Performance Analysis Flowchart

- Barton@VelocitySoftware.com
- HTTP://VelocitySoftware.com

"If you can't Measure it,
I am Just Not Interested ™"

The Analysis Problem

“z” is:

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

The challenge?

What challenge, it is all there!

- 200+ zMON panels (with menus)
- 160+ zMAP reports (with table of contents)
- 3400+ unique variables

Very few companies support full-time performance analysts.

The Analysis Flow Chart

The challenge:

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

My challenge

- Provide a flowchart to resolve problems quickly
- Describe the few panels/reports needed to solve any specific problem

This flowchart is based on decades of analysis

Velocity Tuning Guide – Each web page has more details

The Challenge: z/VM Serves Many Functions (160+ Reports)

```
ESAHDR ESATUNE

*Performance Summary
ESASSUM ESASUM

*Transaction Activity   (5)
ESAUSLA ESAXACT ESARATE
ESACLAS ESAEXCP

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

*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

*TCPIP 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
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
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 ESADSD4 ESADSD5 ESAMDC
ESAVDSK ESATAPE ESA3495
ESASCSI ESASCS2
ESASEEK

*
ESAOPER
```

The Analysis Flow Chart

Analysis starts with “is there a problem?”

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

System Configuration

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

Loads on the system and subsystems

Wait states for those impacted

Subsystem Analysis

- DASD, Storage, Paging, Processor, Network

Tuning Guide – Performance Methodology Tips

The Analysis Flow Chart

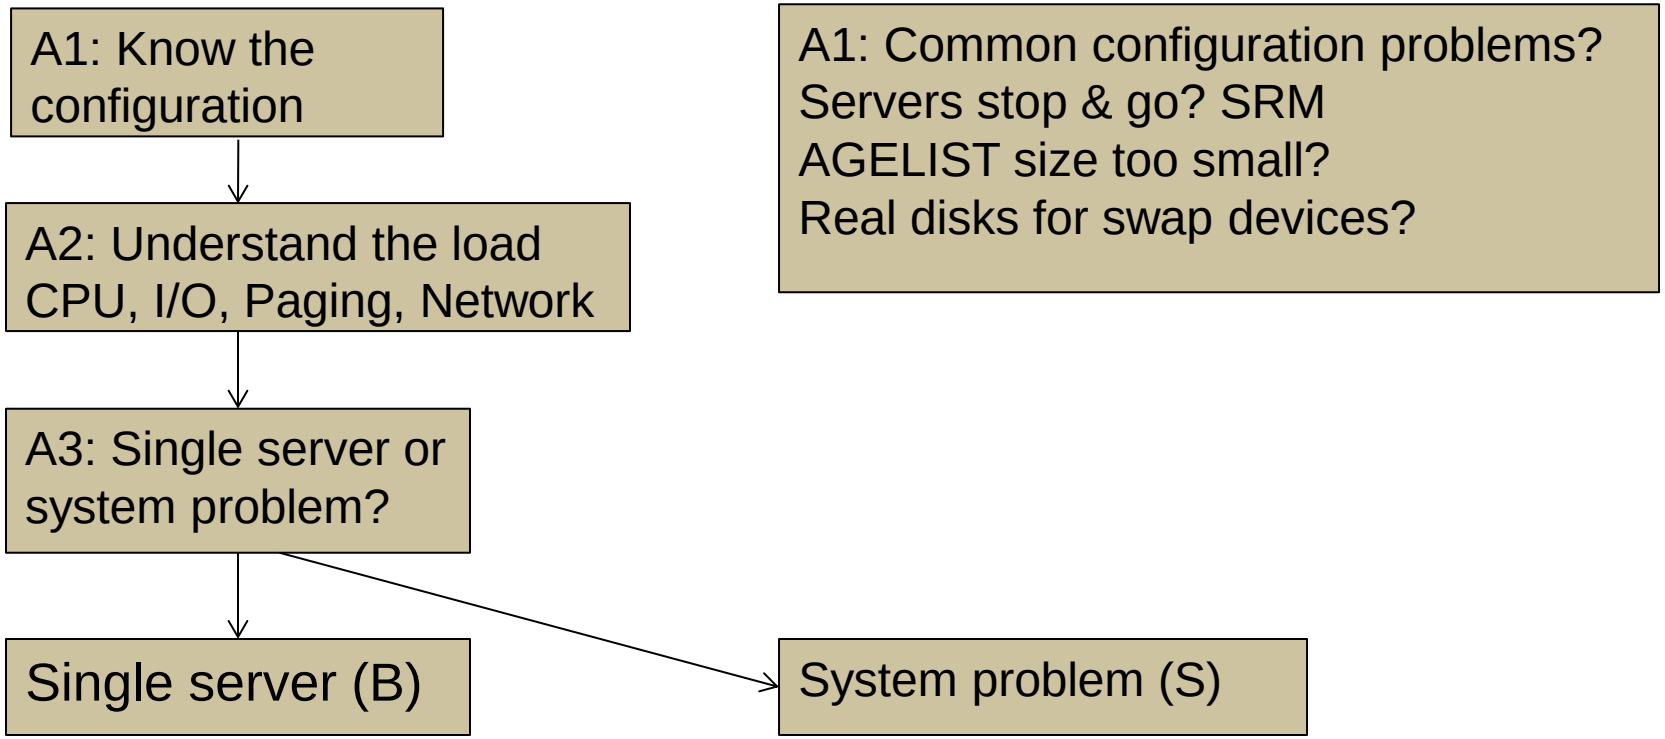

Tuning Guide – Flow Chart Analysis

The Analysis Flow Chart

Single server (B)

B1: Check wait states for constraints

B2: Understand configuration (VM) Share, VMSIZE

B3: Understand Linux configuration

Choose constraint
CPU (C), Paging (D), DASD(E), Network(F)

C1: Check process table, requirements
C2: Check system load (processes)
C3: Validate virtual CPUs
C4: Check “resident/reset”

D1: Check Linux storage/swap sizes
D2: Check paging configuration
D3: Check server page rates
D4: Virtual disk used for swap?

E1: Check data configuration
E2: Check DASD data rates

F1: Check network configuration
F2: Check network data rates

The Analysis Flow Chart

System problem (S)

S1: Check wait states for constraints

S2: Choose constraint:
CPU(T),
Paging(U),
DASD(V),
NETWORK(W)

T1: LPAR Utilization
T2: LPAR overhead
T2: Abusive servers
T3: Cron across multiple servers

U1: Storage requirements
U2: User storage?
U3: Correct virtual disk settings
U4: Page space, block paging
U5: AGELIST settings

V1: Top DASD, Control units?
V2: DASD cache, fast/write
V3: Device configuration

The Analysis Flow Chart

A1: Configuration: ESAHDR
A2: System Load: ESASSUM / ESAMAIN

B1: Check wait states: ESAXACT
B2: Virtual machine config: ESAUSRC / ESAUSR1
B3: Linux configuration: ESALNXS

C1: Process table: ESALNXC
C2: Process Load: ESALNXP
C3: Validate Virtual CPUs: ESAUSP2

D1: Linux Storage: ESAUCD2
D2: Paging configuration: ESAPSDV
D3: Server Paging Rate: ESAUSPG
D4: VDISK for swap: ESAASPC

E1: Data configuration: ESAUSEK / ESAQDIO
E2: DASD Rates: ESADSD2

F1: Network configuration: ESATCPI
F2: Network data rates: ESATCP1/2/4
F3: VSWITCH users: ESANIC
F4: VSWITCH traffic: ESAVSW
F5: OSA traffic: ESAOSA

The Analysis Flow Chart

S1: Wait states:	ESAXACT
T1: LPAR utilization:	ESALPARS
T2: LPAR overhead:	ESALPAR
T3: Abusive Server:	ESAUSP2 / ESAUSR2
T4: Cron across servers:	ESALNXP
U1: Storage requirements:	ESASTR1
U2: User storage:	ESAUSPG
U3: VDISK storage:	ESAVDSK / ESAASPC
U4: Page configuration:	ESAPSDV
U5: Page space:	ESAPSDV / ESABLKP
U6: Expanded storage:	ESAXSTO
V1: Top DASD control units:	ESADSD2
V2: DASD cache, fast/write:	ESADSD5
V3: Device configuration:	ESADSD1

