



VELOCITY
SOFTWARE

Performance Analysis Flowchart

Velocity Software Inc.
196-D Castro Street
Mountain View CA 94041
650-964-8867

Velocity Software GmbH
Max-Joseph-Str. 5
D-68167 Mannheim
Germany
+49 (0)621 373844

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

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

“z” is:

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

The challenge?

What challenge, it is “all” there!

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

Few companies have full time performance analysts.

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

*Performance Summary (4)

ESAHDR ESATUNE ESASSUM ESASUM

*Transaction Activity (5)

ESAUCLA **ESAXACT** ESARATE ESACLAS
ESAEXCP

*User Activity (21)

ESASRVC ESASRV1 **ES AUSRC** 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** **ESADEV**

*Virtual NETWORK Reporting (7)

ESAQDIO ESAQDI2 **ESANIC**
ESA VSWC ESA VSW ESA VSW2
ESAOSA

*Operational Logging

ESAOPER

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

***Shared File System (7)**

ESASFS1 ESASFS2 ESASFS3 ESASFS4
ESASFS5 ESASFS6 ESASFS7

***Byte File System (2)**

ESABFS1 ESABFS2 ESABFS3

***CMS Multitasking (1)**

ESAMTSK

***Web Serving Reports (8)**

ESAWEB1 ESAWEB2 ESAWEB3 ESAWEB4
ESAVWS1 ESAVWS2 ESAVWS3 ESAVWS4

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

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

Network: Source: snmp - VERY efficient, really!

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

ESATCPC **ESATCP1** **ESATCP2** ESATCP3 **ESATCP4**

- Microsoft servers: Host mib

ESAHST1 ESAHST2 ESAHST3 ESAHST4 ESAHSTA

- Linux servers: UCD mib

ESAUCD1 **ESAUCD2** ESAUCD3 **ESAUCD4** ESAUCDD

- More Linux: Velocity mib

ESALNXD ESALNXS ESALNXR **ESALNXP** ESALNXA

ESALNXC **ESALNXF** ESALNXU ESALNXV

ESALNXM ESALNXUP

- Linux Application "vendor" mibs

ESAJVM ESAORAC ESAORAG ESAORAS ESAORAW

ESAGPFN ESAGPFF ESAGPFF ESAGPFD ESAGPFS (Version 5, zvps)

ESAMNG1 ESAMNG2 ESAMNG3 ESAMNG4 ESAMNG5 (Version 5, zvps)

ESADOCK1 ESADOCK2 ESASSC (Version 5, zVPS)

- VSE mibs: IBM, Velocity

ESAVSEC ESAVSES ESAVSEP ESAVSEJ

ESAVSEP ESAVSEJ

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

Configuration: ESAHDR

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

Wait states: ESAXACT

Virtual machine Config: ESAUSRC

CPU:

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

Storage

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

- Configuration: ESAPSDV
- Loads by user: ESAUSPG

DASD

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

Network

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

Analysis starts with “is there a problem?”

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

System Configuration

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

Loads on the system subsystems

Wait states for those impacted

Subsystem Analysis

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

Know the configuration: ESAHDR (B)

Report: ESAHDR z/VM Monitor Analysis
Monitor initialized: 04/25/22 at 00:00:00 on 8561 s
Monitor period: 7200 seconds (2:00:00)

ZMAP Release 5.1.3.1
History Source 5.1.3.0
Key Expiration for: Velocity 04/28/23
Components licensed: TUN
Monitor file created: 04/25/22 00:00:00

z/VM Version: 7 Release 2.0 SLU 2102
TOD clock at termination 20:00:00
Abend code of last termination
TOD clock at last IPL: 03/04/22 19:40:57
System Operator: OPERATOR
Time zone adjustment from GMT: -4 hours

System Identifier VLXPMW08
Checkpoint/Warmstart Volumes V51R01/V51R01
Machine Model/Type Z15:8561/401
Multithreading Status (SMT): Enabled
Core Thread count: 2
Enabled Count: 2

Common configuration problems

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

Know the configuration: ESAHDR (B)

Report: ESAHDR z/VM Monitor Analysis

```
-----  
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 00000000000469C7

