
# *Performance Analysis Flowchart*
# *The Science of Analysis*

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….*

## "z" is:

- Very large (one box easily supports 1,000 servers)
- Very complex and shares resources
- z/VM is very well instrumented
- TECHNOLOGY KEEPS CHANGING…

## The challenge? What challenge, it is all there!

- 200++ zmon panels (with menus), plus graphs, plus…
- 150+ zmap reports (with table of contents)
- 3400++ unique variables

## Few companies have full time performance analysts, even though immense savings

## The challenge – when running high utilization....

- Performance problems are visible,
- "z" applications are often impacted by other applications in shared resource environments

## The performance analysis challenge

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

## This flowchart is based on decades of analysis

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

***Performance Summary (4)**

ESAHDR ESATUNE ESASSUM ESASUM

***Transaction Activity (5)**

ESAUSLA **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 ESAPOOL

***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 ESAEDEV

*Virtual NETWORK Reporting (7)

ESAQDIO ESAQDI2 ESANIC
ESAVSWC ESAVSW ESAVSW2
ESAOSA

*Operational Logging

ESAOPER

# The Challenge: Many subsystems

- z/VM "traditional" Applications (Source: APPLMON)
  - Allows virtual machines to create application unique data
  - (33 reports currently? )

```
*Shared File System (IBM)   (7)

ESASFS1 ESASFS2 ESASFS3 ESASFS4
ESASFS5 ESASFS6 ESASFS7

*Byte File System (IBM)    (2)

ESABFS1 ESABFS2 ESABFS3

*CMS Multitasking (IBM)    (1)

ESAMTSK

*Web Serving Reports (Velocity)   (8)

ESAWEB1 ESAWEB2 ESAWEB3 ESAWEB4
ESAVWS1 ESAVWS2 ESAVWS3 ESAVWS4


*TCPIP / Network Reporting (IBM)   (15)

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

# Data Source: snmp - VERY efficient, really!

- Network: mib ii (first available, from any snmp enabled server
  ESATCPC ESATCP1 ESATCP2 ESATCP3 ESATCP4

- Linux servers: UCD mib – Linux system data
  ESAUCD1 **ESAUCD2** ESAUCD3 **ESAUCD4** ESAUCDD

- More Linux: Velocity mib – Linux process, ram, system data
  ESALNXD ESALNXS ESALNXR **ESALNXP** ESALNXA
  ESALNXC **ESALNXF** ESALNXU ESALNXV
  ESALNXM ESALNXUP

- Linux Application "vendor" mibs
  ESAJVM ESAORAC ESAORAG ESAORAS ESAORAW (Java, Oracle)
  ESAGPFN ESAGPFF ESAGPFF ESAGPFD ESAGPFS (GPFS)
  ESAMNG1 ESAMNG2 ESAMNG3 ESAMNG4 ESAMNG5 (MONGODB)
  ESADOCK1 ESADOCK2 ESASSC (DOCKER, SSC)

# Network: Source: snmp!

- Network: mib ii (first available, from any snmp enabled server
  ESATCPC ESATCP1 ESATCP2 ESATCP3 ESATCP4
- Linux and Microsoft servers: Host mib
  ESAHST1 ESAHST2 ESAHST3 ESAHST4 ESAHSTA
- VSE mibs: IBM, Velocity
  ESAVSEC ESAVSES ESAVSEP ESAVSEJ
  ESAVSEP ESAVSEJ

u

# *The 25 "z/VM" Reports you need in the order you need them*

System Configuration: ESAHDR
System load (z/VM): ESASSUM

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

Paging

- 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)

## System Configuration

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

## Evaluate:

- Loads on the system subsystems
- Wait states for those impacted
- Subsystem Analysis
  - DASD, Storage, Paging, Processor, Network
- Server analysis (linux, vse, z/os….)

# *The Analysis Flow Chart*

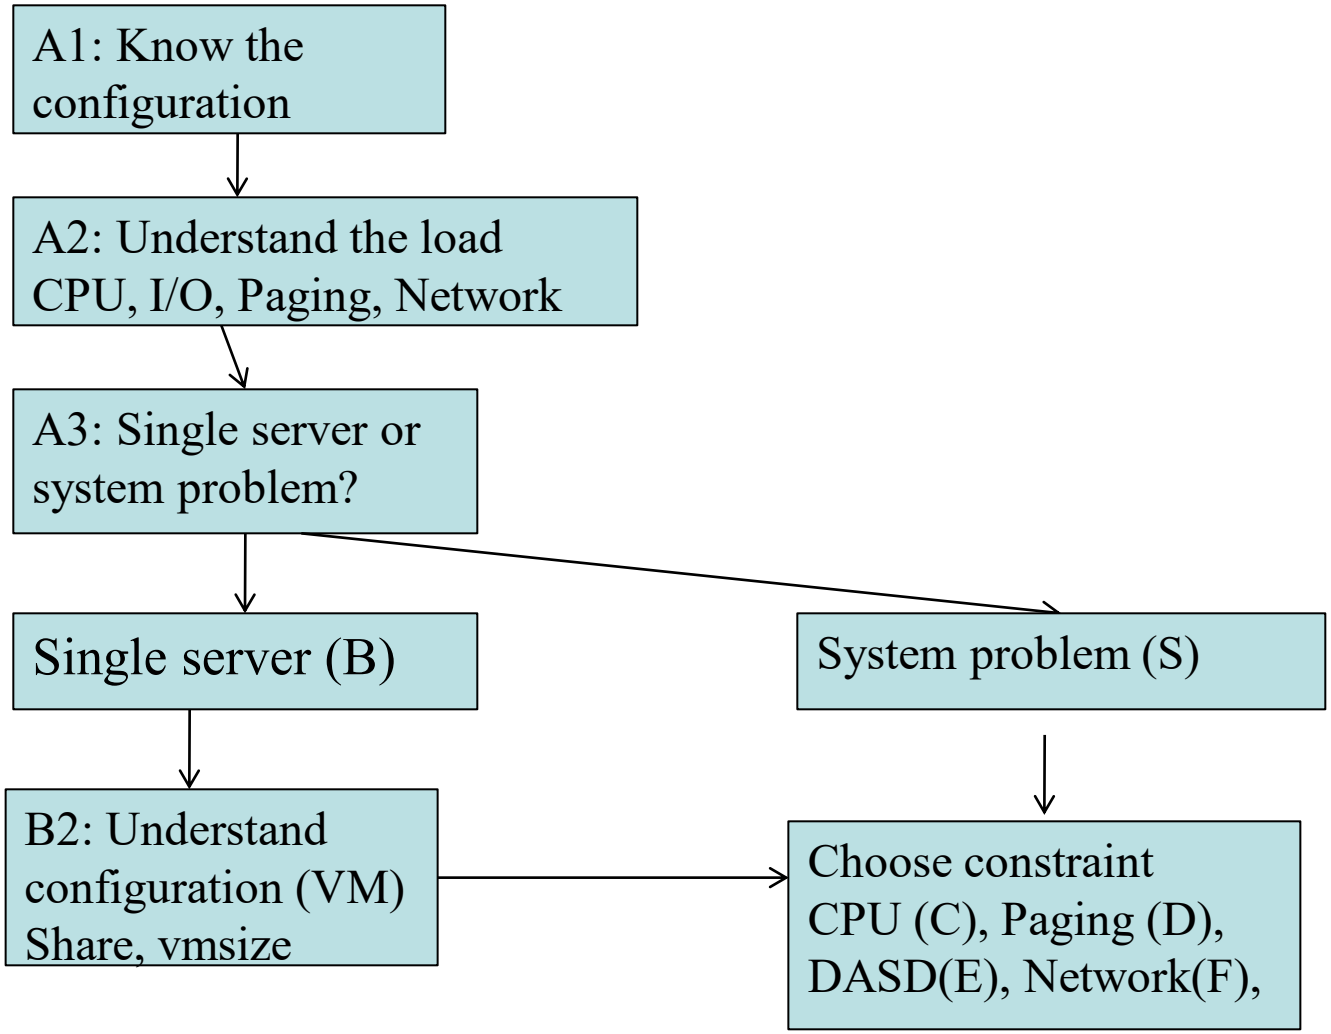




VELOCITY SOFTWARE

# *The System wide Analysis Flow Chart*

System problem (S)

S1: Check Wait states for constraints

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

CPU
T1: LPAR Utilization
T2: LPAR overhead
T2: Abusive servers
T3: cron across multiple servers

Storage:
U1: Storage requirements
U2: User storage?
U3: Correct vdisk settings
U4: Page space, block paging

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

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

V1: top dasd? Control units: ESADSD2
V2: dasd cache, fast/write: ESADSD5
V3: Device configuration: ESADSD1

# The Single user Analysis Flow Chart

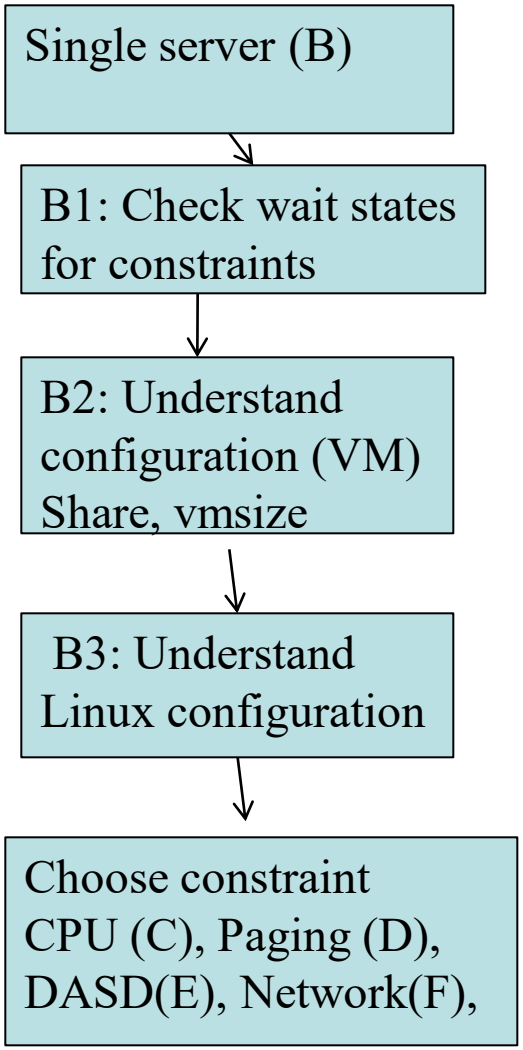


CPU:
C1: Check Linux process table, "magnitude"
C2: check system load (what processes)
C3: Validate virtual cpus

Paging:
D1: check linux storage/swap sizes
D2: check paging configuration
D3: Check server page rates
D4: Vdisk used for swap?

DASD:
E1: check DASD configuration
E2: check DASD Data rates

Network:
F1: check network configuration
F2: check network Data rates

# *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: DASD configuration: ESAUSEK
ESAQDIO
E2: DASD Rates: ESADSD2

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

# *Know the configuration: ESAHDR*

