



VELOCITY
SOFTWARE

Performance Analysis Flowchart

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

Barton Robinson,
barton@velocitysoftware.com
If you can't measure it, I'm just not interested....

Copyright © 2019 Velocity Software, Inc. All Rights Reserved.
Other products and company names mentioned herein may be
trademarks of their respective owners.

“z” is:

- Very large,
- Very complex and
- Very well instrumented
- TECHNOLOGY KEEPS CHANGING...

The challenge?

What challenge, it is “all” there!

- 200 zmon panels (with menus)
- 150 zmap reports (with table of contents)
- 3400 unique variables

Few companies have full time performance analysts.

z/VM Platform, source: Monitor (95+ reports)

*Performance Summary (4)

ESAHDR ESATUNE ESASSUM ESASUM

*Transaction Activity (5)

ESAUCLA **ESAXACT** ESARATE ESACLAS
ESAEXCP

*User Activity (21)

ESASRVC ESASRV1 **ESAUSRC** ESAUSR1
ESAUSR2 ESAUSR3 ESAUSR4 ESAUSR5
ESAUSP2 **ESAUSP3** ESAUSP4
ESAUSCP **ESAUSP5**
ESAUSTR **ESAUSPG** ESAUSEK ESAWKLD
ESAUSRQ ESASCED ESAACCT **ESAPPOOL**

*Processor Subsystem (26)

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

*Storage Subsystem (11)

ESASTRC ESASTOR **ESASTR1** ESASTR2
ESASTR3 ESAME **ESAVDSK**
ESAFREE ESADCSS **ESAASPC** ESASXS

*Paging Subsystem (5)

ESAPSPC ESAPAGE **ESABLKP** ESAXSTO
ESAPSDV

*Input/Output Subsystem (18)

ESADEV1 ESADEV2 ESADSD1 **ESADSD2**
ESADSD6 ESAIOAS ESACHNC ESACHAN
ESACHNH
ESADSDC **ESADSD5** ESAMDC
ESA3495 ESASCSI ESASCS2
ESASEEK **ESAFCP** **ESADEV**

*Virtual NETWORK Reporting (7)

ESAQDIO ESAQDI2 **ESANIC**
ESA\SWC ESA\SWW ESA\SW2
ESAOSA

*Operational Logging

ESAOPER

- z/VM “traditional” Applications (Source: APPLMON)
 - (33 reports currently?)

***Shared File System (7)**

ESASFS1 ESASFS2 ESASFS3 ESASFS4
ESASFS5 ESASFS6 ESASFS7

***Byte File System (2)**

ESABFS1 ESABFS2 ESABFS3

***CMS Multitasking (1)**

ESAMTSK

***Web Serving Reports (8)**

ESAWEB1 ESAWEB2 ESAWEB3 ESAWEB4
ESAVWS1 ESAVWS2 ESAVWS3 ESAVWS4

***TCPIP Reporting (15) (tcpip monitor records, snmp)**

ESATCPC ESATCPI ESATCP1 ESATCP2
ESATCP3 ESATCP4
ESATCP5 ESATCP6 ESATCP7 ESATCP8
ESATCPP ESATCPS ESATCPA **ESATCPU** ESATFTP

Network: Source: snmp - VERY efficient, really!

- Network: mib ii (first available, from any snmp enabled server

ESATCPC ESATCP1 ESATCP2 ESATCP3 ESATCP4

- Microsoft servers: Host mib

ESAHST1 ESAHST2 ESAHST3 ESAHST4 ESAHSTA

- Linux servers: UCD mib

ESAUCD1 ESAUCD2 ESAUCD3 ESAUCD4 ESAUCDD

- More Linux: Velocity mib

ESALNXD ESALNXS ESALNXR ESALNXP ESALNXA

ESALNXC ESALNXF ESALNXU ESALNXV

ESALNXM ESALNXUP

- Linux Application "vendor" mibs

ESAJVM ESAORAC ESAORAG ESAORAS ESAORAW

ESAGPFN ESAGPFF ESAGPFF ESAGPFD ESAGPFS (Version 5, zvps)

ESAMNG1 ESAMNG2 ESAMNG3 ESAMNG4 ESAMNG5 (Version 5, zvps)

ESADOCK1 ESADOCK2 ESASSC (Version 5, zVPS)

- VSE mibs: IBM, Velocity

ESAVSEC ESAVSES ESAVSEP ESAVSEJ

ESAVSEP ESAVSEJ

The 25 “z/VM” Reports you need in the order you need them

Configuration: ESAHDR

System (z/VM LPAR) Load: ESASSUM **Paging**

Wait states: ESAXACT

Virtual machine Config: ESAUSRC

CPU:

- LPAR Summary: ESALPARS
- CPU Consumer: ESAUSP2
- Linux Consumer: ESALNXP
- Linux Processor: ESALNXS
- CPU Cache: ESAMFC

Storage

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

- Configuration: ESAPSDV
- Loads by user: ESAUSPG

DASD

- Configuration: ESADSD1
- Rates: ESADSD2
- CACHE: ESADSD5
- FCP: ESAFCP
- EDEV: ESAEDEV

Network

- Configuration: ESATCPI
- Network management: ESATCP1/2/4
- OSA: ESAOSA/NIC/VSW

Analysis starts with “is there a problem?”

- Describe the problem (what user(s), what time)
- (CPU Utilization too high, CPU Wait)
- Question? How many more IFLs? Wrong question....