```
Processor 0 model/serial 8561-401 /111108 Master  
Processor 1 model/serial 8561-401 /111108  
Processor 2 model/serial 8561-401 /111108  
Processor 3 model/serial 8561-401 /111108  
Processor 4 model/serial 8561-401 /111108  
...  
Processor 38 model/serial 8561-401 /111108  
Processor 39 model/serial 8561-401 /111108
```

Common configuration problems

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

(z/VM perspective, SMT Enabled)

```
CPU(GP) Capability Factor: 3085  
CPU(IFL) Capability Factor: 416  
CPU Cycles/ns: 5200  
CPU Cycles/ns (GP): 781
```

Operating on IFL Processor(s)

Channel Path Measurement Facility(CPMF) Extended is installed

Totals by Processor type:

```
<-----CPU-----> <-Shared Processor busy>  
Type Count Ded shared total assigned Ovhd Mgmt  
-----  
CP 1 0 1 0.4 0.3 0.0 0.1  
IFL 108 0 108 1501 1478.6 21.1 22.2
```

Know the configuration: ESAHDR (B)

Report: ESAHDR z/VM Monitor Analysis
Monitor initialized: 04/25/22 at 00:00:00 on 8561 serial
Monitor period: 3600 seconds (1:00:00)

Number of logical partitions defined: 21

Main Storage installed (MB): 2867199
Main Storage Generated (MB): 2867199
Horizontal/Vertical Scheduling Configuration IFL CPUs

CPU utilization: 616.6 of 4000%
CPU charged to users: 598.3%
System overhead: 18.4%
Capture ratio: 100.00%
DASD I/O rate: 473.3 per second
MDC effectiveness: 0.0%
DASD paging rate: 0.0 per second
Spool utilization: 2%
Communication rate: 392.1 per second

Top users and user classes by CPU consumption:

UserID /Class	<-Relative->		<---Absolute Percent CPU--->			
	<-Pct CPU--> Used	Cum	<Out of 100%> Util	Cum	<Out of 4000%> Util	Cum
1. LINUX3F8	71.6	72	11.0	11	441.4	441
2. LINUX3FD	13.3	85	2.1	13	82.2	524
3. System	3.0	88	0.5	14	18.4	542
4. LINUX3FA	2.2	90	0.3	14	13.7	556
5. LINUX411	1.7	92	0.3	14	10.2	566
6. LINUX3B2	0.8	93	0.1	14	4.8	571
7. LINUX529	0.8	93	0.1	14	4.7	575
8. LINUX400	0.7	94	0.1	15	4.6	580

Common configuration problems

- Real Storage
- High level utilizations
- Top User(s)

Know the overall loads: ESASSUM / ESAMAIN (B)

Report: ESASSUM Subsystem Activity SOMEgr
 Monitor initialized: 04/25/22 at 13:00:00 on 8561 serial 111108 First

<---Users----> <Processor> Storage (MB) <-Paging--> <-----I/O----->

Time	<-avg number->			Utilization		<-Paging-->		<pages/sec>		<-DASD-->		Other
	On	Actv	In Q	Total	Virt.	User	Resid.	XStore	DASD	Rate	Resp	Rate
13:15:00	78	45	191	662	613	74.9	2765K	0	1	165	0.2	0
13:30:00	78	45	190	625	580	74.9	2774K	0	1	168	0.2	0
13:45:00	78	45	196	1303	1248	74.9	2778K	0	0	163	0.2	0
14:00:00	78	46	178	848	811	74.9	2783K	0	0	168	0.2	0
14:15:00	78	45	190	312	287	74.9	2791K	0	0	164	0.2	0
14:30:00	78	45	195	337	308	74.9	2805K	0	0	165	0.2	0
14:45:00	78	45	188	943	891	74.9	2811K	0	2269	187	1.2	0
15:00:00	78	45	189	596	556	74.9	2812K	0	17K	470	2.9	0
15:15:00	78	46	179	374	341	74.9	2812K	0	20K	661	2.4	0
15:30:00	78	45	194	413	381	74.9	2813K	0	8889	489	1.3	0

Look for Spikes, dramatic changes, what time? (14:30-14:45)

- Processor
- Storage for users
- Page rates
- DASD I/O rates

Wait states provide options for improvement

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

Wait state (queue) analysis -> where to focus

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

Normal idle wait states

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

Two types of Wait states are provided by virtual machine

- Monitor frequency: once per minute (**Not interesting**)
- High Frequency (Hi-Freq): once per second (**Interesting**)
 - (60 samples per 1 minute per virtual cpu)

Shown by:

- Summarized for all users (**start here**)
- Summarized for user classes (grouped by installation)
- Servers
- Top users

User class analysis -> where to focus

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

Wait States: ESAXACT (B)

Report: ESAXACT Transaction Delay Analysis SOMEg
Monitor initialized: 04/25/22 at 14:00:00 on 8561 serial 111108 First