Know the configuration: ESAHDR

```
Report: ESAHDR        z/VM Monitor Analysis
Monitor period:          3600 seconds (   1:00:00)
-------------------------------------------------------
z/VM Version:  5                Release 4.0 SLU 1002
TOD clock at termination                     09:49:16
Abend code of last termination
TOD clock at last IPL:             12/26/10  09:49:40
System Operator:                             OPERATOR
Time zone adjustment from GMT:               -7 hours

System Identifier                            ZVM2
Checkpoint/Warmstart Volumes            V2RES1/V2RES1
Machine Model/Type                      z10E:2097/710
System Sequence Code                  00000000000D2655
Processor  0 model/serial             2097-710  /072655 Mast
Processor  1 model/serial             2097-710  /072655
Processor  2 model/serial             2097-710  /072655
Processor  3 model/serial             2097-710  /072655
Processor  4 model/serial             2097-710  /072655

 ESAME (Memory Extension) Nucleus in use
 Power of processor in terms of service Units: 32989
 ESA/370 hardware installed
 Operating on IFL Processor(s)
 Channel Path Measurement Faclity(CPMF) Extended is inst

Main Storage installed (MB):                 70656
Main Storage Generated (MB):                 70656
Number of users in monitor file:                90
Number of DASD in monitor file:                530
Number of non-DASD in monitor file:              2
```

Common configuration problems:

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

Tuning Guide – Configuration

Know the Overall Loads: ESASSUM / ESAMAIN

```
Report: ESASSUM      Subsystem Activity                                Veloci
Monitor initialized: 04/15/11 at 10:00:00 on 2097 serial 72655          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
-------- ---- ---- ---- ------ ----- ----- ----- ----- ------ ------ ----
10:15:00   89   63 61.3  145.1 0.613   262   254  14.4  68662    862  289
10:30:00   89   63 61.3  140.3 0.545   270   261  14.4  68726    886  133
10:45:00   89   63 63.3  134.1 0.563   262   253  14.0  68806   1123  281
11:00:00   89   64 67.4  137.8 0.477   275   259  13.5  68156   2218  665
*********************************************Summary*****************
Average:   89   63 63.3  139.3 0.550   267   257  14.1  68587   1272  342
```

Look for Spikes, dramatic changes, what time?

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

[Tuning Guide – System Load]

Wait State Analysis

Wait states provide options for improvement

- State Sampling – once per minute per user
- Hi-Frequency State Sampling – once per second per vCPU
- (900 samples per vCPU per 15 minute period)

Waits reported by server, class, top user

- Look for what is impacting the users
- Recognize “running” to wait comparison

Tuning Guide – Wait State Analysis

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, SIO -> DASD subsystem
- Loading – special state, loading in working set (LDUBUF)
    - NOT a wait state, indicates thrashing
- Eligible – SRM Settings – has no value with 6.3

Normal idle wait states

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

Wait States: ESAXACT

Report: **ESAXACT** Transaction Delay Analysis Veloc
Monitor initialized: 04/15/11 at 10:00:00 on 2097 serial 72655 First

UserID /Class	<-Samples-> Total	In Q	<-----Percent non-dormant (Wait states)----------> Run	Sim	CPU	SIO	Pag	E- SVM	D- SVM	T- SVM	CF	Tst Idl	<Asynch> I/O	Pag	**Ldg**	Pct Elig
11:00:00	1335	1011	4.0	0.2	0.6	0	0.5	0	0	0.1	0	91	0.1	.	.	0
Hi-Freq:	116K	59208	4.2	0.0	**1.9**	0.0	**0.3**	0	7.9	0.1	0.0	89	**0.4**	0.1	0.2	0
***Key User Analysis ***																
RSCS	893	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
RSCSDNS	893	8	0	0	0	0	0	0	99	100	0	0	0	0	0	0
TCPIP	893	285	0.4	0	2.5	0	0	0	0	0	0	97	0	0	0	0
User Class Analysis																
*Servers	12502	822	0.7	0.1	1.0	0.2	0	0	17	4.5	0	93	0	0	0	0
*System	1786	1437	0.1	0.1	1.1	0	0.2	0	0	0	0	92	0.1	0	0.7	0
*ITM	1786	911	1.5	0.1	2.2	0	0.5	0	0	0	0	78	0.4	0.1	0.2	0
*SOA	35720	31695	7.0	0.0	2.2	0	0.3	0	0	0	0.1	88	0.6	0.0	0.1	0
*ITM	36613	23570	1.1	0.0	1.7	0	0.3	0	0	0	0	91	0.1	0.2	0.4	0
*TheUsrs	24111	480	0.2	0.8	1.3	0	0.6	0	26	5.2	0	91	0.2	0	0.2	0
Top User Analysis																
LNXUWA01	**893**	**893**	**71**	**0**	**2.8**	0	0.1	0	0	0	0	24	1.7	0.4	0	0
LNXUWA03	**1786**	**1786**	**28**	**0.2**	**5.5**	0	1.2	0	0	0	0.6	57	7.2	0.1	0.1	0
LNXUWA02	**1786**	**1786**	**27**	**0.1**	**3.6**	0	0.1	0	0	0	0.4	69	0.1	0	0.1	0
LNXQWA01	1786	1786	4.0	0	2.2	0	0	0	0	0	0	94	0.1	0	0	0
LNXDWA02	1786	1786	6.0	0	2.2	0	0.2	0	0	0	0	91	0.1	0	0	0
LNXDWA04	1786	1786	4.1	0	2.9	0	0	0	0	0	0	93	0	0	0.1	0
V2TPSP02	179	179	35	0	6.1	0	0	0	0	0	0	59	0	0	0	0

User Configuration: ESAUSRC

```
Report: ESAUSRC      User Configuration                          Velocity Software Corporate  ESAMAP 4
Monitor initialized: 04/15/11 at 10:00:00 on 2097 serial 72655   First record analyzed: 04/15/11 10:00:
Monitor period:         3600 seconds (   1:00:00)                Last record:           04/15/11 11:00
-------------------------------------------------------------------------------------------------------
                                          <---------SHARE--------->          <Status> <-MDC> <-Storage->
                    Account  ACI Grp   CPU <Normal> <--MAX-> Lim CPU <Modes>           NO  NO <-VM Size->
UserID   ClassID    Code     Name      Type  Rel Abs Typ Shre -it Cnt VM  STG SVM QDSP FS  INS Dflt    Max
-------- -------- -------- -------- ---- ---- --- --- ---- --- --- --- --- --- ---- --  --- ----- -----
LNXDMS2A *ITM     27482            .  IFL  200   .   .    .   .   2 ESA V=V   N    N  N   N   2.0G  2.0G
LNXDPB02 *System  75113            .  IFL  200   .   .    .   .   2 ESA V=V   N    N  N   N   512M  512M
LNXDWA01 *SOA     03817            .  IFL  400   .   .    .   .   2 ESA V=V   N    N  N   N   6.0G  6.0G
LNXDWA02 *SOA     03817            .  IFL  200   .   .    .   .   2 ESA V=V   N    N  N   N   4.0G  4.0G
LNXDWA03 *SOA     03817            .  IFL  200   .   .    .   .   2 ESA V=V   N    N  N   N   2.0G  2.0G
LNXDWA04 *SOA     03817            .  IFL  200   .   .    .   .   2 ESA V=V   N    N  N   N   7.0G  7.0G
LNXDWA11 *SOA     03817            .  IFL  200   .   .    .   .   2 ESA V=V   N    N  N   N   8.0G  8.0G
LNXQWA01 *SOA     03817            .  IFL  200   .   .    .   .   2 ESA V=V   N    N  N   N   7.0G  7.0G
LNXQWA02 *SOA     03817            .  IFL  200   .   .    .   .   2 ESA V=V   N    N  N   N   2.0G  2.0G
LNXQWA03 *SOA     03817            .  IFL  200   .   .    .   .   2 ESA V=V   N    N  N   N   2.0G  2.0G
LNXQWA04 *SOA     03817            .  IFL  200   .   .    .   .   2 ESA V=V   N    N  N   N   2.0G  2.0G
LNXTWA04 *SOA     03817            .  IFL  400   .   .    .   .   4 ESA V=V   N    N  N   N   5.0G  5.0G
LNXUWA01 *SOA     03817            .  IFL  100   .   .    .   .   1 ESA V=V   N    N  N   N    12G   12G
```

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?)

