

# Performance Case Studies - SMT

# Know the configuration: ESAHDR

```
Report: ESAHDR          z/VM Monitor Analysis
-----
ZMAP Release                4.3.0.0
Monitor file created:      05/27/16 06:55:52

z/VM Version: 6            Release 3.0 SLU 1501
TOD clock at last IPL:    02/21/16 03:41:24
System Operator:          OPERATOR
Time zone adjustment from GMT: -4 hours

System Identifier          C207
Checkpoint/Warmstart Volumes 1E63RS/1E63RS
Machine Model/Type        Z13:2964/400
Multithreading Status:
System Sequence Code      00000000000C3AE7
Processor 0 model/serial   2964-400 /313AE7 Master
Processor 1 model/serial   2964-400 /313AE7
Processor 2 model/serial   2964-400 /313AE7

Processor 10 model/serial  2964-400 /313AE7
Processor 11 model/serial  2964-400 /313AE7

Operating on IFL Processor(s)

Totals by Processor type:
<-----CPU-----> <-Shared Processor busy>
Type Count Ded shared total assigned Ovhd Mgmt
-----
IFL      86   0    86  3085  2895.8  111  189

Number of logical partitions defined:      8

Main Storage installed (MB):              524288
```

## Common configuration problems

- IFLs?
- Real Storage / Expanded
- Release significant
- Master processor significant

# Know the overall loads: ESASSUM / ESAMAIN

```
Report: ESASSUM          Subsystem Activity          Veloci
Monitor initialized: 05/27/16 at 06:55:52 on 2964 serial 13AE7      First
-----
          <---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
-----
06:57:00    128    97   266    37.0 0.184   892   757 215.6 513636      0  232
06:58:00    128    97   264    35.0 0.365   999   907 215.6 513637      0  260
06:59:00    128    95   266    34.0 0.542   857   719 215.6 513636      0  185
07:00:00    128    94   267    36.0 0.352   899   769 215.6 513634      0  209
*****Summary*****
Average:    128    96   266    35.5 0.357   912   788 215.6 513635      0  221
```

Look for Spikes, dramatic changes, what time?

- Processor
- Storage for users
- Page rates
- DASD I/O rates
- (Transactions are for traditional workloads)

# Wait States: ESAXACT

Report: ESAXACT Transaction Delay A  
Monitor initialized: 05/27/16 at 06:55:5

```
-----  
                                     <-----Percent non-d  
UserID    <-Samples->  
/Class    Total    In Q  Run  Sim  CPU  SIO  Pag  
-----  
06:57:00   128     266  3.8  0.4  0.8   0   0  
Hi-Freq: 19920 15807  2.9  0.2  4.7  0.0  0  
***Key User Analysis***  
RSCS        60      60   18  5.0  17    0   0  
TCPIP       60      60   15  43   22    0   0  
***User Class Analysis***  
Servers     900       1    0    0    0    0   0  
ZVPS        480      12  8.3  0    8.3  8.3  0  
TheUsers   18420    15674  2.8  0.1  4.6   0   0  
***Top User Analysis***  
CV52D172   240      240   13   0   9.2   0   0  
CV52D157   240      240   6.7  0.4  7.5   0   0  
CV52D160   240      240   10   0.8  10    0   0  
CV52D151   240      240   6.3  0    6.3   0   0  
CV52D173   240      240   5.0  0    8.3   0   0  
CV52D168   240      240   5.4  0.4  8.8   0   0
```

