3	8GB DP fibre	5	vio2	C3	8GB DP fibre	5	vio1	C3		5		1	
	C4	4PT 10/100/1000	6	vio2	C4	4PT 10/100/1000	6	vio1	C4	8GB DP fibre	2	vio2	2	
	C5	8GB DP fibre	2	vio1	C5	8GB DP fibre	2	vio2	C5	4PT 10/100/1000	4		2	
	C6	4PT 10/100/1000	4	vio1	C6	4PT 10/100/1000	4	vio2	C6	4GB DP fibre	6	lpar1	2	
									C7	4GB DP fibre	7		3	
	D1	146GB disk		vio1	D1	146GB disk		vio1	C8		8		3	
	D4	146GB disk		vio2	D4	146GB disk		vio2	C9		9		3	
									C10		10		3	

10

FORSYTHE

MAKE SURE YOUR PAGE SPACE IS CORRECT

More than one page volume
All the same size including hd6
Page spaces must be on different disks to each other
Do not put on hot disks
Mirror all page spaces that are on internal or non-raided disk

How much page space is needed?

- That depends
- Some software vendors require 2X memory and will not support it without that
- Others are fine with something more sensible

11

FORSYTHE

Help the Hypervisor!

Help the hypervisor cleanly place partitions when they are first defined and activated.

Define dedicated partitions first.

- Define large partitions first.

Within shared pool, define large partitions first.
At system (not partition) IPL, PowerVM will allocate resources cleanly.
Do not set maximum memory setting too high as you will waste memory
Fill your memory dimms to get maximum bandwidth
Don't mix memory dimms of different speeds

12

FORSYTHE

MEMORY USAGE

From HMC

Server-9117-MMA-SN1020AD5

General | Processors | **Memory** | I/O | Migration | Power-On Parameters | Capabilities | Advanced

Details of the managed system's memory are listed below.

Installed memory:	32768 MB
Deconfigured memory:	0 MB
Available memory:	1920 MB
Configurable memory:	32768 MB
Memory region size:	128 MB
Current memory available for partition usage :	30592 MB
System firmware current memory:	2176 MB
Maximum number of memory pools:	1

OK | Cancel | Help

Note firmware Use

Server-8233-E8B-SN0617BFP

General | Processors | **Memory** | I/O | Migration | Power-On Parameters | Capabilities | Advanced

Details of the managed system's memory are listed below.

Installed memory:	131072 MB
Deconfigured memory:	0 MB
Available memory:	0 MB
Configurable memory:	131072 MB
Memory region size:	256 MB
Current memory available for partition usage :	127744 MB
System firmware current memory:	3328 MB
Maximum number of memory pools:	1

OK | Cancel | Help

Also note memory region size

You need to know it for LPM

13

FORSYTHE

PLANNING FOR MEMORY

Don't forget memory overhead
And double if Hypervisor Mirrored

Memory Planning Worksheet
Power7 770

This gives a rough estimate
Assumes LMB size is 256MB

Max RAM Capacity 786432 Ram installed GB 393216 Ram Active 131072 Each active IVE port adds 102 MB
Change the LMB size on this line to match MRO on HMC 384 MB LMB = 256 LMB below in MB Used the largest to show worst possible

Extra high performance ports per VIO 8 NPIV VFCs per VIO 12

LPAR NAME	Desired Memory MB	Maximum Memory MB	Ohead Max Div 64	OH/LMB MB	Roundup OH MB	Actual Ohead (MB) OH * LMB	Memory Needed	Extra high Perf ports	If NPIV
VIOS1	3172	4096	64	0.25	1	256		4096	1680
VIOS2	3172	4096	64	0.25	1	256		4096	1680
LPAR1	12032	16384	256	1.00	1	256			
LPAR2	20224	24576	384	1.50	2	512			
LPAR3	14336	16384	256	1.00	1	256			
LPAR4	16384	24576	384	1.50	2	512			
LPAR5	3072	4096	64	0.25	1	256			
LPAR6	2048	4096	64	0.25	1	256			
LPAR7	17152	17152	268	1.05	2	512			
LPAR8	65536	71680	1120	4.38	5	1280			
LPAR9	32768	36864	576	2.25	3	768			
HYPERVISOR						768			
IVE						102			
I/O drawer (I use 512 per 2)						512			
Safety Net						512			
						MB Total			
MB Total	189896	224000	3500	13.671875	20	7014	196910	8192	3360
GB Total	185					6.85	192	8.00	3.28
							GB Total		

Hypervisor requires 7GB minimum for overhead with these settings — Add High Perf 200
LPARs require 185GB so the total active needed is at least 192GB — Or add NPIV 196
Need to add NPIV and high speed adapter memory needs as well

So if doing both totaloverhead is — Combined New Overhead total 204

8GB and 10GB extra high performance adapters, for each active port DD 512mb
i.e. 20 ports per VIO without NPIV would be 20 * 512 = 10GB plus VIOS base for each VIOS
If NPIV then 140MB per VFC adapter per client
if NPIV then we allocate per client so if there are 20 clients on each VIO then each vio needs 20*140=2.8GB extra

14

FORSYTHE

Memory Tips

Avoid having chips without DIMMs.

Attempt to fill every chip's DIMM slots, activating as needed.

Hypervisor tends to avoid activating cores without "local" memory.

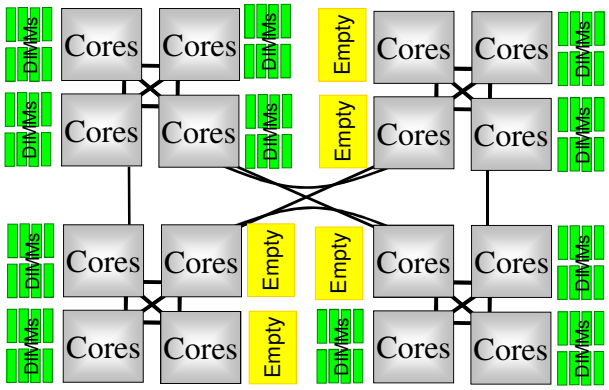

Diagram courtesy of IBM

15

FORSYTHE

TIPS TO KEEP OUT OF TROUBLE

Check the latest performance apars have all been installed

- Yes this means you need to stay current
- See Stephen Nasypany Optimization Presentations

Keep firmware up to date

- In particular, look at the firmware history for your server to see if there are performance problems fixed

Information on the firmware updates can be found at:

- http://www-933.ibm.com/support/fixcentral/
  - i.e. in AL720_064 there is a Hypervisor dispatch bug
  - It is fixed in concurrent update AL720_101

Firmware history including release dates can be found at:

- Power7 Midrange
  - http://download.boulder.ibm.com/ibmdl/pub/software/server/firmware/AM-Firmware-Hist.html
- Power7 High end
  - http://download.boulder.ibm.com/ibmdl/pub/software/server/firmware/AL-Firmware-Hist.html

16

FORSYTHE

More Tips to keep out of trouble

Ensure software stack is current
Ensure compilers are current and that compiled code turns on optimization
To get true MPIO run the correct multipath software
Ensure system is properly architected (VPs, memory, entitlement, etc)

Use the correct tunables for the version you are running

Always read the READMEs for firmware, etc updates
- Sometimes there is a prerequisite action to avoid problems

COLLECT BASELINES when the system is well
DOCUMENTATION

17

FORSYTHE

Have a Plan for making changes

1. Describe the problem.
2. Measure where you're at (baseline).
3. Recreate the problem while getting diagnostic data (perfpmr, your own scripts, etc.).
4. Analyze the data.
5. Document potential changes and their expected impact, then group and prioritize them.
   1. Remember that one small change that only you know about can cause significant problems so document ALL changes
6. Make the changes.
   1. Group changes that go together if it makes sense to do so but don't go crazy
7. Measure the results and analyze if they had the expected impact; if not, then why not?
8. Is the problem still the same? If not, return to step 1.
9. If it's the same, return to step 3.

