



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: Whack a Mole

Case: Use the Force

Case: SMT – turn it on?

“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 Challenge z/VM serves many functions (160+ reports)

ESAHDR ESATUNE

*Performance Summary
ESASSUM ESASUM

*Transaction Activity (5)
ESAUOLA ESAXACT ESARATE
ESACLAS ESAEXCP

*User Activity (21)
ESATUNA
ESASRVC ESASRV1 **ESAUSRC** ESAUSR1
ESAUSR2 ESAUSR3 ESAUSR4 ESAUSR5
ESAUSP2 **ESAUSP3** ESAUSP4 ESAUSCP **ESAUSP5**
ESAUSTR **ESAUSPG** ESAUSEK
ESAWKLD ESAUSRQ ESASCED
ESAACCT
ESAPPOOL

*Multi-Tasking Users
ESAMTSK

*Web Serving Reports (8)
ESAWEB1 ESAWEB2 ESAWEB3 ESAWEB4
ESAVWS1 ESAVWS2 ESAVWS3 ESAVWS4

*Virtual NETWORK Reporting (7)
ESAQDIO ESAQDI2 **ESANIC**
ESAVSWC ESAVSW ESAVSW2
ESAOSA

*TCP/IP Reporting (15)
ESATCPC ESATCPI **ESATCP1** **ESATCP2** ESATCP3 **ESATCP4**
ESATCP5 ESATCP6 ESATCP7 ESATCP8
ESATCPP ESATCPS ESATCPA **ESATCPU** ESATFTP

*LINUX Reporting (20)
ESAUCD1 ESAUCD2 ESAUCD3 ESAUCD4 ESAUCDD ESALNXD
ESAHST1 ESAHST2 ESAHST3 ESAHST4 ESAHSTA
ESALNXS ESALNXR ESALNXP ESALNXA ESALNXC **ESALNXF**
ESALNXU ESALNXV ESALNXM ESALNXUP

*Linux Application Reporting (4)
ESAJVM ESAORAC ESAORAG ESAORAS ESAORAW

*VSE Reporting (4)
ESAVSEC ESAVSES ESAVSEP ESAVSEJ

*Shared File System (7)
ESASFS1 ESASFS2 ESASFS3 ESASFS4
ESASFS5 ESASFS6 ESASFS7

*Byte File System
ESABFS1 ESABFS2 ESABFS3

*Processor Subsystem (24)
ESACPUU **ESACPUA** ESACPUS **ESASMT**
ESADIAG ESAINS ESALCK1 ESALCK2 **ESADIA2**
ESAMFC ESAMFCA ESAMFCC ESACPUV
ESACPU1 ESACPU2
ESAIUCV ESAIUC2 ESAIUER
ESALPARC ESALPAR ESALPARS
ESAPLDV ESAIOP ESACRYPT ESACRY2

*Storage Subsystem (10)
ESASTRC ESASTOR **ESASTR1** ESASTR2 ESASTR3 ESAME
ESAFREE ESADCSS **ESAASPC** ESASXS

*Paging Subsystem (5)
ESAPSPC ESAPAGE **ESABLKP** ESAXSTO
ESAPSDV

*Input/Output Subsystem (23)
ESADEV1 ESADEV2 ESADSD1 ESADSD2
ESADSD6 ESAIOAS ESACHNC ESACHAN ESACHNH
ESADSDC ESADSD5 ESAMDC
ESAVDSK ESA3495
ESASCSI ESASCS2
ESASEEK

*
ESAOPER

The 20 (22) 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 Consumer: ESAUSP2**
- **Linux Consumer: ESALNXP**
- **Linux Processor: ESALNXS**

Storage

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

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

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

Report: ESASSUM		Subsystem Activity						Velocity Software							
Time	<---Users--->			Transactions		<Processor>		Storage (MB)		<Paging-->		<---I/O-->			
	<-avg number->	On	Actv	In	Q	Per	Avg.	Utilization	Fixed	Active	<pages/sec>	<-DASD-->	<---I/O-->		
						Minute	Resp	Total	Virt.	User	Resid.	XStore	DASD	Rate	Resp
13:05:00	128	97	218	183.0	0.174	1019	951	187.1	363467	0	0	179	0.7		
13:06:00	128	95	217	171.0	0.185	938	874	187.1	363455	0	2	178	0.7		
13:07:00	128	96	216	175.0	0.191	898	836	187.1	363457	0	0	170	0.6		
13:08:00	128	97	218	176.0	0.202	881	821	187.1	363461	0	0	207	0.9		
*****Summary*****															
Average:	127	96	217	176.3	0.188	934	871	187.1	363460	0	0	184	0.7		