System Configuration

- Processor model, cpu type
- Number of processors, storage size
- SMT support

Loads on the system subsystems

Wait states for those impacted

Subsystem Analysis

- DASD, Storage, Paging, Processor, Network
- Follow the leads, almost everything is available

Know the configuration: ESAHDR (B)

Report: ESAHDR z/VM Monitor Analysis
Monitor initialized: **06/12/23** at 00:00:00
Monitor period: 86400 seconds (24:00:00)

ZMAP Release 5.1.5.0
History Source 5.1.3.5
Key Expiration for: bdr 07/10/23
Components licensed: IFL ZOS TUN PRO
Monitor file created: 06/12/23 00:00:00

z/VM Version: **7** **Release 2.0** SLU 2201
TOD clock at termination 20:00:00
Abend code of last termination
TOD clock at last IPL: **12/11/22** 00:22:00
System Operator: VMOPER
Time zone adjustment from GMT: -4 hours

System Identifier **VMPROD2**
Checkpoint/Warmstart Volumes MP2RES/MP2RES
Machine Model/Type **Z15**:8561/608
Multithreading Status (SMT): Enabled
Core Thread count: 2
Enabled Count: 2

Validate configuration

- IFLs?
- Release significant
- SMT on, two threads/core

Know the configuration: ESAHDR (B)

Report: ESAHDR z/VM Monitor Analysis
Monitor initialized: **06/12/23** at 00:00:00 on 8561

```
-----  
System Sequence Code          000000000007C278  
Processor 0 model/serial      8561-608 /06C278 Master  
Processor 1 model/serial      8561-608 /06C278  
Processor 2 model/serial      8561-608 /06C278  
Processor 3 model/serial      8561-608 /06C278  
...  
Processor 24 model/serial     8561-608 /06C278  
Processor 25 model/serial    8561-608 /06C278
```

```
CPU(GP) Capability Factor:      716  
CPU(IFL) Capability Factor:     416  
CPU Cycles/ns:                  5200  
CPU Cycles/ns (GP):             3022  
Operating on IFL Processor(s)  
Channel Path Measurement Facility(CPMF)
```

Totals by Processor type:

<-----CPU----->				<-Shared Processor busy>			
Type	Count	Ded	shared	total	assigned	Ovhd	Mgmt
CP	8	0	8	478.4	476.0	2.3	2.5
IFL	36	0	36	3083	3071.4	13.0	11.3
ICF	1	0	1	8.9	6.7	0.2	2.2
ZIIP	4	0	4	115.1	114.1	0.5	1.0

(85% CPU 24 hours)

Number of logical partitions defined: 10

Main Storage installed (MB): **1048576**
Main Storage Generated (MB): 1048576

Validate configuration

- 13 Virtual CPUs (IFLs)?
- 26 threads
- 36 IFLs 85% busy 24 hours
- SMT
- Master processor thread 0
- 1TB

Know the configuration: ESAHDR (B)

Top users and user classes by CPU consumption:

UserID /Class	<-Relative->		<---Absolute Percent CPU--->			
	Used	Cum	<Out of 100%>		<Out of 2600%>	
			Util	Cum	Util	Cum
1. PGFS3J2	6.2	6	2.6	3	68.4	68
2. PGFS1J4	5.6	12	2.4	5	61.7	130
3. PGFS1J2	4.7	17	2.0	7	51.3	181
4. PGFS2J2	4.5	21	1.9	9	49.7	231
5. System	3.1	24	1.3	10	34.0	265
6. PADMIN1	2.5	27	1.1	11	27.6	293
7. PSUBS2J2	2.5	29	1.0	12	27.1	320
8. PSUBS1J2	2.4	31	1.0	13	25.9	346
9. PSUBS4J1	2.2	34	0.9	14	24.4	370
10. PSUBS4J3	2.2	36	0.9	15	24.3	395

Top User CPU Consumption

- As compared to users
- Out of 100 (Relative to 100)
- Out of 26 threads (2600 Absolute)\
- (no single user is major consumer
- Why is “System” so high???

Know the overall loads: ESASSUM / ESAMAIN (B)

Report: ESASSUM Subsystem Activity Velocity S
Monitor initialized: **06/12/23** at 00:00:00 on 856178 First reco

```
-----  
      <---Users----> Transactions <Processor> Pages <-----I/O  
      <-avg number->   Per   Avg. Utilization /Sec <-DASD-->  
te  
Time          On Actv In Q Minute   Resp Total Virt.   DASD Rate Resp  
-----  
06/12/23  
09:00:00    116   103   147    65.5 0.022   1464   1409     342   516   0.5  
09:15:00    116   103   133    65.2 0.027   1418   1364     410   539   0.6  
09:30:00    116   103   159    64.5 0.041   1859   1813     235   443   0.6  
09:45:00    116   103   156    66.1 0.026   1501   1448     308   484   0.5  
10:00:00    116   103   157    65.1 0.026   1772   1717     276   487   0.5  
10:15:00    116   103   133    65.4 0.024   1405   1349     280   699   0.9  
10:30:00    116   103   163    63.9 0.039   1860 1807     237   507   0.6  
10:45:00    116   103   146    65.3 0.028   1442   1384     224   476   0.5  
11:00:00    116   103   149    65.5 0.025   1548   1485     691   600   0.5  
11:15:00    116   102   145    64.9 0.030   1635  1562     506   939   1.5
```

Look for Spikes, changes,

- what time? (10:30)
- Processor **18.6/26**
- Storage for users?
- Page rates?
- DASD I/O rates?

Wait states provide options for improvement

- Sample user status once per second, once per minute
- (900 samples per vcpu per 15 minute period)

Wait state (queue) analysis -> where to focus

- Running / CPU Wait -> CPU Subsystem
- Simulation wait (master processor) -> CPU Subsystem
- Page wait -> Paging/Storage subsystems
- Asynchronous i/o, page -> DASD subsystem

Normal idle wait states

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

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

- Set up user classes FIRST!
- Group Test vs Product
- Group application by application
- Group support servers vs production

Wait States: ESAXACT (B)