This may look like common sense but in an emergency that is the first thing to go out the window

18

FORSYTHE

PARAMETER SETTINGS - SUMMARY

PARAMETER	DEFAULTS AIXv5.3	AIXv6	AIXv7	NEW SET ALL TO	
NETWORK (no)					
rfc1323	0	0	0	1	
tcp_sendspace	16384	16384	16384	262144 (1Gb)	
tcp_recvspace	16384	16384	16384	262144 (1Gb)	
udp_sendspace	9216	9216	9216	65536	
udp_recvspace	42080	42080	42080	655360	
MEMORY (vmo)					
minperm%	20	3	3	3	
maxperm%	80	90	90	90	JFS, NFS, VxFS, JFS2
maxclient%	80	90	90	90	JFS2, NFS
lru_file_repage	1	0	0	0	
lru_poll_interval	?	10	10	10	
Minfree	960	960	960	calculation	
Maxfree	1088	1088	1088	calculation	
page_steal_method	0	0 /1 (TL)	1	1	
JFS2 (ioo)					
j2_maxPageReadAhead	128	128	128	as needed	
j2_dynamicBufferPreallocation	16	16	16	as needed	

19

FORSYTHE

OK I HAVE A PROBLEM – NOW WHAT DO I DO?

20

FORSYTHE

TAKE A DEEP BREATH!

FORSYTHE

RULES OF THE ROAD

Clear the room of unnecessary people so you can think (but be nice)

Always have a baseline to compare to

- In order to know what is bad you have to know what is normal

Take new baselines before and after changes

Don't make so many changes it is impossible to figure out what broke it

Make sure there is good documentation

- Sysplan from HMC or use HMC Scanner (for HMC and Flexmgr)
- LPAR layouts
- Allocation list

Change control

Step 1 – gather information

- Need a clear description of the problem
  - What happened
  - What is the problem
  - Any error messages, etc

FORSYTHE

Google It

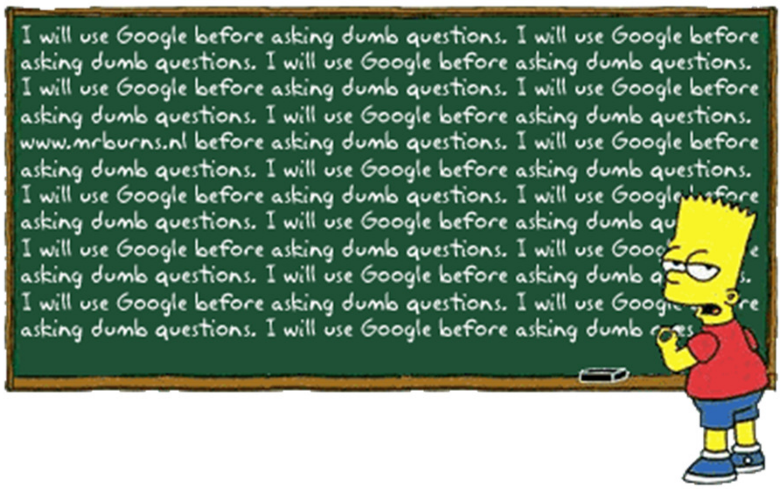

23

FORSYTHE

Baseline

You can use perfpmr, nmon or whatever works for you

- Must be consistent and repeatable

On AIX v5.3 you must download nmon12 – don't use the nmon_topas version

I collect nmon data 7/24 as follows:

Crontab entry:

- 59 23 * * * /usr/local/bin/runnmon.sh >/dev/null 2>&1

runnmon.sh:

```
#!/bin/ksh
#
cd /usr/local/perf
/usr/bin/nmon -ft -A -M -L -^ -s 150 -c 576
```

24

FORSYTHE

LOW HANGING FRUIT

Check change control – was there anything changed?
Do I have any hardware errors in errpt?
Does lsps –a or lsps –s show you have a lot of page space used?
Is my system approaching 100%
- If shared pool am I constantly over entitlement or am I constantly folding/unfolding VPs

Is the ratio of SYS% more than USR%?
Does my batch window extend into my online?
Is there unexplained I/O wait?
Are my CPU's and threads being used fairly evenly?
Is the I/O fairly well spread between disks? / Adapters?
Any full filesystems – especially /var or / or /usr
Error messages
- /etc/syslog.conf will tell you where they are
- Look at errpt – a lot of problems are made clear there

Check at Fix Central in case it is a known bug
- http://www-933.ibm.com/support/fixcentral/

Do the same at the firmware history site in case it is fixed at the next firmware update

Know how to use PerfPMR – before you need to…
http://publib.boulder.ibm.com/infocenter/pseries/v5r3/index.jsp?topic=/com.ibm.aix.prftungd/doc/prftungd/reporting_perf_prob.htm

25

FORSYTHE

WHAT MAKES IT GO SLOW?

Obvious:-
- Not enough CPU
- Not enough memory
- Not enough disk I/O
- Not enough network I/O

Not so obvious:-
- AIX tuning
- Oracle/DB2 parameters log place, SGA, Buffers
- Read vs write characteristics,
- Adapter placement, overloading bus speeds
- Throttling effects – e.g., single-thread dependency
- Application errors
- Background processes (backups, batch processing) running during peak online times?
- Concurrent access to the same files
- Changes in shared resources
- Hardware errors

26

FORSYTHE

How to Measure

Response time is the elapsed time between when a request is submitted and when the response from that request is returned.

- Amount of time for a database query
- Amount of time it takes to echo characters to the terminal
- Amount of time it takes to access a Web page
- How much time does my user wait?

Throughput is a measure of the amount of work that can be accomplished over some unit of time.

- Database transactions per minute
- File transfer speed in KBs per second
- File Read or Write KBs per second
- Web server hits per minute