- ### 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 /Class	<-Samples->		<Asynch>													
	Total	In Q	Run	Sim	CPU	SIO	Pag	E-SVM	D-SVM	T-SVM	CF	Tst Idl	I/O	Pag	Ldg	
13:05:00	128	218	5.0	0	1.4	0	0	0	0	0	0	94	0	.	. 0	
Hi-Freq:	16920	13069	4.7	0.1	1.5	0	0	0	5.1	0.2	0.0	94	0.0	0	0 0	
Key User Analysis																
RSCS	60	1	0	0	0	0	0	0	98	100	0	0	0	0	0 0	
TCPIP	60	37	0	0	0	0	0	0	0	0	0	100	0	0	0 0	
User Class Analysis																
Servers	900	64	0	0	1.6	0	0	0	0.4	1.6	0	97	0	0	0 0	
ZVPS	480	28	0	0	3.6	0	0	0	57	64	0	32	0	0	0 0	
TheUsers	15420	12939	4.8	0.1	1.5	0	0	0	3.4	0.1	0.0	94	0.0	0	0 0	
Top User Analysis																
ILNXA195	180	180	67	0	2.8	0	0	0	0	0	0	30	0	0	0 0	
ILNXA203	180	180	27	0	3.3	0	0	0	0	0	0	69	0	0	0 0	
ILNXA199	180	180	21	0.6	1.7	0	0	0	0	0	0	77	0	0	0 0	
ILNXA204	180	180	21	0	2.8	0	0	0	0	0	0	76	0	0	0 0	
ILNXA198	180	180	13	0.6	2.8	0	0	0	0	0	0	84	0	0	0 0	
ILNXA197	180	180	16	0	3.3	0	0	0	0	0	0	81	0	0	0 0	
ILNXA200	180	180	12	0	3.3	0	0	0	0	0	0	85	0	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 three vcpu has 3 times samples

User Configuration: ESAUSRC

Report: ESAUSRC User Configuration Velocity Software Corporate ZMAP 4.3.1 0

UserID	ClassID	Account Code	<CP POOL> PoolName	CPU Type	<-----SHARE----->				<---CPU---><Status>		<-MDC>				<Storage>		
					<Normal> Rel	<--MAX--> Abs	Lim	Shre	<Count> Def	On	Mode	SVM	Dsp	FS	INS	<VM Size> Dflt	Max
ILNXA140	TheUsers	WFN40000	. IFL	200	3	3	ESA	N	N	N	N	6.0G	6.0G
ILNXA141	TheUsers	WFN40000	. IFL	200	3	3	ESA	N	N	N	N	6.0G	6.0G
ILNXA142	TheUsers	WFN40000	. IFL	200	3	3	ESA	N	N	N	N	6.0G	6.0G
ILNXA143	TheUsers	WFN40000	. IFL	200	3	3	ESA	N	N	N	N	6.0G	6.0G
ILNXA144	TheUsers	WFN40000	. IFL	200	3	3	ESA	N	N	N	N	6.0G	6.0G
ILNXA145	TheUsers	WFN40000	. IFL	200	3	3	ESA	N	N	N	N	6.0G	6.0G
ILNXA146	TheUsers	WFN40000	. IFL	200	3	3	ESA	N	N	N	N	6.0G	6.0G
ILNXA147	TheUsers	WFN40000	. IFL	200	2	2	ESA	N	N	N	N	2.0G	2.0G
ILNXA148	TheUsers	WFN40000	. IFL	200	3	3	ESA	N	N	N	N	6.0G	6.0G
ILNXA149	TheUsers	WFN40000	. IFL	200	3	3	ESA	N	N	N	N	6.0G	6.0G
ILNXA150	TheUsers	WFN40000	. IFL	200	3	3	ESA	N	N	N	N	6.0G	6.0G

Look for “Interesting configurations”

- Large relative shares / absolute shares
- CPU Counts, matching shares (100 Rel / vcpu)
- CPU “defined”, and CPU “OnLine”
- 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 (Part 1)