Report: ESAXACT Transaction Delay Analysis Veloc
Monitor initialized: 06/12/23 at 00:00:00 on 8561 serial 06C278 First

```
-----Percent non-dormant (Wait states)-----
UserID  <-Samples->
/Class  Total  In Q  Run Sim  CPU  SIO  Pag  E-  D-  T-  Tst <Asynch>
-----  -----  ---  ---  ---  ---  ---  ---  ---  ---  ---  ---  ---  ---  ---
10:30:00 1740 2441 14 0 38 0 0 0 0 0.1 0.3 48 0.0 . .
Hi-Freq: 298K 143K 12 0.1 24 0.0 0.0 0 0.9 0.0 0.1 64 0.1 0 0.0
***User Class Analysis***
zVPS      9900  748 1.9 2.4 2.0 0.7 0 0 7.5 1.9 0 91 0 0 0
Other    288K 142K 12 0.1 24 0.0 0.0 0 0.7 0.0 0.1 64 0.1 0 0.0
***Top User Analysis***
PGFS3J2 13440 12961 18 0.0 31 0 0 0 0 0 0 51 0 0 0
PGFS1J4 14400 12934 17 0.1 31 0 0 0 0 0 0.1 52 0 0 0
PGFS1J2 13440 12518 14 0.1 24 0 0 0 0 0 0.2 62 0.0 0 0
PGFS2J2 14400 13077 13 0.1 27 0 0 0 0 0 0 60 0 0 0
PFRS3J2 1800 1550 27 0 38 0 0 0 0 0 0.1 35 0 0 0
PADMIN1 1440 1436 17 0.1 37 0 0 0 0 0 0 45 0 0 0
PFRS3J4 720 672 36 0.4 37 0 0 0 0 0 0 25 0.9 0 0
PSUBS4J3 8400 1590 13 0.2 21 0 0 0 0 0 0.1 65 0.1 0 0
PSUBS1J2 4320 1548 11 0.2 23 0 0 0 0 0 0.3 66 0.1 0 0
PSUBS2J2 3360 1432 9.2 0.1 22 0 0 0 0 0 0 69 0 0 0
```

What is the wait state? Identifiable? A problem? Why so many samples?

- Processor wait?
- Page wait / IO Wait non existent? **NO PROBLEMS Except CPU**

User Configuration: ESAUSRC (B)

Report: ESAUSRC User C Velocity Softporate vsZMAP 5
 Monitor initialized: 06/12/ 8561 serial 06C278 First record : 06/12/23 00:0
 Monitor period: 86400:00) Last record: 06/12/23 24:0

UserID	ClassID	Account Code	> CPU Type	<-----SHARE----->				<---CPU-us>			<-MDC>		<Sto <VM Dflt	
				<Normal> Rel	<--MAX--> Abs	Typ	Shre	Lim -it	<Count> Def	Qck On	MDsp	FS		INS
PADMIN1	Other	PADMIN1	IFL	200	2	2	Y	N	N	20G
PFRS4J2	Other	PFRS4J2	IFL	200	2	2	Y	N	N	12G
PFRS4J4	Other	PFRS4J4	IFL	200	2	2	Y	N	N	20G
PFRS6	Other	PFRS6	IFL	200	2	2	Y	N	N	20G
PGFS1J2	Other	PGFS1J2	IFL	1600	16	16	Y	N	N	85G
PGFS1J4	Other	PGFS1J4	IFL	1600	16	16	Y	N	N	85G
PGFS2J2	Other	PGFS2J2	IFL	1600	16	16	Y	N	N	85G
PGFS3J2	Other	PGFS3J2	IFL	1600	16	16	Y	N	N	85G
ZADMIN	zVPS	ZADMIN	IFL	100	1	1	Y	N	N	128M
ZMAP	zVPS	ZMAP	IFL	100	1	1	N	N	N	64M
ZSERVE	zVPS	ZSERVE	IFL	300	1	1	Y	N	N	32M
ZTCP	zVPS	ZTCP	IFL	.	3.0	.	.	.	1	1	N	N	N	64M
ZVPS	zVPS	ZVPS	IFL	100	1	1	N	N	N	64M
ZWEBLOG	zVPS	ZWEBLOG	IFL	100	1	1	N	N	N	48M
ZWRITE	zVPS	ZWRITE	IFL	.	3.0	.	.	.	1	1	Y	N	N	32M

Look for “Interesting configurations”

- Large relative shares / absolute shares
- Shares matter most when there is CPU contention
- CPU Counts, **matching shares (100 Rel / vcpu)**
- CPU Type (IFL, CP)
- Virtual machine storage sizes (too large?, largest?)
- MDC enabled where important? **(ZWRITE!!)**
- Quickdsp? No-op after 6.3

Report: **ESAUSP2** User Resource Rate Re
Monitor initialized: 06/12/23 at 00:00:00

```
-----<---CPU time---> <----Main Storag
UserID <(Percent)> T:V -<Resident> Lock
/Class Total Virt Rat Totl Activ -ed
-----
PGFS3J2 222.4 221.2 1.01 20M 20.0M 1519
PGFS1J2 224.0 222.7 1.01 16M 15.8M 1506
PGFS3J2 226.5 224.6 1.01 19M 18.7M 1660
PGFS3J2 227.4 225.6 1.01 20M 20.0M 1745
PGFS1J4 227.4 226.1 1.01 17M 16.9M 1638
PGFS1J4 228.3 226.7 1.01 16M 15.7M 1687
PGFS1J4 241.8 240.9 1.00 16M 15.8M 1497
PGFS3J2 244.3 242.2 1.01 20M 20.0M 1620
PGFS1J2 245.2 244.2 1.00 16M 15.8M 1507
PGFS3J2 247.2 246.1 1.00 19M 18.7M 1523
PGFS3J2 248.5 247.4 1.00 20M 20.0M 1623
PGFS1J2 255.9 254.3 1.01 17M 16.8M 1555
PGFS1J4 276.5 274.5 1.01 16M 15.8M 1604
PSUBS4J1 280.8 279.7 1.00 7.3M 7253K 1152
```

Look for “Interesting configurations”

- Sort cpu for whole day (cpu by thread)
- Pick largest samples (PGFS1J2/3J2/1J4,PSUB4J1)
- How many vcpu really needed?
- How much extra overhead in cpu management?
- Probably should implement **zVRM**....