27

FORSYTHE

Performance Support Flowchart

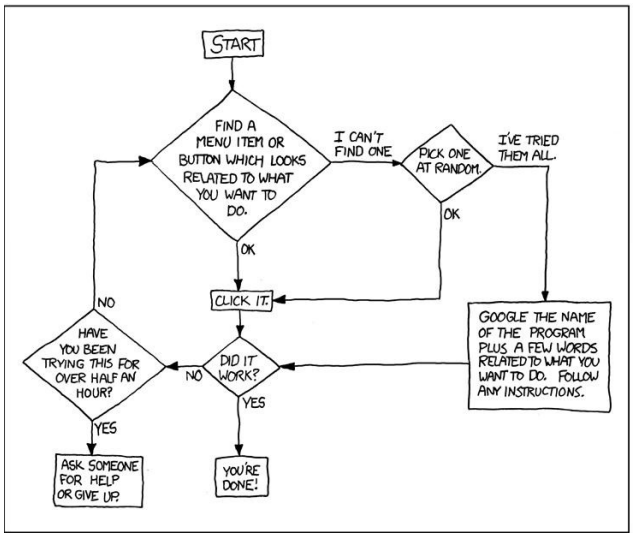

Courtesy of XKCD

28

FORSYTHE

Performance Analysis Flowchart

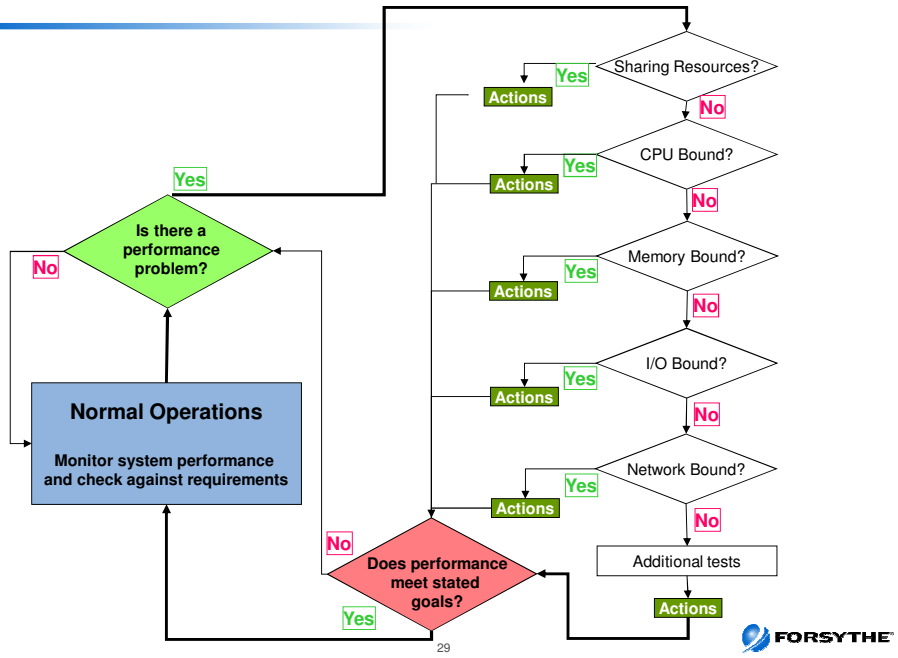

Monitoring CPU

User, system, wait and idle are fine for dedicated LPARs

They are not fine for SPLPAR or dedicated donating LPARs

You need to measure and charge back based on used CPU cycles

Moral of the story – use Physc (Physical consumed)

lparstat

- Use with no flags to view partition configuration and processor usage

30

FORSYTHE

TERMS TO UNDERSTAND

Process

- A process is an activity within the system that is started with a command, a shell script, or another process.

Run Queue

- Each CPU has a dedicated run queue. A run queue is a list of runnable threads, sorted by thread priority value. There are 256 thread priorities (zero to 255). There is also an additional global run queue where new threads are placed.

Time Slice

- The CPUs on the system are shared among all of the threads by giving each thread a certain slice of time to run. The default time slice of one clock tick is 10 ms

FORSYTHE

USING SAR –P ALL TO SHOW SMT THREADS

SMT2 Example

```
sar -P ALL 1 1
AIX sys01a 3 5 00CDAF6F4C00 ent=0.80
System Configuration: lcpu=4 ent=0.80
```

12:18:01	cpu	%usr	%sys	%wio	%idle	%physc	%entc
12:18:01	0	0	7	0	93	0.03	3.3
	1	100	0	0	0	0.37	46.8
	2	100	0	0	0	0.38	46.9
	3	0	1	0	99	0.02	3.1
	-	**94**	**0**	**0**	**6**	**0.80**	**100**

System is clearly busy – now map this to the mpstat command

```
mpstat –s 1 1
System configuration: lcpu=4 ent=0.80
```

Proc0		Proc1	
39.99%		39.76%	
cpu0	cpu1	cpu2	cpu3
2.55%	37.45%	37.57%	2.19%

Oracle tends to really like SMT and to take advantage of it

FORSYTHE

USING SAR –P ALL (POWER7 & SMT4)

AIX bpicnim 1 7 00F6934B4C00 10/05/11 (1 core and 2 VPs)

System configuration: lcpu=8 ent=1.00 mode=Uncapped

19:40:49 cpu	%usr	%sys	%wio	%idle	physc	%entc
19:40:50	0	7	88	0	5	0.01	1.4
	1	0	0	0	100	0.00	0.3
	2	0	1	0	99	0.00	0.3
	3	0	0	0	100	0.00	0.3
	7	0	59	0	41	0.00	0.0
	U	-	-	0	98	0.98	97.5
	-	0	1	0	99	0.02	2.5

In the above cpu4-6 are missing as they are 0 so sar did not print them to save space

mpstat –s 1 1
System configuration: lcpu=8 ent=1.0 mode=Uncapped

Proc0				Proc4			
2.26%				0.01%			
cpu0	cpu1	cpu2	cpu3	cpu4	cpu5	cpu6	cpu7
1.33%	0.31%	0.31%	0.31%	0.00%	0.00%	0.00%	0.01%

33

FORSYTHE

SAR –MU –P ALL

System configuration: lcpu=64 ent=7.00 mode=Uncapped

CPU	%usr	%sys	%wio	%idle	physc	%entc
Average						
0	66	32	0	2	0.47	6.8
1	48	15	2	34	0.20	2.9
2	0	4	0	96	0.09	1.3
3	0	4	0	96	0.09	1.3
4	80	16	0	3	0.43	6.2
5	74	6	1	18	0.28	4.0
6	0	4	0	95	0.08	1.1
7	0	4	0	95	0.08	1.1
8	78	19	0	2	0.45	6.4
9	54	13	2	32	0.21	3.0
10	0	4	0	96	0.09	1.2
11	0	4	0	96	0.09	1.2
12	77	20	0	3	0.42	6.0
13	63	10	2	25	0.23	3.3
14	0	5	0	95	0.08	1.2
15	1	4	0	95	0.08	1.2
.....						
60	78	19	0	3	0.42	6.0
61	56	16	2	27	0.22	3.1
62	0	4	0	96	0.08	1.1
63	0	6	0	94	0.08	1.2
AVE -	**52**	**17**	**1**	**30**	**12.90**	**184.2**

On average exceeding entitlement
Need to adjust it

34

FORSYTHE

VIO CONSTANTLY EXCEEDING ENTITLEMENT

System configuration: lcpu=16 ent=1.10 mode=Uncapped

	cpu	%usr	%sys	%wio	%idle	physc	%entc	
	0	0	96	0	4	0.48	43.9	
	1	0	35	0	65	0.14	12.5	
	2	0	20	0	80	0.11	10.0	
	3	0	16	0	84	0.11	9.7	.83
	4	1	67	0	33	0.10	9.1	
	5	0	33	0	67	0.05	5.0	
	6	0	31	0	69	0.05	4.7	
	7	0	31	0	69	0.05	4.6	.25
	8	0	71	0	29	0.11	9.9	
	9	0	32	0	68	0.06	5.0	
	10	0	31	0	69	0.05	4.8	
	11	0	31	0	69	0.05	4.8	.27
	12	0	82	0	18	0.18	16.4	
	13	0	27	0	73	0.07	6.1	
	14	0	25	0	75	0.06	5.8	
	15	0	25	0	75	0.06	5.8	.37
AVE	-	0	57	0	43	1.74	158.0	

4 VPs
Ent=1.1
Using 1.74 on ave

35

FORSYTHE

VMSTAT EXAMPLE

System configuration: lcpu=64 mem=163840MB ent=7.00

kthr			memory		page					faults			cpu					
r	b	p	avm	fre	fi	fo	pi	po	fr	sr	in	sy	cs	us	sy	id	wa	pc	ec
13	0	0	34899243	3335256	13	191	0	0	0	0	2001	179983	36609	36	32	31	1	8.20	117.1
15	0	0	34897063	3337424	10	12	0	0	0	0	2422	185355	28888	36	32	31	1	8.21	117.3
13	0	0	34899746	3334982	9	9	0	0	0	0	2114	273146	41908	36	33	29	1	7.90	112.8
13	0	0	34899376	3335340	5	27	0	0	0	0	2173	256828	41451	35	32	32	1	7.83	111.9
14	0	0	34899005	3335696	4	14	0	0	0	0	2079	258935	46276	35	30	34	1	8.34	119.1
11	0	0	34896462	3338226	8	29	0	0	0	0	4535	321194	77820	34	28	37	1	8.15	116.4
12	0	0	34895235	3339441	2	18	0	0	0	0	1653	180720	25206	34	30	35	1	8.41	120.2
13	0	0	34899626	3335032	4	93	0	0	0	0	1996	252956	44036	35	28	36	1	8.23	117.6