Report: ESALPARS Logical Partition Summary Velocity Software Corporat

Time	<--Complex-->		<-----Logical Partition----->						<-Assigned Shares---->				Cap- ped	Wait Comp
	Phys CPUs	Dispatch Slice	Name	Nbr	Virt CPUs	CPU Type	<%Assigned> Total	Ovhd	<---LPAR--> Weight	<VCPU Pct> Pct	/SYS	/CPU		
13:05:00	55	Dynamic	Totals:	00	71	IFL	1055	18.7	1001	100				
			LNX5	05	20	IFL	828.6	12.6	475	47.5	2.37	54.6	No	No
			LNXX	0F	1	IFL	0.5	0.1	50	5.0	4.99	115	No	No
			LNX2	02	12	IFL	1201	0.1	Ded	21.8	0	0	No	Yes
			LNX3	03	20	IFL	2000	0.1	Ded	36.4	0	0	No	Yes
			LNX8	08	10	IFL	224.7	5.7	475	47.5	4.75	109	No	No
			ZS02	0E	8	IFL	1.3	0.3	1	0.1	0.01	0.29	No	No

The MOST important CPU Report!!!

Look for “Shared processors”

- First LPAR is “us”, z/vm where data collected (LNX5)
- IFLs shared between LPARs
- Check weights
- Assigned pct/CPU > 100 ??? -> excess share? (LNXX)
- Assigned “Pct/CPU” is rated speed of IFL when system running at 100%

Understanding SMT: ESALPARS (Part 2)

```
Report: ESALPARS          Logical Partition Summary          ZMAP 4.3.1.0
-----
      <--Complex--> <-----Logical Partition----->
      Phys Dispatch      Virt CPU <%Assigned> <-Thread->
Time   CPUs   Slice Name      Nbr CPUs Type Total  Ovhd  Idle  cnt
-----
13:05:00   55  Dynamic Totals:    00   71 IFL   1055  18.7
                LNX5     05   20 IFL   828.6  12.6  594.5   2
                LNXXX    0F    1 IFL    0.5   0.1    0    1
                LNX2     02   12 IFL  1201   0.1    0    1
                LNX3     03   20 IFL  2000   0.1    0    1
                LNX8     08   10 IFL  224.7   5.7    0    1
                ZS02     0E    8 IFL    1.3   0.3    0    1

Totals by Processor type:
<-----CPU-----> <-Shared Processor busy->
Type Count Ded shared  Total  Logical Ovhd Mgmt
-----
IFL      55  32    23 1073.7  1036.5 18.7 18.4
```

SMT – Two threads, one CPU

- 20 IFLs has 40 threads
- Total time an IFL was assigned to LPAR (2 threads) 828%
- Or about 1750% of thread time
- Of which 595% was considered “available”

IFL Utilization Summary, 32 Dedicated, 23 Shared, 1036% assigned

Consumers within LPAR: ESAUSP2

Report: ESAUSP2 User Resource Rate Report Velocity Software

UserID	<---CPU time-->			<----Main Storage (pages)----->					<-----Paging (pages)----->					
/Class	<(Percent)>	T:V	<Resident>	Lock	<-----WSS----->		<---Allocated--->			<Pgs/Secnd>				
	Total	Virt	Rat	Totl	Activ	-ed	Totl	Activ	Avg	Total	ExStg	Disk	Read	Write
13:05:00	972.1	951.3	1.0	93M	93.0M	24K	93M	93.0M	727K	9814	0	9814	0	0
***Key User Analysis ***														
RSCS	0.00	0.00	1.4	3125	3125	1	3124	3124	3124	159	0	159	0	0
TCPIP	0.34	0.11	3.1	3946	3946	670	3255	3255	3255	111	0	111	0	0
User Class Analysis														
Servers	0.06	0.04	1.8	34K	16374	508	33K	15833	1056	3592	0	3592	0	0
ZVPS	0.70	0.60	1.2	28K	27981	19K	8843	8843	1105	0	0	0	0	0
TheUsers	971.0	950.6	1.0	93M	93.0M	4044	93M	93.0M	903K	5952	0	5952	0	0
Top User Analysis														
ILNXA195	202.8	202.3	1.0	1.3M	1310K	54	1.3M	1310K	1M	24	0	24	0	0
ILNXA203	77.36	76.77	1.0	1.3M	1310K	54	1.3M	1310K	1M	23	0	23	0	0
ILNXA199	67.44	66.75	1.0	1.3M	1310K	54	1.3M	1310K	1M	21	0	21	0	0
ILNXA204	57.35	56.22	1.0	1.3M	1310K	54	1.3M	1310K	1M	21	0	21	0	0
ILNXA198	49.73	48.74	1.0	1.3M	1310K	54	1.3M	1310K	1M	20	0	20	0	0
ILNXA197	40.01	39.35	1.0	1.3M	1310K	54	1.3M	1310K	1M	21	0	21	0	0
ILNXA200	38.80	38.03	1.0	1.3M	1310K	54	1.3M	1310K	1M	21	0	21	0	0
ILNXA142	27.03	26.12	1.0	1.6M	1572K	54	1.6M	1572K	2M	12	0	12	0	0
ILNXA146	25.43	24.57	1.0	1.6M	1572K	54	1.6M	1572K	2M	31	0	31	0	0