```
-----Percent non-dormant (Wait states)-----
UserID  <-Samples->
/Class  Total  In Q  Run Sim CPU SIO Pag SVM SVM SVM  CF  Tst <Asynch>
-----
14:45:00 1170 2815 7.1 0.1 2.1  0  0  0  0  0  0  90  0  .  .
Hi-Freq: 302K 174K 4.6 0.1 1.2 0.0 0.0  0 1.0 0.0  0  94  0 0.0 0.0
***User Class Analysis***
Servers 10800  459  0 0.2 0.4  0  0  0  10 8.9  0  90  0  0  0
Velocity 9000  115 1.7 0.9  0  0  0  0  0.2 8.7  0  89  0  0  0
CATech  4500   1  0  0  0  0  0  0  0  0  0  0  0  0  0
TheUsrs 278K 174K 4.6 0.1 1.2 0.0 0.0  0  0.7 0.0  0  94  0 0.0 0.0
***Top User Analysis***
LINUX3F8 28800 25036 24 0.4 3.2  0  0  0  0  0  73  0  0  0
LINUX3FD 28800 27916 4.3 0.2 2.5  0 0.0  0  0  0  0  93  0  0  0
LINUX538 28800 15500 1.9 0.0 0.3  0  0  0  0  0  0  98  0  0  0
LINUX3FA 28800 15945 1.1 0.1 0.4  0  0  0  0  0  0  98  0  0  0
LINUX411  3600  3600 2.9 0.1 3.2  0  0  0  0  0  0  94  0 0.0 0.1
-----
15:00:00 1170 2840 4.8 0.1 1.3  0 0.2  0  0 0.0  0  93  0  .  .
Hi-Freq: 302K 173K 2.9 0.1 0.6  0 0.1  0  1.0 0.0  0  96 0.0 0.0 0.7
***Top User Analysis***
LINUX3F8 28800 21011 11 0.1 1.0  0 0.1  0  0  0  0  88  0  0 0.1
LINUX3FD 28800 27855 4.4 0.3 1.6  0 0.2  0  0  0  0  93  0  0 0.1
LINUX538 28800 17749 4.8 0.0 0.4  0 0.1  0  0  0  0  95  0  0 0.0
LINUX3FA 28800 15107 1.1 0.0 0.4  0  0  0  0  0  0  88  0  0 0.0
```

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

- Processor wait trivial?
- Page wait non existent?

Wait States: ESAXACT (A)

Report: ESAXACT Transaction Delay Analysis Velocity S
 Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669C7 First reco

