



VELOCITY
S O F T W A R E

Performance Case Studies – Performance Analysis

Barton Robinson,
barton@velocitysoftware.com

Velocity Software Inc.
196-D Castro Street
Mountain View CA 94041
650-964-8867

Velocity Software GmbH
Max-Joseph-Str. 5
D-68167 Mannheim
Germany
+49 (0)621 373844



Flow chart process for Performance Analysis

State sampling and how to use it

Sample Analysis

Case Parking

Case SMT (How do I read this?)

Case SMT2 – a good one?

Case Tyrone (is java application a problem?)

Case Tyrone2 (Why do I have an Eligible list?)

“z” is:

- Very large,
- Very complex and
- Very well instrumented

The challenge?

- Performance problems are visible,
- “z” applications are often impacted by other applications

What challenge, it is all there!

- 200 zmon panels (with menus)
- 150 zmap reports (with table of contents)
- 5,000 unique variables

The 20 Reports you need in the order you need them

Configuration: ESAHDR

System (z/VM LPAR) Load: ESASSUM

Wait states: ESAXACT

Virtual machine Config: ESAUSRC

CPU:

- **Lpar Configuration: ESALPARS**
- **CPU Utilization: ESACPUU**
- **CPU Consumer: ESAUSP2**
- **Linux Consumer: ESALNXP**

Storage

- **z/VM Requirements: ESASTR1**
- **User Storage: ESAUSPG**
- **Linux Storage: ESAUCD2**
- **VDISK for swap: ESAASPC**

Paging

- **Configuration: ESAPSDV**
- **Loads by user: ESAUSPG**
- **Blocking: ESABLKP**

DASD

- **Configuration: ESADSD1**
- **Rates: ESADSD2**
- **CACHE: ESADSD5**

Network

- **Configuration: ESATCPI**
- **Network management: ESATCP1/2/4**

Analysis starts with “is there a problem?”

- Describe the problem (what user(s), what time)

System Configuration

- Processor model, cpu type
- Number of processors, storage size
- SMT support? (Increased level of complexity)
- Linux using IFLs in mixed mode LPAR?

High level Loads on the subsystems

Wait states for those impacted

Subsystem Analysis

- DASD, Storage, Paging, Processor, Network

Report: ESAHDR z/VM Monitor Analysis
Monitor period: 240 seconds (4:00)

z/VM Version: 6 Release 3.0 SLU 1401
TOD clock at last IPL: 08/02/15 00:56:05
System Operator: OPERATOR
Time zone adjustment from GMT: -4 hours

System Identifier VS0P2
Checkpoint/Warmstart Volumes RESP26/RESP26
Machine Model/Type EC12:2827/712
System Sequence Code 0000000000082BA7
Processor 0 model/serial 2827-712 /3xxxA7
Processor 1 model/serial 2827-712 /3xxxA7
Processor 2 model/serial 2827-712 /3xxxA7
Processor 3 model/serial 2827-712 /3B2BA7 Master

Processor 14 model/serial 2827-712 /3xxxA7
Processor 15 model/serial 2827-712 /3xxxA7
Processor 16 model/serial 2827-712 /3xxxA7
Processor 17 model/serial 2827-712 /3xxxA7

Configuration Errors just happen Common configuration problems

- z/VM Release significant
- IFLs?
- Real Storage Size
- Master processor can be significant
- Vertical problematic

Power of processor in terms of service Units: 59701
CPU Capability Factor: 552
CPU(GP) Capability Factor: 552
Operating on IFL Processor(s)
Channel Path Measurement Facility(CPMF) Extended is installed
Main Storage installed (MB): 421888
Main Storage Generated (MB): 421888
Horizontal/Vertical Scheduling Configuration IFL CPUs
Confidence Percent5F
Algorithm 02
Confidence 5A

Know the z/VM LPAR overall loads: ESASSUM / ESAMAIN

Report: ESASSUM Subsystem Activity

```

-----
          <---Users----> Transactions <Processor> Storage (MB) <-Paging-->
          <-avg number->   Per   Avg. Utilization Fixed Active <pages/sec>
Time      On Actv In Q   Minute   Resp Total Virt.   User Resid. XStore DASD
-----
17:05:00   85   50 85.0   247.0 0.008   271   249 186.1 381973     0    0
17:06:00   85   50 86.0   249.0 0.008   269   244 186.1 381973     0    0
17:07:00   85   53 89.0   250.0 0.008   266   243 186.1 381989     0    0
17:08:00   85   56 84.0   244.0 0.007   181   161 186.1 381996     0    0
*****Summary*****
Average:   84   52 86.0   247.5 0.008   247   224 186.1 381981     0    0
    
```

- z/VM LPAR Loads**
- Processor utilization
 - Storage consumed by users
 - Page rates to DASD, Expanded Storage
 - DASD I/O rates
 - (Transactions are for traditional workloads)

- Any changes? Spikes?
what time?**
- Processor
 - Storage for users
 - Page rates
 - DASD I/O rates

Wait states provide direction for investigation

Wait state (queue) analysis -> “Interesting”

- **Running / CPU Wait** -> CPU Subsystem
- **Simulation wait** -> Master Processor (CPU)
- **Page wait / asynchronous** -> Paging/Storage subsystems
- **SIO / Asynchronous i/o** -> DASD subsystem
- **Eligible / Loading (LDG)** -> special state, loading working set

Normal idle wait states, “not interesting”

- TCPIP, Linux: test idle
- Traditional servers: SVM (service machine wait)
- Traditional users: idle (not in queue)

Compare “running” vs “interesting”

- Anything “Interesting” more than 1% is Interesting,
- **ELIGIBLE ALWAYS Interesting**

Two types of Wait states are provided by virtual machine

- Monitor frequency: once per minute (**Not interesting**)
- High Frequency (Hi-Freq): once per second (**Interesting**)
 - (60 samples per 1 minute per virtual cpu)

Shown by:

- Summarized for all users (**start here**)
- Summarized for user classes (grouped by installation)
- Servers
- Top users

User class analysis -> where to focus

- Group Test vs Product
- Group application by application
- Group support servers vs production