Ent = 7
Using last line PC is 8.23 cores or 117.6% of entitlement
US + SY = 35+28= 63%

36

FORSYTHE

VMSTAT –I OUTPUT

```
vmstat -I  2 10

System Configuration: lcpu=22 mem=90112MB

 kthr    memory              page              faults        cpu
-------- ----------- ------------------------ ------------ -----------
 r   b    p     avm        fre          fi        fo     pi   po       fr          sr           in           sy           cs       us   sy  id wa
70 309  0 8552080   9902   75497 9615   9   3    84455 239632 18455 280135 91317  42  37  0 20
79 285  0 8537038   9371   83963 7568  44   2    84266 230503 19400 406846 77938  58  37  0   5
56 301  0 8540516   8895   91385 8912  12   3  101110 253980 17943 388340 86999  52  38  0 10
48 306  0 8544771   9565 101529 9966  14   3  112865 277552 16930 358515 82444  50  41  0   9
73 285  0 8544667   8763   94305 5915  25   3    95071 277963 19299 438769 83214  49  35  0 16
23 317  0 8547888   9846   91608 5481  12   1    97364 235613 19148 393468 74293  55  34  0 11
16 352  0 8541280   8845   92946 5246  14   0    93028 244146 18471 448516 87874  44  37  0 19
```

fre is meaningless if you do not know the minfree, maxfree and mempools values
SR:FR should be <= 4:1
244146: 93028 is around 2.61 : 1

```
System configuration: lcpu=32 mem=122880MB ent=8.00
 kthr    memory              page              faults              cpu
-------- ----------- ------------------------ ------------ ----------------------
 r  b  p       avm          fre              fi    fo   pi  po  fr   sr  in       sy       cs     us    sy  id    wa    pc        ec
 1  0  0 16760611 13937801      36   0     0   0   0     0  405   4413  1165   5      1   94   0      0.79     9.9   .79 is
9.9% of entitlement
 1  0  0 16760407 13938004       0    0     0   0   0     0  357   4445    979   5      1   93   0      0.81   10.1
```

37

SVMON –G OUTPUT TRANSLATED TO GB

MEMORY	VALUE	GB	MB
size	46137344	176.00	180224.00
inuse	23832056	90.91	93093.97
free	22305116	85.09	87129.36
pin	3470487	13.24	13556.59
virtual	16886019	64.42	65961.01
page sz	4194304	16.00	16384.00
page inuse	106961	0.41	417.82
pin work	2128407	8.12	8314.09
pin persist	0	0.00	0.00
pin client	0	0.00	0.00
pin lpage	1342080	5.12	5242.50
inuse work	16885847	64.41	65960.34
inuse persist	0	0.00	0.00
inuse client	6946209	26.50	27133.63

38

LSSRAD -AV

Large LPAR on a 770

REF1	SRAD	MEM	CPU
0			
	0	171529.00	0-15 20-23 28-31 36-39 44-47 52-55 60-63
	1	114771.50	16-19 24-27 32-35 40-43 48-51 56-59

Smaller LPAR

REF1	SRAD	MEM	CPU
0			
	0	88859.50	0-7
	2	36354.00	
1			
	1	42330.00	8-11
	3	20418.00	

REF1 indicates where
REF1=0 SRAD=0 is local
REF1=0 SRAD=1 is near
Other REF values are far
This is relative to the process home
SRAD = CPU + Memory group
MEM = Mbytes
CPU = LCPU number, assuming SMT4

39

FORSYTHE

WHAT IS IOWAIT?

- **iowait** is a form of idle time

The **iowait** statistic is simply the percentage of time the CPU is idle AND there is at least one I/O still in progress (started from that CPU)

The **iowait** value seen in the output of commands like **vmstat**, **iostat**, and **topas** is the **iowait** percentages across all CPUs averaged together

High I/O wait does not mean that there is definitely an I/O bottleneck

Zero I/O wait does not mean that there is not an I/O bottleneck

A CPU in I/O wait state can still execute threads if there are any runnable threads

40

FORSYTHE

IO WAIT AND WHY IT IS NOT NECESSARILY USEFUL

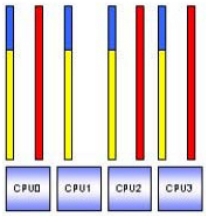

Diagram courtesy of IBM

System has 3 threads blocked (red threads)
SMT is turned on
There are 4 threads ready to run so they get dispatched and each is using 80% user and 20% system

Metrics would show:
%user = .8 * 4 / 4 = 80%
%sys = .2 * 4 / 4 = 20%
Idle will be 0% as no core is waiting to run threads
IO Wait will be 0% as no core is idle waiting for IO to complete as something else got dispatched to that core
SO we have IO wait
BUT we don't see it

41

FORSYTHE

KERNEL I/O LAYERS

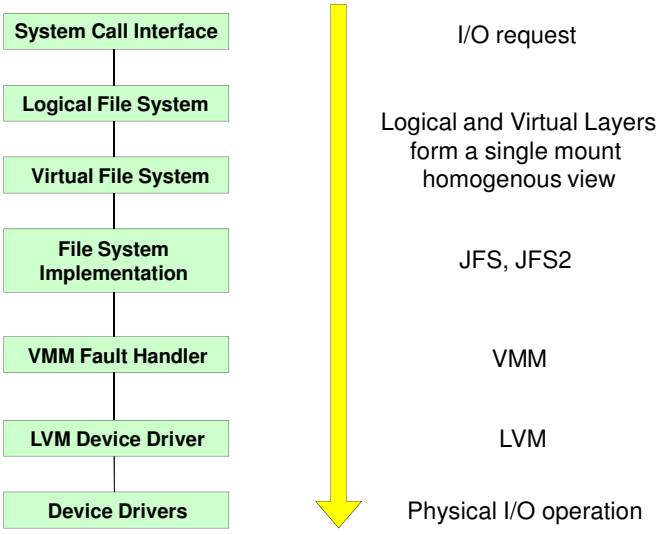

Diagram courtesy of IBM

42

FORSYTHE

ROUGH ANATOMY OF AN I/O

LVM requests a PBUF
- Pinned memory buffer to hold I/O request in LVM layer

Then placed into an FSBUF
- 3 types
- These are also pinned
- Filesystem — JFS
- Client — NFS and VxFS
- External Pager — JFS2

If paging then need PSBUFs (also pinned)
- Used for I/O requests to and from page space

Then queue I/O to hdisk (queue_depth)
Then queue it to adapter (num_cmd_elems)
Adapter queues it to the disk subsystem

43

FORSYTHE

VMSTAT –V OUTPUT

3.0 minperm percentage
90.0 maxperm percentage
45.1 numperm percentage
45.1 numclient percentage
90.0 maxclient percentage

1468217 pending disk I/Os blocked with no pbuf	pbufs
11173706 paging space I/Os blocked with no psbuf	pagespace
2048 file system I/Os blocked with no fsbuf	JFS
238 client file system I/Os blocked with no fsbuf	NFS/VxFS
39943187 external pager file system I/Os blocked with no fsbuf	JFS2

numclient=numperm so most likely the I/O being done is JFS2 or NFS or VxFS
Based on the blocked I/Os it is clearly a system using JFS2
It is also having paging problems
pbufs also need reviewing