Tuning Guide - Configuration

Processor Subsystem Analysis

Top down:

- CEC / LPAR
- LPAR / z/VM
- Virtual machine
- Linux process

CPU Capture ratio 100% down to process

Tuning Guide – LPAR Level

LPAR Configuration: ESALPARS

```
Report: ESALPARS     Logical Partition Summary                    Velocity Software Corporate
Monitor initialized: 04/15/11 at 10:00:00 on 2097 serial 72655    First record analyzed: 04/1
-------------------------------------------------------------------------------------------
          <--Complex--> <-------Logical Partition---> <-Assigned Shares----->            Proce
          Phys Dispatch              Virt <%Assigned> <---LPAR--> <VCPU Pct> Cap-  Wait Type
Time      CPUs    Slice Name     Nbr CPUs Total  Ovhd  Weight  Pct /SYS /CPU  ped   Comp
--------  ---- -------- -------- --- ---- -----  ----  ------ ---- ---- ----  ----  ---- -----
04/15/11

10:15:00    18  Dynamic Totals:    0   34 968.7   4.9    1080 88.9
                        SYS4N3     7    5 263.5   1.2      80  6.6 1.32 23.7    No    No IFL
                        SYS4P1     3    3  22.9   0.4      60  4.9 1.65 29.6    No    No CP
                        SYS4N1     1    8 323.3   1.6     590 48.6 6.07  109    No    No CP
                        SYS4N2     2    2  17.1   0.4      60  4.9 2.47 44.4    No    No CP
                        SYS4D1     4    7  98.3   0.8     160 13.2 1.88 33.9    No    No CP
                        SYS4D2     5    5  35.9   0.4     100  8.2 1.65 29.6    No    No CP
                        SYS4D3     6    2   9.0   0.2      30  2.5 1.23 22.2    No    No CP
                        SYS4D4     8    1 100.0   0.0     Ded  5.6 5.56  100    No   Yes ICF
                        SYS4D5     9    1  98.6   0.0     Ded  5.6 5.56  100    No   Yes ICF
```

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

Tuning Guide – LPAR Level

Already Know the Overall Loads: ESASSUM / ESAMAIN

```
Report: ESASSUM      Subsystem Activity                                  Veloci
Monitor initialized: 04/15/11 at 10:00:00 on 2097 serial 72655           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
-------- ---- ---- ---- ------ ----- ----- ----- ----- ------ ------ ----
10:15:00   89   63 61.3  145.1 0.613   262   254  14.4  68662    862  289
10:30:00   89   63 61.3  140.3 0.545   270   261  14.4  68726    886  133
10:45:00   89   63 63.3  134.1 0.563   262   253  14.0  68806   1123  281
11:00:00   89   64 67.4  137.8 0.477   275   259  13.5  68156   2218  665
*******************************************Summary*****************
Average:   89   63 63.3  139.3 0.550   267   257  14.1  68587   1272  342
```

Look for Spikes, dramatic changes, what time? (11:00)

- What else changed at that time?
- Processor (Also, ESACPUU, ESACPUA)

Tuning Guide – System Load

LPAR Overhead - 2: ESALPARS

Report: **ESALPARS** Logical Partition Summary

Totals by Processor type:

Type	Count	Ded	shared	Total	Logical	Ovhd	Mgmt
	<---------CPU------->			<-Shared Processor busy->			
CP	1	0	1	21.8	21.7	0.1	0.1
IFL	11	0	11	**180.1**	**167.6**	**5.4**	**7.1**
ICF	3	2	1	100.0	99.6	0.0	0.3
ZIIP	2	0	2	0.0	0.0	0.0	0.0

Tuning Guide – LPAR Level

Look for processor type busy

- IFLs shared between LPARs (4 LPARs)
- TOTAL IFL Busy: 167% out of 1100
- Check overheads – high overhead result of too many vCPU
  - Logical overhead part of LPAR assigned
  - Physical overhead is CEC Management

LPAR Configuration - 2: ESALPARS

Report: **ESALPARS** Logical Partition Summary Velocity Softw

	<--Complex-->		<--------Logical Partition------->			<%Assigned>		<-Assigned Shares----			
	Phys	Dispatch			Virt	CPU			<---LPAR-->	<VCPU	Pct
Time	CPUs	Slice	Name	Nbr	CPUs	Type	Total	Ovhd	Weight	Pct	/SYS	/CPU
11:20:00	17	Dynamic	Totals:	0	2	CP	21.7	0.1	167	100		
			Totals:	0	18	**IFL**	173.0	5.4	100	100		
			VT4	44	7	**IFL**	112.4	3.2	60	60.0	8.57	94.3
			CFED2	15	1	ICF	100.0	0.0	Ded	5.9	0	0
			CFEH2	13	1	ICF	12.5	0.0	90	9.0	9.00	9.00
			CFEN2	14	1	ICF	100.0	0.0	Ded	5.9	0	0
			CFEA2	31	1	ICF	74.7	0.0	820	82.0	82.0	82.0
			CFEI2	30	1	ICF	12.5	0.0	90	9.0	9.00	9.00
			ITKP	21	1	CP	0.8	0.0	50	29.9	29.9	29.9
			VTT	47	2	**IFL**	3.0	0.4	2	2.0	1.00	11.0
			VT3	43	2	**IFL**	2.9	0.3	8	8.0	4.00	44.0
			VT8	45	7	**IFL**	54.7	1.6	30	30.0	4.29	47.1
			DRITE4	29	1	CP	0	0	50	29.9	29.9	29.9
			DRITE1	28	2	CP	20.9	0.0	50	29.9	15.0	15.0

Look for “Shared processors”

- IFLs shared between LPARs (4 LPARs)
- Check weights
- Assigned pct/CPU > 100 ??? -> excess share?

Tuning Guide – LPAR Weights/Overhead Analysis

LPAR Overhead - 3: ESALPAR

Report: **ESALPAR** Logical Partiti
Monitor initialized: 04/15/11 at 10:

Physical CPU Management time

CPU	Percent	Type
0	**3.838**	**CP**
1	4.412	CP
2	3.134	CP
3	2.222	CP
4	4.429	CP
5	3.924	CP
11	0.132	ZAP
13	0.068	ZAP
14	0.311	ZAP
15	1.070	ZIIP
17	1.391	ZIIP
18	0.945	ZIIP
19	1.298	IFL
24	0.121	ZAP
30	3.111	CP
33	0.408	ZAP
37	0.293	ZAP
40	1.903	IFL
41	1.786	IFL
42	1.687	IFL
43	1.161	IFL
44	1.176	IFL
45	1.158	IFL
46	1.178	IFL

Look for processor type overhead

- CPs shared between LPARs (13 LPARs)
- Check overheads:
  - High overhead result of too many vCPUs
- Total CP Utilization 835 / 900 = 93%

ESALPARS

Totals by Processor type:

	<---------CPU------->			<-Shared Processor busy->			
Type	Count	Ded	shared	Total	Logical	Ovhd	Mgmt
CP	**9**	**0**	**9**	**835.8**	**779.4**	**12.5**	**31.4**
ZAP	9	2	7	214.8	208.9	1.5	2.9
IFL	31	0	31	1778.5	1669.4	28.4	52.2
ICF	3	0	3	300.2	292.4	0.2	7.3
ZIIP	6	0	6	328.8	311.5	4.2	9.0