Report: ESAXACT Transaction Delay Analysis

```

-----
<-----Percent non-dormant (Wait states)-----
UserID  <-Samples->          E-  D-  T-      Tst <Asynch>
/Class  Total  In Q  Run Sim CPU SIO Pag SVM SVM SVM  CF Idl I/O Pag Ldg
-----
UserID  <-Samples->          E-  D-  T-      Tst <Asynch>
/Class  Total  In Q  Run Sim CPU SIO Pag SVM SVM SVM  CF Idl I/O Pag Ldg
-----
17:05:00      85      85  8.2  2.4  5.9    0    0    0    0    0    0    84    0    .    .
Hi-Freq:  9000   4902  3.2  0.2  0.4    0    0    0    5.3  0.4    0    96  0.1    0    0
***Key User Analysis ***
TCPIP        60      23    0    0    0    0    0    0    0    0    0    0  100    0    0    0
***User Class Analysis***
Servers     1020      23    0  4.3    0    0    0    0  3.7  70    0    26    0    0    0
ZVPS        600      13    0    0    0    0    0    0    0    0    0    100    0    0    0
TheUsers    7080   4843  3.2  0.2  0.4    0    0    0  5.0  0.1    0    96  0.1    0    0
***Accounting Code Analysis**
LNXSAP     4920   4829  3.2  0.1  0.4    0    0    0    0    0    0    0    96  0.1    0    0
***Top User Analysis***
LNXP22     480     479   15  1.5  0.6    0    0    0    0    0    0    0    82  0.2    0    0
LNXP14     480     430   9.3    0  0.9    0    0    0    0    0    0    0    90    0    0    0
LNXP12     480     468   1.1    0  0.2    0    0    0    0    0    0    0    99  0.2    0    0
LNXPMD11   480     475   1.3    0  0.2    0    0    0    0    0    0    0    99    0    0    0
LNXP41     120     120   1.7    0    0    0    0    0    0    0    0    0    98    0    0    0
LNXP3P2    240     240   1.7    0  0.4    0    0    0    0    0    0    0    98    0    0    0
LNXP3P4    240     237   1.7    0    0    0    0    0    0    0    0    0    98    0    0    0
LNXPBWP4   360     360   1.7    0    0    0    0    0    0    0    0    0    98    0    0    0
  
```

“HI-Freq:” - Wait states for all users
 Key users – as defined by installation
 User Class – as defined by installation
 Top “n” users – Note two vcpu has 2 times samples

User Configuration: ESAUSRC

Report: ESAUSRC User Configuration Velocity Software

UserID	ClassID	Account Code	CPU Type	<-----SHARE----->				<---CPU <Count>		<Storage>		
				<Normal> Rel	Abs	<--MAX--> Typ	Shre	Lim -it	Def On	Dflt	Max	
FTPSERVE	Servers	FTPSERVE	IFL	100	1	1	32M	32M
LNXBWP2	TheUsers	LNXSAP	IFL	100	6	6	43G	43G
LNXBWP4	TheUsers	LNXSAP	IFL	100	6	6	43G	43G
LNXBWP6	TheUsers	LNXSAP	IFL	100	6	6	42G	42G
LNXEP12	TheUsers	LNXSAP	IFL	100	8	8	44G	44G
LNXEP14	TheUsers	LNXSAP	IFL	100	8	8	32G	32G
LNXEP32	TheUsers	LNXSAP	IFL	100	4	4	34G	34G
LNXEP34	TheUsers	LNXSAP	IFL	100	4	4	34G	34G
LNXEP42	TheUsers	LNXSAP	IFL	100	2	2	16G	16G
LNXGFT01	TheUsers	LNXSAP	IFL	100	4	4	16G	16G
LNXGFT02	TheUsers	LNXSAP	IFL	100	4	4	16G	16G
LNXGP12	TheUsers	LNXSAP	IFL	100	2	2	6.0G	6.0G
LNXMD11	TheUsers	LNXSAP	IFL	100	8	8	48G	48G
LNXPP41	TheUsers	LNXSAP	IFL	100	2	2	16G	16G
LNXR3P2	TheUs											
LNXR3P4	TheUs											
LNXP22	TheUs											
LNXT42	TheUs											
OPERATNS	Server											

Look for “Interesting configurations”

- Large relative shares / absolute shares
- CPU Counts, matching shares (100 Rel / vcpu)
- CPU Type (IFL, CP)
- Virtual machine storage sizes (too large?, largest?)

From State Sampling (ESAXACT), If CPU Wait, or “running” is dominant state, then:

Top down:

- CEC / LPAR (ESALPARS)
- z/VM (ESACPUU)
- Virtual machine (ESAUSP2)
- Linux process (ESALNXP)

CPU Capture ratio 100% down to process

LPAR Configuration: ESALPARS

Report: ESALPARS Logical Partition Summary Velocity Software Corporate

```

-----
<--Complex--> <-----Logical Partition-----> <-Assigned Shares---->
Time      Phys Dispatch      Virt CPU  <%Assigned>  <---LPAR--> <VCPU Pct>  Cap-  Wait
          CPUs      Slice Name    Nbr CPUs  Type  Total  Ovhd  Weight  Pct  /SYS /CPU  ped  Comp
-----
17:05:00   69   Dynamic
Totals:    00   33   CP   364.1   8.2   1016  100
Totals:    00   50  IFL   822.8  20.3  1000  100
FR9VS0P2   3B   18  IFL   280.5   7.0   405  40.5  2.25  56.2  No   No
FR9IWAS    0E    1  ICF   100.0   0.0   Ded   1.4    0    0    No   Yes
FR9WAS4    0B    5  IFL    0      0     Ded   7.2    0    0    No   Yes
FR9WAS6    01    2   CP    18.1    0.7   131  12.9  6.45  77.4  No   No
FR9WAS6    01    3  ZAP    46.1    0.2   117  27.2  9.07  118   No   No
FR9S0PB    3A    9  ZIP    36.2    0.3   273  33.2  3.69  40.5  No   No
FR9VS0P0   38   19  IFL   408.3   7.1   460  46.0  2.42  60.5  No   No
FR9VS0Q0   36   13  IFL   134.1   6.2   135  13.5  1.04  26.0  No   No
FR9WAS2    05    2   CP    7.6    0.3    16   1.6   0.79  9.45  No   No
FR9WAS2    05    2  ZAP   10.7    0.1    26   6.0   3.02  39.3  No   No
FR9ISRV    13    1  ICF   97.5    0.2    50  50.0  50.0  50.0  No   No
    
```

Look for “Shared processors”

- First LPAR is “us”, z/vm where data collected (FR9VSOP2)
- Check weights (405/1000 = 40.5% of IFLs)
- Dedicated / non-shared

LPAR Configuration: ESALPARS

```
Report: ESALPARS          Logical Partition Summary          Velocity Software Corporate
-----
<--Complex--> <-----Logical Partition-----> <-Assigned Shares---->
      Phys Dispatch          Virt CPU  <%Assigned>  <---LPAR--> <VCPU Pct>  Cap-  Wait
Time    CPUs    Slice Name      Nbr CPUs Type  Total  Ovhd  Weight  Pct  /SYS  /CPU  ped   Comp
-----
17:05:00    69  Dynamic Totals:    00   33  CP  364.1   8.2   1016  100
```

```

Totals by Processor type:
<-----CPU-----> <-Shared Processor busy->
Type Count Ded shared  Total  Logical Ovhd Mgmt
-----
CP      12   0    12  377.6   355.9   8.2  13.5
ZAP     13   0    13  186.4   183.4   1.1   1.9
IFL    30   5    25  847.4   802.5  20.3  24.5
ICF     3   2     1  100.0    97.5   0.2   2.3
ZIIP    11   0    11  168.0   163.9   1.4   2.8
```