Look for consumers, in percent of cpu

- **By class (TheUsers - default)**
- **Abusive servers (ILNXA195)?**
- **Correct per expected? Not a performance question**

Understanding Capacity, chargeback with SMT: ESAUSP5

```

Report: ESAUSP5      User SMT CPU Consumption Analysis
-----
UserID      <-----CPU Percent Consumed (Total)-----> <-CPU PCT Primar
/Class     <Traditional> <MT-Equivalent> <MT Prorated> <MT-Equivalent>
Total     Virt     Total     Virtual     Total     Virtual     Total     Virtual
-----
13:05:00  972.1   951.3   830.8     813.0     0         0         830.8     813.0
***Key User Analysis***
RSCS      0.00    0.00    0.00      0.00      0         0         0.00      0.00
TCPIP     0.34    0.11    0.29      0.09      0         0         0.29      0.09
***User Class Analysis***
Servers   0.06    0.04    0.06      0.03      0         0         0.06      0.03
ZVPS      0.70    0.60    0.61      0.53      0         0         0.61      0.53
TheUsers  971.0   950.6   829.9     812.3     0         0         829.9     812.3
***Top User Analysis***
ILNXA195  202.8   202.3   176.6     176.1     0         0         176.6     176.1
ILNXA203  77.36   76.77   64.13     63.63     0         0         64.13     63.63
ILNXA199  67.44   66.75   56.32     55.73     0         0         56.32     55.73
ILNXA204  57.35   56.22   49.20     48.22     0         0         49.20     48.22
ILNXA198  49.73   48.74   43.41     42.55     0         0         43.41     42.55
  
```

What is CPU consumed?

- Traditional is “assigned to core”
- MT-Equivalent is time “using core”
- (How do all linux tools measure cpu?)

Understanding Capacity, chargeback with SMT: ESAUSP5

Report: ESAUSR5 User SMT CPU Consumption Analysis

```
-----  
      <----Raw CPU Seconds Consumed (Total)---->  
UserID  <Traditional> <MT-Equivalent> <MT Prorated>  
/Class  Total      Virt      Total      Virtual      Total      Virtual  
-----  
10:32:00 660.4    641.7    476.0     462.5    432.0     420.0  
  ***User Class Analysis***  
TheUsers 660.2    641.6    475.9     462.4    431.9     419.9  
  ***CPU POOL User Analysis***  
DB2      15.63    15.42    12.13     11.97    12.23     12.09  
EEMSCSP  9.03     8.97     6.91      6.87     6.59     6.55  
IIB      498.7    488.6    360.4     353.2    321.8     315.4
```

What is CPU consumed?

- **Traditional:** Time assigned and dispatched on a thread
- **Time Would take if non-SMT (MT-Equivalent)**
- **Cycles really used (approximately, prorated)**
- **(How do all linux tools measure cpu?)**

Linux Process Load: ESALNXP

Report: ESALNXP LINUX HOST Process Statistics Report Veloc

node/ Name	<-Process ID	Ident-> PPID	GRP	Nice Valu	PRTY Valu	<-----CPU Tot	Percents-----> sys	user	syst	usrt	<----- Size	RSS
ILNXa195	0	0	0	0	0	203	1.66	201	0.15	0.10	7396	4592
splunkd	2782	1	2781	0	20	0.20	0.10	0.10	0	0	329	19
java	2940	1	2910	0	20	0.85	0.02	0.83	0	0	619	43
perl	3027	2940	2910	0	20	0.23	0.02	0.02	0.13	0.07	7	1
java	37087	1	0	0	20	175	0.45	175	0	0	1546	1392
java	40406	1	0	0	20	0.10	0.03	0.07	0	0	731	337
java	40513	1	0	0	20	25.3	0.82	24.5	0	0	1691	1431
java	40514	1	0	0	20	0.55	0.08	0.47	0	0	1587	1335
snmpd	57526	1	0	-10	10	0.07	0.05	0.02	0	0	25	5

Look for processes within Linux, in percent of CPU

- **By relevant server (ILNXa195)**
- **Correct? Relevant? Cron? Agents?**
- (Traditional CPU – must be prorated for chargeback)

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)