44

FORSYTHE

ADAPTER QUEUE PROBLEMS

Look at BBBF Tab in NMON Analyzer or run fcstat command

Adapter device drivers use DMA for IO
From fcstat on each fcs
NOTE these are since boot

FC SCSI Adapter Driver Information
No DMA Resource Count: 0
No Adapter Elements Count: 2567
No Command Resource Count: 34114051

No DMA resource – adjust max_xfer_size
No adapter elements – adjust num_cmd_elems
No command resource - adjust num_cmd_elems

If using NPIV make changes to VIO and client, not just VIO

45

FORSYTHE

NETWORK

If 10Gb network check out Gareth's Webinar

- https://www.ibm.com/developerworks/wikis/download/attachments/153124943/7_PowerVM_10Gbit_Ethernet.pdf?version=1

netstat –v

- Look for overflows and memory allocation failures
  - Max Packets on S/W Transmit Queue: 884
  - S/W Transmit Queue Overflow: 9522
- "Software Xmit Q overflows" or "packets dropped due to memory allocation failure"
  - Increase adapter xmit queue
  - Use lsattr –EL ent? To see setting
- Look for receive errors or transmit errors
- dma underruns or overruns
- mbuf errors

Lparstat 2

- Look for high vcsw – indicator that entitlement may be too low

tcp_nodelay (or tcp_nodelayack)

- Disabled by default
- 200ms delay by default as it waits to piggy back acks on packets

Also check errpt – people often forget this

46

FORSYTHE

ENTSTAT -V

ETHERNET STATISTICS (ent18) :
Device Type: Shared Ethernet Adapter
Elapsed Time: 44 days 4 hours 21 minutes 3 seconds

Transmit Statistics:	Receive Statistics:
Packets: 94747296468	Packets: 94747124969
Bytes: 99551035538979	Bytes: 99550991883196
Interrupts: 0	Interrupts: 22738616174
Transmit Errors: 0	Receive Errors: 0
Packets Dropped: 0	Packets Dropped: 286155
	Bad Packets: 0

Max Packets on S/W Transmit Queue: 712
S/W Transmit Queue Overflow: 0
Current S/W+H/W Transmit Queue Length: 50

Elapsed Time: 0 days 0 hours 0 minutes 0 seconds

Broadcast Packets: 3227715	Broadcast Packets: 3221586
Multicast Packets: 3394222	Multicast Packets: 3903090
No Carrier Sense: 0	CRC Errors: 0
DMA Underrun: 0	DMA Overrun: 0
Lost CTS Errors: 0	Alignment Errors: 0
Max Collision Errors: 0	**No Resource Errors: 286155 check those tiny, etc Buffers**
Late Collision Errors: 0	Receive Collision Errors: 0
Deferred: 0	Packet Too Short Errors: 0
SQE Test: 0	Packet Too Long Errors: 0
Timeout Errors: 0	Packets Discarded by Adapter: 0
Single Collision Count: 0	Receiver Start Count: 0

Multiple Collision Count: 0
Current HW Transmit Queue Length: 50

47

FORSYTHE

ENTSTAT –V VIO

SEA

Transmit Statistics:	Receive Statistics:
Packets: 83329901816	Packets: 83491933633
Bytes: 87482716994025	Bytes: 87620268594031
Interrupts: 0	Interrupts: 18848013287
Transmit Errors: 0	Receive Errors: 0
Packets Dropped: 0	**Packets Dropped: 67836309**
	Bad Packets: 0

Max Packets on S/W Transmit Queue: 374
S/W Transmit Queue Overflow: 0
Current S/W+H/W Transmit Queue Length: 0

Elapsed Time: 0 days 0 hours 0 minutes 0 seconds

Broadcast Packets: 1077222	Broadcast Packets: 1075746
Multicast Packets: 3194318	Multicast Packets: 3194313
No Carrier Sense: 0	CRC Errors: 0
DMA Underrun: 0	DMA Overrun: 0
Lost CTS Errors: 0	Alignment Errors: 0
Max Collision Errors: 0	**No Resource Errors: 67836309**

Virtual I/O Ethernet Adapter (l-lan) Specific Statistics:

Hypervisor Send Failures: 4043136
Receiver Failures: 4043136
Send Errors: 0
Hypervisor Receive Failures: 67836309

"No Resource Errors" can occur when the appropriate amount of memory can not be added quickly to vent buffer space for a workload situation.

48

FORSYTHE

BUFFERS

Virtual Trunk Statistics
Receive Information
Receive Buffers

Buffer Type	Tiny	Small	Medium	Large	Huge
Min Buffers	512	512	128	24	24
Max Buffers	2048	**2048**	256	64	64
Allocated	513	2042	128	24	24
Registered	511	506	128	24	24
History					
Max Allocated	532	2048	128	24	24
Lowest Registered	502	354	128	24	24

"Max Allocated" represents the maximum number of buffers ever allocated
"Min Buffers" is number of pre-allocated buffers
"Max Buffers" is an absolute threshhold for how many buffers can be allocated

chdev –l <veth> -a max_buf_small=4096 –P
chdev –l <veth> -a min_buf_small=2048 –P
Above increases min and max small buffers for the virtual ethernet adapter configured for the SEA above

49

FORSYTHE

LPARSTAT

lparstat 30 2 output

System configuration: type=Shared mode=Uncapped smt=4 lcpu=12 mem=6144MB psize=4 ent=1.50

%user	%sys	%wait	%idle	physc	%entc	lbusy	app	vcsw	phint
0.1	9.6	0.0	90.3	0.38	25.5	3.4	3.55	16678	5
0.1	8.2	0.0	91.7	0.33	21.8	3.5	3.59	13922	5

FORSYTHE

DEMOTED I/O IN ORACLE

CIO write fails because IO is not aligned to FS blocksize

- i.e app writing 512 byte blocks but FS has 4096

Ends up getting redone

- Demoted I/O consumes more kernel CPU
- And more physical I/O

To find demoted I/O (if JFS2)

trace –aj 59B,59C ; sleep 2 ; trcstop ; trcrpt –o directio.trcrpt

grep –i demoted directio.trcrpt

Look in the report for:

51

FORSYTHE

REDO LOGS AND DEMOTED I/O

lsfs -a output

Name	Nodename	Mount Pt	VFS	Size	Options	Auto	Accounting
/dev/hd4	--	/	jfs2	524288	rw	yes	no
/dev/hd1	--	/home	jfs2	20971520	rw	yes	no
/dev/hd2	--	/usr	jfs2	8912896	rw	yes	no

lsfs -q output

Name	Nodename	Mount Pt	VFS	Size	Options	Auto	Accounting
/dev/hd4	--	/	jfs2	524288	rw	yes	no

(lv size: 524288, fs size: 524288, **block size: 4096**, sparse files: yes, inline log: no, inline log size: 0, EAformat: v1, Quota: no, DMAPI: no, VIX: yes, EFS: no, ISNAPSHOT: no, MAXEXT: 0, MountGuard: no)

NOTE THE BLOCKSIZE ABOVE IS 4096 – redo log should be 512
So look for that with redo logs (usually /u99 or some such – check with DBA)

52

FORSYTHE

CHECKING I/O BANDWIDTH

- **Taken from an NMON report and totaled here**

	AVE	MAX	MB AVE	MB MAX
Disk read KB/S	19374.7	67151	18.92	65.58
Disk write KB/s	6259.5	40462.5	6.11	39.51
BOTH	25634.2	107613.5	25.03	105.09
FCS0 read KB/s	4710.1	22204.7	4.60	21.68
FCS0 write KB/s	1412.5	9903.9	1.38	9.67
BOTH	6122.6	32108.6	5.98	31.36
FCS1 read KB/s	4710.1	20129.6	4.60	19.66
FCS1 write KB/s	1591.9	14330.6	1.55	13.99
BOTH	6302	34460.2	6.15	33.65
FCS2 read KB/s	4988	17924.6	4.87	17.50
FCS2 write KB/s	1666.2	13539.8	1.63	13.22
BOTH	6654.2	31464.4	6.50	30.73
FCS3 read KB/s	4953.9	21645.9	4.84	21.14
FCS3 write KB/s	1528.4	9945.1	1.49	9.71
BOTH	6482.3	31591	6.33	30.85
ALL FCS	13136.5	63055.4	**12.83**	**61.58**

53

FORSYTHE

OPENING A PMR

Do yourself a favor and open via the web
- Yes you need a login but it is worth it!

