

Table des matières

Ipar2rrd

- <http://www.ipar2rrd.com/>

VP-to-entitlement ratio

Ideally the ratio should be 2.5 or less. Anything above 4.0 is performance unfriendly, especially on multi-node systems (770 and above).

How to estimate the number of virtual processors per uncapped shared LPAR:

The first step would be to monitor the utilization of each partition and for any partition where the average utilization is ~100%, then add one virtual processor. (use capacity of the already configured virtual processors before adding more)

If the peak utilization is well below 50%, then look at the ratio of virtual processors to configured entitlement and if the ratio is > 1, then consider reducing the ratio. (In any case if there are too many virtual processors configured, AIX can "fold" those processors.)

AIX monitors the utilization of each virtual processor and the utilization of an PLPAR, and if utilization goes below 50%, AIX will start folding down the virtual CPUs so that fewer virtual CPUs will be dispatched. (If utilization goes beyond 50% AIX starts unfolding virtual CPUs.)

Considerations for Virtual Processor (VP) and Entitled Capacity:

- Lpars that require high performance (such as critical database) can be forced to get the best resources by activating the critical LPAR first prior to activating any other LPARs including VIO Server.
- The best practice for LPAR entitlement would be setting entitlement close to average utilization and let the peak addressed by additional uncapped capacity. (exceptions could be LPARs with higher priority)
- For each shared LPAR the number of VPs must be less than (or equal) to the number of cores of the shared pool
- Shared uncapped LPARS with too low VPs will not cover Production Need (VP number is a limit for uncapped LPARS)
- When AIX folding is turned off it can happen that PhysC (physical cores used) is high, but AIX shows high percentage of idle time. (This is because unused Virtual Processors are also assigned to cores, but they are not doing any work at all.)

Checking how many Virtual Processors are active:

```
root@bb_lpar:/ # lparstat -i | grep Virt
Online Virtual CPUs : 2                                <--we have 2 virtual processors configured
Maximum Virtual CPUs : 8
Minimum Virtual CPUs : 1
Desired Virtual CPUs : 2

root@bb_lpar:/ # bindprocessor -q
The available processors are: 0 1 2 3 4 5 6 7          <--this shows smt=4 active (4 threads/virtual processor)

root@bb_lpar:/ # echo vpm | kdb
...
0 0 ACTIVE 0 AWAKE 0000000000000000 00000000 00
1 0 ACTIVE 0 AWAKE 0000000000000000 00000000 00
2 0 ACTIVE 0 AWAKE 0000000000000000 00000000 00
3 0 ACTIVE 0 AWAKE 0000000000000000 00000000 00
4 0 DISABLED 0 AWAKE 0000000000000000 00000000 00      <--4 lines are DISABLED, so 1 Virt. proc. is inactive (folding)
5 11 DISABLED 0 SLEEPING 00000000515B4478 29DBE3CA 02
6 11 DISABLED 0 SLEEPING 00000000515B4477 2C029174 02
7 11 DISABLED 0 SLEEPING 00000000515B4477 2C0292A1 02
```

SMT

threads = VP x (SMT threads par processeur) = logical CPUs

Soit la partoche ci-dessous :

```
root@partoche:/root # lparstat -i |grep Virtual
Online Virtual CPUs : 3
Maximum Virtual CPUs : 6
Minimum Virtual CPUs : 1
Desired Virtual CPUs : 3
```

```
root@partoche:/root # smtctl
```

```
This system is SMT capable.
This system supports up to 4 SMT threads per processor.
SMT is currently enabled.
```

SMT boot mode is not set.
 SMT threads are bound to the same virtual processor.

proc0 has 4 SMT threads.
 Bind processor 0 is bound with proc0
 Bind processor 1 is bound with proc0
 Bind processor 2 is bound with proc0
 Bind processor 3 is bound with proc0

proc4 has 4 SMT threads.
 Bind processor 4 is bound with proc4
 Bind processor 5 is bound with proc4
 Bind processor 6 is bound with proc4
 Bind processor 7 is bound with proc4

proc8 has 4 SMT threads.
 Bind processor 8 is bound with proc8
 Bind processor 9 is bound with proc8
 Bind processor 10 is bound with proc8
 Bind processor 11 is bound with proc8

Topas Monitor for host: partoche						EVENTS/QUEUES	FILE/TTY
Fri Nov 27 15:50:05 2015 Interval: 2						Cswitch	1323 Readch 1815.1K
						Syscall	4913 Writech 612.6K
CPU	User%	Kern%	Wait%	Idle%	Physc	Reads	574 Rawin 0
0	81.9	16.7	1.2	0.2	0.41	Writes	363 Ttyout 356
2	0.0	1.0	0.0	99.0	0.08	Forks	6 Igets 0
3	0.0	1.0	0.0	99.0	0.08	Execs	7 Namei 469
4	0.0	43.4	0.0	56.6	0.00	Runqueue	1.0 Dirblk 0
5	0.0	31.7	0.0	68.3	0.00	Waitqueue	0.0
1	0.0	0.9	0.0	99.1	0.08	MEMORY	
6	0.0	0.3	0.0	99.7	0.00	PAGING	Real, MB 24576
11	0.0	0.0	0.0	100.0	0.01	Faults	1554 % Comp 90
7	0.0	0.3	0.0	99.7	0.00	Steals	0 % Noncomp 1
8	0.0	74.6	0.0	25.4	0.01	PgspIn	0 % Client 1
9	0.0	2.3	0.0	97.7	0.01	PgspOut	0
10	0.0	0.0	0.0	100.0	0.01	PageIn	0 PAGING SPACE
						PageOut	0 Size, MB 25600
Network	KBPS	I-Pack	O-Pack	KB-In	KB-Out	Sios	0 % Used 2
Total	218.6	346.0	329.9	92.1	126.5		% Free 98
NFS (calls/sec)							
Disk	Busy%	KBPS	TPS	KB-Read	KB-Writ	Serv2	0 WPAR Activ 0
Total	2.4	2126.9	226.0	1640.6	486.4	Cliv2	0 WPAR Total 0
						Serv3	0 Press: "h"-help
FileSystem	KBPS	TPS	KB-Read	KB-Writ	Cliv3	0	"q"-quit
Total	2.2K	331.3	1.7K	486.3	Serv4	0	
					Cliv4	0	
Name	PID	CPU%	PgSp	Owner			
oracle	9830502	15.1	6.7	orair3			
oracle	14483686	13.9	14.0	orair3			
oracle	26411184	10.0	10.6	orair3			
oracle	6684822	8.6	6.7	orair3			
enserver	29425668	1.3	56.3	ir3adm			
oracle	16580828	0.6	26.8	ir3adm			
oracle	27132004	0.4	8.0	ir3adm			
bgscolle	11403318	0.2	3.3	bmcpor			
sapstart	16187400	0.2	22.3	ir3adm			
init	1	0.1	0.8	root			
PatrolAg	8454164	0.0	15.7	patrol			
syncd	3211376	0.0	0.6	root			
lrud	262152	0.0	0.6	root			
gil	1769526	0.0	0.9	root			
getty	4194474	0.0	0.6	root			
nfssync_	3604592	0.0	0.7	root			
random	4587558	0.0	0.4	root			
vmmd	458766	0.0	0.8	root			
nfsd	4915360	0.0	1.8	root			
bdaemon	7471354	0.0	1.8	root			

From:
<https://unix.ndlp.info/> - Where there is a shell, there is a way

Permanent link:
https://unix.ndlp.info/doku.php/informatique:nix:ibm:ibm_aix_powervm:configuration

Last update: 2015/11/27 14:55