```
Report: ESAHDR       z/VM Monitor Analysis
Monitor initialized: 04/15/21 at 00:00:00 on 8562 serial 040F78
Monitor period:        86400 seconds (  24:00:00)
------------------------------------------------------------
ZMAP Release                                    5.1.2.1
History Source                                  5.1.2.0
Key Expiration for: VELOCITY                   01/01/22
Components licensed:  IFL  ZOS  TUN  PRO
Monitor file created:            04/15/21  00:00:00

z/VM Version:  7              Release 1.0 SLU 2001
TOD clock at termination                       20:00:00
Abend code of last termination
TOD clock at last IPL:           09/24/20  06:35:42
System Operator:                               OPERATOR
Time zone adjustment from GMT:                 -4 hours

System Identifier                              VSIVM4
Checkpoint/Warmstart Volumes                VM4RS2/VM4RS2
Machine Model/Type                          Z15s:8562/A02
Multithreading Status(SMT): Enabled
  Core Thread count:        2
  Enabled Count:            2

System Sequence Code                 0000000000040F78
Processor  0 model/serial            8562-A02  /040F78 Master
Processor  1 model/serial            8562-A02  /040F78
Processor  3 model/serial            8562-A02  /040F78
Processor  4 model/serial            8562-A02  /040F78
Processor  5 model/serial            8562-A02  /040F78

 CPU(GP) Capability Factor:     8563
 CPU(IFL) Capability Factor:     478
 CPU Cycles/ns:                 4500
 Operating on IFL Processor(s)
 Main Storage installed (MB):                  14336
```

Common configuration problems

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

# *Sample Issue, start with LPAR configuration*

```
Report: ESAHDR       z/VM Monitor Analysis
Monitor initialized: 08/29/18 at 20:00:35 on 2827
Monitor period:         3600 seconds (   1:00:00)
-------------------------------------------------------------
ZMAP Release                                    4.4.0.0
History Source                                  4.3.4.0
Monitor file created:              08/29/18  20:00:35

z/VM Version:  6                Release 4.0 SLU 1601
TOD clock at last IPL:             01/21/18  01:10:34
System Operator:                             OPERATOR
Time zone adjustment from GMT:               -4 hours

System Identifier                              VML1
Checkpoint/Warmstart Volumes            LV1RES/LV1RES
Machine Model/Type                      EC12:2827/702

Apar installed: VM65918 VSW Aggreg Load Balance
Apar installed: VM65925 NICDEF Security Controls
Apar installed: VM65942 Add user diagnose tables
Apar installed: VM65943 eav minidisks (large)
Apar installed: VM65985 System Hang with mdc on
Apar installed: VM65988 Processor Scalability
Apar installed: VM66026 HYPERPAV/PAV/ZHPF Monitr
Apar installed: VM66063 High PR/SM LPAR Mgt Time
Apar installed: VM66083 Wait, CUWait monitor rec

System Sequence Code                 00000000000469X7
Processor  0 model/serial          2827-702  /0669X7 Master
Processor  1 model/serial          2827-702  /0669X7

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

Useful information

- Performance APARs
- Release significant
- Master processor significant
- Running on IFLs….