Gather documentation
- IBM needs a clear description of the problem
  - Include any error messages
- They also need serial numbers, etc for entitlement

Software
- https://www.ibm.com/support/servicerequest

Hardware
- http://www-947.ibm.com/usrsrvc/support/esc/signin.jsp

Determine and set severity correctly
- Sev 4 is documentation and seems to be default
- Sev 1 means hard down and you are on call 7/24

Uploading files to IBM
- http://www-05.ibm.com/de/support/ecurep/index.html

54

FORSYTHE

PERFORMANCE TOOLS

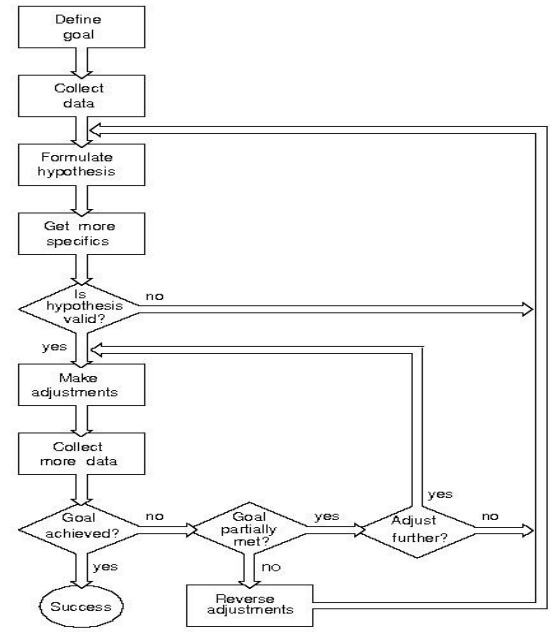

Flowchart illustrating the methodology for system performance tuning

AIX Performance Monitoring Tools (lots of options)

Tools	Monitor status and stats	Trace	Tune
Virtualization	lparstat, mpstat, schedo, hpmcount, hpmstat, VIOS and HMC commands	VIOS and HMC commands	schedo, VIOS commands, HMC commands
Processor	vmstat, topas, nmon, iostat, ps, lparstat, mpstat, sar, time, emstat, netpmon, wlmstat, xmperf, procmon	tprof, curt, splat, trace, trcpt	schedo, fdpr, bindprocessor, nice/renice, setpri, smtctl
Memory	vmstat, sar, topas, nmon, ps, lsps, ipcs, svmon, netpmon, filemon, xmperf, wlmstat, pagesize	trace, trcpt	vmo, rmss, fdpr, chps/mkps
Network	netstat, topas, nmon, nfsstat, atmstat, entstat, tokstat, fddstat, nfsstat, ifconfig, netpmon tcpdump, wlmstat, iperf, netperf, jperf	iptrace, tcpdump, ipreport, trace, trcpt	no, nfso, chdev, ifconfig
I/O, LVM, JFS2	vmstat, sar, topas, nmon, iostat, fcstat, lvmstat, lsps, lsdev, lsattr, lspv, lsvg, lslv, fileplace, trcpt, filemon, ncheck, xmperf, wlmstat	trace, trcpt	ioo, lvmo, chdev, nfso, migratepv, chlv, reorgvg, chps
Kernel	ps, pstat, topas, nmon, ipcs, emstat, svmon, truss, kdb, dbx, gprof, fuser, prof, ncheck, procmon	truss, prof, curt, splat, trace, trcpt	chdev, fdpr, schedo, schedtune, tunchange, tuncheck, tunrestore, tunsave, tundefault, raso
Application	emstat, gprof, trpof, truss, probevue, prof, time	emstat, gprof, trpof, truss, probevue, prof, time	emstat, gprof, trpof, truss, probevue, prof, time

57

FORSYTHE

TOOLS

https://www.ibm.com/developerworks/wikis/display/WikiPtype/Other+Performance+Tools

topas
- New –L flag for LPAR view

nmon

nmon analyzer
- Windows tool so need to copy the .nmon file over in ascii mode
- Opens as an excel spreadsheet and then analyses the data
- Also look at nmon consolidator

sar
- sar -A -o filename 2 30 >/dev/null
- Creates a snapshot to a file – in this case 30 snaps 2 seconds apart
- Must be post processed on same level of system

errpt

Check for changes from defaults

ioo, vmo, schedo, vmstat –v

lvmo

lparstat, mpstat

iperf, jperf, netperf

iostat

Check out Alphaworks for the Graphical LPAR tool

Ganglia - http://ganglia.info

Nmonrrd and nmon2web and pGraph

Commercial IBM
- PM for AIX
- Performance Toolbox
- Tivoli ITM

58

FORSYTHE

OTHER TOOLS

filemon
- filemon -v -o filename -O all
- sleep 30
- trcstop
- Most active LVs, PVs and files

pstat to check async I/O in 5.3
- pstat -a | grep aio | wc –l

perfpmr to build performance info for IBM if reporting a PMR
- /usr/bin/perfpmr.sh 300

FORSYTHE

NMON

nmon -ft -A -s 15 -c 120
- Grabs a 30 minute nmon snapshot with async I/O

nmon -ft -A -M -L -^ -s 15 -c 120
- Same as above but includes large pages and some other features

Must be running nmon12e or higher
Nmon comes with AIX at 5.3 tl09 or 6.1 tl01 and higher BUT on 5.3 I download the latest version from the web so I get the latest v12 for sure
Creates a file in the working directory that ends .nmon
This file can be transferred to your PC and interpreted using nmon analyser or other tools
nmon –f –O – now gets seastats for VIO server
nmon –f –K - dump libperfstat structures

http://www.ibm.com/developerworks/wikis/display/WikiPtype/nmon
http://www.ibm.com/developerworks/wikis/display/WikiPtype/nmonanalyser
http://www.ibm.com/developerworks/wikis/display/WikiPtype/nmonconsolidator

FORSYTHE

NMON

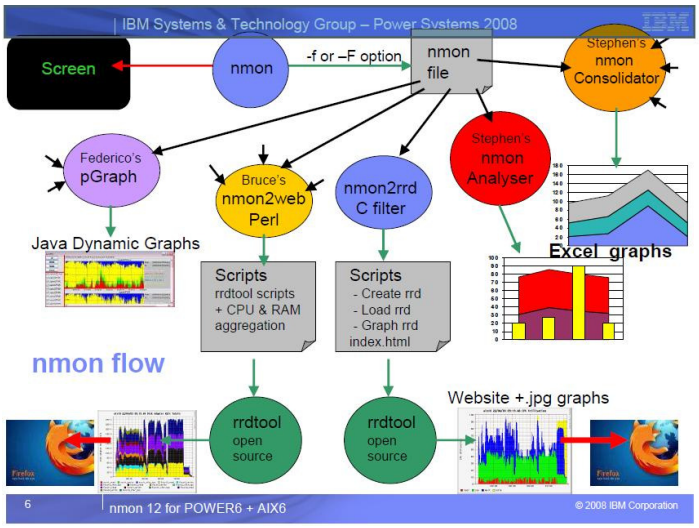

Courtesy Nigel Griffiths - IBM

61

FORSYTHE

NMON ON POWER6 & AIX6 + - NEW FEATURES FOR V12

Disk Service Times
Selecting Particular Disks
Time Drift
Multiple Page Sizes
Timestamps in UTC & no. of digits
More Kernel & Hypervisor Stats *
High Priority nmon