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
-------- ----- ----- --- ---- ----- ---- ---- ----- ---- ----- ----- ----- ----- -----
11:00:00 262.6 259.3 1.0  17M   17M  234  19M   19M 213K   13M 4346K 8891K 166.3 391.8
 ***Key User Analysis ***
TCPIP     0.12  0.05 2.4 1286  1286   79  316   316  316  5005   736  4269   0.0   0.0
 ***User Class Analysis***
*Servers  0.40  0.36 1.1  957   951    3 1704  1067   76 16285  2162 14123   0.1   0.5
*SOA     239.2 236.7 1.0  15M   15M   39  17M   17M 843K 5138K 2431K 2707K  79.1 184.0
*ITM     22.47 21.83 1.0   2M 1971K    7   2M 2117K  96K 7686K 1761K 5925K  74.7 126.4
*TheUsrs  0.21  0.18 1.2 2869  2862   17 4372  3688  135  185K 82382  102K   2.5   2.1
 ***Top User Analysis***
LNXUWA01 67.65 67.32 1.0   3M 2889K    1   3M 3146K   3M  324K 65398  259K  15.3   0.1
LNXUWA03 54.43 53.29 1.0   4M 3848K    1   4M 3855K   4M 72353 63975  8378   7.5   0.3
LNXUWA02 50.18 49.92 1.0 685K  685K    0 855K  855K 855K  381K  296K 84613   2.2   2.7
LNXQWA01 12.23 12.11 1.0   1M 1246K    7   1M 1334K   1M  592K  541K 51075   3.1   3.0
LNXDWA02 11.73 11.64 1.0 713K  713K    6 844K  844K 844K  205K 56215  148K   2.0   0.7
LNXDWA04 10.18 10.10 1.0   1M 1152K    1   1M 1248K   1M  689K  593K 96720   1.0  70.8
```

Look for consumers in percent of CPU

- By class (SOA)
- Abusive servers (LNXUWA*)?
- Correct per expected? Not a performance question

Tuning Guide – User Level

Linux Process Load: ESALNXP

Report: **ESALNXP** LINUX HOST Process Statistics Report Velocity Software Corporate ESAMAP 4.1.1 0

node/ Name	<-Process Ident-> ID	PPID	GRP	Nice Valu	<------CPU Percents----> Tot	sys	user	syst	usrt	<------CPU Seconds------> Total	sys	user	syst	usrt	<Stg (k)> Size	RSS	<-Faults/Second-> min	maj	mint	majt
LNXQWA01	0	0	0	0	11.9	1.72	7.91	1.42	0.88	107.4	15.5	71.2	12.8	7.88	11M	6M	21	0	7530	0
java	1235	1	1235	0	1.11	0.19	0.92	0	0	10.0	1.68	8.32	0	0	894K	470K	0	0	0	0
java	7124	1	7124	0	0.86	0.15	0.71	0	0	7.7	1.37	6.36	0	0	720K	415K	0	0	0	0
kcawd	8853	1	4390	0	2.24	0.01	0.02	1.38	0.83	20.1	0.10	0.14	12.4	7.49	38K	5428	2	0	7392	0
java	10522	1	10522	0	1.08	0.17	0.91	0	0	9.8	1.57	8.19	0	0	758K	437K	0	0	0	0
java	15498	1	15498	0	1.09	0.19	0.90	0	0	9.8	1.72	8.07	0	0	763K	523K	0	0	0	0
LNXUWA01	**0**	**0**	**0**	**0**	**67.0**	**5.98**	**59.0**	**1.20**	**0.81**	**601.9**	**53.8**	**531**	**10.8**	**7.29**	**13M**	**9M**	**88**	**0**	**7566**	**0**
java	4444	1	4444	0	1.10	0.07	1.03	0	0	9.9	0.65	9.25	0	0	1M	801K	0	0	0	0
kd4agent	5576	1	4362	0	4.71	1.68	3.03	0	0	42.4	15.1	27.3	0	0	99K	64K	0	0	0	0
kynagent	9569	1	4362	0	2.48	0.07	2.41	0	0	22.3	0.63	21.7	0	0	314K	212K	5	0	0	0
kcawd	9634	1	4362	0	1.92	0.01	0.01	1.14	0.75	16.4	0.06	0.13	10.3	6.78	37K	6936	1	0	7200	0
java	10547	1	10547	0	0.82	0.07	0.75	0	0	7.4	0.64	6.74	0	0	870K	743K	1	0	0	0
java	11751	4877	4877	0	0.57	0.07	0.50	0	0	5.2	0.67	4.49	0	0	617K	98K	6	0	0	0
java	11837	1	11837	0	3.28	0.12	3.16	0	0	29.5	1.10	28.4	0	0	3M	1M	1	0	0	0
java	**21374**	**15199**	**21374**	**0**	**46.3**	**3.07**	**43.2**	**0**	**0**	**416.9**	**27.6**	**389**	**0**	**0**	**3M**	**3M**	**34**	**0**	**0**	**0**
java	24567	1	24567	0	2.27	0.18	2.09	0	0	20.4	1.59	18.8	0	0	1M	831K	0	0	0	0
java	28060	1	28060	0	1.23	0.09	1.14	0	0	11.1	0.82	10.3	0	0	1M	821K	0	0	0	0
java	32428	1	32428	0	1.17	0.10	1.07	0	0	10.5	0.87	9.7	0	0	810K	538K	5	0	0	0

Look for processes within Linux, in percent of cpu

- By relevant server (LNXUWA01)
- Correct? Relevant? Cron?

Tuning Guide – Linux Process Usage

Top down:

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

CPU Capture ratio 100% down to server

Tuning Guide – Storage Level

Storage Utilization: ESASTR1

```
Report: ESASTR1      Main Storage Analysis                          Velocity Software Corporate   ESAMAP 4.1.1 01/21/
Monitor initialized: 04/15/11 at 10:00:00 on 2097 serial 72655     First record analyzed: 04/15/11 10:00:00
-------------------------------------------------------------------------------------------------------------------
         Users <-----------------------------Pages-------------------------------------------------------------->
         Loggd System  Fixed Non-  Free Frame <Available> Systm  User  NSS/DCSS  <-AddSpace> VDISK <MDC> Diag
Time        On Storage Store Pgble Stor Table <2gb  >2gb  ExSpc Resdnt Resident  Systm User  Rsdnt Rsdnt  98
-------- ----- ------- ----- ----- ---- ----- ----- ----- ----- ------ --------  ----- ----- ----- ----- ----
10:15:00    89  18088K  2252  3691  700  141K    79  1032  4710 17577K     4771   226K     0 26852 81157 1126
10:30:00    89  18088K  2252  3683  700  141K    89  1193  4686 17594K     4769   226K     0 30182 61307 1126
10:45:00    89  18088K  2252  3583  700  141K    78  1050  4681 17614K     4769   225K     0 46189 25812 1126
11:00:00    89  18088K  2252  3455  700  141K    82  1062  4688 17448K     4775   223K     0  237K  1418 1126
```

Total storage analysis (in pages)

- MDC? 300mb? SET MDC MAX/MIN
- VDISK Spike (1gb) ? Which server?
- User resident should be a large percent

Tuning Guide – System Storage Analysis

Virtual Machine Storage: ESAUSPG

```
Report: ESAUSPG      User Storage Analysis                          Velocity Software Corporate
Monitor initialized: 04/15/11 at 10:00:00 on 2097 serial 72655     First record analyzed: 04/1
-------------------------------------------------------------------------------------------
         <---Storage occupancy in pages---> <--Main Storage page Read/Write--> Pages  <Address
UserID   <---Main Storage---> <--Paging---> <-Page Writes to:--> <Page Reads:> Moved  <pages R
/Class   Total   >2gb   <2GB  Xstor   DASD   Xsto   Disk   Migr  Xstor  Disk   <2GB   VirtDisk
-------- ------ ------ ------ ------ ------ ------ ------ ------ ------ ------ ------ --------
11:00:00 17448K 16943K 504640  4346K  8891K  1120K 352582 320630 822546 149628      0   237286
***Top User Analysis***
LNXUWA01  2889K  2798K  90725  65398 258675  10999    112      0   5390  13806      0        0
LNXUWA03  3848K  3762K  85186  63975   8378  21875    277      0 221201   6714      0   223173
LNXUWA02 685385 648345  37040 296256  84613  36427   2443      0  22943   1983      0        0
LNXQWA01  1246K  1218K  28190 541178  51075  35529   2727      0  14094   2787      0     1428
LNXDWA02 713091 672702  40388  56215 148406  16314    649      0    451   1828      0        0
LNXDWA04  1152K  1120K  31859 592756  96720  13708  63725  63261   1189    942      0        0
LNXDWA03 330601 324021   6581   4194  39207   3926   5601   5345    120    734      0        8
LNXTWA04 883228 860363  22865  90734 129722   7768     31      0    182     66      0     1889
LNXUWA15 693689 664995  28694  53516 137150  10556   1382      0    553    457      0        0
```

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

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

Tuning Guide – User Storage Analysis

VDISK for Swap: ESAVDSK

Report: **ESAVDSK** VDISK Analysis Report Velocity Software Corporate

Owner	Space Name	AddSpc Pages	VDSK Blks	Cre- ates	Del- etes	<AddSpce> or Shrd	Priv rate /sec	VIO User Links	Resi- dent	<--pages--- Lock- ed	Sto- len	DASD Read
10:45:00												
LNXQWA01	VDISK$LNXQWA01$0206$0530	64256	512K	0	0	Shrd	0.00	1	122	0	0.7	0.0
LNXQWA01	VDISK$LNXQWA01$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$0520	513K	4102K	0	0	Shrd	0.28	1	4751	0	0.4	0.4
System Totals:		7805K	125M	0	0	.	5.09	204	46K	0	7.3	8.1
11:00:00												
LNXQWA01	VDISK$LNXQWA01$0206$0530	64256	512K	0	0	Shrd	0	1	46.9	0	0.1	0
LNXQWA01	VDISK$LNXQWA01$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**
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 - PRIORITIZED!!!

Tuning Guide – VDISK Analysis

Storage Utilization (By Megabyte): ESASTR1

