

# Basic Analysis

## Getting Started with z/VM System Performance

- [Barton@VelocitySoftware.com](mailto:Barton@VelocitySoftware.com)
- [HTTP://VelocitySoftware.com](http://VelocitySoftware.com)

“If you can’t Measure it,  
I am Just Not Interested™”

# Agenda to Basics

- **Know where to start**
- **Know what is important**
- **Know what is the configuration**

**(Not necessarily in that order)**

# The Analysis Problem

## “z” is:

- Very large,
- Very complex and
- Very well-instrumented
- Very suitable for “KI” based analysis (zTUNE...)

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

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

## How long does it take to learn all of this?

- You probably don't have the time

# The Challenge - z/VM serves many functions (190+ reports)

## ESAHDR ESATUNE

\*Performance Summary  
ESASSUM ESASUM

\*Transaction Activity (5)  
ESAUOLA **ESAXACT** ESARATE  
ESACLAS ESAEXCP

\*User Activity (21)  
ESATUNA  
ESASRVC ESASRV1 **ESAUSRC** ESAUSR1  
ESAUSR2 ESAUSR3 ESAUSR4 ESAUSR5  
**ESAUSP2** **ESAUSP3** ESAUSP4 ESAUSCP  
**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**

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

\*LINUX Reporting (20)  
ESAUCD1 **ESAUCD2** ESAUCD3 ESAUCD4 ESAUCDD ESALNXD  
ESAHST1 ESAHST2 ESAHST3 ESAHST4 ESAHSTA  
ESALNXS ESALNXR **ESALNXP** ESALNXA ESALNXC  
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 based on wait states

- DASD, Storage, Paging, Processor, Network

## Tuning Guide – Performance Methodology Tips

# 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 Facility(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 (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 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 system, class, top user

- System: What impacts the whole system?
- User classes: Does one class stand out?
- Users: Is there something specific?
- Recognize “running” to wait comparison

## Tuning Guide – Wait State Analysis

# Wait States: ESAXACT

## Determine what servers waiting for what?

- If "run" is dominant, more cycles are required

```

Report: ESAXACT Transaction Delay Analysis Veloc
Monitor initialized: 04/15/11 at 10:00:00 on 2097 serial 72655 First
-----
<-----Percent non-dormant (Wait states)----->
UserID <-Samples->
/Class Total In Q Run Sim CPU SIO Pag E- D- T- Tst <Asynch> Pct
----- Total In Q Run Sim CPU SIO Pag SVM SVM SVM CF Idl I/O Pag Ldg 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
***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***
LN XUWA01 893 893 71 0 2.8 0 0.1 0 0 0 0 24 1.7 0.4 0 0
LN XUWA03 1786 1786 28 0.2 5.5 0 1.2 0 0 0 0.6 57 7.2 0.1 0.1 0
LN XUWA02 1786 1786 27 0.1 3.6 0 0.1 0 0 0 0.4 69 0.1 0 0.1 0
LN XQWA01 1786 1786 4.0 0 2.2 0 0 0 0 0 0 94 0.1 0 0 0
LN XDWA02 1786 1786 6.0 0 2.2 0 0.2 0 0 0 0 91 0.1 0 0 0
LN XDWA04 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

```

# Wait States: ESAXACT

## Compare wait against "run"

- CPU problem, master processor too?

```

Report: ESAXACT      Transaction Delay Analysis
-----
                                <-----Percent non-dormant (Wait states)-
UserID  <-Samples->
/Class  Total  In Q  Run  Sim  CPU  SIO  Pag  E-  D-  T-  Tst
-----  -----  ---  ---  ---  ---  ---  ---  ---  ---  ---  ---
14:06:00  95   189   19  3.7   60   0   0   0   0   0   0.5  16
Hi-Freq: 14220 11367  19  3.6   30   0.0  0   0   1.3  0.1  0.4  47
ZSERVE   60    1    0    0    0    0   0   0   0   0   0   100
ZWRITE   60   49   2.0  0    0   2.0  0   0   0   0   0   96
ZWEB03   60    1    0    0    0    0   0   0   0   0   0   0
ZALERT   60    2   50    0   50   0   0   0   0   0   0   0
ZTCP     60    8    0   50   25   0   0   0   0   0   0   25
TCPIP    60   44    0  9.1  4.5   0   0   0   0   0   0   86
SMTP     60    1    0    0    0    0   0   0   0   0   0   0
***User Class Analysis***
zVMgmt   2580  142  3.5  7.0  11  0.7  0   0  7.0  7.0  2.8  68
CLS1     3120 3120  36  4.9  47  0.0  0   0   0   0  1.0  12
TPMS     480   480  79  4.6  14  0   0   0   0   0  0.6  2.1
Breakerr 480   480  18  3.8  40  0   0   0   0   0   0   39
MyMQ     780   779  6.8  4.4  41  0   0   0   0   0  0.6  47
Others   4920 4511  3.0  2.4  17  0   0   0   0   0  0.1  78
***Top User Analysis***
LNx081   480   480  48  5.8  40  0   0   0   0   0  3.8  2.5
LNx082   480   480  45  5.4  44  0   0   0   0   0   0  5.8
  
```

# 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

UserID	ClassID	Account Code	ACI Grp Name	CPU Type	<-----SHARE----->					CPU <Modes>	<Status>		<-MDC>		<-Storage->		
					<Normal> Rel	Abs	<--MAX-> Typ	Shre	Lim		-it	Cnt	NO	NO	<-VM Size-> 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

- **Understand where to start**
- **Validate the configuration**
- **Wait state (queue) analysis**  
-> **Where to focus**