Linux CPU Details by program: ESAHSTA

Report: **ESAHSTA** LINUX HOST Application Report Ve
 Monitor initialized: **06/12/23** at 00:00:00 on 8561 serial 06C278 Fi

Node/ Date Time	Process/ Application name	<Application Status Counts>				<-----Processor----->			
		Total	Actv	Run-	Res	Load	<---Utilization--->		
				ning	Wait	-ed	Percent	seconds	Avg
10:30:00									
Node	Groups								
Other	*Totals*	10967	1523	106	11K	1.2	783.3	7050.1	0.1
	cpuplugd	56.0	56.0	0.2	55.6	0.2	4.4	40.0	0.1
	gpfs.gan	4.7	4.7	0	4.7	0	2.6	23.5	0.6
	java	280.5	279	0	280	0	579.5	5215.7	2.1
	ksoftirq	237.5	101	5.7	232	0	3.0	26.7	0.0
	kworker/	819.7	51.3	1.1	818	0.1	0.7	6.5	0.0
	mmfsd	57.4	57.4	0	57.4	0	87.3	786.0	1.5
	mmsysmon	57.9	57.9	0	57.9	0	13.1	118.1	0.2
	oneagent	150.0	150	0.5	150	0	34.7	312.6	0.2
	pmcollec	1.0	1.0	0	1.0	0	2.6	23.3	2.6
	pmsensor	58.1	58.1	0	58.1	0	8.1	72.5	0.1
	python	177.0	97.1	0	177	0	18.2	163.8	0.1
	rcu_sche	38.1	38.1	19.8	18.3	0	0.5	4.7	0.0
	snmpd	59.0	59.0	59.0	0	0	7.4	66.7	0.1
	splunkd	117.7	61.5	0.5	117	0	13.1	118.2	0.1
	systemd	59.0	59.0	0.7	58.3	0	1.5	13.4	0.0
	tuned	49.7	49.7	0	49.7	0	0.7	6.0	0.0
	xfssaild/	284.0	124	5.5	278	0	1.6	14.1	0.0

What is running? Process data by program name

- What is workload? (java)
- Anything interesting? (cpuplugd,oneagent?)
- Several servers not instrumented

Report: ESAUSCP Virtual Machine VCPU Analysis
 Monitor initialized: 06/12/23 at 00:00:00 on 8561 serial 06C278

UserID	<---CPU time-->			<-SHARE-->	CPU	<---Percent					
	CPUvadd	<-Percent-->				<-Samples-->	Run	Sim	CPU		
	Cnt	TOT	Virt	Type	Value	TYPE	Total	In Q			
11:15:00	0	1577	1562	.	.	.	297K	135K	11	0.1	13
PSUBS4J1	10	280.8	279.7	REL	2000	IFL	6600	5814	44	0.2	26
CPU-00		28.27	28.15	REL	1000	IFL	660	656	40	0	25
CPU-01		30.18	30.06	REL	1000	IFL	660	620	42	0.2	27
CPU-02		27.99	27.88	REL	1000	IFL	660	602	43	0.3	24
CPU-03		27.55	27.44	REL	1000	IFL	660	593	43	0.3	26
CPU-04		27.32	27.21	REL	1000	IFL	660	581	42	0.2	24
CPU-05		28.82	28.71	REL	1000	IFL	660	573	45	0.2	28
CPU-06		27.22	27.12	REL	1000	IFL	660	564	46	0.5	24
CPU-07		28.21	28.10	REL	1000	IFL	660	553	48	0.2	26
CPU-08		28.79	28.68	REL	1000	IFL	660	546	49	0.2	24
CPU-09		26.40	26.30	REL	1000	IFL	660	526	45	0	29

Look for “Interesting configurations”

- Look at virtual CPU utilization
- 10 VCPU
- Less than 30% busy
- Linux balances CPU, ALL are used
- Each vcpu waits 1/3rd time for cpu
- Each vcpu idle 25% of the time

Linux CPU Details: ESALNXS (B)

Report: **ESALNXS** LINUX VSI System Analysis Report
Monitor initialized: **06/12/23** at 00:00:00 on 8561 ser

```
-----  
Node/      <---Load Numbers--> CPU <Processor Pct Util>  
Time       Users Procs MaxProc NBR Total  Syst  User  Idle  
-----  
10:30:00  
PGFS3J2      0   266      0 Tot 246.2 10.5  232 1077  
              1   16.3  0.9  15.0 75.5  
              2   16.6  0.8  15.5 77.7  
              3   16.1  0.7  15.0 76.0  
              4   16.1  0.7  15.1 79.0  
              5   16.3  0.7  15.3 76.0  
              6   15.6  0.7  14.7 75.1  
              7   16.1  0.7  15.1 73.6  
              8   15.9  0.6  15.0 72.0  
              9   13.9  0.6  13.0 74.0  
             10   14.4  0.6  13.6 73.4  
             11   15.0  0.6  14.2 70.9  
             12   15.4  0.7  14.5 69.7  
             13   15.0  0.7  14.1 69.2  
             14   15.2  0.7  14.3 67.7  
             15   16.3  0.7  15.3 65.0  
             16   14.8  0.6  14.0 63.0
```

Linux balances across all CPUs

- Spin locks are expensive
- More use of locks with more CPUs
- What if vcpu holds lock, but paged out?

Report: **ESAUSCP** Virtual Machine VCPU Analysis
 Monitor initialized: **06/12/23** at 00:00:00 on 8561 serial 06C278

UserID	<---CPU time-->			<---Percent							
	CPUvadd	<-Percent->		<-SHARE-->	CPU	<-Samples->					
	Cnt	TOT	Virt	Type	Value	TYPE	Total	In Q	Run	Sim	CPU
10:30:00	0	1818	1807	.	.	.	298K	143K	12	0.1	24
PGFS3J2	16	247.2	246.1	REL	1600	IFL	13440	12961	18	0.0	31
CPU-00		15.23	15.15	REL	100	IFL	840	836	17	0.1	31
CPU-01		15.95	15.88	REL	100	IFL	840	822	20	0	32
CPU-02		15.77	15.69	REL	100	IFL	840	821	18	0	31
CPU-03		15.87	15.80	REL	100	IFL	840	819	19	0	33
CPU-04		16.21	16.13	REL	100	IFL	840	813	16	0	31
CPU-05		15.57	15.50	REL	100	IFL	840	811	19	0	32
CPU-06		16.11	16.04	REL	100	IFL	840	819	17	0	30
CPU-07		15.70	15.63	REL	100	IFL	840	816	17	0	32
CPU-08		13.95	13.88	REL	100	IFL	840	814	16	0.1	31
CPU-09		14.47	14.41	REL	100	IFL	840	806	17	0	28
CPU-10		15.12	15.06	REL	100	IFL	840	809	17	0	31
CPU-11		15.52	15.46	REL	100	IFL	840	800	18	0	30
CPU-12		15.08	15.02	REL	100	IFL	840	814	18	0	31
CPU-13		15.35	15.28	REL	100	IFL	840	795	18	0.1	29
CPU-14		16.40	16.33	REL	100	IFL	840	804	18	0.1	32
CPU-15		14.92	14.86	REL	100	IFL	840	762	18	0	31

Too many VCPU cause **spin locks** – wasted cpu

- Look at largest consumer during the day
- 16 VCPU
- Less than 15% busy
- Linux balances CPU
- Cpuplug not helping – needs zVRM

Linux Configuration Guideline Summary

Virtual machine size

- Minimize until some swap
- **Minimize vcpu counts to avoid overhead**

Swapping

- Swap to virtual disk
- Define 2 virtual disks,
 - One to meet the average requirement
 - Second one for overflow - Insurance
- Use DIAG driver instead of FBA
 - Reduces I/O by factor of 8

Virtual processors

- **Minimize to meet the workload/application requirement**
- Ensure diag 9c, not 44

Infrastructure costs

- Minimize – shared resource architecture

Why Linux Configuration Guidelines?