Storage Utilization: ESASTR1

Report: **ESASTR1** **Main Storage Analysis** Velocity Software Corporate

```
----->
Users <-----MegaBytes----->
Loggd System Frame <Available> User NSS/DCSS <-AddSpace> VDISK <MDC>
Time On Storage Table <2gb >2gb Resdnt Resident System User Rsdnt Rsdnt
-----
13:05:00 128 524288 4096 588 140K 363635 551 3262 0 11419 233
13:06:00 128 524288 4096 588 140K 363634 551 3262 0 11419 235
13:07:00 128 524288 4096 588 140K 363634 551 3262 0 11420 237
13:08:00 128 524288 4096 588 140K 363635 551 3262 0 11423 238
*****Summary*****
Average: 128 524288 4096 588 140K 363634 551 3262 0 11420 236
```

Total storage analysis (in megabytes – was pages)

- System Storage (512GB)
- MDC 236 mb? ALWAYS “SET MDC MAX/MIN”
- User resident should be large percent of “System Storage”
- VDISK ? Which server?
- Using VDISK allows server size to be greatly reduced without impact

Virtual Machine Storage : ESAUSPG

Report: ESAUSPG		User Storage Analysis					Velocity Software Corporate					
UserID /Class	<-Storage Occupancy in MegaBytes->			<--Paging-->		<--Main Storage page Read/Write-->			<Page Reads:>		<Address <pages R	
	Total	>2gb	<2GB	Xstor	DASD	Xsto	Disk	Migr	Xstor	Disk	VirtDisk	
13:05:00	363635	362385	1250	0	38	0	0	0	0	0	11419	
***Key User Analysis ***												
RSCS	12	12	0	0	1	0	0	0	0	0	0	
TCPIP	15	15	0	0	0	0	0	0	0	0	0	
User Class Analysis												
Servers	131	131	1	0	14	0	0	0	0	0	0	
ZVPS	109	109	0	0	0	0	0	0	0	0	0	
TheUsers	363367	362119	1249	0	23	0	0	0	0	0	11419	
Top User Analysis												
ILNXA195	5118	5100	18	0	0	0	0	0	0	0	399	
ILNXA203	5118	5100	18	0	0	0	0	0	0	0	1138	
ILNXA199	5118	5099	18	0	0	0	0	0	0	0	376	
ILNXA204	5118	5099	18	0	0	0	0	0	0	0	350	
ILNXA198	5118	5099	18	0	0	0	0	0	0	0	379	

Total storage analysis (in “megabyte” option)

- Largest consumer(s) resident storage
- Largest consumer - which virtual disk?
- VDISK large (1gb) ? Which server? (ILNXA203)

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	
13:05:00													
ILNXA141	VDISK\$ILNXA141\$0160\$0043	25088	200K	0	0	Priv	0	1	25K	0	0	0	0
ILNXA141	VDISK\$ILNXA141\$0161\$0044	267K	2133K	0	0	Priv	0	1	27K	0	0	0	0
ILNXA142	VDISK\$ILNXA142\$0160\$0045	25088	200K	0	0	Priv	0	1	25K	0	0	0	0
ILNXA142	VDISK\$ILNXA142\$0161\$0046	267K	2133K	0	0	Priv	0	1	38K	0	0	0	0
ILNXA143	VDISK\$ILNXA143\$0160\$0047	25088	200K	0	0	Priv	0	1	25K	0	0	0	0
ILNXA143	VDISK\$ILNXA143\$0161\$0048	267K	2133K	0	0	Priv	0	1	23K	0	0	0	0
ILNXA144	VDISK\$ILNXA144\$0160\$0049	25088	200K	0	0	Priv	0	1	25K	0	0	0	0
ILNXA144	VDISK\$ILNXA144\$0161\$004A	267K	2133K	0	0	Priv	0	1	46K	0	0	0	0
ILNXA145	VDISK\$ILNXA145\$0160\$004B	25088	200K	0	0	Priv	0	1	25K	0	0	0	0
ILNXA145	VDISK\$ILNXA145\$0161\$004C	267K	2133K	0	0	Priv	0	1	56K	0	0	0	0

Virtual Disk Analysis

- Any virtual disk spiked?
- Are there multiple vdisks, and PRIORITIZED!!!
- Different sizes?
- If prioritized, expect 1st to have more pages resident
- Recommendation – 1st vdisk is normally sufficient

VDISK for Swap: “growth case study”

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!!!

What changed

What resource is constrained

What configuration changes would be recommended

The problem:

- The User's complain about “steal time”
- Performance is “slow”

The (common “linux admin expert”) problem???