Look for cpu configuration

- 30 IFLs, 847% busy,
 - logical overhead 20%,
 - physical overhead 24%

Consumers within LPAR: ESAUSP2

Report: ESAUSP2		User Resource Rate Report								Velocity Software C				
<---CPU time-->			<----Main Storage (pages)----->					<-----Paging (pages)----->						
UserID	<(Percent)>	T:V	<Resident>	Lock	<-----WSS----->	<---Allocated--->	<Pgs/Secnd>							
/Class	Total	Virt	Rat	Totl	Activ	-ed	Totl	Activ	Avg	Total	ExStg	Disk	Read	Write
17:05:00	258.2	249.4	1.0	98M	97.8M	90	98M	97.8M	1M	91942	0	91942	0	0
***Key User Analysis ***														
TCPIP	0.14	0.06	2.5	4519	4519	10	3518	3518	3518	2416	0	2416	0	0
User Class Analysis														
Servers	0.05	0.03	1.4	33K	16535	19	33K	16401	965	26914	0	26914	0	0
ZVPS	0.28	0.25	1.1	13K	6698	0	13K	6696	670	338	0	338	0	0
TheUsers	257.7	249.0	1.0	98M	97.8M	20	98M	97.7M	2M	56437	0	56437	0	0
***Accounting Code Analysis**														
LNXSAP	256.4	247.8	1.0	98M	97.7M	12	98M	97.7M	6M	9800	0	9800	0	0
Top User Analysis														
LNXP22	120.5	114.3	1.1	11M	10.8M	0	11M	10.8M	11M	1	0	1	0	0
LNXP14	69.47	69.14	1.0	5.2M	5247K	0	5.2M	5245K	5M	5664	0	5664	0	0
LNXP12	9.86	9.56	1.0	7.9M	7889K	0	7.9M	7887K	8M	4121	0	4121	0	0
LNXP11	9.31	9.18	1.0	13M	12.5M	3	13M	12.5M	13M	1	0	1	0	0
LNXP41	6.85	6.67	1.0	4.2M	4185K	0	4.2M	4184K	4M	1	0	1	0	0

Look for consumers, in percent of cpu

- **By class (SAP Workload, defaults TheUsers)**
- **Abusive servers ?(LNXP22)?**
- Correct per expected? Not a performance question

Linux Process Load: ESALNXP

Report: ESALNXP LINUX HOST Process Statistics Report Veloc
Monitor initialized: 10/31/16 at 17:03:00 on 2827 serial 3B2BA7 First

```
-----  
node/      <-Process Ident->  Nice PRTY <-----CPU Percents----->  
Name       ID      PPID  GRP  Valu Valu  Tot  sys user syst usrt  
-----  
lplsp22      0        0     0    0    0   103 11.9 90.5 0.34 0.45  
cybAgent    3904     1   3904    0   20  0.12 0.08 0.03  0    0  
BESCLien    3939     1   3938    0   20  0.39 0.27 0.12  0    0  
SP2_01_D    4123  15839  4123    0   20  4.55 0.37 4.18  0    0  
saphoste    4316     1   4315    0   20  0.40  0    0  0.12 0.29  
sapstart    4636     1   4636    0   20  0.35 0.05 0.30  0    0  
snmpd       5041     1   5040   -10  10  0.10 0.05 0.05  0    0  
SP2_01_D    8092  15839  8092    0   20  1.92 0.20 1.72  0    0  
SP2_01_B    8784  15839  8784    0   20  20.5 1.42 19.0  0    0  
SP2_01_D    9578  15839  9578    0   20  3.94 2.59 1.35  0    0  
jstart     10761   5211   5211    0   20  0.35 0.05 0.30  0    0  
java       15136     1  15120    0   20  21.7 0.71 21.0  0    0  
SP2_01_D   15839  15699  15839    0   20  0.37 0.02 0.03  0.19 0.13  
jlaunch    25866  15871  15839    0   20  3.29 0.39 2.90  0    0  
SP2_01_D   31890  15839  31890    0   20  38.5 4.26 34.3  0    0
```

Look for processes within Linux, in percent of CPU

- By relevant server (LNXLSP22)
- Correct? Relevant? Cron? Agents?

Report: ESALNXP

LINUX HOST Process Statistics Report

node/ Name	<-Process ID	Ident-> PPID	GRP	<-----CPU Tot	Percents-----> sys	user	syst	usrt	<Stg (k)> Size	RSS
LNXUWA01	0	0	0	67.0	5.98	59.0	1.20	0.81	13M	9M
java	4444	1	4444	1.10	0.07	1.03	0	0	1M	801K
kd4agent	5576	1	4362	4.71	1.68	3.03	0	0	99K	64K
kynagent	9569	1	4362	2.48	0.07	2.41	0	0	314K	212K
kcawd	9634	1	4362	1.92	0.01	0.01	1.14	0.75	37K	6936
java	10547	1	10547	0.82	0.07	0.75	0	0	870K	743K
java	11751	4877	4877	0.57	0.07	0.50	0	0	617K	98K
java	11837	1	11837	3.28	0.12	3.16	0	0	3M	1M
java	21374	15199	21374	46.3	3.07	43.2	0	0	3M	3M
java	24567	1	24567	2.27	0.18	2.09	0	0	1M	831K
java	28060	1	28060	1.23	0.09	1.14	0	0	1M	821K
java	32428	1	32428	1.17	0.10	1.07	0	0	810K	538K

Look for processes within Linux, in percent of CPU

- Look for rogue administrators running agents
 - 10% per server adds up to an IFL every 10 servers just for agents....
- Correct? Relevant? Cron? Agents?

Spike in paging? Look at Storage.

Top down:

- z/VM (ESASTR1)
- Virtual machines (ESAUSPG)
- VDISK / MDC / Address Space (ESAASPC)
- Linux server (ESAUCD2)
- Linux process (ESALNXP)

Report: ESASTR1 Main Storage Analysis Velocity Software Corporate ZMAP 4.