Reduce Virtual machine size

- Large servers own storage but do not use it.
- Are they using the storage? (ESAUCD2)
- Use Option to dynamically add to virtual storage instead?

Swapping

- Linux eliminates unused storage
- Intel/x admins do not like swapping.
- Swapping to virtual disk VERY VERY fast

Virtual processors

- Minimize spin locks held by process (that might be paged out?)
- Is there diagnose activity?

Case Study

- Validate storage requirements
- Validate diagnose rates

Virtual machine VCPU count higher than available?

- Determine real requirements?

z/VM With Hiperdispatch parks engines (shared IFLs)

- 20 Physical LPAR VCPU
- SMT gives 40 threads
- Parking based on entitlement

Virtual processors

- Minimize to meet the workload/application requirement
- **Ensure diag 9c**, not 44

Infrastructure costs

- Minimize – shared resource architecture

```
Report: ESAOPER      Operator/System Log
Monitor initialized: at on
10:15:00 PGFS3J2 vcpu stopped: 1
10:15:00 PGFS3J2 vcpu stopped: 2
10:15:00 PGFS3J2 vcpu stopped: 3
10:15:00 PGFS3J2 vcpu stopped: 4
10:15:00 PGFS3J2 vcpu stopped: 5
10:17:00 PGFS3J2 vcpu started: 1
10:17:00 PGFS3J2 vcpu started: 2
10:17:00 PGFS3J2 vcpu started: 3
10:17:00 PGFS3J2 vcpu started: 4
10:17:00 PGFS3J2 vcpu started: 5
10:17:00 PGFS3J2 vcpu started: 6
10:17:00 PGFS3J2 vcpu started: 7
10:17:00 PGFS3J2 vcpu started: 8
10:17:00 PGFS3J2 vcpu started: 9
10:17:00 PGFS3J2 vcpu started: 16
10:17:00 PGFS3J2 vcpu started: 17
10:17:00 PGFS3J2 vcpu started: 18
10:17:00 PGFS3J2 vcpu started: 19
10:17:00 PGFS3J2 vcpu started: 20
10:17:00 PGFS3J2 vcpu started: 21
10:29:00 PGFS3J2 vcpu stopped: 20
10:29:00 PGFS3J2 vcpu stopped: 21
10:30:00 PGFS3J2 vcpu started: 20
10:30:00 PGFS3J2 vcpu started: 21
10:36:00 PGFS3J2 vcpu stopped: 19
```

Operational changes are logged

- Evaluated at monitor start
- Vcpu start/stops?

Cpuplugd at work

- Is it effective?

Virtual Machine Diagnose Counts: ESAUSRD (B)

Report: ESAUSRD Virtual Machine Diagnose Analysisporate vsZ
 Monitor initialized: 06/12/23 at 00:00:00 on 8561 seri: 06/12/23

UserID /ClassID	Total rate	diag counts / second									
		000	004	008	00C	010	014	024	044	09C	0A0
10:30:00	5820	0.1	0	1.1	0.2	0.1	0.0	0.1	3191	2441	.
User Class Analysis											
zVPS	78.8	0.0	.	0.4	0.1	.	0	0.0	0	482	.
Other	5741	0.1	0	0.7	0.1	0.1	0.0	0.1	3191	2441	.
PFRS4J4	15.0	13.8	1.1	.
PFRS6	33.2	33.0	0.2	.
PGFS1J2	499	201	298	.
PGFS1J4	984	319	664	.
PGFS2J2	487	149	337	.
PGFS3J2	887	253	634	.
PLAF1J2	65.1	36.2	28.8	.
PLEGESS2	13.9	13.7	0.1	.
PLEGESS4	2.7	2.6	0.0	.
PLEGNON2	10.3	10.1	0.1	.
PMET2J1	15.2	14.7	0.4	.
PSERDMGR	40.7	39.6	1.0	.
PSERFW2	37.7	36.6	1.0	.
PSER2	25.7	5.5	20.2	.
PSUBS1J2	94.2	87.0	6.5	.
PSUBS2J2	313	301	11.1	.
PSUBS3J2	42.4	39.8	2.2	.
PSUBS4J1	38.3	34.3	3.5	.
PSUBS4J3	109	93.6	14.9	.
PSUBS5J2	76.3	61.8	14.0	.

Spin locks normal

- Changes?
- Diag 44 vs Diag 9C.

If owning process paged out,

- long delay

Impacted by vcpu count

z/VM CPU Details: ESACPUA (B)

Report: **ESACPUU** CPU Utilization Report
Monitor initialized: **06/12/23** at 00:00:00 on 8561 serial 06C278