VELOCITY SOFTWARE

# *Know the CEC configuration: ESAHDR*

```
Report: ESAHDR       z/VM Monitor Analysis
Monitor initialized: 08/29/18 at 20:00:35 on 2827
Monitor period:         3600 seconds (   1:00:00)
-------------------------------------------------------------------
Totals by Processor type:
<---------CPU-------> <-Shared Processor busy>
Type Count Ded shared total assigned Ovhd Mgmt
---- ----- --- ------ ----- -------- ---- ----
CP       2   0      2 165.4    162.7  1.6  2.8
IFL      4   4      0   0.0        0    0  0.0
ICF      2   1      1 100.4     99.8  0.0  0.6
ZIIP     1   0      1   6.1      5.6  0.2  0.5

Number of logical partitions defined:           20
```

Configuration overview

- IFLs? Not used….

Number of LPARs

# *Know the overall loads: ESASSUM / ESAMAIN*

```
Report: ESASSUM       Subsystem Activity
Monitor initialized: 08/29/18 at 20:00:35 on 2827
-------------------------------------------------------------------------------
         <---Users----> Transactions <Processor> Storage (MB) <-Paging--> <----
         <-avg number->    Per  Avg. Utilization Fixed Active <pages/sec> <-DAS
Time       On Actv In Q Minute  Resp Total Virt.  User Resid. XStore DASD Rate
-------- ---- ---- ---- ------ ----- ----- ----- ----- ------ ------ ---- ----
08/29/18
20:01:35   48   31 35.0   27.0 0.404   200   193 162.2 143708      0  931  381
20:02:35   48   31 34.0   31.0 0.222   200   192 162.2 143706      0  165  168
20:03:35   48   37 35.0   29.0 0.202   199   191 162.2 143717      0  272  124
20:04:35   48   33 35.0   31.0 0.236   200   193 162.3 143709      0   29   99
20:05:35   48   35 35.0   30.0 0.013   200   190 162.2 143711      0  118  120
20:06:35   49   39 35.0   34.0 0.266   200   188 162.2 143721      0  175  164
20:07:35   49   31 35.0   32.0 0.206   200   188 162.2 143708      0   40  103
```

Look for Spikes, dramatic changes, what time? (20:01 is worst?

- 2 IFL Processors fully utilized
- Storage for users
- Page rates
- DASD I/O rates
- (Transacations are for traditional workloads)

```
13
14
18
67
20
98
```



VELOCITY SOFTWARE

## When performance is a problem:

- Complaints about performance
- Production requirements not met
- Improvements requested

## Start with ESAXACT

- What resource is critical to workload performance?
- System level
- User class / workload level
- Virtual machine level

## Wait states provide options for improvement

## 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
- Loading -> thrashing (not a wait state, operational state)

## Normal idle wait states

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

# *Wait State Analysis directs further Analysis*

## Two types of Wait state samples provided by virtual cpu

- 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*

```
Report: ESAXACT       Transaction Delay Analysis                    Velocity S
Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669X7     First reco
-------------------------------------------------------------------------------
                      <-----Percent non-dormant (Wait states)-----------------
UserID   <-Samples->                              E-  D-  T-      Tst <Asynch>
/Class    Total  In Q Run Sim CPU SIO Pag SVM SVM SVM  CF Idl I/O Pag Ldg Oth
-------- ----- ----- --- --- --- --- --- --- --- --- --- --- --- --- --- ---
20:01:35    48    35 5.7   0  40   0   0   0   0   0   0  54   0   .   .   0
Hi-Freq:  3720  2114 5.4 0.0  41   0 0.7   0 2.8 0.0   0  52   0 0.2 0.7 0.0
 ***Key User Analysis ***
TCPIP       60    55   0   0 3.6   0   0   0   0   0   0  96   0   0   0   0
 ***User Class Analysis***
Servers    600     3   0  33   0   0   0   0   0   0   0  67   0   0   0   0
Velocity   600    15 6.7   0 6.7   0   0   0   0   0   0  87   0   0   0   0
*Prod     1860  1860 5.8   0  45   0 0.4   0   0   0   0  49   0   0 0.1 0.1
*Util      120   120   0   0  15   0   0   0   0   0   0  85   0   0   0   0
TheUsrs    480    61  10   0  15   0  11   0  21 1.6   0  36   0 6.6  20   0
 ***Top User Analysis***
L24BP      120   120  18   0  77   0   0   0   0   0   0 5.0   0   0   0   0
L233P      120   120  18   0  74   0   0   0   0   0   0 7.5   0   0   0 0.8
L200P      120   120  10   0  80   0 4.2   0   0   0   0 5.8   0   0   0   0
L239P      120   120 6.7   0  51   0   0   0   0   0   0  43   0   0   0   0
L203P      120   120 5.0   0  75   0 0.8   0   0   0   0  19   0   0   0   0
L20BP      120   120 4.2   0  38   0   0   0   0   0   0  58   0   0   0   0
L244P      120   120 3.3   0  28   0   0   0   0   0   0  69   0   0   0   0
L208P      120   120 3.3   0  51   0   0   0   0   0   0  46   0   0   0   0
L20DP      120   120  13   0  44   0   0   0   0   0   0  43   0   0   0   0
L224P      120   120 4.2   0  27   0   0   0   0   0   0  69   0   0   0   0
```

System is in CPU wait, and needs more CPU, focus on CPU

Users impacted very slightly by paging

# *Wait States: ESAXACT*

```
Report: ESAXACT      Transaction Delay Analysis
Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669X7
-------------------------------------------------------------------
                              <Percent non-dormant (Wait states)
UserID    <-Samples->                                        E-
/Class    Total  In Q  Run  Sim  CPU  SIO  Pag  SVM
--------  -----  -----  ---  ---  ---  ---  ---  ---
20:01:35     48     35  5.7    0   40    0    0    0
Hi-Freq:   3720   2114  5.4  0.0   41    0  0.7    0

***User Class Analysis***
Velocity    600     15  6.7    0  6.7    0    0    0
*Prod      1860   1860  5.8    0   45    0  0.4    0
*Util       120    120    0    0   15    0    0    0
TheUsrs     480     61   10    0   15    0   11    0
```

Production "*Prod":

- Running 5.8%
- CPU Wait 45%  (90% of work time)
- Remainder is test-idle

Eliminate CPU Wait improves performance by 10x

## How to improve CPU Wait?

- Faster CPU
- More engines for LPAR
- Higher weighting for LPAR
- Virtual machine share of LPAR (user share)