# User Configuration: ESAUSRC

Report: ESAUSRC      User Configuration      city Software

```

-----SHARE- <---CPU---><
<Normal> <--MA <Count>
UserID   ClassID  Account  ACI Grp  <CP POOL> CPU <Normal> <--MA <Count>
-----
CV52D176 TheUsers EV500000 . . IFL 2000 . . 4 4 ESA
CV52D177 TheUsers EV500000 . . IFL 2000 . . 4 4 ESA
CV52D178 TheUsers EV500000 . . IFL 2000 . . 4 4 ESA
OPMGRM1  TheUsers 5697-J10 . . IFL 5000 . . 1 1 XC
OPMGRS1  TheUsers 5697-J10 . . IFL 5000 . . 1 1 XC
OPMGRS2  TheUsers 5697-J10 . . IFL 5000 . . 1 1 XC
OP1      Servers 3 . . IFL 100 . . 1 1 ESA
PROPGT   TheUsers PROPGT . . IFL 100 . . 1 1 ESA
RACFVM   Servers SYSTEMS . . IFL 100 . . 1 1 ESA
RSCS     KeyUser 1 . . IFL 1500 . . 1 1 ESA
SFSZVPS  TheUsers VMGRP001 . . IFL 1000 . . 1 1 XC
SNMPD    TheUsers SNMPD . . IFL 100 . . 1 1 ESA
TCPIP   KeyUser TCPIP . . IFL 3000 . . 1 1 ESA
ZSERVE   ZVPS     ZSERVE . . IFL 3000 . . 1 1 ESA
ZTCP     ZVPS     ZTCP . . IFL . 3.0 . . 1 1 ESA
ZVPS     TheUsers VMGRP001 . . IFL 2000 . . 1 1 ESA
ZWEB01   ZVPS     ZWEB01 . . IFL 2100 . . 1 1 ESA
ZWEB02   ZVPS     ZWEB02 . . IFL 2100 . . 1 1 ESA
ZWEB03   ZVPS     ZWEB03 . . IFL 2100 . . 1 1 ESA
ZWEB04   ZVPS     ZWEB04 . . IFL 2100 . . 1 1 ESA
ZWEB05   ZVPS     ZWEB05 . . IFL 2100 . . 1 1 ESA
ZWRITE   ZVPS     ZWRITE . . IFL 2100 . . 1 1 ESA
* * * End of File * *

```

## 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?)

# User Resource Requirements: ESAUSP2

```
Report: ESAUSP2           User Resource
Monitor initialized: 05/27/16 at 06
-----
      <---CPU time--> <---Main
UserID <(Percent)> T:V <Resident>
/Class  Total   Virt  Rat  Totl  Activ
-----
06:57:00 831.6 757.3 1.1 131M 131M
***Key User Analysis***
RSCS      15.57 10.15 1.5 144 144.0
TCPIP    47.29 4.65 10 7880 7880
***User Class Analysis***
Servers    0.00 0.00 1.8 1824 1630
ZVPS       1.65 1.50 1.1 26K 26135
TheUsers   767.1 741.0 1.0 131M 131M
```

Look for “Resource Requirements”

- CPU reasonable?

# LPAR Configuration: ESALPARS

```
Report: ESALPARS          Logical Partition Summary          Veloc
-----
```

Time	Phys CPUs	Dispatch Slice	Complex Name	Nbr	Virt CPUs	CPU Type	<%Assigned> Total	Ovhd	<-Assigned S LPAR--> Weight	Pct
06:57:00	86	Dynamic	Totals:	00	124	IFL	2870	114	1000	100
			<b>C207</b>	<b>31</b>	<b>12</b>	<b>IFL</b>	<b>930.3</b>	<b>24.8</b>	<b>70</b>	<b>7.0</b>
			P113	FF	0					
			P115	FF	0					
			P213	33	24	IFL	401.6	19.7	220	22.0
			P215	35	24	IFL	502.2	18.8	220	22.0
			P217	37	24	IFL	583.2	25.5	220	22.0
			P219	39	28	IFL	300.4	13.1	220	22.0
			V2N3	32	12	IFL	152.1	11.8	50	5.0

```
Totals by Processor type:
<-----CPU-----> <-Shared Processor busy->
Type Count Ded shared Total Logical Ovhd Mgmt
-----
```

<b>IFL</b>	<b>86</b>	<b>0</b>	<b>86</b>	<b>3065.1</b>	<b>2756.2</b>	<b>114</b>	<b>195</b>
------------	-----------	----------	-----------	---------------	---------------	------------	------------

## Look for “Shared processors”

- IFLs shared between LPARs (none)
- Check weights
- Assigned pct/CPU > 100 ??? -> excess share?
- First LPAR is “us”, z/vm where data collected