Time	<----Load---->			CPU	CPU Type	<-----CPU (percentages)---				
	<-Users-> Actv	In	Q /sec			Total util	Emul time	User ovrhd	Sys ovrhd	Idle time
10:30:00	103	163	1.1	0	IFL	85.5	82.2	0.5	2.7	14.0
				1	IFL	85.6	83.5	0.4	1.7	13.8
				2	IFL	86.0	83.9	0.4	1.7	13.6
				3	IFL	85.5	83.5	0.4	1.6	14.1
				4	IFL	86.3	84.3	0.4	1.6	13.2
				5	IFL	85.9	84.0	0.4	1.5	13.7
				6	IFL	86.2	84.1	0.4	1.7	13.4
				7	IFL	85.7	83.8	0.4	1.6	13.9
				8	IFL	83.4	80.8	0.6	2.0	16.1
				9	IFL	82.8	80.4	0.5	1.9	16.7
				10	IFL	83.4	80.6	0.6	2.2	16.1
				11	IFL	83.0	80.4	0.5	2.0	16.5
				12	IFL	86.1	83.8	0.4	1.9	13.4
				13	IFL	85.7	83.6	0.4	1.7	13.8
				14	IFL	86.1	83.6	0.6	1.9	13.4
				15	IFL	85.8	83.5	0.5	1.8	13.7
				16	IFL	86.2	83.8	0.5	1.9	13.3
				17	IFL	85.8	83.5	0.5	1.8	13.8
				18	IFL	86.3	83.8	0.5	1.9	13.2
				19	IFL	85.7	83.4	0.5	1.8	13.8
				20	IFL	52.8	50.6	0.4	1.8	16.8
				21	IFL	52.1	50.1	0.4	1.6	17.5
				22	IFL	23.9	23.0	0.1	0.7	5.8
				23	IFL	23.7	22.9	0.1	0.7	6.0
				24	IFL	0.0	0	0	0.0	0
				25	IFL	0.0	0	0	0.0	0
						-----	-----	-----	-----	-----
						1860	1807	10.6	41.8	329.5

System:

LPAR Configuration: ESALPARS

Report: ESALPARS Logical Partition Summary
 Monitor initialized: **06/12/23** at 00:00:00 on 8561 serial 06

Time	<-----Logical Partition----->						<-Assigned Shares----->				Entitl CPU Cn
	Name	Nbr	Virt CPUs	CPUs Type	<%Assigned> Total	Ovhd	<---LPAR--> Weight	Pct	<VCPU Pct> /SYS /CPU		
10:30:00	Totals:	00	16	CP	468.8	2.1	966	100			
	Totals:	00	44	IFL	3540	5.1	420	100			
	VMPROD2	06	13	IFL	1010	2.7	128	30.5	2.34	84.4	10.97
	ATPCF1	0E	1	ICF	7.5	0.2	50	50.0	50.0	50.0	0.50
	ATPCF2	0F	1	ICF	0.0	0.0	50	50.0	50.0	50.0	0.50
	ATPMVSA	01	7	CP	179.1	0.9	403	41.7	5.96	47.7	3.34
	ATPMVSA	01	4	ZIP	84.3	0.3	403	42.9	10.7	42.9	1.71
	ATPMVSB	04	5	CP	170.9	0.6	338	35.0	7.00	56.0	2.80
	ATPMVSB	04	3	ZIP	17.2	0.1	338	36.0	12.0	47.9	1.44
	ATPMVSC	02	1	CP	4.2	0.1	26	2.7	2.69	21.5	0.22
	ATPMVSD	03	3	CP	114.6	0.5	199	20.6	6.87	54.9	1.65
	ATPMVSD	03	2	ZIP	26.9	0.1	199	21.2	10.6	42.3	0.85
	VMDEV1	07	9	IFL	744.9	1.1	82	19.5	2.17	78.1	7.03
	VMDEV2	08	9	IFL	716.0	0.3	82	19.5	2.17	78.1	7.03
	VMPROD1	05	13	IFL	1069	0.9	128	30.5	2.34	84.4	10.97

Totals by Processor type:

<-----CPU----->		<-Shared Processor busy->					
Type	Count	Ded	shared	Total	Logical	Ovhd	Mgmt
CP	8	0	8	471.2	466.8	2.1	2.3
IFL	36	0	36	3544.7	3534.9	5.1	4.7
ICF	1	0	1	9.6	7.3	0.2	2.1
ZIIP	4	0	4	129.2	127.8	0.5	0.9

Look for “Shared processors”

- IFLs shared (few)
- CPU HIGH utilization
- **(98%)**
- **VMPROD2: 1010%**

LPAR Configuration: ESALPARS

Report: ESALPARS Logical Partition Summary
 Monitor initialized: **06/12/23** at 00:00:00 on 8561 serial 06

Time	<-----Logical Partition----->						<-Assigned Shares----->				Entitl CPU Cn
	Name	Nbr	Virt CPUs	CPUs Type	<%Assigned> Total	Ovhd	<---LPAR--> Weight	Pct	<VCPU Pct> /SYS /CPU		
10:30:00	Totals:	00	16	CP	468.8	2.1	966	100			
	Totals:	00	44	IFL	3540	5.1	420	100			
	VMPROD2	06	13	IFL	1010	2.7	128	30.5	2.34	84.4	10.97
	ATPCF1	0E	1	ICF	7.5	0.2	50	50.0	50.0	50.0	0.50
	ATPCF2	0F	1	ICF	0.0	0.0	50	50.0	50.0	50.0	0.50
	ATPMVSA	01	7	CP	179.1	0.9	403	41.7	5.96	47.7	3.34
	ATPMVSA	01	4	ZIP	84.3	0.3	403	42.9	10.7	42.9	1.71
	ATPMVSB	04	5	CP	170.9	0.6	338	35.0	7.00	56.0	2.80
	ATPMVSB	04	3	ZIP	17.2	0.1	338	36.0	12.0	47.9	1.44
	ATPMVSC	02	1	CP	4.2	0.1	26	2.7	2.69	21.5	0.22
	ATPMVSD	03	3	CP	114.6	0.5	199	20.6	6.87	54.9	1.65
	ATPMVSD	03	2	ZIP	26.9	0.1	199	21.2	10.6	42.3	0.85
	VMDEV1	07	9	IFL	744.9	1.1	82	19.5	2.17	78.1	7.03
	VMDEV2	08	9	IFL	716.0	0.3	82	19.5	2.17	78.1	7.03
	VMPROD1	05	13	IFL	1069	0.9	128	30.5	2.34	84.4	10.97

Totals by Processor type:

<-----CPU----->		<-Shared Processor busy->					
Type	Count	Ded	shared	Total	Logical	Ovhd	Mgmt
CP	8	0	8	471.2	466.8	2.1	2.3
IFL	36	0	36	3544.7	3534.9	5.1	4.7
ICF	1	0	1	9.6	7.3	0.2	2.1
ZIIP	4	0	4	129.2	127.8	0.5	0.9

Look for “Shared processors”