# User Configuration: ESAUSRC

**Report: ESAUSRC** User Configuration Velocity Software Corporate ZMAP 4.4.0
Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669X7 First record analyzed: 08/29/18 20:00:35

| | | | | | | <-------SHARE-------> | | | | | <---CPU---> | | <Status> | | <-MDC> | | | <Storage> | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | Account | ACI Grp | <CP POOL> | CPU | <Normal> | | <--MAX-> | | Lim | <Count> | | | Qck | NO | NO | | <VM Size> | |
| UserID | ClassID | Code | Name | PoolName | Type | Rel | Abs | Typ | Shre | -it | Def | On | Mode | SVM | Dsp | FS | INS | Dflt | Max |
| DISKACNT | Servers | 10 | . | . | IFL | 100 | . | . | . | . | 1 | 1 | ESA | N | N | N | N | 32M | 32M |
| DTCVSW1 | TheUsrs | DTCVSW1 | . | . | IFL | 3000 | . | . | . | . | 1 | 1 | ESA | N | Y | N | N | 128M | 128M |
| DTCVSW2 | TheUsrs | DTCVSW2 | . | . | IFL | 3000 | . | . | . | . | 1 | 1 | ESA | N | Y | N | N | 128M | 128M |
| EREP | Servers | EREP | . | . | IFL | 100 | . | . | . | . | 1 | 1 | ESA | N | N | N | N | 32M | 32M |
| FTPSERVE | Servers | FTPSERVE | . | . | IFL | 100 | . | . | . | . | 1 | 1 | XC | N | Y | N | N | 32M | 32M |
| L10CU | *Util | L10CU | . | . | IFL | 100 | . | . | . | . | 2 | 2 | ESA | N | N | N | N | 16G | 16G |
| **L203P** | ***Prod** | **L203P** | . | . | **IFL** | 120 | . | . | . | . | 2 | 2 | **ESA** | **N** | **N** | **N** | **N** | **6.0G** | **6.0G** |
| **L208P** | ***Prod** | **L208P** | . | . | **IFL** | 200 | . | . | . | . | 2 | 2 | **ESA** | **N** | **N** | **N** | **N** | **6.0G** | **6.0G** |
| **L215P** | ***Prod** | **L215P** | . | . | **IFL** | . | 8 | Abs | 8.0 | Hrd | 2 | 2 | **ESA** | **N** | **N** | **N** | **N** | **8.0G** | **8.0G** |
| **L216P** | ***Prod** | **L216P** | . | . | **IFL** | 100 | . | . | . | . | 1 | 1 | **ESA** | **N** | **N** | **N** | **N** | **4.0G** | **4.0G** |
| **L222P** | ***Prod** | **L222P** | . | . | **IFL** | 100 | . | . | . | . | 2 | 2 | **ESA** | **N** | **N** | **N** | **N** | **6.0G** | **6.0G** |
| **L224P** | ***Prod** | **L224P** | . | . | **IFL** | 200 | . | . | . | . | 2 | 2 | **ESA** | **N** | **N** | **N** | **N** | **12G** | **12G** |
| **L225P** | ***Prod** | **L225P** | . | . | **IFL** | 70 | . | . | . | . | 2 | 2 | **ESA** | **N** | **N** | **N** | **N** | **2.0G** | **2.0G** |
| **L23AC** | **TheUsrs** | **L23AC** | . | . | **IFL** | 100 | . | . | . | . | 1 | 1 | **ESA** | **N** | **N** | **N** | **N** | **2.0G** | **2.0G** |
| **L23BP** | ***Prod** | **L23BP** | . | . | **IFL** | 50 | . | . | . | . | 1 | 1 | **ESA** | **N** | **N** | **N** | **N** | **6.0G** | **6.0G** |
| **L233P** | ***Prod** | **L233P** | . | . | **IFL** | 200 | . | . | . | . | 2 | 2 | **ESA** | **N** | **N** | **N** | **N** | **18G** | **18G** |
| **L239P** | ***Prod** | **L239P** | . | . | **IFL** | 200 | . | . | . | . | 2 | 2 | **ESA** | **N** | **N** | **N** | **N** | **32G** | **32G** |
| **L24BP** | ***Prod** | **L24BP** | . | . | **IFL** | 200 | . | . | . | . | 2 | 2 | **ESA** | **N** | **N** | **N** | **N** | **27G** | **27G** |
| **L24FP** | ***Prod** | **L24FP** | . | . | **IFL** | 100 | . | . | . | . | 1 | 1 | **ESA** | **N** | **N** | **N** | **N** | **2.0G** | **2.0G** |
| **L244P** | ***Prod** | **L244P** | . | . | **IFL** | 200 | . | . | . | . | 2 | 2 | **ESA** | **N** | **N** | **N** | **N** | **12G** | **12G** |
| **L25BP** | | | | | | | | | | | | | | | | | **N** | **6.0G** | **6.0G** |

Best practices, 100 relative share per virtual cpu

- No Large relative shares, use absolute shares
- CPU Counts, matching shares (100 Rel / vcpu)
- CPU Type (IFL, CP)
- Virtual machine storage sizes (too large?, largest?)

## Top down:

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

## CPU Capture ratio 100%

- By lpar
- By virtual machine
- By Linux process

# LPAR Configurations, Utilizations: ESALPARS

**Report: ESALPARS** Logical Partition Summary Velocity Softw
Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669X7 First record a

| | <--Complex--> | | <--------Logical Partition-------> | | | | | <-Assigned Shares---- | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | Phys | Dispatch | | | Virt | CPU | <%Assigned> | | <---LPAR--> | | <VCPU Pct |
| Time | CPUs | Slice | Name | Nbr | CPUs | Type | Total | Ovhd | Weight | Pct | /SYS | /CPU |
| 20:01:35 | 9 | Dynamic | Totals: | 00 | 12 | CP | 198.7 | 0.9 | 1480 | 100 | | |
| | | | **VML1** | **06** | **2** | **IFL** | **200.0** | 0.0 | Ded | 22.2 | 0 | 0 |
| | | | VML3 | 02 | **2** | **IFL** | 200.1 | 0.0 | Ded | 22.2 | 0 | 0 |
| | | | CER2 | 18 | **2** | **CP** | **8.9** | 0.1 | 60 | 4.1 | 2.03 | 4.05 |
| | | | CER2 | 18 | 1 | ZIP | 1.1 | 0.0 | 100 | 6.2 | 6.21 | 6.21 |
| | | | ESN1 | 19 | **2** | **CP** | **13.4** | 0.1 | 105 | 7.1 | 3.55 | 7.09 |
| | | | ESN1 | 19 | 1 | ZIP | 0.1 | 0.0 | 485 | 30.1 | 30.1 | 30.1 |
| | | | ICF11 | 1E | 1 | ICF | 99.4 | 0.0 | Ded | 11.1 | 0 | 0 |
| | | | ICF12 | 0D | 1 | ICF | 49.9 | 0.0 | 10 | 50.0 | 50.0 | 50.0 |
| | | | ICF13 | 1F | 1 | ICF | 50.0 | 0.0 | 10 | 50.0 | 50.0 | 50.0 |
| | | | PRD1 | 11 | **2** | **CP** | **100.3** | 0.1 | **666** | 45.0 | 22.5 | 45.0 |
| | | | PRD1 | 11 | 1 | ZIP | 0.8 | 0.0 | 530 | 32.9 | 32.9 | 32.9 |
| | | | PRD3 | 1B | **2** | **CP** | **61.7** | 0.3 | 433 | 29.3 | 14.6 | 29.3 |
| | | | PRD3 | 1B | 1 | ZIP | 2.0 | 0.1 | 345 | 21.4 | 21.4 | 21.4 |

Look for "Shared processors"

- First LPAR is "us", z/vm where data collected, 2 dedicated IFLs
- IFLs shared between LPARs (none)
- Check weights
- Assigned pct/CPU > 100 ??? -> excess share?