- Advanced, POWER6 and AIX6 items

Virtual I/O Server SEA
Partition Mobility (POWER6)
WPAR & Application Mobility (AIX6)
Dedicated Donating (POWER6)
Folded CPU count (SPLPAR)
Multiple Shared Pools (POWER6)
Fibre Channel stats via entstat

62

FORSYTHE

NMON SUMMARY

Samples	120	First	10:00:32	Last	10:30:18

System Summary xxxxxxx 9/25/2012

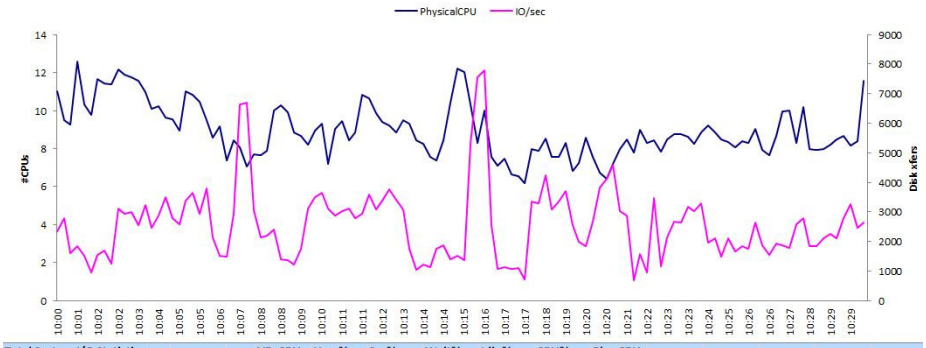

Total System I/O Statistics		VP_CPU:	User%	Sys%	Wait%	Idle%	CPU%	PhysCPU
Avg tps during an interval:	2,582	Avg	26.8	7.3	0.1	21.7	55.9	9.0
Max tps during an interval:	7,784	Max	41.6	9.8	0.6	29.6	78.6	12.6
Max tps interval time:	10:16:18	Max:Avg	1.6	1.3	4.7	1.4	1.4	1.4
Total number of Mbytes read:	434,480							
Total number of Mbytes writte	12,142							
Read/Write Ratio:	35.8							

63

FORSYTHE

PCPU_ALL TAB

PCPU Total xxxxxxx 9/25/2012

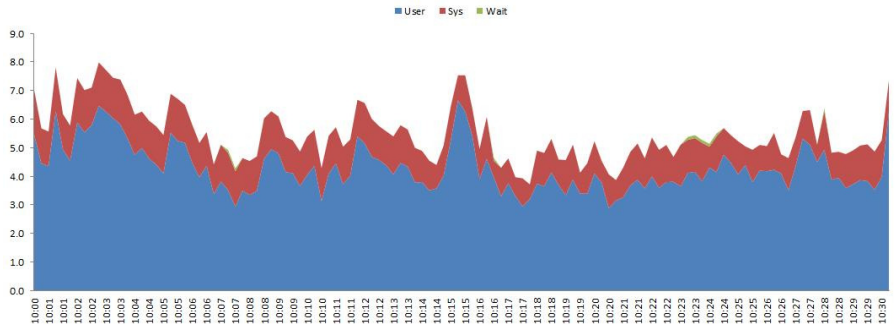

64

FORSYTHE

NMON BBBL TAB

	A	B	C	D	E	F	G	H	I	J
1	lparno	3								
2	lparname	xxxxxx								
3	CPU in sys	24								
4	Virtual CPU	16								
5	Logical CPU	64								
6	Pool CPU	16								
7	smt threads	4								
8	capped	0								
9	min Virtual	8								
10	max Virtual	20								
11	min Logical	8								
12	max Logical	80								
13	min Capacity	8								
14	max Capacity	16								
15	Entitled Capacity	10								
16	Weight	150								
17	min Memory MB	131072								
18	max Memory MB	327680								
19	online Memory	294912								
20	pool id	2								
21	Flags	LPARed DRable SMT Shared UnCapped PoolAuth Migratable Not-Donating AMSable.								

65

FORSYTHE

LPAR TAB

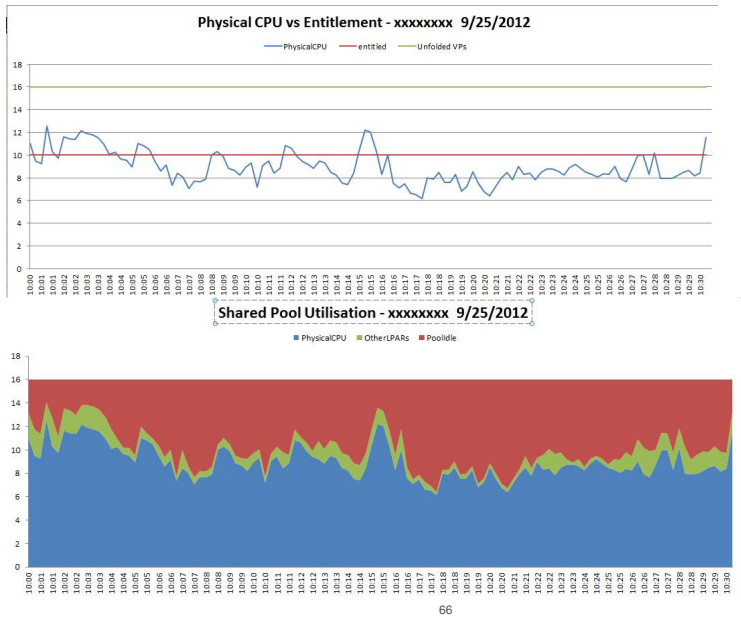

66

FORSYTHE

USEFUL WEB LINKS

Sign up for Storage and System Notifications
- https://www14.software.ibm.com/webapp/set2/subscriptions/onvdq

Article on this topic
- http://www.ibmsystemsmag.com/aix/administrator/performance/performance_updates/

POWER Firmware Code Matrix
- http://www-304.ibm.com/webapp/set2/sas/f/power5cm/power7.html

Perfpmr
- http://www-01.ibm.com/support/docview.wss?uid=aixtools-27a38cfb
- ftp://ftp.software.ibm.com/aix/tools/perftools/perfpmr

Fix Level Recommendation Tool (FLRT)
- http://www-304.ibm.com/support/customercare/flrt/home

Nigel's AIXPert Blog
- https://www.ibm.com/developerworks/mydeveloperworks/blogs/aixpert/?lang=en

IBM Performance Tools
- http://www.ibm.com/developerworks/wikis/display/WikiPtype/Other+Performance+Tools
- Includes new advisors for Java, VIOS, Virtualization

67

FORSYTHE

USEFUL LINKS

AIX Wiki
- https://www.ibm.com/developerworks/wikis/display/WikiPtype/AIX

HMC Scanner
- http://www.ibm.com/developerworks/wikis/display/WikiPtype/HMC+Scanner

Workload Estimator
- http://ibm.com/systems/support/tools/estimator

Performance Tools Wiki
- http://www.ibm.com/developerworks/wikis/display/WikiPtype/Performance+Monitoring+Tools

Performance Monitoring
- https://www.ibm.com/developerworks/wikis/display/WikiPtype/Performance+Monitoring+Documentation