Users are worried about STEAL TIME (SAR)

12:00:02	AM	CPU	%user	%nice	%system	%iowait	%steal	%idle
01:00:02	AM	all	1.95	0.00	2.52	0.67	3.75	91.12
01:10:02	AM	all	4.41	0.00	10.55	6.88	53.24	24.92
01:20:02	AM	all	3.54	0.00	12.59	2.84	52.32	28.70
01:30:04	AM	all	3.86	0.00	13.71	2.94	55.24	24.24
01:40:07	AM	all	4.02	0.00	13.84	2.54	62.26	17.33
01:50:04	AM	all	4.05	0.00	12.77	2.89	73.78	6.50

Server performance?

- Nothing can be determined at Linux Level

Report: ESAHDR z/VM Monitor Analysis
Monitor initialized: 07/01/17 at 18:00:00 on 2827 serial 0168A7
Monitor period: 3600 seconds (1:00:00)

ZMAP Release 4.3.1.1
History Source 4.3.0.5
Monitor file created: 07/01/17 18:00:00

z/VM Version: 6 Release 3.0 SLU 1501
TOD clock at last IPL: 02/19/17 07:16:17
System Operator: OPERATOR
Time zone adjustment from GMT: -4 hours

System Identifier VML2
Checkpoint/Warmstart Volumes LV2RES/LV2RES
Machine Model/Type **EC12:2827/702**
System Sequence Code 0000000000468A7
Processor 0 model/serial 2827-702 /0168A7 Master
Processor 1 model/serial 2827-702 /0168A7
Processor 2 model/serial 2827-702 /0168A7
Processor 3 model/serial 2827-702 /0168A7

Power of processor in terms of service Units: 73394
Operating on IFL Processor(s)
Channel Path Measurement Facility(CPMF) Extended is installed

IFLs, 4 in LPAR
z/VM 6.3,
EC12, IFLs

CASE WAM. Subsystem Activity

Report: ESASSUM SubActivity Velocity Softwar
 Monitor initialized: 01/t 13:00:00 on 2964 serial 06EC47 First record ana

Time	<---Users---> tions <Processor> Storage (MB) <-Paging--> <-----I/O				Avg. Utilization		Fixed Active		<pages/sec> <-DASD-->		Rate Resp		
	On	Actv	In	Q	Resp	Total	Virt.	User	Resid.	XStore	DASD	Rate	Resp
07/01/17													
18:01:00	55	39	44.0		1.297	231	218	166.1	163692	0	0	150	1.1
18:02:00	55	38	45.0		1.858	255	229	166.1	163690	0	0	103	0.4
18:03:00	55	38	41.0		1.717	228	211	166.1	163690	0	0	95	0.4
18:04:00	55	38	42.0		1.284	232	214	166.1	163690	0	0	120	0.9
18:05:00	55	44	43.0		3.303	222	206	166.1	163715	0	0	109	0.5
18:06:00	55	39	43.0		1.957	224	209	166.1	163692	0	0	110	0.5
18:07:00	55	38	44.0		2.202	223	208	166.1	163690	0	0	121	0.7
18:08:00	55	38	43.0		1.254	221	199	166.1	163690	0	0	100	0.4
18:09:00	55	38	43.0		1.688	229	210	166.1	163690	0	0	125	0.5
18:10:00	55	44	44.0		1.619	233	215	166.1	163715	0	0	97	0.3
18:11:00	55	39	42.0		1.619	221	211	166.1	163692	0	0	122	0.5
18:12:00	55	38	44.0		1.508	226	213	166.1	163690	0	0	95	0.3
18:13:00	55	38	43.0		1.597	257	241	166.1	163690	0	0	101	0.3

CPU Spikes? at 18:02? (255 out of 400?)
Nothing can be determined at the LPAR level