- IFLs shared (few)
- CPU HIGH utilization
- **(98%)**
- VMPROD2: 1010%

LPAR Configuration: ESALPARS

Report: ESALPAR Logical Partition Analysis
Monitor initialized: 06/12/23 at 00:00:00 on 8561 serial 06C278

Time	CEC <-Logical Partition->				<-----Logical Processor-->					
	Phys CPUs	Name	No	Pool Name	VCPU Addr	<%Assigned> Total	VCPU Ovhd	Weight/ TYPE	Polar	
10:30:00	49	VMPROD2	06	.	0	92.4	0.2	IFL	128	VHi
					1	92.8	0.2	IFL	128	VHi
					2	93.1	0.2	IFL	128	VHi
					3	93.0	0.2	IFL	128	VHi
					4	90.8	0.3	IFL	128	VHi
					5	90.9	0.3	IFL	128	VHi
					6	92.9	0.2	IFL	128	VHi
					7	92.8	0.2	IFL	128	VHi
					8	92.9	0.2	IFL	128	VHi
					9	93.0	0.2	IFL	128	VHi
					10	59.4	0.3	IFL	128	VMe
					11	25.7	0.1	IFL	128	VLo
					12	0.0	0.0	IFL	128	VLo
					LPAR	1010	2.7			

Processor availability?

- Using 10 / 13
- Add weight?

Linux Process Load: ESALNXP

Report: ESALNXP LINUX HOST Process Statistics Report
Monitor initialized: 06/12/23 at 00:00:00 on 8561 serial 06C

```
-----  
node/      <Process Ident> Nice PRTY <-----CPU Percents----->  
Name       ID      PPID Valu Valu Tot  sys user syst usrt  
-----  
10:30:00  
PADMIN1    0        0    0    0 15.0 2.59 6.41 2.49 3.48  
pmcollec   2176     1    0   20 2.58 0.49 2.10 0    0  
java       2533     1    0   20 1.97 0.07 0.33 0.46 1.10  
python     5466     1    0   20 0.66 0.06 0.16 0.28 0.16  
mmfsd      10365    9794 -20   0 0.83 0.20 0.63 0    0  
mmcesmon   17856    17855 0   20 1.58 0.01 0.03 0.52 1.02  
mmsnmpag   20841    1    0   20 0.72 0.13 0.08 0.14 0.36  
mmsysmon   46467    1    0   20 5.27 1.13 2.39 1.04 0.71  
PGFS1J2    0        0    0    0 99 4.71 93.7 0.53 0.47  
oneagent   1332     1328 0   20 0.80 0.39 0.33 0.07 0.02  
python     9368     1    0   20 0.86 0.06 0.19 0.30 0.31  
mmfsd      21161    20651 -20   0 4.04 0.76 3.28 0    0  
java       31562    1    0   20 21.3 0.49 20.8 0    0  
java       31751    1    0   20 22.4 0.51 21.9 0    0  
java       32133    1    0   20 10.6 0.26 10.3 0    0  
java       32390    1    0   20 12.3 0.47 11.9 0    0  
java       32828    1    0   20 13.7 0.53 13.2 0    0  
java       33251    1    0   20 4.93 0.23 4.70 0    0  
java       33670    1    0   20 3.66 0.19 3.47 0    0  
java       34163    1    0   20 1.11 0.07 1.03 0    0  
java       34486    1    0   20 0.94 0.07 0.87 0    0  
java       35054    1    0   20 0.72 0.07 0.65 0    0
```

Look for processes within Linux, in percent of cpu

- By relevant server (L233P)
- Correct? Relevant? Cron? Init?
- Process data adds up to same as virtual machine
- Spectrum scale being used?

Top down:

- z/VM
- Virtual machines
- VDISK / MDC / Address Space
- Linux server
- Linux process

CPU Capture ratio 100% down to server

Storage Utilization: ESASTR1

Report: ESASTR1 Masis Velocity Software Corporate vsZMA
Monitor initialized: 0600 on 8561 serial 0 First record analyzed: 06/12/23 0

Time	Users	<-GB-->	---MegaBytes---		Loggd	System	Frame	<Available>	User	NSS/	<-AddSpace>	VDISK	<MDC>	
			Table	<2gb	>2gb				Resdnt	DCSS	System	User	Rsdnt	Rsdnt
06/12/23														
08:00:00	115	1024.0	24576	55	203	1029K	67	9880	0	119	39			
08:15:00	115	1024.0	24576	56	200	1029K	67	9872	0	126	18			
08:30:00	115	1024.0	24576	56	201	1029K	67	9870	0	131	49			
08:45:00	115	1024.0	24576	58	199	1029K	68	9870	0	99	60			
09:00:00	116	1024.0	24576	59	185	1103K	69	9869	0	91	32			
09:15:00	116	1024.0	24576	59	197	1029K	69	9868	0	160	39			
09:30:00	116	1024.0	24576	59	198	1029K	69	9866	0	194	38			
09:45:00	116	1024.0	24576	59	198	1029K	69	9865	0	198	62			
10:00:00	116	1024.0	24576	59	198	1029K	69	9864	0	224	53			
10:15:00	116	1024.0	24576	57	200	1029K	69	9863	0	254	52			
10:30:00	116	1024.0	24576	56	202	1029K	69	9863	0	265	67			
10:45:00	116	1024.0	24576	56	200	1029K	69	9862	0	259	45			
11:00:00	116	1024.0	24576	57	185	1103K	69	9860	0	255	28			
11:15:00	116	1024.0	24576	57	199	1029K	69	9858	0	350	14			
11:30:00	115	1024.0	24576	57	199	1029K	69	9858	0	352	10			
11:45:00	115	1024.0	24576	57	198	1029K	69	9859	0	330	13			
12:00:00	115	1024.0	24576	57	200	1029K	69	9860	0	307	17			

Total storage analysis

- MDC? minimal? SET MDC MAX/MIN
- VDISK in use, growing?
- User resident should be large percent
- Storage constrained? (zVRM again?)