Time	Users	Loggd System		name <Available>		System	User	NSS/DCSS	<-AddSpace>		VDISK	<MDC>
	<-----MegaBytes----->	On	Storage	able	<2gb	>2gb	ExSpc	Resdnt	Resident	System	User	Rsdnt
17:05:00	85	421888	3296	335	23523	55	382107	562	4043	0	7420	236
17:06:00	85	421888	3296	335	23524	55	382106	562	4043	0	7420	236
17:07:00	85	421888	3296	335	23523	55	382106	562	4043	0	7420	237

Total storage analysis (in Megabytes – previously pages)

- System Storage (420GB)
- MDC 236mb? SET MDC MAX/MIN
- User resident should be large percent of “System Storage”
- VDISK stable?
- Large available means unconstrained

Virtual Machine Storage : ESAUSPG

Report: ESAUSPG		User Storage Analysis						Velocity Software Corporate				
		<---Storage occupancy in pages--->			<--Main Storage page			Read/Write-->		<Address		
UserID	<---Main Storage---	<---Paging---		<-Page Writes to:-->		<Page Reads:>		<pages R				
/Class	Total	>2gb	<2GB	Xstor	DASD	Xsto	Disk	Migr	Xstor	Disk	VirtDisk	
11:00:00	17448K	16943K	504640	4346K	8891K	1120K	352582	320630	822546	149628	237286	
Top User Analysis												
LNXUWA01	2889K	2798K	90725	65398	258675	10999	112	0	5390	13806	0	
LNXUWA03	3848K	3762K	85186	63975	8378	21875	277	0	221201	6714	223173	
LNXUWA02	685385	648345	37040	296256	84613	36427	2443	0	22943	1983	0	
LNXQWA01	1246K	1218K	28190	541178	51075	35529	2727	0	14094	2787	1428	
LNXDWA02	713091	672702	40388	56215	148406	16314	649	0	451	1828	0	
LNXDWA04	1152K	1120K	31859	592756	96720	13708	63725	63261	1189	942	0	
LNXDWA03	330601	324021	6581	4194	39207	3926	5601	5345	120	734	8	
LNXTWA04	883228	860363	22865	90734	129722	7768	31	0	182	66	1889	
LNXUWA15	693689	664995	28694	53516	137150	10556	1382	0	553	457	0	

Example from storage constrained system:

Total storage analysis (in pages, new “megabyte” option)

- Largest consumer(s) resident storage
- Largest consumer - which virtual disk?
- VDISK Spike (1gb) ? Which server?
- (LNXUWA03 was in Asynchronous I/O Wait)

VDISK for Swap: ESAVDSK (Or ESAASPC)

Report: ESAVDSK VDISK Analysis Report Velocity Software Corporate

Owner	Space Name	-----<--Size--->	<AddSpce>	Priv	VIO	<--pages-->							
		AddSpc	VDSK	Cre-	Del-	or	rate	User	Resi-	Lock-	Sto-	DASD	
		Pages	Blks	ates	etes	Shrd	/sec	Links	dent	ed	len	Read	
-----<--Size--->													
10:45:00													
LNXQWA01	VDISK\$LNQWA01\$0206\$0530	64256	512K	0	0	Shrd	0.00	1	122	0	0.7	0.0	
LNXQWA01	VDISK\$LNQWA01\$0207\$0531	64256	512K	0	0	Shrd	0.04	1	2565	0	3.5	0.2	
LNXTWA04	VDISK\$LNXTWA04\$0206\$051C	131K	1049K	0	0	Shrd	1.28	1	11K	0	0	0.0	
LNXUWA03	VDISK\$LNXUWA03\$0206\$051E	250K	2002K	0	0	Shrd	0.65	1	14K	0	1.6	6.7	
LNXUWA03	VDISK\$LNXUWA03\$0207\$051F	375K	3002K	0	0	Shrd	0.29	1	4980	0	0.4	0.7	
LNXUWA03	VDISK\$LNXUWA03\$0208\$0520	513K	4102K	0	0	Shrd	0.28	1	4751	0	0.4	0.4	
-----<--Size--->													
System Totals:		7805K	125M	0	0	.	5.09	204	46K	0	7.3	8.1	
-----<--Size--->													
11:00:00													
LNXQWA01	VDISK\$LNQWA01\$0206\$0530	64256	512K	0	0	Shrd	0	1	46.9	0	0.1	0	
LNXQWA01	VDISK\$LNQWA01\$0207\$0531	64256	512K	0	0	Shrd	0	1	1381	0	0.3	0	
LNXTWA04	VDISK\$LNXTWA04\$0206\$051C	131K	1049K	0	0	Shrd	0	1	3984	0	11.7	0	
LNXUWA03	VDISK\$LNXUWA03\$0206\$051E	250K	2002K	0	0	Shrd	10.1	1	46K	0	12.9	58.4	
LNXUWA03	VDISK\$LNXUWA03\$0207\$051F	375K	3002K	0	0	Shrd	16.2	1	88K	0	6.1	19.7	
LNXUWA03	VDISK\$LNXUWA03\$0208\$0520	513K	4102K	0	0	Shrd	16.1	1	88K	0	5.8	20.2	
-----<--Size--->													
System Totals:		7805K	125M	0	0	.	84.6	204	237K	0	37.2	98.3	

Virtual Disk Analysis

- Which virtual disk spiked?
- Are there multiple vdisks, and PRIORITIZED!!!

If Paging increased

- Check Virtual disk resident storage increase
- Associated with one linux server. (LNXUWA03)
- 2nd largest server in terms of cpu

CPU utilization?

- Understand agents using 10% per server
- SAP / Java applications

The problem???

Servers waiting for CPU, CPU is available?

CASE Parking. Subsystem Activity

Report: ESASSUM	Subsystem Activity						Velocity Software				
Time	<---Users---> <-avg number-> On Actv In Q			<Processor> Utilization Total Virt.		Storage (MB) Fixed Active User Resid.		<-Paging--> <pages/sec> XStore DASD		<-----I/O- <-DASD--> Rate Resp	
12:05:00	146	107	220	1459	1262	214.4	491672	0	0	350	0.5
12:06:00	146	106	226	1488	1285	214.4	491672	0	0	349	0.5
12:07:00	146	107	218	1424	1245	214.4	491676	0	0	357	0.5
12:08:00	146	104	225	1396	1196	214.4	491664	0	0	345	0.5
12:09:00	146	103	221	1456	1260	214.4	491660	0	0	353	0.5
12:10:00	146	107	223	1396	1165	214.4	491674	0	0	348	0.4
12:11:00	146	104	223	1397	1217	214.4	491663	0	0	347	0.6
12:12:00	146	105	222	1322	1134	214.4	491671	0	0	353	0.4
12:13:00	146	103	225	1315	1112	214.4	491661	0	0	363	0.5
12:14:00	146	103	226	1258	1061	214.4	491662	0	0	346	0.4
12:15:00	146	106	224	1434	1244	214.4	491674	0	0	347	0.5
12:16:00	146	104	226	1459	1258	214.4	491666	0	0	380	0.6
12:17:00	146	105	221	1443	1248	214.4	491674	0	0	348	0.6