# CEC Configuration,Utilization: ESALPARS

```
Report: ESALPARS      Logical Partition Summary
Monitor initialized: 08/29/18 at 20:00:35 on 2827
-------------------------------------------------------------
Totals by Processor type:
<---------CPU-------> <-Shared Processor busy->
Type Count Ded shared  Total  Logical Ovhd Mgmt
---- ----- --- ------ ------ -------- ---- ----
CP       2   0      2  200.0    197.7  0.9  1.3
IFL      4   4      0    0.0        0    0  0.0
ICF      2   1      1  100.5     99.9  0.0  0.6
ZIIP     1   0      1    4.9      4.4  0.1  0.4
```

At "CEC Level", Look for "Shared processors"

- Know capacity for each processor type
- Dedicated show up as 100% "assigned"
- Detail is on ESALPAR

# *Already Know the overall local LPAR load: ESASSUM / ESAMAIN*

```
Report: ESASSUM      Subsystem Activity                              Velocity Software Corporate   ZMAP 4.4.0 08/31/18   Page   34
Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669X7      First record analyzed: 08/29/18 20:00:35
-------------------------------------------------------------------------------------------------------------------------------
         <---Users----> Transactions <Processor> Storage (MB) <-Paging--> <-----I/O-----> <MiniDisk> Spool Communications Captur
         <-avg number->    Per  Avg. Utilization Fixed Active <pages/sec> <-DASD--> Other <-Cache--> Page <-per second->  Ratio
Time       On Actv In Q Minute  Resp Total Virt.  User Resid. XStore DASD Rate Resp  Rate  Rate %Hit Rate   IUCV    VMCF (pct)
-------- ---- ---- ---- ------ ----- ----- ----- ----- ------ ------ ---- ---- ---- ----- ----- ---- ----- ------ ------ ------
20:01:35   48   31 35.0   27.0 0.404   200   193 162.2 143708      0  931  381  1.6   1.1   3.4 56.9     0    453      0 100.00
20:02:35   48   31 34.0   31.0 0.222   200   192 162.2 143706      0  165  168  0.4   1.2   0.6 51.4     0    362      0 100.00
20:03:35   48   37 35.0   29.0 0.202   199   191 162.2 143717      0  272  124  0.5   1.1   0.4 38.1     0    341      0 100.00
```

Next step, "Top Down" CPU in LPAR

- Validate capture ratio

```
Report: ESASSUM      Subsystem Activity                                Velocity Software Corporate ZMAP
-------------------------------------------------------------------------------------------------------
         <---Users----> Transactions <Processor> Storage (MB) <-Paging--> <-----I/O-----> <MiniDisk>
         <-avg number->    Per  Avg. Utilization Fixed Active <pages/sec> <-DASD--> Other <-Cache-->
Time       On Actv In Q Minute  Resp Total Virt.  User Resid. XStore DASD Rate Resp  Rate  Rate %Hit
-------- ---- ---- ---- ------ ----- ----- ----- ----- ------ ------ ---- ---- ---- ----- ----- ----
20:01:35   48   31 35.0   27.0 0.404   200   193 162.2 143708      0  931  381  1.6   1.1   3.4 56.9
20:02:35   48   31 34.0   31.0 0.222   200   192 162.2 143706      0  165  168  0.4   1.2   0.6 51.4
20:03:35   48   37 35.0   29.0 0.202   199   191 162.2 143717      0  272  124  0.5   1.1   0.4 38.1
20:04:35   48   33 35.0   31.0 0.236   200   193 162.3 143709      0   29   99  0.2   1.2   0.8 52.2
20:05:35   48   35 35.0   30.0 0.013   200   190 162.2 143711      0  118  120  0.3   1.1   0.4 36.4
20:06:35   49   39 35.0   34.0 0.266   200   188 162.2 143721      0  175  164  0.6   2.4   6.0 19.8
20:07:35   49   31 35.0   32.0 0.206   200   188 162.2 143708      0   40  103  0.2   1.1   0.4 36.4
```

Look for Spikes, dramatic changes, what time?

- No "cpu spike", paging maybe interesting
- Processor (Also, ESACPUU, ESACPUA)



VELOCITY SOFTWARE

# *LPAR Overhead Shared example: 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 |

Screen: ESALPMGS Velocity Software ESAMON 5.1

1 of 1 Physical CPU Utilization by CPU Type CPU TOTALS

| Time | CPU Type | <- Count -> Tot | Ded | Shr | <--- Shared Processor Busy ---> %CPU | Total | Assign | Ovhd | Mgmt |
|---|---|---|---|---|---|---|---|---|---|
| 13:01:00 | IFL | 2 | 0 | 2 | 63.2 | 126.3 | 125.5 | 0.4 | 0.9 |
| | CP | 2 | 0 | 2 | 85.6 | 171.2 | 170.9 | 0.2 | 0.3 |

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 Overhead example: 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)
- TOTAL IFL Busy: 167% out of 1100
- Check overheads – high overhead result of too many vcpu
  - 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 25.0 31.4
ZAP      9   2      7  214.8    208.9  3.0  2.9
IFL     31   0     31 1778.5   1669.4 56.9 52.2
ICF      3   0      3  300.2    292.4  0.2  7.3
ZIIP     6   0      6  328.8    311.5  8.4  9.0
```



VELOCITY SOFTWARE

# *Consumers (Virtual Machines) within LPAR: ESAUSP2*

```
Report: ESAUSP2      User Resource Rate Report                          Velocity Software C
Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669X7    First record analyzed:
-------------------------------------------------------------------------------------------
         <---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
-------- ----- ----- --- ---- ----- ---- ---- ----- ---- ----- ----- ----- ----- -----
20:01:35 196.6 192.8 1.0  37M 36.8M 6065  42M 41.9M 872K   23M     0   23M 454.7 443.1
 ***Key User Analysis ***
TCPIP     0.29  0.14 2.1 1639  1639  671  968 968.0  968  3174     0  3174     0     0
 ***User Class Analysis***
Servers   0.01  0.00 2.2  563 212.0    4  629 211.0   21 14476     0 14476     0     0
Velocity  0.75  0.68 1.1 6491  3705    2 6575  3702  370  9040     0  9040     0     0
*Prod     188.3 185.2 1.0  36M 36.5M 4624  42M 41.6M   2M   22M     0   22M  54.4 395.8
*Util     1.83  1.78 1.0  50K 50484  238  54K 53913  54K  263K     0  263K   0.1   1.6
TheUsrs   5.37  5.07 1.1 235K  235K  526 237K  237K  30K  443K     0  443K 400.3  45.6
 ***Top User Analysis***
L24BP     30.45 30.37 1.0 6.5M 6548K  238 7.1M 7078K   7M 3187K     0 3187K   0.0  46.6
L233P     30.04 29.58 1.0 3.7M 3738K  287 4.5M 4486K   4M 2295K     0 2295K   6.4  18.3
L200P     28.25 28.01 1.0 2.2M 2174K  497 2.6M 2610K   3M 1990K     0 1990K  15.8   3.0
L239P     16.68 16.35 1.0 5.7M 5685K  287 6.8M 6822K   7M 1276K     0 1276K   0.2  57.9
L203P     13.92 13.74 1.0 1.4M 1405K  312 1.6M 1573K   2M  889K     0  889K  11.6   9.2
L20BP     12.91 12.83 1.0 3.1M 3121K  239 3.1M 3146K   3M     1     0     1     0     0
```

Look for consumers, in percent of cpu

- By class (Prod)
- Abusive servers (not really)?
- Correct per expected? Not a performance question



VELOCITY SOFTWARE

# *Linux Process Load (CPU by Process): ESALNXP*