Case WaM User Configuration: ESAUSRC

Report: ESAUSRC User ocity SoftwareMAP 4.3.1

UserID	ClassID	Account Code	CPU Type	<-----S>		<---CPU--->		<Storage>		
				<Normal> Rel	<Abs> Abs	<Count> m	<Def On Mode> t	<VM Size> Dflt	Max	
L10AC	TheUsrs	L10AC	IFL	100	. .	2	2	ESA	4.0G	4.0G
L10DU	*Util	L10DU	IFL	100	. .	2	2	ESA	6.0G	6.0G
L20EP	*Prod	L20EP	IFL	200	. .	2	2	ESA	4.0G	4.0G
L20FP	*Prod	L20FP	IFL	200	. .	2	2	ESA	4.0G	4.0G
L21EP	*Prod	L21EP	IFL	200	. .	2	2	ESA	10G	10G
L211P	*Prod	L211P	IFL	100	. .	1	1	ESA	4.0G	4.0G
L212P	*Prod	L212P	IFL	100	. .	1	1	ESA	4.0G	4.0G
L214P	*Prod	L214P	IFL	100	. .	1	1	ESA	2.0G	2.0G
L217P	*Prod	L217P	IFL	100	. .	1	1	ESA	8.0G	8.0G
L218P	*Prod	L218P	IFL	100	. .	1	1	ESA	7.5G	7.5G
L22AP	*Prod	L22AP	IFL	200	. .	2	2	ESA	6.0G	6.0G
L22CP	*Prod	L22CP	IFL	200	. .	2	2	ESA	8.0G	8.0G
L220P	*Prod	L220P	IFL	200	. .	2	2	ESA	21G	21G
L223P	*Prod	L223P	IFL	10	. .	2	2	ESA	2.0G	2.0G
L229P	*Prod	L229P	IFL	100	. .	1	1	ESA	2.0G	2.0G
L23CP	*Prod	L23CP	IFL	200	. .	1	1	ESA	6.0G	6.0G
L241C	TheUsrs	L241C	IFL	100	. .	1	1	ESA	2.0G	2.0G
L245P	*Prod	L245P	IFL	200	. .	2	2	ESA	12G	12G
L246C	TheUsrs	L246C	IFL	40	. .	2	2	ESA	2.0G	2.0G
L25AP	*Prod	L25AP	IFL	100	. .	2	2	ESA	16G	16G
L25CP	*Prod	L25CP	IFL	100	. .	1	1	ESA	6.0G	6.0G
L25EP	*Prod	L25EP	IFL	100	. .	1	1	ESA	16G	16G
L251C	TheUsrs	L251C	IFL	30	. .	1	1	ESA	2.0G	2.0G
TCPIP	KeyUser	TCPIP	IFL	3000	. .	1	1	ESA	128M	128M

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?)

Servers “prioritized”!

CASE WaM. Wait State Analysis

Report: ESAXACT		Transaction Delay Analysis											Veloc			
		<-----Percent non-dormant (Wait states)----->														
UserID /Class	<-Samples->		Run	Sim	CPU	SIO	Pag	E-SVM	D-SVM	T-SVM	CF	Tst Idl	<Asynch>		Ldg	
	Total	In Q											I/O	Pag		
18:01:00	55	44	6.8	0	57	0	0	0	0	0	0	36	0	.	.	
Hi-Freq:	4200	2581	6.9	0.2	37	0	0	0	2.7	0.2	0.0	55	0.1	0	0	
User Class Analysis																
Servers	540	1	0	0	0	0	0	0	0	0	0	0	0	0	0	
Velocity	600	20	5.0	0	5.0	0	0	0	0	10	0	80	0	0	0	
*Prod	2040	2013	8.4	0.2	39	0	0	0	0	0	0.0	52	0.1	0	0	
*Util	120	120	0	0	18	0	0	0	0	0	0	83	0	0	0	
TheUsrs	780	368	2.2	0	40	0	0	0	14	0.5	0	57	0	0	0	
Top User Analysis																
L22AP	120	120	20	0	69	0	0	0	0	0	0	11	0	0	0	
L217P	60	60	32	3.3	65	0	0	0	0	0	0	0	0	0	0	
L256P	120	120	18	0	62	0	0	0	0	0	0	21	0	0	0	
L220P	120	120	10	0.8	61	0	0	0	0	0	0	28	0	0	0	
L21EP	120	120	13	0	57	0	0	0	0	0	0	28	2.5	0	0	
L23CP	60	60	27	0	38	0	0	0	0	0	0	35	0	0	0	
L218P	60	60	15	0	65	0	0	0	0	0	0	20	0	0	0	
L258P	60	60	15	0	42	0	0	0	0	0	0	43	0	0	0	
L245P	120	120	5.8	0.8	51	0	0	0	0	0	0	43	0	0	0	
L22CP	120	120	6.7	0	34	0	0	0	0	0	0.8	58	0	0	0	

Any obvious waits? (CPU....)
Nothing else significant

Totals by Processor:

		CPU		Shared Processor busy->			
Type	Count	Ded	shar	Total	Logical	Ovhd	Mgmt
CP	2	0		193.3	184.5	2.4	6.4
IFL	4	0		394.3	393.2	0.4	0.7
ICF	3	3		0.5	0	0	0.5
ZIIP	1	0		8.5	7.9	0.1	0.5