CPU Consistent, no CPU Spikes
No paging, DASD I/O static
Nothing interesting

User Configuration: ESAUSRC

Report: ESAUSRC

User Configuration

Velocity SoftwareMAP 4.3.0 0

```
<-SHARE-><--CPU--->< <Storage>
```

UserID	ClassID	Account Code	CPU Type	<Normal> Rel	<Normal> Abs	Count ef	On	Mode	<VM Size> Dflt	<VM Size> Max
DTCSMAPI	TheUsers	DTCSMAPI	IFL	3000	.	1 1	ESA		32M	32M
DTCVSW1	Servers	DTCVSW1	IFL	5000	.	1 1	ESA		32M	32M
DTCVSW2	Servers	DTCVSW2	IFL	5000	.	1 1	ESA		32M	32M
TSTDD002	TheUsers	EV500000	IFL	2000	.	5 5	ESA		16G	16G
TSTDD004	TheUsers	EV500000	IFL	2000	.	5 5	ESA		16G	16G
TSTDD006	TheUsers	EV500000	IFL	2000	.	5 5	ESA		16G	16G
TSTDD008	TheUsers	EV500000	IFL	2000	.	5 5	ESA		16G	16G
TSTDD010	TheUsers	EV500000	IFL	2000	.	5 5	ESA		16G	16G
FTPSEVE	Servers	FTPSEVE	IFL	100	.	1 1	XC		32M	32M
HTTPD	TheUsers	12345678	IFL	1000	.	1 1	ESA		32M	32M
IPGATE	TheUsers	IPGATE	IFL	100	.	1 1	ESA		256M	256M
LNx1A004	TheUsers	HWSE0000	IFL	2000	.	4 4	ESA		6.0G	6.0G
LNx1A006	TheUsers	HWSE0000	IFL	1200	.	4 4	ESA		7.5G	7.5G
LNx1A008	TheUsers	HWSE0000	IFL	2000	.	4 4	ESA		4.0G	4.0G
LNx1A160	TheUsers	HWSE0000	IFL	2000	.	4 4	ESA		6.0G	6.0G
LNx1A168	TheUsers	HWSE0000	IFL	2000	.	4 4	ESA		6.0G	6.0G
LNx1A170	TheUsers	HWSE0000	IFL	1200	.	4 4	ESA		7.5G	7.5G
LNx1A172	TheUsers	HWSE0000	IFL	2000	.	4 4	ESA		6.0G	6.0G
LNx1A178	TheUsers	HWSE0000	IFL	2000	.	4 4	ESA		6.0G	6.0G
LNx1A186	TheUsers	HWSE0000	IFL	2000	.	4 4	ESA		6.0G	6.0G
LNx1A188	TheUsers	HWSE0000	IFL	1200	.	4 4	ESA		7.5G	7.5G
LNx1A196	TheUsers	HWSE0000	IFL	2000	.	4 4	ESA		6.0G	6.0G
LNx1A200	TheUsers	HWSE0000	IFL	2000	.	4 4	ESA		6.0G	6.0G

Look for “Interesting configurations”

- Norm is 500/vcpu
- Large relative shares / absolute shares
- CPU Counts, matching shares (500 Rel / vcpu)
- CPU Type (IFL, CP)
- Virtual machine storage sizes?

CASE Parking. Wait State Analysis

Report: ESAXACT

Transaction Delay Analysis

Veloc

-----Percent non-dormant (Wait states)-----															
UserID /Class	<-Samples->											Tst <Asynch>			
	Total	In Q	Run	Sim	CPU	SIO	Pag	E- SVM	D- SVM	T- SVM	CF	Idl	I/O	Pag	Ldg
12:05:00	146	220	8.2	0	4.1	0	0	0	0	0.5	0	87	0	.	.
Hi-Freq:	17820	13385	6.6	0.1	4.1	0.0	0	0	3.3	0.3	0	89	0.1	0	0
Hi-Freq:	17820	13323	6.6	0.2	5.3	0.0	0	0	3.4	0.4	0	87	0.0	0	0
Hi-Freq:	17820	13397	6.2	0.2	4.1	0.0	0	0	3.3	0.6	0	89	0.1	0	0
Hi-Freq:	17820	13369	5.7	0.2	2.7	0.0	0	0	3.6	0.5	0	91	0.1	0	0
Hi-Freq:	17820	13369	5.9	0.2	3.9	0.0	0	0	3.7	0.5	0	89	0.1	0	0
Hi-Freq:	17820	13359	5.5	0.2	3.1	0.0	0	0	3.3	0.6	0	90	0.0	0	0
Hi-Freq:	17820	13328	5.8	0.1	3.9	0.0	0	0	3.5	0.7	0	89	0.1	0	0
Hi-Freq:	17820	13311	5.1	0.2	2.1	0.0	0	0	3.7	0.5	0	92	0.1	0	0
Hi-Freq:	17820	13344	5.9	0.2	2.8	0.0	0	0	3.5	0.6	0	90	0.1	0	0
Hi-Freq:	17820	13385	5.5	0.1	3.3	0.0	0	0	3.5	0.7	0	90	0.1	0	0
Hi-Freq:	17820	13388	6.2	0.2	4.4	0.0	0	0	3.2	0.6	0	89	0.0	0	0
Hi-Freq:	17820	13418	6.1	0.2	4.1	0.0	0	0	3.0	0.5	0	89	0.1	0	0

Edit the ESAXACT report, say "ALL /Hi-F/", see totals

Any obvious waits?

Running is good,

cpu wait is only on significant

Report: ESALPARS Logical Partition Summary Velocity Software
 Monitor initialized: 01/2/25/17 at 12:03:27 on 2964

Time	Phys CPUs	<-----Logical Partition----->						<-Assigned Shares---->				
		Name	Nbr CPUs	Virt CPU Type	<%Assigned> Total	Ovhd	<---LPAR--> Weight	Pct	<VCPU Pct> /SYS	/CPU		
12:05:00	45	Totals:	00	84	IFL	3787	125	1000	100			
		V1TST1	3F	28	IFL	1519	40.5	370	37.0	1.32	59.5	
		P1PRD1	36	20	IFL	726.4	28.1	200	20.0	1.00	45.0	
		V1TST2	3E	8	IFL	430.6	17.8	60	6.0	0.75	33.7	SMT
		V1TST3	3C	28	IFL	1111	38.6	370	37.0	1.32	59.5	

Look for “Shared processors”

- First LPAR is “us”, z/vm where data collected (V1TST1)
 - Using 15 out of 28 VCPU
- IFLs shared between LPARs
- Check weights
- “pct/CPU” is how fast cpu will run when all are busy
- V1TST1 is guaranteed by weights 37% of 45, 16 IFLs.