```
Report: ESASTR1      Main Storage Analysis                        Velocity Software Corporate   ZMAP 4.2.3
Monitor initialized: 01/24/14 at 00:00:00 on 2827 serial 55AB7   First record analyzed: 01/24/14 00:00:00
-------------------------------------------------------------------------------------------------------
         Users <--------------------------MegaBytes-------------------------------------------------->
         Loggd System  Fixed Non-  Free Frame <Available> Systm  User  NSS/DCSS  <-AddSpace> VDISK <MDC>
Time        On Storage Store Pgble Stor Table <2gb  >2gb  ExSpc Resdnt Resident  Systm User  Rsdnt Rsdnt
-------- ----- ------- ----- ----- ---- ----- ----- ----- ----- ------ --------  ----- ----- ----- -----
00:05:00   114   10240    11    55    1    80  1993  2656    22   4474       97     93     0   362   241
00:10:00   115   10240    11    55    1    80  1993  2649    22   4484       97     96     0   362   242
00:15:00   114   10240    11    56    1    80  1992  2644    22   4480      103     97     0   362   243
00:20:00   113   10240    11    56    1    80  1992  2658    22   4474       98     97     0   362   242
```

Total storage analysis (“megabyte” option)

- uspg_byMB = '1'b (Impacts ESASTR1, ESAUSPG) – now default
- MDC? 240mb? SET MDC MAX/MIN
- VDISK normal?
- User resident should be a large percent
- System “oversized”

Invalid but Resident Storage Analysis

```
Report: ESAUSTR      User Storage Analysis
Monitor initialized: 07/07/15 at 13:03:48 on 2964 serial 5C2A7
-------------------------------------------------------------------
         <-------Virtual Server Storage (Pages)---------> <Resident>
UserID   Size  Alloc Resi- UFO   <-----IBR----> <AgeList> <Unreferd>
/Class               dent Activ TOT  <2gb >2gb <2gb >2gb <2gb >2gb
-------- ----- ----- ----- ----- ---- ---- ---- ---- ---- ---- ----
13:08:00  109M 93.1M   93M 93.0M 4405 1368 3037  316 123K    0    0
 ***User Class Analysis***
Servers   186K 33583 33583  8730  568  107  461 54.0  24K    0    0
ZVPS      420K 27906 27906 27906    0    0    0    0    0    0    0
TheUsers  108M 93.0M   93M 92.9M 3530 1135 2395  241  95K    0    0
 ***Top User Analysis***
LINXA195 1311K 1310K 1310K 1309K  3.0  3.0    0  3.0 1066    0    0
LINXA203 1311K 1310K 1310K 1309K  2.0  2.0    0  3.0 1072    0    0
LINXA204 1311K 1310K 1310K 1309K  3.0  1.0  2.0  3.0 1072    0    0
LINXA198 1311K 1310K 1310K 1309K  4.0  4.0    0  3.0 1072    0    0
LINXA199 1311K 1310K 1310K 1309K  4.0  4.0    0  3.0 1072    0    0
LINXA197 1311K 1310K 1310K 1309K 49.0 49.0    0  3.0 1069    0    0
LINXA155 1573K 1572K 1572K 1571K 23.0 12.0 11.0  3.0 1076    0    0
LINXA146 1573K 1572K 1572K 1571K  6.0  5.0  1.0  3.0 1073    0    0
LINXA148 1573K 1572K 1572K 1571K 17.0  3.0 14.0  3.0 1094    0    0
LINXA150 1573K 1572K 1572K 1571K  158  128 30.0  3.0 1075    0    0
```

Invalid but Resident (IBR)

- Are correct servers losing pages? (Yes)

Tuning Guide – IBR Analysis

Linux Storage - 2: ESAUCD2