Linux LPAR is using 4 virtual CPUs

IFLs running at 95-100%

Nothing else significant

Report: ESALPAR Logical Partition Summary

Velocity Software Co

Time	<-----Logical Partition----->						<--Assigned Shares---->				Cap- ped
	Name	Nbr	Virt CPUs	CPU Type	<%Assigned> Total	Ovhd	<---LPAR--> Weight	<VCPU Pct> Pct	/SYS	/CPU	
18:01:00	Totals:	00	10	CP	186.9	2.4	1400	100			
	Totals:	00	7	IFL	393.6	0.4	135	100			
	VML2	01	4	IFL	231.0	0.1	75	55.6	13.9	55.6	No
	VML4	03	3	IFL	162.7	0.3	60	44.4	14.8	59.3	No
	CER1	17	2	CP	8.8	0.1	60	4.3	2.14	4.28	No
	CER1	17	1	ZIP	0.6	0.0	100	4.5	4.48	4.48	No

Look for “Shared processors”

- First LPAR is “us”, z/vm where data collected (LINUX)
- IFLs shared between two LPARs
- Check weights $(75 / 135) = 55.6\%$ of 4 IFLs (220%)
- “pct/CPU” is how fast cpu will run when all are busy
- VML2 getting more that what was set

CASE WaM. Virtual Machine Consumption Analysis

```
Report: ESAUSP2          Use
-----
<---CPU time-->
UserID  <(Percent)>  T:V
/Class  Total    Virt  Rat
-----  -
18:01:00 225.6  218.1  1.0
***Key User Analysis**
TCPIP      0.21   0.11  1.9
***User Class Analysis*
Servers    0.00   0.00  2.6
Velocity   0.74   0.69  1.1
*Prod     213.5  206.4  1.0
*Util     0.76   0.72  1.1
TheUsrs   10.38  10.11  1.0
***Top User Analysis***
L22AP     28.80  28.27  1.0
L217P     23.70  20.68  1.1
L256P     22.43  22.22  1.0
L220P     21.18  20.93  1.0
L21EP     20.90  19.91  1.0
L23CP     20.06  19.66  1.0
L218P     13.94  13.85  1.0
L258P     11.10  10.92  1.0
```

No major consumer

User classification provided,

- **Production is all of it**

CASE WAM. Linux Process consumption

Report: ESALNXP

LINUX HOST Process Statistics Report

node/ Name	<-Process ID	Ident-> PPID	GRP	Nice Valu	PRTY Valu	<-----CPU Tot	Percents-----> sys	user	syst	usrt
L22AP	0	0	0	0	0	27.3	7.56	13.9		
init	1	1	1	0	20	5.36	0	0		
timestam	3299	1	2263	0	20	0.29	0.04	0.11		
ohasd.bi	5431	1	5431	0	20	0.34	0.09	0.24		
oraroota	6139	1	6139	0	20	0.17	0.04	0.13		
oraagent	6207	1	6207	0	20	0.17	0.02	0.15		
evmd.bin	6225	1	6225	0	20	0.47	0.06	0.41		
gipcd.bi	6269	1	6269	0	20	0.77	0.26	0.50		
ocssd.bi	6345	1	6345	0	-100	0.50	0.17	0.34		
octssd.b	6587	1	6587	0	20	0.28	0.06	0.22		
crsd.bin	6632	1	6632	0	20	0.82	0.11	0.71		
oraagent	6709	1	6709	0	20	0.39	0.06	0.26		
oraroota	6728	1	6728	0	20	0.52	0.09	0.43		
asm_vktm	7201	1	7201	0	-2	0.22	0.11	0.11		
asm_dia0	7233	1	7233	0	20	0.21	0.02	0.19		
asm_lmon	7240	1	7240	0	20	0.21	0.04	0.17		
asm_lmd0	7245	1	7245	0	20	0.13	0.04	0.09		
asm_lms0	7250	1	7250	0	-2	0.13	0.02	0.11		
oraagent	7560	1	7560	0	20	0.32	0.04	0.28		
ora_vktm	7693	1	7693	0	-2	0.22	0.11	0.11		
ora_dia0	7742	1	7742							
ora_lmon	7747	1	7747							
ora_lmd0	7752	1	7752							
ora_lms0	7757	1	7757							

Oracle workload, nothing significant,
 "init" is short term processes terminating (oracle?)

Issues – Not enough CPU for workload

Solution:

- Prioritize work (whack a mole)
- Use less
- Buy more

Analysis is NOT rocket science

- **Few screens needed for most problems**

Send your problems to barton@VelocitySoftware.com