Totals by Processor type:

```
<-----CPU-----> <-Shared Processor busy->
Type Count Ded shared Total Logical Ovhd Mgmt
-----
IFL      45   0      45 4039.2 3662.3 125 252
```

V1TST1 using 15 virtual CPUs out of 28 available

45 REAL IFLs available, using 40.3

IFLs running at 90%

If there is CPU available, vcpu available,

- **why are users waiting for CPU**

CASE Parking. Vertical scheduling analysis

Report: Report: ESACPUU CPU Utilization Report Velocity Software
 Monitor 1/25/17 at 12:03:27 on 2964 serial

```

-----
<-----CPU (percentages)----->d <-Vertical-->
Time      CPU   CPU   Total  Emul   User   Sys   Idle  Steal  Entitle- Park
-----  ---  ----  ----  ----  ----  ----  ----  ----  -----  -----
12:05:00  0   IFL   80.3  71.5   2.8   6.0   17.2   2.5   1.00    0
          1   IFL   78.2  67.9   3.0   7.3   19.0   2.8   1.00    0
          2   IFL   78.6  68.8   2.9   6.9   18.6   2.8   1.00    0

          16  IFL   62.0  49.6   3.0   9.5   18.9   19.1   0.65    0
          17  IFL   62.4  50.4   2.9   9.1   18.7   18.8   0        0
          18  IFL   38.8  22.2   3.4   13.3  42.5   18.6   0        2.0
          19  IFL   23.2  15.5   1.5   6.2   29.6   47.2   0       23.2
          20  IFL   12.3   8.8   0.7   2.8   16.6   71.1   0       40.0
          21  IFL    5.5   4.3   0.2   1.0   6.1   88.4   0       52.0
          22  IFL    0.0    0     0     0.0   0    100.0   0       60.0

          27  IFL    0.0    0     0     0.0   0    100.0   0       60.0
-----  ---  ----  ----  ----  ----  ----  ----  ----  -----  -----
System:   1459  1262  53.4  143.0  438.5  902.9  16.65  477.1
  
```

LPAR Weight guarantees 16 cpu
 Hiperdispatch parks the rest

- **EVEN IF CPU AVAILABLE**



CASE Parking. Virtual Machine Consumption Analysis

Report: ESAOPER Operator/System Log
Monitor initialized: 01/25/17 at 12:03:27 o

```
-----  
12:03:31 CPU Park from 20 to 18 CPUUtil= "14.1", Projected= "15.6"  
12:03:33 CPU Unpark from 18 to 19 CPUUtil= "14.5", Projected= "15.6"  
12:03:35 CPU Park from 19 to 17 CPUUtil= "14.9", Projected= "15.6"  
12:03:37 CPU Unpark from 17 to 19 CPUUtil= "12.9", Projected= "15.9"  
12:03:41 CPU Unpark from 19 to 20 CPUUtil= "15.5", Projected= "17.2"  
12:03:43 CPU Park from 20 to 17 CPUUtil= "15.7", Projected= "16.1"  
12:03:45 CPU Unpark from 17 to 19 CPUUtil= "13.3", Projected= "16.2"  
12:04:15 CPU Unpark from 19 to 20 CPUUtil= "13.5", Projected= "16.7"  
12:04:21 CPU Unpark from 20 to 21 CPUUtil= "15.1", Projected= "16.6"  
12:04:23 CPU Unpark from 21 to 22 CPUUtil= "15.1", Projected= "16.6"  
12:04:29 CPU Park from 22 to 20 CPUUtil= "16.5", Projected= "18.0"  
12:04:31 CPU Unpark from 20 to 22 CPUUtil= "15.2", Projected= "16.2"  
12:04:33 CPU Park from 22 to 18 CPUUtil= "14.7", Projected= "16.2"  
12:04:35 CPU Unpark from 18 to 21 CPUUtil= "14.0", Projected= "16.1"
```

Parking monitor record is every 2 seconds

A lot of parking when not wanted?

CASE Parking. Virtual Machine Consumption Analysis

Report: ESAMFC MainFrame Cache Magnitudes

Monitor initialized: 01/25/17 at 12:03:27 on 296

```

-----
                <CPU Busy> <-----Processor----->
                <percent>  Speed/<-Rate/Sec->
Time           CPU Totl User  Hertz Cycles Instr Ratio
-----
12:05:02      0 80.3 71.5  5000M  4017M 1976M 2.033
                1 78.2 67.9  5000M  3912M 1741M 2.247
                2 78.6 68.8  5000M  3930M 1744M 2.253
                3 76.7 67.4  5000M  3837M 1983M 1.934

                16 62.0 49.6  5000M  3099M  585M 5.301
                17 62.4 50.4  5000M  3119M  606M 5.151
                18 38.8 22.2  5000M  1962M  483M 4.062
                19 23.2 15.5  5000M  1166M  322M 3.620
                20 12.3  8.8  5000M   616M  182M 3.377
                21  5.5  4.3  5000M   278M  101M 2.748
                22  0.0  0  5000M   700.7  8.15 85.97
                23  0.0  0  5000M   648.6  8.15 79.58

                27  0.0  0  5000M   564.0  8.15 69.20
-----
System:        1459 1262  5000M  68.0G 31.3G 2.172
    
```

Quick check on MFC
- 2.1 cycles per instruction is "OK"

- 5.3 is NOT OK

Vertical supposed to help by localizing work

CPU BUSY is "cycles used per time", NOT instructions executed

No issues

System doing what you told it to,

- Just not what you wanted

Validate LPAR Weights???

Use HORIZONTAL Scheduling.....

The question:

- Is SMT good?
- Is there CPU capacity issue?

Report: ESAHDR z/VM Monitor Analysis
Monitor initialized: 02/23/17 at 13:46:37 on 2964 serial 3
Monitor period: 300 seconds (5:00)

Totals by Processor type:

Type	Count	Ded	shared	total	assigned	Ovhd	Mgmt
IFL	88	0	88	7006	5477.4	339	1528

Running on IFLs

88 IFLs, 80% utilized on average

LPAR Overhead HIGH

- **15 of 88 for physical management???**
- Normally less than 5%
- Leaves 79% of CPU left for real work
- No cpu or thread can have more than 79%

Report: ESASSUM Subsystem Activity Velocity Software

Time	<---Users--->			Transactions		<Processor>		Storage (MB)		<-Paging-->	
	<-avg number-> On	Actv	In Q	Per Minute	Avg. Resp	Utilization Total	Virt.	Fixed User	Active Resid.	<pages/sec> XStore	DASD
13:48:00	234	196	390	229.0	0.907	503	402	208.8	756863	0	2
13:49:00	234	186	396	223.0	0.547	483	383	208.8	756855	0	3
13:50:00	234	188	390	244.0	0.522	460	361	208.8	756858	0	0
13:51:00	234	191	387	205.0	0.450	468	368	208.8	756872	0	4
13:52:00	234	189	395	201.0	1.485	496	390	208.8	756863	0	2
*****Summary*****											
Average:	233	190	392	220.4	0.769	482	381	208.8	756861	0	2

CPU in LPAR about 25% used (500 / 2000)

- 500% utilization over 20 threads

Paging low, no real spikes

No real spikes – problem is more static?