```
Report: ESAUCD2       LINUX UCD Memory Analysis Report                Velocity Softwar
-------------------------------------------------------------------------------------
Node/     <-------------------------Storage Sizes (in MegaBytes)------------------->
Time/     <--Real Storage--> <-----SWAP Storage----> Total <----Storage in Use----->
Date      Total  Avail Used  Total Avail Used  MIN   Avail CMM    Buffer Cache Ovrhd
--------- ------ ----- ----- ----- ----- ----- ----- ----- ------ ------ ----- -----
*** Nodes *******
LINUXVM2   495.2   7.2 488.1  63.5  63.5   0.0  15.6  70.7      0   63.9 283.2 141.0
LNXDPB02   493.0  52.5 440.5     0     0     0  15.6  52.5      0   89.6 278.8  72.1
V2TPSP01  1992.8  28.7  1964 269.5  84.9 184.6  16.4 113.6      0  218.3 669.7  1076
V2TPSP06  1895.4 757.1  1138 256.3 256.3     0  15.6  1013      0  126.9 901.2 110.2
V2TPSP04  1895.5 756.9  1139 256.3 256.3     0  15.6  1013      0  127.0 901.1 110.4
V2TPSP05  1895.5 756.8  1139 256.3 256.3     0  15.6  1013      0  126.6 901.3 110.8
V2TPSP03  1895.4 723.4  1172 256.3 201.8  54.5  15.6 925.2      0  109.0 655.7 407.2
V2TMSP04  1501.1   8.3  1493 256.3 256.3   0.0  15.6 264.7      0   82.0 599.3 811.5
V2TMSP05  1501.1 121.7  1379 256.3 256.3   0.0  15.6 378.0      0   84.0 269.2  1026
V2TMSP02  1501.1  65.3  1436 256.3 256.3   0.0  15.6 321.6      0  105.9 599.5 730.3
V2TMSP03  1501.1  64.2  1437 256.3 256.3   0.0  15.6 320.5      0   80.4 270.3  1086
```

Linux Storage Map

- Opportunities?
  - High available (greater than 5%)
  - High buffer (greater than 20mb)
- Issues? Swap
- If swap used, but also a large buffer, CMM?

Tuning Guide – Linux Storage Analysis

Top down:

- z/VM
- Configuration
- Rates
- Space full
- Device busy

Paging rules change in 6.3

Tuning Guide – Storage Level

Paging Subsystem: ESAPSDV

Report: **ESAPSDV** Page And Spool Device Activity Velo

Dev No.	Serial	Paging Slots Avail	Paging Slots Used	Paging Slots %Use	Paging Slots Max	Paging per sec Read	Paging per sec Write	Spooli Slots Avail	Spooli Slots Used	Spooli Slots %Use	
11:00:00											
E92F	V2PAG1	1803K	1121K	62	1129K	25.2	35.1	.	.	.	
E93F	V2PAG2	1803K	1114K	62	1122K	24.1	35.2	.	.	.	
E930	V2PAG3	1803K	1117K	62	1123K	22.5	31.2	.	.	.	
E940	V2PAG4	1803K	1081K	60	1089K	21.0	35.8	.	.	.	
E933	V2PAG5	1803K	904950	50	913775	23.2	37.2	.	.	.	
E934	V2PAG6	1803K	894360	50	903958	23.7	39.4	.	.	.	
E935	V2PAG7	1803K	840048	47	848995	23.8	37.2	.	.	.	
E937	V2PAG8	1803K	709086	39	718015	24.4	37.1	.	.	.	
E93C	V2PAG9	1803K	726428	40	734888	24.8	36.1	.	.	.	
E938	V2PA10	1803K	596028	33	604582	25.0	37.4	.	.	.	
E93B	V2PA11	1803K	594606	33	603738	26.7	38.9	.	.	.	
EA4A	V2SPL1	.	.	.	.	0	0	5897K	546231	9	54
Total:		19832K	9697K	**49**	9791K	264.6	400.5	5897K	546231	9	54

Paging Configuration:

- How many devices (11)
- Equal sizes?
- How full? (50%)
- Rates reasonable? Device type dependent

Tuning Guide – System Page Analysis

Page Device Busy: ESADSD2

Report: **ESADSD2** DASD Performance Analysis Velocity Sof

Dev No.	Serial	Device Type	<--SSCH--> Total	ERP	<%DevBusy> Avg	Peak	<SSCH/sec> avg	peak	<--------DASD Response tim Resp	<--Service times--> Serv	Pend	Disc	Conn
11:00:00													
Top DASD by Device busy													
E95C	**V2U019**	**3390-9**	23344	0	**10.6**	**44.6**	**26.4**	**116.6**	4.8	4.0	0.3	1.4	2.2
E930	V2PAG3	3390-9	9170	0	6.2	19.5	10.4	29.3	5.9	5.9	0.3	0.0	5.6
E93F	V2PAG2	3390-9	9759	0	5.9	15.8	11.0	**31.7**	5.3	5.3	0.3	0.0	5.0
E93C	V2PAG9	3390-9	8101	0	5.8	17.1	9.2	29.3	6.3	6.3	0.3	0.0	6.0
E92F	V2PAG1	3390-9	10137	0	5.7	15.6	11.5	31.4	5.0	5.0	0.3	0.0	4.6
E940	V2PAG4	3390-9	8869	0	5.2	14.8	10.0	29.9	5.2	5.2	0.3	0.0	4.8
E933	V2PAG5	3390-9	8418	0	5.1	12.8	9.5	28.9	5.3	5.3	0.3	0.0	5.0
E934	V2PAG6	3390-9	7858	0	5.0	13.4	8.9	26.9	5.6	5.6	0.3	0.0	5.3
E937	V2PAG8	3390-9	7568	0	5.0	13.3	8.6	28.9	5.8	5.8	0.3	0.0	5.5
E935	V2PAG7	3390-9	8284	0	4.9	13.1	9.4	30.8	5.2	5.2	0.3	0.0	4.9
End Top DASD by Device busy													

Page Device Analysis – DASD Subsystem

- Page Devices are usually in “Top Ten DASD”
- Device busy > 20% cause for concern
- Device busy > 50% serious
- Minute by minute analysis would show a 30% “Peak”

Paging Analysis: ESABLKP

Report: **ESABLKP** Block Paging Analysis Velocity Software Corporate
Monitor initialized: 04/15/11 at 10:00:00 on 2097 serial 72655 First record analyzed: 04/15/1

	<----Load---->			Serv	<-Block->		<-Blocks Formed By->				Block	<--Block Exceptions/sec→			
	<-Users->		Tran	Time	<-Reads->		<-Steal->		<Migrate>		Fault	<Single Read>		<No Refers>	
Time	Actv	In Q	/sec	(ms)	/sec	Size	/sec	Size	/sec	Size	/sec	User	System	Migr	Steal
10:15:00	63	61.3	2.4	45.9	19.9	7.0	0.0	31.0	10.2	13.2	9.0	8.8	0.0	0.8	50.0
10:30:00	63	61.3	2.3	47.1	10.3	7.0	0.0	25.1	3.7	13.7	4.7	5.6	0.0	0	45.1
10:45:00	63	63.3	2.2	33.0	18.8	7.0	0.0	29.4	6.0	20.9	8.4	11.1	0.0	0	57.2
11:00:00	64	67.4	2.3	57.8	27.1	7.7	1.0	33.3	**26.0**	**13.6**	**11.0**	**34.6**	**0.1**	**12.9**	**176.8**

Block Paging Analysis

- Block page read – optimal 10 pages
- Steal should be zero prior to 6.3
- Migrate should be zero with 6.3 and beyond
- Pages stolen, unreferenced – Storage stress
- Single page read – goes up with 6.3

Paging Analysis: ESABLKP

Report: **ESABLKP** Block Paging Analysis TEST MAP

Time	<----Load----> <-Users-> Actv	In Q	Tran /sec	Serv Time (ms)	<-Block-> <-Reads-> /sec	Size	<-Blocks Formed By-> <-Steal-> /sec	Size	<Migrate> /sec	Size	Block Fault /sec	<--Block Exceptions/sec--> <Single Read> User	System	<No Refers> Migr	Steal
07:49:00	83	262	0.7	.	65.6	**5.6**	31.4	**18.8**	0	0	25.4	**291.2**	1.7	0	0

Block Paging Analysis for 6.3+

- Block page read – optimal 5 pages??
- Migrate should be zero (No expanded storage)
- Pages stolen, unreferenced – zero with 6.3
- Single page read – goes up with 6.3
- Faster paging devices? (new market for SSD)

DASD Subsystem Analysis

Top down:

- Configuration
- DASD I/O for system
- Rates by control unit
- Rates by device
- Rates by minidisk (by user)
- Cache

Tuning Guide – DASD Level

DASD Configuration: ESADSD1

Report: **ESADSD1** DASD Configuration Velocity Software Corporate

Dev No.	Sys ID	Serial	Device Type	SHR	<CHPIDS OnLn> 01	02	03	04	MDisk Links	<----Extent----> Type	Start	Size	<--MDC St Elig	Def
E92F	1B89	V2PAG1	3390-9	NO	**7A**	**7B**	**78**	**79**	0	Page	1	10K	Yes	On
E930	1B8A	V2PAG3	3390-9	NO	7A	7B	78	79	**0**	Page	1	10K	Yes	On
E931	1B8B	540RES	3390-9	NO	7A	7B	78	79	0	.	.	.	No	On
E933	1B8D	V2PAG5	3390-9	NO	7A	7B	78	79	0	Page	1	10K	Yes	On
E934	1B8E	V2PAG6	3390-9	NO	7A	7B	78	79	0	Page	1	10K	Yes	On
E935	1B8F	V2PAG7	3390-9	NO	7A	7B	78	79	0	Page	1	10K	Yes	On
E936	1B90	V4SPL2	3390-9	NO	7A	7B	78	79	0	.	.	.	No	On
E937	1B91	V2PAG8	3390-9	NO	7A	7B	78	79	0	Page	1	10K	Yes	On
E938	1B92	V2PA10	3390-9	NO	7A	7B	78	79	0	Page	1	10K	Yes	On
E939	1B93	VME939	3390-9	NO	7A	7B	78	79	0	.	.	.	No	On
E93B	1B95	V2PA11	3390-9	NO	7A	7B	78	79	0	Page	1	10K	Yes	On
E93C	1B96	V2PAG9	3390-9	NO	7A	7B	78	79	0	Page	1	10K	Yes	On
E93E	1B98	VME93E	3390-9	NO	7A	7B	78	79	0	.	.	.	No	On
E93F	1B99	V2PAG2	3390-9	NO	7A	7B	78	79	0	Page	1	10K	Yes	On
E940	1B9A	V2PAG4	3390-9	NO	7A	7B	78	79	0	Page	1	10K	Yes	On
E958	1BB2	V2U011	3390-9	NO	7A	7B	78	79	113	.	.	.	**Yes**	On
E959	1BB3	V2U013	3390-9	NO	7A	7B	78	79	15	.	.	.	Yes	On
E95A	1BB4	V2U015	3390-9	NO	7A	7B	78	79	39	.	.	.	Yes	On
E95B	1BB5	V2U017	3390-9	NO	7A	7B	78	79	29	.	.	.	Yes	On

DASD Configuration
- Multi channels to devices
- No minidisks on page devices
- MDC enabled appropriately

Tuning Guide – DASD Configuration Analysis

Control Unit Data Rates: ESADSD2

```
Report: ESADSD2       DASD Performance Analysis                      Velocity Sof
Monitor initialized: 04/15/11 at 10:00:00 on 2097 serial 72655        First record
-------------------------------------------------------------------------------
                                                    <--------DASD Response tim
 Dev          Device <--SSCH--> <%DevBusy> <SSCH/sec->          <--Service times-->
 No. Serial   Type    Total  ERP  Avg  Peak   avg  peak  Resp   Serv Pend Disc Conn
---- ------ ------ ------ --- ---- ---- ----- ----- -----  ---- ---- ---- ----
11:00:00