```
Report: ESALNXP      LINUX HOST Process Statistics Report          Velocity Software Corporate   ZMAP 4.4.0 08/
Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669X7    First record analyzed: 08/29/18 20:00:35
```

| node/ | <-Process Ident-> | | | Nice | PRTY | <------CPU Percents----> | | | | | <---------------Storage Metrics (MB)----------> | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Name | ID | PPID | GRP | Valu | Valu | Tot | sys | user | syst | usrt | Size | RSS | Peak | Swap | Data | Stk | EXEC | Lib | Lck | PTbl |
| **L233P** | 0 | 0 | 0 | 0 | 0 | **29.1** | 4.25 | 21.0 | 1.08 | 2.75 | 3M | 559K | 3.1M | 0.14 | 7438 | 94.2 | 72K | 7K | 950 | 1874 |
| init | 1 | 1 | 1 | 0 | 20 | 2.99 | 0 | 0.02 | 0.55 | **2.43** | 2 | 1 | 2.4 | 0 | 0.19 | 0.1 | 0.0 | 2.0 | 0 | 0.01 |
| ksoftirq | 3 | 2 | 0 | 0 | 20 | 0.12 | 0.12 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| udevd | 403 | 1 | 403 | -4 | 16 | 0.13 | 0.02 | 0 | 0.06 | 0.05 | 3 | 1 | 3.1 | 0 | 0.57 | 0.1 | 0.1 | 2.0 | 0 | 0.01 |
| timestam | 3769 | 1 | 2724 | 0 | 20 | 0.26 | 0.03 | 0.10 | 0.08 | 0.05 | 4 | 2 | 4.3 | 0 | 1.23 | 0.1 | 0.6 | 2.2 | 0 | 0.01 |
| ohasd.bi | 6293 | 1 | 6293 | 0 | 20 | 0.27 | 0.06 | 0.21 | 0 | 0 | 317 | 72 | 381 | 0 | 218 | 0.3 | 30.8 | 65 | 0 | 0.32 |
| oraroota | 7002 | 1 | 7002 | 0 | 20 | 0.14 | 0.03 | 0.11 | 0 | 0 | 402 | 57 | 466 | 0 | 323 | 0.1 | 10.6 | 65 | 0 | 0.28 |
| oraagent | 7065 | 1 | 7065 | 0 | 20 | 0.13 | 0.02 | 0.11 | 0 | 0 | 268 | 40 | 327 | 0 | 189 | 0.2 | 10.7 | 65 | 0 | 0.22 |
| evmd.bin | 7083 | 1 | 7083 | 0 | 20 | 0.29 | 0.03 | 0.26 | 0 | 0 | 300 | 30 | 364 | 0 | 221 | 0.1 | 1.4 | 65 | 0 | 0.17 |
| gipcd.bi | 7126 | 1 | 7126 | 0 | 20 | 0.45 | 0.18 | 0.27 | 0 | 0 | 239 | 32 | 302 | 0 | 151 | 0.1 | 0.5 | 65 | 0 | 0.16 |
| ocssd.bi | 7202 | 1 | 7202 | 0 | -100 | 0.40 | 0.14 | 0.26 | 0 | 0 | 286 | 164 | 350 | 0 | 207 | 0.1 | 1.7 | 65 | 286 | 0.34 |
| octssd.b | 7638 | 1 | 7638 | 0 | 20 | 0.18 | 0.03 | 0.14 | 0 | 0 | 236 | 27 | 300 | 0 | 159 | 0.1 | 0.3 | 65 | 0 | 0.15 |
| crsd.bin | 7687 | 1 | 7687 | 0 | 20 | 0.45 | 0.06 | 0.39 | 0 | 0 | 350 | 74 | 410 | 0 | 241 | 0.3 | 30.7 | 65 | 0 | 0.36 |
| oraagent | 7743 | 1 | 7743 | 0 | 20 | 0.30 | 0.05 | 0.19 | 0.03 | 0.03 | 359 | 56 | 359 | 0 | 280 | 0.2 | 10.7 | 65 | 0 | 0.27 |
| oraroota | 7752 | 1 | 7752 | 0 | 20 | 0.22 | 0.03 | 0.19 | 0 | 0 | 321 | 31 | 321 | 0 | 242 | 0.1 | 10.6 | 65 | 0 | 0.18 |
| asm_vktm | 8298 | 1 | 8298 | 0 | -2 | 0.13 | 0.06 | 0.06 | 0 | 0 | 1348 | 19 | 1348 | 0 | 3.65 | 0.1 | 237 | 18 | 0 | 0.41 |
| asm_dia0 | 8330 | 1 | 8330 | 0 | 20 | 0.13 | 0.02 | 0.11 | 0 | 0 | 1363 | 42 | 1363 | 0 | 18.1 | 0.2 | 237 | 18 | 0 | 0.50 |
| oraagent | 8997 | 1 | 8997 | 0 | 20 | 0.16 | 0.03 | 0.13 | 0 | 0 | 249 | 35 | 292 | 0 | 170 | 0.2 | 10.7 | 65 | 0 | 0.20 |
| ora_vktm | 9100 | 1 | 9100 | 0 | -2 | 0.14 | 0.06 | 0.08 | 0 | 0 | 13K | 21 | 13K | 0 | 3.70 | 0.1 | 261 | 18 | 0 | 0.44 |

Look for processes within Linux, in percent of cpu

- Total of all processes valid?
- By relevant server (L233P)
- Correct? Relevant? Cron? Init (short term proceses)?



VELOCITY SOFTWARE

## 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      Main Storage Analysis                          Velocity Software Corporate   ESAMAP 4.1.1
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   . Systm  User  NSS/DCSS  <-AddSpace> VDISK <MDC> Diag
Time         On Storage  . ExSpc Resdnt Resident  Systm User  Rsdnt Rsdnt  98
-------- ----- -------  . ----- ------ --------  ----- -----  ----- ----- ----
10:15:00    89  18088K  .  4710 17577K     4771   226K     0 26852 81157 1126
10:30:00    89  18088K  .  4686 17594K     4769   226K     0 30182 61307 1126
10:45:00    89  18088K  .  4681 17614K     4769   225K     0 46189 25812 1126
11:00:00    89  18088K  .  4688 17448K     4775   223K     0  237K  1418 1126
```

Total storage analysis (in pages, vs megabytes)

Best Practices?

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



VELOCITY SOFTWARE

# *Virtual Machine Storage (by Virtual Machine) : 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

| UserID /Class | <---Storage occupancy in pages---> <---Main Storage---> Total | >2gb | <2GB | <--Paging---> Xstor | DASD | <--Main Storage page Read/Write--> <-Page Writes to:--> Xsto | Disk | Migr | <Page Reads:> Xstor | Disk | Pages Moved <2GB | <Address <pages R 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?