Report: ESAXACT Transaction Delay Analysis Velocity Software C
 Monitor initialized: 01/13/17 at 13:00:00C57 First record analyzed

<-----Percent non-do----->

UserID /Class	<-Samples->		<Asynch>					Lim Lst	Pct Elig				
	Total	In Q	Run	Sim	CPU	SIO	Pag			I/O	Pag	Ldg	Oth
Hi-Freq:	29978	23180	1.3	0.1	0.8	0.0	0	0.0	0	0	0.0	0	0
Hi-Freq:	30000	23202	1.5	0.1	0.7	0.0	0	0.0	0	0	0.0	0	0
Hi-Freq:	30000	23291	1.0	0.0	0.9	0	0	0.0	0	0	0.0	0	0
Hi-Freq:	30000	23328	1.1	0.0	0.3	0	0	0	0	0	0.0	0	0
Hi-Freq:	30000	23286	1.3	0.0	1.0	0.0	0	0.0	0	0	0.0	0	0
Hi-Freq:	150K	116K	1.2	0.0	0.7	0.0	0	0.0	0	0	0.0	0	0
Hi-Freq:	150K	116K	1.2	0.0	0.7	0.0	0	0.0	0	0	0.0	0	0

NO real problems
Some CPU Wait, page wait nill
no ELIGIBLE

SMT – High Level LPAR Perspective

Report: ESALPARS Logical Partition **Summary** Velocity Software
 Report: Logical Partition Summary ZMAP 4.3.0 03/04
 Monitor 2/23/17 at 13:46:37 on 2964

```

-----
                <-----Logical Partition----->
                Virt CPU  <%Assigned>  <LPAR-  <-Thread->
Time           Name      Nbr  CPUs  Type  Total  Ovhd  Weight  Idle  cnt
-----
13:48:00 Totals:      00   120  IFL   5450   332    1000
          V2TEST      32    10  IFL   397.2  28.4     50  215.9    2
          C2TEST      34    16  IFL   616.0  25.4     70  149.9    2
          C2QA         36    10  IFL   535.4  31.4     70  142.7    2
          P2PROD1      38    28  IFL   1608   79.6    270     0     1
          P2PROD2      3A    28  IFL   1166   90.5    270     0     1
          P2PROD3      3C    28  IFL   1127   76.3    270     0     1
  
```

Totals by Processor type:

```

<-----CPU-----> <-Shared Processor busy->
Type Count Ded shared Total Logical Ovhd Mgmt
-----
IFL      88    0      88 6925.4 5118.3 332 1475
  
```

Mgmt Overhead REALLY BAD (1475/8800 = 17%)

- **SMT RELATED????**

IFL TOTAL UTILIZATION: 6925/8800 = 79%

SMT – Logical/Physical CPU: ESALPAR

Report: ESALPAR: Logical Partition Analysis
Monitor 02/23/17 at 13:46:37 on 2964 serial 3

```
-----  
                <--Logical-> <-----Logical Pro  
                <-Partition> VCPU <%Assigned> VCPU  
Time           Name      No. Addr Total  Ovhd TYPE  
-----  
13:48:00      V2TEST      32   0  38.7   2.8 IFL  
                1  40.9   2.8 IFL  
                2  39.5   3.0 IFL  
                3  39.0   2.8 IFL  
  
                7  39.0   2.6 IFL  
                8  38.5   2.9 IFL  
                9  40.6   3.3 IFL  
                -----  
                LPAR 397.2  28.4
```

System total logical partition

Physical CPU Management time:

CPU	Percent	Type
63	23.181	IFL
65	16.287	IFL
66	25.658	IFL
67	23.745	IFL
68	22.041	IFL

Logical CPUs assigned: 40% busy
Physical CPUs overhead VERY HIGH (23%)
(cpu numbers don't start at 0/1)

SMT: Assigned vs z/VM Metrics

Report: **ESALPAR** Logical Partition Analysis

Time	<-Logical CPU-->			<-----CPU (percentages)----->						<Multi-thread>	
	VCPU Addr	<%Assigned> Total	Ovhd	Total util	Emul time	User ovrhd	Sys ovrhd	Idle time	Stl Pct	Idle Time	cp1/cp2
13:48:00	0	38.7	2.8	49.0	35.2	2.2	11.5	144.3	6.67	22.16	0 / 1
	1	40.9	2.8	51.6	40.7	2.3	8.6	141.8	6.67	23.22	2 / 3
	2	39.5	3.0	49.6	38.1	2.3	9.2	143.3	7.17	22.00	4 / 5
	3	39.0	2.8	49.9	39.8	2.2	7.9	133.4	16.6	20.42	6 / 7
	4	41.7	2.7	54.0	44.4	2.2	7.4	129.5	16.4	22.01	8 / 9
	5	41.5	3.1	52.5	42.7	2.3	7.6	130.3	17.1	22.03	10 /11
	6	37.8	2.6	47.9	38.8	2.0	7.0	136.1	16.0	20.53	12 /13
	7	39.0	2.6	49.7	40.5	2.1	7.0	134.6	15.8	21.15	14 /15
	8	38.5	2.9	48.0	39.0	2.0	7.1	135.6	16.3	21.25	16 /17
	9	40.6	3.3	51.4	41.6	2.3	7.6	130.7	17.9	21.14	18 /19
	LPAR	397.2	28.4	503.6	400.8	21.9	80.9	1360	137	215.9	0 / 0

ESALPAR: The IFL from the “HMC” perspective vs z/VM

Hardware: 38.7 % assigned to LPAR

z/VM: 2 threads total 49.0%

Emulation time

User overhead

System overhead

Report: **ESALPAR** Logical Partition Analysis

Time	VCPU Addr	<%Assigned> Total	<Logic> Ovhd	<CPU (percentages)> Total util	Emul time	User ovrhd	Sys ovrhd	Idle time	Stl Pct	Idle Time	cp1/cp2
13:48:00	0	38.7	2.8	49.0	35.2	2.2	11.5	144.3	6.67	22.16	0 / 0
	1	40.9	2.8	51.6	40.7	2.3	8.6	141.8	6.67	23.22	2 / 3
	2	39.5	3.0	49.6	38.1	2.3	9.2	143.3	7.17	22.00	4 / 5
	3	39.0	2.8	49.9	39.8	2.2	7.9	133.4	16.6	20.42	6 / 7
	7	39.0	2.6	49.7	40.5	2.1	7.0	134.6	15.8	21.15	14 /15
	8	38.5	2.9	48.0	39.0	2.0	7.1	135.6	16.3	21.25	16 /17
	9	40.6	3.3	51.4	41.6	2.3	7.6	130.7	17.9	21.14	18 /19
	LPAR	397.2	28.4	503.6	400.8	21.9	80.9	1360	137	215.9	0 / 0