VIOS Advisor
- https://www.ibm.com/developerworks/wikis/display/WikiPtype/Other+Performance+Tools#OtherPerformanceTools-VIOSPA

68

FORSYTHE

USEFUL LINKS

Service and support best practices
- http://www14.software.ibm.com/webapp/set2/sas/f/best/home.html

Fix Central - HMC, SDMC, Firmware, AIX Updates
- http://www-933.ibm.com/support/fixcentral/

IBM Prerequisite Tool
- https://www-912.ibm.com/e_dir/eserverprereq.nsf

IBM System Planning Tool
- http://www-947.ibm.com/systems/support/tools/systemplanningtool/

IBM Systems Workload Estimator
- http://www-947.ibm.com/systems/support/tools/estimator/index.html

nmon wiki
- http://www.ibm.com/developerworks/wikis/display/WikiPtype/nmon

nmon analyser wiki
- http://www.ibm.com/developerworks/wikis/display/WikiPtype/nmonanalyser

nmon consolidator wiki
- http://www.ibm.com/developerworks/wikis/display/WikiPtype/nmonconsolidator

IBM Redbooks
- http://www.redbooks.ibm.com

FORSYTHE

DON'T FORGET!

When you're right
no-one remembers

When you're wrong
no-one forgets

FORSYTHE

PERFORMANCE WIKI

Overview | New to | Forums | Wikis

Other Performance Tools

View | Attachments (39) | Info | Browse Space

Added by nag, last edited by nagger on Feb 14, 2012 (view change)
Labels: (None)

Performance - Other Tools (non AIX commands)

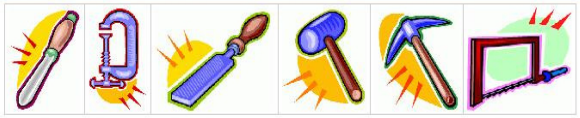

Quick Links to this page:
NEW: Java Performance Advisor - VIOS Performance Advisor - Virtualization Performance Advisor
nmon based: nmon - nmon-Analyser - nmon-Consolidator - pGraph - topas-CEC-nmon-Analyser - nmon2web - nmon2rrd
Regular: qmon v71n - Visual Performance Analyzer - Roll-Your-Own - nstress-Tools
LPAR-Monitor - SEA-Monitor - HMC-LPAR-data-2-rrd - WLM - rPerf - ryaacct - nworms
Products and larger tools: PM - Ganglia - Munin - Galileo - Performance-Toolbox

VIOS Performance Advisor

The VIOS advisor is an application that runs within the customer's VIOS for a user specified amount of time (hours), which polls and collects key performance metrics before analyzing results and providing a health check report and proposes changes to the environment or areas to investigate further. The goal of the VIOS advisor is not to provide another monitoring tool, but instead have an expert system view performance metrics already available to the customer and make assessments and recommendations based on the expertise and experience available within the IBM systems performance group.
Take this link to the VIOS Performance Advisor

71

FORSYTHE

VIOS ADVISOR

https://www.ibm.com/developerworks/wikis/display/WikiPtype/VIOS+Advisor

- Application that collects performance metrics and does a health check
- Following slides run on a production VIO during a regular production day

72

FORSYTHE

VIOS ADVISOR

The ratings and recommendations in the table below were chosen with the following information:

Hostname : vio1. '.com
PartitionID: 2
Monitoring Start Time : 03/09 11:45:19
Monitoring Stop Time : 03/09 13:45:19 **Duration :** 120 min
IBM Systems Workload Estimator link: http://ibm.com/systems/support/tools/estimator *(VIOS Sizings)*

SYSTEM - CONFIGURATION

Name	Value
Processor Family	POWER6
Server Model	IBM,9117-MMA
Server Frequency	4.208 GHz
Server - Online CPUs	10 cores
Server - Maximum Supported CPUs	16 cores
VIOS Level	2.2.0.13-FP24 SP-03
VIOS Advisor Release	121211B

73

FORSYTHE

VIOS ADVISOR

VIOS - CPU

Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1=lowest 5=highest	Impact 1=lowes 5=highes
CPU Capacity	1.0 ent	-	03/09 11:45:19	-	n/a	n/a
CPU Consumption	avg:5.4% (cores:0.1) high:40.2% (cores:0.5)	-	-	-	n/a	n/a
Processing Mode	Shared CPU, (UnCapped)	-	03/09 11:45:19	-	n/a	n/a
Variable Capacity Weight	200	-	03/09 11:45:19	-	n/a	n/a
Virtual Processors	2 vCPUs	-	03/09 11:45:19	-	n/a	n/a
SMT Mode	SMT2	-	03/09 11:45:19	-	n/a	n/a

74

FORSYTHE

VIOS ADVISOR

SYSTEM - SHARED PROCESSING POOL

	Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1=lowest 5=highest	Impact 1=lowes 5=highes
✓	Shared Pool Monitoring	enabled	-	03/09 11:45:19	-	n/a	n/a
🔒	Shared Processing Pool Capacity	10.0 ent.	-	03/09 11:45:19	-	n/a	n/a
✓	Free CPU Capacity	avg_free:9.4 ent. lowest_free:7.7 ent.	-	-	-	n/a	n/a

75

FORSYTHE

VIOS ADVISOR

VIOS - I/O ACTIVITY

	Name	Value
🔒	Disk I/O Activity	avg: 229 iops @ 32KB peak: 1916 iops @ 137KB
🔒	Network I/O Activity	[avgSend: 0 iops 0.0MBps , avgRcv: 0 iops 0.0MBps] [peakSend: 0 iops 0.0MBps , peakRcv: 0 iops 0.0MBps]

76

FORSYTHE

VIOS ADVISOR

VIOS - DISK ADAPTERS

Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1=lowest 5=highest	Impact 1=lowest 5=highest
FC Adapter Count	3	-	03/09 11:45:19	-	n/a	n/a
FC Avg IOps	avg: 77 iops @ 32KB	-	03/09 11:45:19	03/09 13:45:19	n/a	n/a
FC Adapter Utilization	optimal	-	-	-	n/a	n/a
FC Port Speeds	running at speed	-	-	-	n/a	n/a

77

FORSYTHE

VIOS ADVISOR

VIOS - DISK DRIVES

Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1=lowest 5=highest	Impact 1=lowest 5=highest
Physical Drive Count	93	-	03/09 11:45:19	-	n/a	n/a
I/Os Blocked	optimal	-	-	-	n/a	n/a
Long I/O Latency (hdisk3)	avg:9.7ms (9.7 + 0.0) high:11.5ms (11.5 + 0.0)	Range: 8-12ms	03/09 12:35:58	03/09 13:44:02	n/a	n/a

78

FORSYTHE

VIOS ADVISOR

VIOS - MEMORY

Name	Measured Value	Recommended Value	First Observed	Last Observed	Risk 1=lowest 5=highest	Impact 1=lowest 5=highest
Real Memory	4.000 GB	7.000 GB	03/09 11:45:19	-	1	2
Available Memory	0.346 GB	1.5 GB Avail.	03/09 11:45:39	03/09 13:45:05	n/a	n/a
Paging Rate	0.2 MB/s pg rate	-	-	-	n/a	n/a
Paging Space Size	8.000 GB	-	03/09 11:45:19	-	n/a	n/a
Free Paging Space	7.923 GB free	-	-	-	n/a	n/a
Pinned Memory	1.262 GB pinned	-	-	-	n/a	n/a

79

FORSYTHE

THANK YOU FOR YOUR TIME

If you have questions please email me at:
lynchj@forsythe.com

Movie replay of this session is at:
http://www.circle4.com/movies

This presentation at:
http://www.circle4.com/papers/common-perfprobs.pdf

80

FORSYTHE