VELOCITY SOFTWARE

# VDISK for Swap: ESAVDSK

Report: ESAVDSK VDISK Analysis Report Velocity Software Corporate

| Owner | Space Name | AddSpc Pages | VDSK Blks | Creates | Deletes | AddSpce or Shrd | Priv rate /sec | VIO User Links | Resident | Locked | Stolen | 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 |
| 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!!!



VELOCITY SOFTWARE

# *Linux Storage - 2: ESAUCD2*

```
Report: ESAUCD2      LINUX UCD Memory Analysis Report                Velocity Software Corpor
-------------------------------------------------------------------------------------------
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 Shared
-------- ------ ----- ----- ----- ----- ----- ----- ----- ------ ------ ----- ----- ------
20:01:35
***Node Groups***
*Prod    179728  9206  167K  150K  147K  2314 281.3  156K      0 4769.9  129K 33148      0
*Util     15871 15055 815.9 18828 18828     0  15.6 33883      0  111.9 411.9 292.1      0
TheUsrs  2005.3   6.7  1999  1173  1173     0  15.6  1180      0  282.7  1347 369.2      0
*** Nodes *******
L200P     10064  53.3 10010  1669  1669     0  15.6  1723      0   15.7  6693  3301      0
L203P    6036.0 103.7  5932  1669  1523 145.9  15.6  1627      0   72.1  3944  1916      0
L210P    4022.1  26.9  3995  2843  2525 317.7  15.6  2552      0   96.5  2737  1161      0
L215P    8055.9 488.8  7567  1669  1501 168.1  15.6  1990      0   73.5  5999  1495      0
L23AC    2005.3   6.7  1999  1173  1173     0  15.6  1180      0  282.7  1347 369.2      0
L23BP    6036.0  37.2  5999  2843  2656 186.5  15.6  2693      0    9.8  4825  1164      0
L233P     18119 548.7 17571 35611 35495 116.3  15.6 36043      0  352.0 11915  5304      0
l239p     31753  6101 25653 35611 35611     0  15.6 41711      0  532.8 21944  3176      0
L24BP     27209 289.6 26919  2843  2795  48.1  15.6  3084      0  614.5 23608  2696      0
L24FP    2005.3   8.4  1997  2843  2067 775.4  15.6  2076      0    9.2  1350 637.8      0
```

Linux Storage Map

- Opportunities (Oversized servers consume storage)?
  - High available (greater than 5%)
  - High write buffer (greater than 20mb)
- Issues? Swap
- Swap used vs available….

## Top down:

- z/VM
- Configuration (ESAPSDV)
- Rates
- Space full
- Device busy

# *Paging Subsystem: ESAPSDV*

```
Report: ESAPSDV      Page And Spool Device Activity                  Veloc
Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669X7      First
-------------------------------------------------------------------------
            <----Paging/Spooling--------->       </Sec><Device->
 Dev        <------Slots-------> <per sec>        SSCH Serv Resp %Alloc
 No. Serial Avail Used %Use  Max Read Writ Queue +RSCH Time Time Select
---- ------ ----- ---- ---- ---- ---- ---- ----- ----- ---- ---- ------
20:01:35
Page Devices
2A9A LV1P0A 1803K 547K   30 547K 10.0  8.3     0   4.9  1.3  1.3  100.0
2AD6 LV1P0B 1803K 551K   31 551K  9.0  8.6     0   4.5  1.3  1.3  100.0
2A23 LV1P0C 1803K 549K   30 549K 11.9  8.9     0   5.8  1.0  1.0  100.0
2A5D LV1P0D 1803K 551K   31 551K 11.4 17.2     0   7.6  0.4  0.4  100.0
2A9B LV1P0E 1803K 549K   30 549K 10.8  8.4     0   5.9  1.7  1.7  100.0
...
2A21 LV1P05 1803K 551K   31 551K 12.0  8.7     0   5.9  0.6  0.6  100.0
2A9F LV1P29 1803K 547K   30 547K 15.4 17.4     0   9.7  0.3  0.3  100.0
            ----- ---- ---- ---- ---- ---- ----- ----- ---- ---- ------
Total Page    78M  24M   31  24M  473  458
Spool Devices
            ----- ---- ---- ---- ---- ---- ----- ----- ---- ---- ------
Total Spl   1803K 986K   55 986K    0    0     0 268.9  395  395  100.0
```

Paging Configuration:

- How many devices (11)
- Equal sizes?
- How full? (70% target?)
- Rates reasonable? Device type dependent

# *Page Device Busy: ESADSD2*

```
Report: ESADSD2      DASD Performance Analysis                      Velocity Software Corporate
Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669X7     First record analyzed: 08/2
-------------------------------------------------------------------------------------------------
                                                          <--------DASD Response times (ms)------>
 Dev         Device <--SSCH--> <%DevBusy> <SSCH/sec->          <--Service times--> <--Queueing->
 No. Serial  Type   Total  ERP  Avg  Peak   avg  peak  Resp   Serv Pend Disc Conn DASD Cntl THR
---- ------ ------ ------ ---  ---- ---- ----- ----- -----   ---- ---- ---- ---- ---- ---- ---
20:01:35
***Top DASD by Device busy***
2A8C LV1P22 3390-9    392   0   1.6  1.6   6.6   6.6   2.5    2.5  0.1  2.2  0.2    0    0   0
2A9D LV1P16 3390-9    434   0   1.6  1.6   7.4   7.4   2.1    2.1  0.1  2.0  0.1    0    0   0
2A89 LV1P19 3390-9    387   0   1.5  1.5   6.6   6.6   2.3    2.3  0.1  2.1  0.1    0    0   0
2ADA LV1P2A 3390-9    298   0   1.5  1.5   5.1   5.1   3.0    3.0  0.1  2.7  0.2    0    0   0
2AC4 LV1P1A 3390-9    343   0   1.4  1.4   5.8   5.8   2.4    2.4  0.1  2.2  0.1    0    0   0
```

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 30% “Peak”



VELOCITY SOFTWARE

# *Paging Analysis: ESABLKP*

```
Report: ESABLKP      Block Paging Analysis                          Velocity Software Corporate
Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669X7     First record analyzed: 08/29/1
```