# Server Requirement Case Study

TCPIP used 47% of a processor at peak

LPAR has 12 processors

TCPIP has a requirement of 4% of the system

Calculate normalizeShare =

$$(\text{RelShare} / \text{SRMRELDL}) * (100 - \text{SRMABSDL}) = 3000 / 141388 = 2\%$$

```
Report: ESASUM           System Summ
Monitor initialized: 05/27/16 at
Monitor period:         240 sec
```

```
-----
Variable Average Minimum Maximum
-----
SRMABSDL          0          0          0
SRMRELDL  141388  139000  146000
```



# Dispatch rates – workload suitable for SMT?

Report: ESAPLDV Processor Local Dispatch VectoSoftware Corporate ZMAP 4  
 Monitor initialized/27/16 at 06:55:52 on 2964 sord analyzed: 05/27/16 06:5

Time	<---Load-		CPU	<VMDBK Moves/sec>		Avg	Dispatcher		<-CPU Steals fr		
	Actv	In Q		Steals	To Master		Long Paths	<-From Nesting	Same	NL1	NL2
06:57:00	97	266	0	9094.5	8.0	2.0	25011.3	6547	2548	0	
			1	9316.9	0	1.8	25456.0	7052	2265	0	
			2	9108.4	0	1.2	23481.9	7177	1932	0	
			3	8870.0	0	0.8	21525.3	7148	1722	0	
			4	8700.2	0	0.6	20082.6	7135	1565	0	
			5	16591	0	0.9	37073.5	0	17K	0	
			6	8278.3	0	0.3	18392.6	6924	1355	0	
			7	9370.5	0	0.6	22102.4	7204	2167	0	
			8	7208.0	0	1.1	21210.2	5371	1837	0	
			9	7410.2	0	1.1	20587.6	5716	1694	0	
			10	7610.1	0	1.0	19939.7	6030	1580	0	
			11	7752.8	0	1.0	19363.9	6256	1497	0	
System:				109310	8.0	12.5	274226.8	73K	37K	0	

Very high dispatch rate

- Not suitable to SMT?
- Interesting to analyze
- Research needed

# Dispatch rates – workload suitable for SMT?

Report: ESAUSR3 User Rorporate ZMAP 4.3.0 06/21/16 Page 62  
 Monitor initialized: 05/27/ed: 05/27/16 06:56:00

UserID /Class	DASD		MDisk	<--Messages Queued-->				<Message>		<Dispatch>	
	DASD I/O	Block I/O	Cache Hits	<----IUCV----->				<Errors-->		<Rate/Sec>	
				Send	Recv	Reply	VMCF	IUCV	VMCF	Disp	Waits
06:57:00	2787	9	1047	1	86	0	3	0	0	<b>194K</b>	194K
***Key User Analysis***											
RSCS	0	0	0	0	0	0	0	0	0	15K	15019
<b>TCPIP</b>	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>	<b>86</b>	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>	<b>15K</b>	<b>14805</b>
***User Class Analysis***											
Servers	20	9	8	0	0	0	1	0	0	1	1
ZVPS	296	0	85	0	0	0	0	0	0	299	299
TheUsers	2471	0	954	1	0	0	2	0	0	164K	164K
***Top User Analysis***											
CV52D172	85	0	0	0	0	0	0	0	0	3095	3095
CV52D157	60	0	0	0	0	0	0	0	0	3787	3787
CV52D160	40	0	0	0	0	0	0	0	0	4684	4684
CV52D151	26	0	0	0	0	0	0	0	0	4533	4533
CV52D173	76	0	0	0	0	0	0	0	0	3088	3088