1800 Control Unit    3000   0   0.0  0.0   3.4   3.4   0.3   0.3  0.3    0  0.0

1880 Control Unit    3000   0   0.0  0.0   3.4   3.4   0.3   0.3  0.2    0  0.0

E900 Control Unit  186192   0   0.7  1.8 210.4 530.4   3.9   3.8  0.3  0.4  3.1

E980 Control Unit    1500   0   0.0  0.0   1.7   1.7   0.4   0.4  0.4    0  0.1

EA00 Control Unit   42722   0   0.1  0.5  48.3  93.2   2.1   2.1  0.3  0.2  1.5

EA80 Control Unit    1500   0   0.0  0.0   1.7   1.7   0.4   0.4  0.3    0  0.1
                    ------ ---  ---- ---- ----- ----- -----  ---- ---- ---- ----
System:             237914   0   0.2  0.5 268.8 633.7   3.4   3.4  0.3  0.3  2.7
```

DASD Control Units Rates, Performance:

- Shows activity by control unit
- Pend - indication of cache problems
- Compare control units to determine normality

Tuning Guide – DASD Utilization Analysis

Data Rates, Device Performance: ESADSD2

Report: **ESADSD2** DASD Performance Analysis Velocity Sof

Dev No.	Serial	Device Type	<--SSCH--> Total	ERP	<%DevBusy> Avg	Peak	<SSCH/sec-> avg	peak	<--------DASD Response tim Resp	<--Service times--> Serv	Pend	Disc	Conn
11:00:00													
Top DASD by Device busy													
E95C	**V2U019**	**3390-9**	23344	0	**10.6**	**44.6**	**26.4**	**116.6**	4.8	4.0	0.3	1.4	2.2
E930	V2PAG3	3390-9	9170	0	6.2	19.5	10.4	29.3	5.9	5.9	0.3	0.0	5.6
E93F	V2PAG2	3390-9	9759	0	5.9	15.8	11.0	31.7	5.3	5.3	0.3	0.0	5.0
E93C	V2PAG9	3390-9	8101	0	5.8	17.1	9.2	29.3	6.3	6.3	0.3	0.0	6.0
End Top DASD by Device busy													
1880	Control Unit		3000	0	0.0	0.0	3.4	3.4	0.3	0.3	0.2	0	0.0
E900	Control Unit		186192	0	0.7	1.8	**210.4**	**530.4**	3.9	3.8	0.3	0.4	3.1
E980	Control Unit		1500	0	0.0	0.0	1.7	1.7	0.4	0.4	0.4	0	0.1
EA00	Control Unit		42722	0	0.1	0.5	48.3	93.2	2.1	2.1	0.3	0.2	1.5
System:			237914	0	0.2	0.5	268.8	633.7	3.4	3.4	0.3	0.3	2.7

DASD Rates, Performance:

- System: rate, average service/response time
- Pend, disconnect low -> else DASD cache
- Connect low -> else faster channels
- Response = service, else queueing
- Peak busy for device (1 minute peak)

DASD Cache: ESADSD5

```
Report: ESADSD5       3990-3 Cache Analysis                          Velocity Software Corporate   ES
Monitor initialized: 04/15/11 at 10:00:00 on 2097 serial 72655     First record analyzed: 04/15/11
-------------------------------------------------------------------------------------------------------
             Pct. <--------------------per second--------------------> <--------Write activity per se
 Dev         Actv <------Total--------> <----Read----> <--Seq Read--> Total   DFW  DFW SEQ         NVS
 No.  Serial Samp  I/O Hits Hit% Read%  I/O Hits Hit%  I/O Hits Hit%   I/O   I/O Hits  I/O Hit% Full
---- ------ ---- ---- ---- ---- ----- ---- ---- ---- ---- ---- ---- ----- ---- ---- ---- ---- ----
11:00:00
***Top DASD by Device busy***
E95C  V2U019  100 25.9 21.3 82.0  62.5 16.2 11.5 71.3    0    0    0   9.7  9.7  9.7    0  100    0
E930  V2PAG3  100 10.1  7.6 75.9  58.6  5.9  3.5 58.9    0    0    0   4.2  4.2  4.2    0  100    0
E93F  V2PAG2  100 10.9  8.5 77.3  58.2  6.4  3.9 61.1    0    0    0   4.6  4.6  4.6    0  100    0
E93C  V2PAG9  100  8.9  6.3 70.0  65.8  5.9  3.2 54.5    0    0    0   3.1  3.1  3.1    0  100    0
E92F  V2PAG1  100 11.2  8.5 76.3  59.2  6.6  4.0 60.1    0    0    0   4.6  4.6  4.6    0  100    0
***End Top DASD by Device busy***

1800 CtlUnit  100  220  219  100   4.6 10.1  9.7 96.7    0    0    0 209.6  210  210    0  100    0
1880 CtlUnit  100  1.8  1.8  100 100.0  1.8  1.8  100    0    0    0     0    0    0    0    0    0
E900 CtlUnit  100  368  331 89.8  27.3  101 63.3 62.9    0    0    0 267.8  268  268    0  100    0
EA00 CtlUnit  100 73.0 72.3 99.1   6.9  5.0  4.4 86.8    0    0    0  68.0 68.0 68.0    0  100    0
---- ------ ---- ---- ---- ---- ----- ---- ---- ---- ---- ---- ---- ----- ---- ---- ---- ---- ----
System:       100  663  624 94.2  17.7  118 79.2 67.4    0    0    0 545.3  545  545    0  100    0
```

DASD Cache:

- Hit percent (read, write)
- Low hit% -> need more cache or batch (backups)
- NVS full -> fast write stops
- Data shows activity from all LPARS to device/ctl unit

Tuning Guide – DASD Cache Analysis

Data activity by user: ESASEEK, ESAUSEK

```
Report: ESAUSEK      User DASD Seeks Report                          Velocity
Monitor initialized: 04/15/11 at 10:00:00 on 2097 serial 72655       First re
Monitor period:         3600 seconds (   1:00:00)                    Last rec
-------------------------------------------------------------------------------
 Userid    Dev Volume <--Minidisk-> <Cylinder-> Total <---Non-zero---> Read
 /Time     No. Serial Ownerid  Addr Start Stop  Seeks Seeks Pct. Dist. Pct.
--------  ---- ------ -------- ---- ----- ----- ----- ----- ---- ----- ----
**********************************************Summary*********************
Average:
LNXUWA01  E95C V2U019 LNXUWA01 0233 40591 40722  2389  1699 71.1  9685    0
          EA59 V2U016 LNXUWA01 0210     1 16698 14762  9854 66.8  2220    0
          E903 V2U034 LNXUWA01 021F 15207 32689  7542  4394 58.3  1578 16.6
          E903 V2U034 LNXUWA01 0220 32986 33350    63    63  100 10459    0
          E95A V2U015 LNXUWA01 0209     1 12084 10345  4849 46.9  4981 28.4
          E95A V2U015 LNXUWA01 020A 12085 19617  2608  2024 77.6  8521    0
          E95A V2U015 LNXUWA01 020F 52329 53478    24    16 66.7 33363    0
          E926 V2U041 LNXUWA01 0232  6062  7598  2239  1544 69.0  4294    0
          E95B V2U017 LNXUWA01 021E 26231 28597    42    36 85.7 10207    0
          E95E V2U023 LNXUWA01 0204 63268 63850   675   327 48.4 21376    0
          EA58 V2U014 LNXUWA01 0205  3029  3033     3     2 66.7 31999    0
```

DASD activity by virtual machine: ESAUSEK

DASD activity by minidisk/volume: ESASEEK

- Correlate activity to poor performing disks
- Note read percent for Linux minidisks

Network Activity

- Configuration
- Rates
- Errors
- vSwitch/Guest LAN

Tuning Guide – Network Level

Network Configuration: ESATCPI

```
Report: ESATCPI      TCPIP Interface Configuration Report          Velocity Sof
-------------------------------------------------------------------------------
NODE      Idx        Speed <-Status-> Up   <-------------------Interface---------
          Nbr MTU    (Est) Oper Admin Time MACAddress         Description Type
--------  --- -----  ----- ---- ----- ---- ----------------- ----------- --------