Virtual Machine Storage : ESAUCD2 (again)

Report: ESAUCD2 LINUX UCD Memory Analysis Report
 Monitor initialized: 06/12/23 at 00:00:00 on 8561 serial 06C278

Node/ Time/ Date	-----Storage Sizes (in MegaBytes)-----									
	<---Real Storage-->			<-----SWAP Storage----->			Total <-----			
	Total	Avail	Used	Total	Avail	Used	MIN	Avail	CMM	B
10:30:00										
PADMIN1	19933	201.6	19731	2810	2253	557.0	15.6	2454		0
PDDS2	19933	815.5	19117	2810	2530	279.3	15.6	3346		0
PDDS4	19933	1304	18629	2810	2580	229.4	15.6	3884		0
PDIST1J2	19933	1683	18250	2810	2810	0	15.6	4493		0
PDIST1J4	11894	171.0	11723	2810	2804	5.4	15.6	2975		0
PDMNT1J2	11894	1619	10275	2810	2739	70.5	15.6	4358		0
PFFWDMGR	11894	1484	10409	2810	2809	0.5	15.6	4293		0
PFMGAD2	19933	3317	16616	2810	2810	0	15.6	6126		0
PFMGFW1	19933	1852	18081	2810	2760	50.0	15.6	4612		0
PFMGFW3	11894	644.2	11249	2810	2482	327.4	15.6	3126		0
PFRS1J2	11894	114.3	11779	2810	1944	866.0	15.6	2058		0
PFRS1J4	11894	216.3	11677	2810	2084	725.9	15.6	2300		0
PFRS2	19933	2011	17921	2810	2783	26.6	15.6	4794		0
PFRS2J2	19933	1668	18264	2810	2751	58.2	15.6	4420		0
PFRS2J4	19933	5744	14189	2810	2810	0	15.6	8553		0
PFRS3J2	19933	967.0	18966	2810	2810	0	15.6	3777		0
PFRS3J4	11894	131.7	11762	2810	2361	448.2	15.6	2493		0
PFRS4	19933	1730	18203	2810	2782	28.0	15.6	4511		0
PFRS4J2	11894	1151	10743	2810	2274	535.7	15.6	3425		0
PFRS4J4										
PFRS6										
PGFS1J2										
PGFS1J4										
PGFS2J2										
PGFS3J2	85304	31377	53927	2810	2810	0	15.6	34187		0
PLAF1J2	19933	4021	15912	2810	2810	0	15.6	6830		0

Linux storage analysis ("megabyte")

- Swap to vdisk - good

Virtual Machine Storage : ESAVDSK

Report: ESAVDSK VDISK Analysis Report 3 Pg
Monitor initialized: 06/12/23 at 00:00:00 on

Maximum VDISK: Blocks (MB)
System storage: --No Limit--
Storage per user: --No Limit--

Owner	Space Name	<---Size---> AddSpc Pages	VDSK> Blksr	Page Slots
06/12/23 00:15:00				
PADMIN1	VDISK\$PADMIN1\$\$0300\$0036	65792	524K	65470
PADMIN1	VDISK\$PADMIN1\$\$0301\$003C	131K	1049K	131K
PDDS2	VDISK\$PDDS2\$\$\$\$0300\$0037	65792	524K	65470
PDDS2	VDISK\$PDDS2\$\$\$\$0301\$003E	131K	1049K	131K
PDDS4	VDISK\$PDDS4\$\$\$\$0300\$0042	65792	524K	65470
PDDS4	VDISK\$PDDS4\$\$\$\$0301\$004A	131K	1049K	123K
PDIST1J2	VDISK\$PDIST1J2\$0300\$003D	65792	524K	4906
PDIST1J2	VDISK\$PDIST1J2\$0301\$0043	131K	1049K	1346
PDIST1J4	VDISK\$PDIST1J4\$0300\$0038	65792	524K	24656
PDIST1J4	VDISK\$PDIST1J4\$0301\$0048	131K	1049K	1346
PFMGAD2	VDISK\$PFMGAD2\$\$0300\$0057	65792	524K	826
PFMGAD2	VDISK\$PFMGAD2\$\$0301\$0059	131K	1049K	1346
PFMGFW1	VDISK\$PFMGFW1\$\$0300\$005A	65792	524K	56941
PFMGFW1	VDISK\$PFMGFW1\$\$0301\$005E	131K	1049K	1346
PFMGFW3	VDISK\$PFMGFW3\$\$0300\$005B	65792	524K	65470
PFMGFW3	VDISK\$PFMGFW3\$\$0301\$0060	131K	1049K	131K
PFRS1J2	VDISK\$PFRS1J2\$\$0300\$0040	65792	524K	65376
PFRS1J2	VDISK\$PFRS1J2\$\$0301\$0047	131K	1049K	130K
PFRS1J4	VDISK\$PFRS1J4\$\$0300\$003F	65792	524K	65429
PFRS1J4	VDISK\$PFRS1J4\$\$0301\$0046	131K	1049K	131K
PFRS2	VDISK\$PFRS2\$\$\$\$0300\$0044	65792	524K	65085
PFRS2	VDISK\$PFRS2\$\$\$\$0301\$004B	131K	1049K	1346

Swap Guideline:

- 2 vdisks
- **PRIORITIZED**
- **Do NOT format disks**

Paging: setting best practices – when storage constrained

cp query agelist **(DEFAULT)**
Target size = 37888K (37M) 2.0% of pageable storage **(5% option)**
In use = 20480K
Pending writes = 18452K
Early writes = Yes **(Write out unused pages early)**
Sizing = Variable
Keep slot = Yes

q agelist **(installation “B”)**
Target size = 74562560K (72815M) **2.0%** of pageable storage **(Needed10%)**
In use = 8616968K
Pending writes = 8616928K
Early writes = No **(So everything waits until we write out pages)**
Sizing = Variable
Keep slot = No **(saves disk, but degrades performance of writing)**

Storage Case Study Recap

From Linux perspective, ensure snmp is operational

- **Ensure linux vdisks NOT formatted, prioritized**

From z/VM perspective, manage storage with CMM

z/VM configuration parameters

- suited for SMALL LPARS (1-2GB)
- Change AGELIST to maximum 5%
- Enable Early writes
- Keep slot – Yes if storage constrained and paging

How many CPUs needed?

- Target 80% at peak period for production
- Target 90% at peak period for dev

There IS logic to doing performance analysis

- Check ESATUNE if available
- Follow the flow chart
- See the tuning guide
- Call Velocity Software for help