## Very high dispatch rate for TCPIP?

- Not suitable to SMT?
- Interesting to analyze
- Research needed

# Dispatch rates – workload suitable for SMT?

Report: ESAMFC MainFrame Cache Magnitudes  
 Monitor initialized: 05/27/16 at 06:55:52 on 296

```

-----
                <CPU Busy> <-----Processor----->
                <percent>  Speed/<-Rate/Sec->
Time          CPU Totl User  Hertz Cycles Instr Ratio
-----
06:57:02     0  77.4  64.5  5000M  3866M 1328M 2.911
              1  75.6  63.6  5000M  3776M 1175M 3.213
              2  75.5  64.5  5000M  3773M 1272M 2.967
              3  76.1  66.2  5000M  3802M 1406M 2.703
              4  76.8  67.5  5000M  3836M 1579M 2.429
              5  50.2  35.3  5000M  2528M   668M 3.784
              6  75.3  65.7  5000M  3768M 1149M 3.278
              7  77.5  66.3  5000M  3876M 1206M 3.214
              8  77.3  65.6  5000M  3868M 1201M 3.222
              9  77.0  65.5  5000M  3850M 1227M 3.139
             10  76.9  66.0  5000M  3844M 1241M 3.098
             11  77.0  66.5  5000M  3849M 1287M 2.990
-----
System:           892   757  5000M  41.6G 13.7G 3.028
  
```

Cache, CPI 3.0. High for z13

- Why?

# Dispatch rates – workload suitable for SMT?

Report: ESAMFCA      Mais      Velocity Software Corporate      ZMAP 4  
 Monitor initialized: 05/4 serial 13AE7      First record analyzed: 05/27/16 06:5

Time	CPU	<CPU Busy>		<-----Rate per 100 Instructions----->							<TLB>	
		Totl	User	<-----Data source read from----->							Inst	Miss
				L1	L2	L3	L4L	L4R	MEM			Cost%
06:57:02	0	77.4	64.5	2.146	1.154	0.607	0.161	0.012	0.103	0.208	0.108	51.16
	1	75.6	63.6	2.451	1.282	0.723	0.191	0.015	0.112	0.236	0.129	54.92
	2	75.5	64.5	2.232	1.179	0.651	0.169	0.013	0.105	0.210	0.114	53.32
	3	76.1	66.2	1.991	1.078	0.564	0.144	0.011	0.096	0.191	0.097	50.48
	4	76.8	67.5	1.745	0.958	0.486	0.124	0.009	0.084	0.165	0.082	47.89
	5	50.2	35.3	2.675	1.279	0.694	0.352	0.136	0.094	0.291	0.147	45.71
	6	75.3	65.7	2.176	1.031	0.511	0.322	0.085	0.102	0.222	0.116	55.77
	7	77.5	66.3	2.338	1.114	0.649	0.285	0.083	0.092	0.234	0.128	52.11
	8	77.3	65.6	2.323	1.124	0.629	0.278	0.084	0.093	0.232	0.128	51.57
	9	77.0	65.5	2.284	1.111	0.621	0.271	0.080	0.088	0.226	0.124	51.47
	10	76.9	66.0	2.211	1.062	0.606	0.265	0.081	0.088	0.220	0.121	51.13
	11	77.0	66.5	2.139	1.044	0.577	0.252	0.078	0.084	0.212	0.115	50.22
System:		892	757	2.193	1.108	0.603	0.225	0.052	0.095	0.216	0.115	51.46

TLB Miss

- 50% of cycles used in DAT
- z/OS less than 20%
- High rate of turn over

## Dispatch rate analysis

**25,000 dispatch per second (per processor)**

**CPU 77% utilization**

**.1 memory access per 100 cycles**

**4,000,000,000 (4B) cycles used per second (rounded)**

**4,000,000 (4M) cache lines from memory per second**

**4M \* 256 bytes = 1GB per second loaded from memory**

## Cache Sizes – z13

**L1: 96K Instruction, 128K Data**

**L2: 2MB Instruction, 2MB data**

**L3: 64MB (Chip, Shared over 8 CPUS)**

**L4: 480MB + 224M NIC (per node)**

# SMT Analysis summary

**Dispatch rate high**

**25,000 dispatch per second (per processor)**

**DAT takes 50% of cycles per thread**

**ONLY ONE DAT PER CORE**

**SMT value**

**Share cycles – not many to share**

**Share L1/L2 cache – already bad, would be worse**

**If 77% busy now, with 3.0 cycles per instruction**

**Can only get worse with SMT;**