*****************************************Summary************************
Average:
TCPIP       1  1500 1000M    .     .    .                    . ETHERNET viETHERNET-
VMLOCAL     1  1500 1000M   UP    UP    0 00:20:20:20:20:20 ETHERNET viETHERNET-
LINUXVM2    2  1500  100M   UP    UP    0 02:00:00:00:00:30 eth0       ETHERNET-
LNXDPB02    3  1492  100M   UP    UP    0 02:00:00:00:00:04 eth0       ETHERNET-
V2TPSP01    1 16436   10M   UP    UP    0 00:20:20:20:20:20 lo         Software
            2  1500  100M   UP    UP    0 02:00:00:00:00:15 eth0       ETHERNET-
V2TMSP05    1 16436   10M   UP    UP    0 00:20:20:20:20:20 lo         Software
            2  1500  100M   UP    UP    0 02:00:00:00:00:09 eth0       ETHERNET-
V2TMSP02    1 16436   10M   UP    UP    0 00:20:20:20:20:20 lo         Software
            2  1500  100M   UP    UP    0 02:00:00:00:00:06 eth0       ETHERNET-
V2TMSP03    1 16436   10M   UP    UP    0 00:20:20:20:20:20 lo         Software
            2  1500  100M   UP    UP    0 02:00:00:00:00:07 eth0       ETHERNET-
LNXUWA01    1 16436   10M   UP    UP    0 00:20:20:20:20:20 lo
            4  1492  100M   UP    UP    0 02:00:00:00:00:22 eth0
```

Interface configuration

- Ethernet adapter
- Loop back
- MTU check

Tuning Guide – Network Configuration Analysis

Network Data Rates: ESATCP4

```
Report: ESATCP4       TCPIP Hardware Layer/Interfaces Report          Ve
-------------------------------------------------------------------------
Date/          <Total Octets> Avg  <-Subnet packets / Sec-> <-----Pack
Time           <-Per second->   Q  <-Unicast-> <NonUnicast> <In Error>
Node       IFT Input   Output  Len  Input Outpt Input Output Inpt Outpt
-------- ---- -----  ------  ---- ----- ----- ----- ------ ---- -----
11:00:00
*** Nodes *******
TCPIP    -   1 16897  6231.9     0 25.74  21.3     0      0    0     0
VMLOCAL  -   1 16859  6223.3     0 25.70  21.3     0      0    0     0
LINUXVM2 -   2 93.06  208.92     0  0.38   0.4     0      0    0     0
LNXDPB02 -   3 293.8  590.32     0  2.25   2.4     0      0    0     0
V2TPSP01 -   1 418.3  418.26     0  1.54   1.5     0      0    0     0
         -   2 188.6  666.61     0  0.95   1.2     0      0    0     0
V2TMSP05 -   1 323.6  323.61     0  6.16   6.2     0      0    0     0
         -   2  1517  2481.8     0  4.70   4.5     0      0    0     0
LNXDMS2A -   3 103.4  299.74     0  0.47   0.6     0      0    0     0
LNXUWA01 -   1 21167   21167     0 57.81  57.8     0      0    0     0
         -   4  109K    122K     0 236.9 268.5     0      0    0     0
LNXDWA02 -   1 920.2  920.23     0  5.03   5.0     0      0    0     0
         -   4  9112   10306     0 25.84  24.3     0      0    0     0
```

Network activity, server, by interface
Understand rates
Check for errors

[Tuning Guide – Network Data Rate]

QDIO Data Rates: ESAQDIO

```
Report: ESAQDIO      Queued I/O Report                                  Velocity Software Corpor
Monitor initialized: 04/15/11 at 10:00:00 on 2097 serial 72655     First record analyzed: 0
-------------------------------------------------------------------------------------------------
                                  Number <QDIO SIGA Instructions/Sec-> <-Throughtput / sec->
Date/     Dev.          Virt QDIO Queues <---Guest----> <----CP------> <Buffers> <--Bytes→
Time      Nmbr owner    DevN Fmt  In Out Read Writ "s"  Read Writ "s"  Sent From  Sent  From
-------- ---- -------- ---- ---- -- --- ---- ---- ---- ---- ---- ---- ---- ---- ----- -----
-------------------------------------------------------------------------------------------------
11:00:00 0000 Totals   0000 QDIO  0   0    0    0    0    0  693    0 1066  676  644K  422K
         F3D8 VSWCTRL2 F3D8 QDIO  1   1    0    0    0    0  573    0  895  535  527K  306K
         F3E0 VSWCTRL2 F3E0 QDIO  1   1    0    0    0    0  119    0  171  141  118K  117K
         F53E LNXUWA02 7002 HPER  1   4    0    0    0    0  0.6    0    1    0    89     0
*************************************************Summary*************************************
Average: 0000 Totals   0000 QDIO  0   0    0    0    0    0  639    0 1040  621  615K  441K
         F3C8 VSWCTRL1 F3C8 QDIO  1   1    0    0    0    0    0    0    0    0     0     0
         F3D8 VSWCTRL2 F3D8 QDIO  1   1    0    0    0    0  530    0  891  491  529K  322K
         F3E0 VSWCTRL2 F3E0 QDIO  1   1    0    0    0    0  108    0  149  130 85716  119K
         F3F0 VSWCTRL1 F3F0 QDIO  1   1    0    0    0    0    0    0    0    0     0     0
         F515 LNXDPB02 7002 HPER  1   4    0    0    0    0    0    0    0    0     0     0
         F518 LNXDWA01 7002 HPER  1   4    0    0    0    0    0    0    0    0     0     0
         F53B LNXUWA01 7002 HPER  1   4    0    0    0    0    0    0    0    0     0     0
         F53E LNXUWA02 7002 HPER  1   4    0    0    0    0  0.6    0    1    0    92     0
         F542 LNXUWA03 7002 HPER  1   4    0    0    0    0    0    0    0    0     0     0
         F545 LNXUWA04 7002 HPER  1   4    0    0    0    0    0    0    0    0     0     0
         F548 LNXDMS2A 7002 HPER  1   4    0    0    0    0    0    0    0    0     0     0
```

QDIO activity

- Hipersockets
- Virtual switch

Tuning Guide – QDIO Utilization Analysis

Guest LAN / Virtual Switch Data Rates: ESANIC / ESATCP4

```
Screen: ESANIC   Velocity Software - VSIVM4
1 of 3  Virtual NIC Activity

                                      <-- Data Th
          VSWITCH/ <Virtual NIC> <Bytes/Sec>
Time      GuestLAN Userid   Addr   Sent  Rcvd
--------- -------- -------- ---- ----- -----
15:24:00 VSIINT   TIML2    0600  4048 11059
                  SLES11X3 0600  1160   628
                  RKS2LV   0600   481   839
                  REDHAT71 0600   573   376
                  REDHAT64 0600  1818   846
                  REDHAT56 0600  2415   964

F1=Help                    PF3=Quit      PF4=S
            PF8=Forward    PF9=Sort      PF10=
```

```
Screen: ESATCP4  Velocity Software - VSIVM4
1 of 2  TCPIP Hardware Layer / Interfaces

                              <Total Octets>
          Node/               <-Per second->
Time      Group    Interface    Input Output
--------- -------- ---------  ------ ------
15:24:00 redhat71 enccw0.0.   390.87 584.07
          redhat71 lo              0      0
          redhat64 eth0        918.03   1908
          redhat64 lo              0      0
          redhat6x eth0        818.33   1900
          redhat6x eth1          0.47      0
          redhat6x lo            3059   3059
          redhat6  eth0          1862   4660
          redhat6  lo               0      0
```

Guest LAN / Virtual Switch activity

- ESANIC: CP Monitor data
- ESATCP4: SNMP data
- Compare “received to input”

- Redhat7 renamed eth0

Tuning Guide – Network User Data Rate Analysis

OSA Adapter: ESAOSA

```
Report: ESAOSA          OSA System Configuration Report
-------------------------------------------------------
Collector <---------OSA Configuration--> MacAddress
Node      Idx  Name  Nbr  Type Level Shrd Active
--------  ---  ----- ---  ---- ----- ---- ------------
00:15:00
OSA178      2  OSA1    0 1G Eth 6.00  Yes 6CAE8B483FD4

redhat6x    3  OSA1    0 1G Eth 6.00  Yes 6CAE8B483FD4
```

OSA data collected via SNMP

- Configuration data
- Total data
- Data by LPAR if shared
- (New with 4.3)

```
Report: ESAOSA                        Velocity Software Corporate
-----------------------------------------------------------------
Collector <-------- LPAR Bus  CPHID  KBytes/Sec Packets/sec
Node      Idx  Name   NBR Util Util    IN   OUT    In   OUT
--------  ---  ------  --- ---- ----  ----- ----- ----- -----

OSA178      2  OSA1   Tot    0   15    4.0   8.1  25.5  16.7
                        2    0    .     53    15
                        4    0    .    288   291
                        5    0    .     59    55

redhat6x    3  OSA1   Tot    0   15   12.7   5.3  26.8  16.8
                        1    0    .      2    56
                        2    0    .     61    15
                        4    0    .    312   400
                        5    0    .     59    55
```

Tuning Guide – OSA Utilization Analysis