Adding up the numbers

IFL assigned 38.7% , subtract Logical overhead (2.8%) = 36%
 - (two threads = 72% available)

Two threads (z/VM) utilization total 49.0%

LPAR assigned 2 threads concurrently even if only one used

One thread idle time: 22%

Total 71%, so numbers are close with rounding

Report: ESACPUU CPU Utilization Report .0 03/04/17
 Monitor 2/23/17 at 13:46:37 on 2964 serial 0

Time	CPU	CPU Type	Total util	Emul time	User ovrhd	Sys ovrhd	MThread Core/Thread	<-Vertical--> Entitlement	Park seconds
13:48:00	0	IFL	24.5	17.6	1.1	5.8	0/ 0	1.00	0
	1	IFL	23.7	18.5	1.1	4.1	0/ 1	1.00	0
	2	IFL	25.9	20.3	1.2	4.5	1/ 0	1.00	0
	3	IFL	25.6	20.4	1.1	4.1	1/ 1	1.00	0
	4	IFL	25.6	19.5	1.2	4.9	2/ 0	1.00	0
	5	IFL	24.0	18.6	1.1	4.3	2/ 1	1.00	0
	15	IFL	24.1	19.7	1.0	3.3	7/ 1	0	0
	16	IFL	24.6	19.7	1.0	3.8	8/ 0	0	0
	17	IFL	23.5	19.3	1.0	3.2	8/ 1	0	0
	18	IFL	26.0	20.8	1.2	4.0	9/ 0	0	0
	19	IFL	25.4	20.7	1.1	3.6	9/ 1	0	0
System:			502.8	401.7	21.9	79.2	0/55	1.10	0

z/VM "Threads" 25% busy
No parking

SMT : MFC – Hardware effectiveness

Report: ESAMFC MainFrame Cache Magnitudes R
Monitor initialized: 02/23/17 at 13:46:37 on 2964

Time	CPU	<CPU Busy> <percent>		<-----Processor-----> Speed/<-Rate/Sec->			
-----	---	Totl	User	Hertz	Cycles	Instr	Ratio
-----	---	----	----	-----	-----	-----	-----
13:48:02	0	24.5	17.6	5000M	1231M	410M	3.004
	1	23.7	18.5	5000M	1189M	407M	2.923
	2	25.9	20.3	5000M	1303M	517M	2.521
	3	25.6	20.4	5000M	1286M	515M	2.496
	4	25.6	19.5	5000M	1283M	464M	2.762
	5	24.0	18.6	5000M	1207M	421M	2.869
	6	25.8	20.4	5000M	1298M	278M	4.671
	7	24.1	19.4	5000M	1217M	249M	4.892
	8	27.2	22.2	5000M			
	9	26.8	22.2	5000M			
	10	26.7	21.6	5000M			
	11	25.8	21.1	5000M			
	12	24.7	19.8	5000M			
	13	23.2	19.0	5000M			
	14	25.6	20.8	5000M	1287M	262M	4.911
	15	24.1	19.7	5000M	1211M	235M	5.157
	16	24.6	19.7	5000M	1235M	260M	4.747
	17	23.5	19.3	5000M	1183M	248M	4.762
	18	26.0	20.8	5000M	1306M	251M	5.201
	19	25.4	20.7	5000M	1281M	260M	4.919
		----	----	-----	-----	-----	-----
System:		503	402	5000M	23.6G	6454M	3.919

**Z13, expect 2.0 cycles/instruction
With SMT, get 4 cycles / instruction**

Has SMT Provided more capacity?

20% of both threads used by physical management

2 times as many cycles / instruction required

Net capacity: negative

Threads can run 160%

Thread effectiveness 50%

Total: 80% of thruput

NOTE: Perceived benefit of SMT is less users in CPU wait

- They run ½ speed, but nobody measures

The question:

- Is SMT good?
- Is there CPU capacity issue?

Report: ESAHDR z/VM Monitor Analysis
Monitor initialized: 02/23/17 at 13:46:37 on 2964 serial 3
Monitor period: 300 seconds (5:00)

Totals by Processor type:

Type	Count	CPU		Shared Processor busy			
		Ded	shared	Total	Logical	Ovhd	Mgmt
CP	8	0	8	794.1	786.6	3.1	4.4
IFL	10	10	0	0.1	0	0	0.1
ICF	2	2	0	0.0	0	0	0.0
ZIIP	2	0	2	51.6	49.1	0.9	1.6

Running on IFLs - **DEDICATED**

88 IFLs, 80% utilized on average

LPAR Overhead HIGH

- 15 of 88 for physical management???

- No management overhead

CASE SMT2. MFC Analysis

Report: ESAMFC MainFrame Cache Magnitudes R
Monitor initialized: 01/18/17 at 10:30:00 on 2964

Time	CPU	<CPU Busy> <percent>		<-----Processor-----> Speed/<-Rate/Sec->			
		Totl	User	Hertz	Cycles	Instr	Ratio
10:32:00	0	48.4	45.5	5000M	2415M	1295M	1.864
	1	50.1	47.8	5000M	2496M	1420M	1.757
	2	44.0	40.9	5000M	2185M	1175M	1.859
	3	48.6	45.6	5000M	2409M	1386M	1.738
	4	58.6	55.9	5000M	2913M	1569M	1.857
	5	58.5	56.3	5000M	2905M	1545M	1.881
	6	60.7	58.2	5000M	3022M	1511M	1.999
	7	62.7	60.4	5000M	3121M	1623M	1.924
	8	65.7					
	9	66.1					
	10	57.5					
	11	58.3					
	12	60.6					
	13	61.0					
	14	65.6					
	15	66.9					
	16	64.2					
	17	66.6	64.1	5000M	3312M	1690M	1.959
	18	61.5	58.9	5000M	3066M	1606M	1.909
	19	62.5	60.1	5000M	3115M	1657M	1.880
System:		1188	1140	5000M	55.1G	29.5G	1.864

Threads 60% utilized

- 1200% utilization over 20 threads

Cycles / Instruction 1.8



SMT may cost large physical management time
SMT may greatly increase cycles per instruction
Parking/Vertical doesn't help

Analysis is NOT rocket science

- **Few screens needed for most problems**

Send your problems to barton@VelocitySoftware.com