| | <----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 |
| 20:01:35 | 31 | 35.0 | 0.4 | 10.9 | 21.3 | 9.6 | 29.4 | 15.1 | 0 | 0 | 21.3 | **202.2** | 0 | 0 | 202.0 |
| 20:02:35 | 31 | 34.0 | 0.5 | 10.9 | 1.3 | 8.6 | 7.3 | 16.3 | 0 | 0 | 1.3 | 9.9 | 0 | 0 | 49.7 |
| 20:03:35 | 37 | 35.0 | 0.5 | 10.9 | 3.4 | 10.2 | 11.4 | 16.3 | 0 | 0 | 3.4 | 23.3 | 0 | 0 | 81.9 |
| 20:04:35 | 33 | 35.0 | 0.5 | 10.9 | 0.9 | 7.3 | 1.8 | 15.4 | 0 | 0 | 0.9 | 5.1 | 0 | 0 | 0.7 |
| 20:05:35 | 35 | 35.0 | 0.5 | 10.9 | 1.2 | 7.6 | 6.9 | 16.3 | 0 | 0 | 1.2 | 6.3 | 0 | 0 | 12.9 |
| 20:06:35 | 39 | 35.0 | 0.6 | 10.9 | 2.7 | 7.9 | 7.4 | 18.4 | 0 | 0 | 2.8 | 36.4 | 6.8 | 0 | 16.2 |
| 20:07:35 | 31 | 35.0 | 0.5 | 10.9 | 1.5 | 7.7 | 3.6 | 17.4 | 0 | 0 | 1.4 | 2.0 | 0 | 0 | 16.6 |
| 20:08:35 | 37 | 35.0 | 0.5 | 10.9 | 1.0 | 3.3 | 3.3 | 17.4 | 0 | 0 | 1.0 | 1.1 | 0 | 0 | 19.2 |
| 20:09:35 | 35 | 37.0 | 0.6 | 10.9 | 0.8 | 3.8 | 1.6 | 17.9 | 0 | 0 | 0.8 | 0.7 | 0 | 0 | 7.1 |
| 20:10:35 | 40 | 35.0 | 0.6 | 10.9 | 1.5 | 4.0 | 3.8 | 16.6 | 0 | 0 | 1.5 | 2.3 | 0 | 0 | 20.5 |

Block Paging Analysis

- Block page read – optimal 10 pages
- 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
-------------------------------------------------------------------------------
-
          <----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
--------- ---- ---- ----  ----  ---- ----  ---- ----  ---- ----  -----  ------ ------  ----- -----
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 - z/VM Block pages "working sets"

- 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)

## Top down:

- Configuration (ESADSDC)
- DASD I/O for system (ESADSD2/ESADSD6)
- Rates by control unit
- Rates by device
- Rates by minidisk (by user) (ESAUSEK)
- Cache (ESADSD5)

# *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

# *Control Unit Data Rates for LPAR: 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 ESADSD2

- By control unit shows where activity is
- Pend, indication of cache problems
- Compare control units to determine normality

# *Data Rates, Device Performance: ESADSD2*

```
Report: ESADSD2      DASD Performance Analysis                      Velocity Sof
-------------------------------------------------------------------------------
                                                    <--------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
***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 ESADSD2

- 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)



VELOCITY SOFTWARE

# *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



VELOCITY SOFTWARE

## Network Activity

- Configuration
- Rates
- Errors
- Vswitch/guest lan

# *Network Configuration: ESATCPI*

```
Report: ESATCPI      TCPIP Interface Configuration Report          Velocity Sof
Monitor initialized: 04/15/11 at 10:00:00 on 2097 serial 72655     First record
Monitor period:         3600 seconds (   1:00:00)                  Last record:
-------------------------------------------------------------------------------
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 (note virtual mac addresses)

- Ethernet adapter
- Loop back
- MTU check



VELOCITY SOFTWARE

# *Network Configuration: ESATCPI*

```
Report: ESATCPI      TCPIP Interface Configuration Report          Velocity Software Corpo
Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669X7    First record analyzed:
Monitor period:         3600 seconds (   1:00:00)                 Last record:
-------------------------------------------------------------------------------------------
NODE     Idx       Speed <-Status-> Up   <-------------------Interface--------------------
         Nbr MTU   (Est) Oper Admin Time MACAddress        Description Type
-------- --- ----- ----- ---- ----- ---- ----------------- ----------- -------------------

*****************************************Summary*********************************
Average:
L10CU      1 16436   10M   UP    UP    0 00:20:20:20:20:20 lo          Software LoopBack
           2  1492   10G   UP    UP    0 02:78:C1:01:0C:00 eth0        ETHERNET-CSMACD
           3  1492   10G   UP    UP    0 02:78:C1:01:0C:01 eth1        ETHERNET-CSMACD
L20BP      1 16436   10M   UP    UP    0 00:20:20:20:20:20 lo          Software LoopBack
           2  1492   10G   UP    UP    0 02:78:C1:02:0B:00 eth0        ETHERNET-CSMACD
           3  1492   10G   UP    UP    0 02:78:C1:02:0B:01 eth1        ETHERNET-CSMACD
L20DP      1 16436   10M   UP    UP    0 00:20:20:20:20:20 lo          Software LoopBack
           2  1492   10G   UP    UP    0 02:78:C1:02:0D:00 eth0        ETHERNET-CSMACD
           3  1492   10G   UP    UP    0 02:78:C1:02:0D:01 eth1        ETHERNET-CSMACD
L24BP      1 16436   10M   UP    UP    0 00:20:20:20:20:20 lo          Software LoopBack
           2  1492   10G   UP    UP    0 02:78:C1:02:4B:00 eth0        ETHERNET-CSMACD
           3  1492   10G   UP    UP    0 02:78:C1:02:4B:01 eth1        ETHERNET-CSMACD
```

Interface configuration (non-virtual / assigned mac address)

- Ethernet adapter
- Loop back
- MTU check



VELOCITY SOFTWARE

# Network Data Rates: ESATCP4

```
Report: ESATCP4       TCPIP Hardware Layer/Interfaces Report        Velocity Software Corp
Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669X7    First record analyzed:
-------------------------------------------------------------------------------------------
Date/          <Total Octets> Avg  <-Subnet packets / Sec-> <-----Packets Discarded------>
Time           <-Per second->   Q  <-Unicast-> <NonUnicast> <In Error> <NonError> Unknown
Node      IFT  Input  Output  Len  Input Outpt Input Output Inpt Outpt Inpt Outpt Protocol
-------- ---- -----  ------  ---- ----- ----- ----- ------ ---- ----- ---- ----- --------
08/29/18
20:01:35
***Node Groups***
KeyUser  -   0 61803   12635     0 100.0  58.7     0   1.25    0  1.33    0     0        0
*Prod    -   0  728K  829279     0  1267  1221 520.4      0    0     0 9.19     0        0
*Util    -   0  3884  1851.0     0 11.70   2.7 36.43      0    0     0    0     0        0
TheUsrs  -   0  3537  1911.6     0  2.91   3.6 37.01      0    0     0    0     0        0
*** Nodes *******
L24BP    -   1 96.61   96.61     0  0.63   0.6     0      0    0     0    0     0        0
         -   2  5210   14818     0 24.11  29.3     0      0    0     0    0     0        0
         -   3  2874       0     0     0     0 34.08      0    0     0    0     0        0
```

Snmp data source:

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

# 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

# OSA Adapter: ESAOSA

```
Report: ESAOSA        OSA System Configuration Report
Monitor initialized: 06/15/16 at 00:00:00 on 2828 seria
-----------------------------------------------------
Collector <---------OSA Configuration--> MacAddress
Node      Idx  Name  Nbr  Type Level Shrd Active
--------  ---  ----  ---  ---- ----- ---- ------------
06/15/16
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   Z
Monitor initializedal 314X7       First record analyzed: 06/15/16
-------------------------------------------------------------------
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
```



VELOCITY SOFTWARE

## "z" is complex:

- z/VM is very well instrumented
- zVPS enhances instrumentation for networks, Linux

## Analysis flowchart simple:

- Wait states
- Configuration
- Performance details

## Next step – Subsystem Analysis and tuning:

- CPU
- Storage
- DASD