```

-----
<-----Percent non-dormant (Wait states)-----
UserID  <-Samples->
/Class  Total  In Q  Run Sim  CPU  SIO  Pag  SVM  SVM  SVM  CF  Tst <Asynch>
-----  -----  -----  -----  -----  -----  -----  -----  -----  -----  -----  -----  -----  -----  -----  -----  -----
08/29/18
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
  
```

Sometimes it is easy: The primary wait state is CPU, 77% CPU WAIT

User Configuration: ESAUSRC (B)

Report: ESAUSRC SOMEgroup - MAP 5.1.3 0
 Monitor initializ 8561 serial 111108 First record 00:00:00
 Monitor period: 0:00) Last record: 01:00:00

UserID	ClassID	Type	<-----SHARE----->				<---CPU		<Storage>		
			> CPU	<Normal>	<--MAX-->	Lim	<Count>	<VM Size>			
			Rel	Abs	Typ	Shre	-it	Def	On	Dflt	Max
LINUX3B2	TheUsrs	IFL	1600	16	16	244G	244G
LINUX3FA	TheUsrs	IFL	3200	32	32	488G	488G
LINUX3FD	TheUsrs	IFL	3200	32	32	488G	488G
LINUX3F5	TheUsrs	IFL	3200	32	32	488G	488G
LINUX3F8	TheUsrs	IFL	3200	32	32	488G	488G
LINUX40A	TheUsrs	IFL	400	4	4	61G	61G
LINUX40D	TheUsrs	IFL	400	4	4	61G	61G
LINUX400	TheUsrs	IFL	400	4	4	61G	61G
LINUX404	TheUsrs	IFL	200	2	2	31G	31G
LINUX408	TheUsrs	IFL	400	4	4	61G	61G
LINUX41B	TheUsrs	IFL	800	8	8	122G	122G
LINUX411	TheUsrs	IFL	400	4	4	61G	61G
LINUX51E	TheUsrs	IFL	3200	32	32	488G	488G
LINUX529	TheUsrs	IFL	3200	32	32	488G	488G
LINUX538	TheUsrs	IFL	3200	32	32	488G	488G

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?)

Very Bad Performance reported by user (B)

Users want virtualized Linux “just like x”

Virtual machines are not waiting – or are they?

- 32 virtual cpus are samples
- 32 * 60 samples per minute * 15 minutes = 28800 samples
- Why are 88% samples “in queue”, but not doing anything?
- If 11% are using CPU, how many CPU’s needed by workload?

Performance problem is “internal” to Linux

```
Report: ESAXACT          Transaction Delay Analysis          C
Monitor initialized: 04/25/22 at 14:00:00 on 8561 serial 111108  F
-----
                <-----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
-----
15:00:00    1170   2840  4.8 0.1 1.3   0 0.2   0   0 0.0   0  93   0   .
Hi-Freq:    302K   173K  2.9 0.1 0.6   0 0.1   0  1.0 0.0   0  96  0.0 0.0
***Top User Analysis***
LINUX3F8   28800 21011  11 0.1 1.0   0 0.1   0   0  0   0  88   0   0
LINUX3FD  28800 27855  4.4 0.3 1.6   0 0.2   0   0  0   0  93   0   0
```

Linux Configuration Guideline Summary

Virtual machine size

- Minimize until some swap

Swapping

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

Virtual processors

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

Infrastructure costs

- Minimize – shared resource architecture

Why Linux Configuration Guidelines?

Reduce Virtual machine size

- Large servers own storage but do not use it.
- 7 servers * 500GB barely fit in 2.7 TB LPAR
- Are they using the storage? (ESAUCD2)
- Use Option to dynamically add to virtual storage instead?

Swapping

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

Virtual processors

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

Case Study

- Validate storage requirements
- Validate diagnose rates

Linux Storage: ESAUCD2 (B)

Report: ESAUCD2 LINUX UCD Memory Analysis Repo
Monitor initialized: 04/25/22 at 00:00:00 on 8561 s

```
-----  
Node/      <-----Storage Sizes (in MegaBytes)-----  
Time/      <--Real Storage--> <-----SWAP Storage----->  
Date       Total  Avail Used   Total Avail Used   MIN  
-----
```

```
04/25/22  
00:15:00  
***Node Groups***  
TheUsrs    4044K 2477K 1566K 480K 480K 7.8 234.4  
*** Nodes *****  
LINUX3B2  245373 172K 69731 32768 32765 2.8 15.6  
LINUX3FA  490922 137K 342K 32768 32768 0 15.6  
LINUX3FD  490922 9564 470K 32768 32763 4.8 15.6  
LINUX3F5  490922 459K 20737 32768 32768 0 15.6  
LINUX3F8  490922 22709 457K 32768 32768 0 15.6  
LINUX40A   61210 31079 30131 32768 32768 0 15.6  
LINUX40D   61210 52677 8533 32768 32768 0 15.6  
LINUX400   61210 46124 15086 32768 32768 0 15.6  
LINUX404   30516 22208 8309 32768 32768 0 15.6  
LINUX408   61210 53121 8089 32768 32768 0 15.6  
LINUX41B  122598 99777 22820 32768 32768 0 15.6  
LINUX411   61210 21646 39564 32768 32768 0 15.6  
LINUX51E  490861 465K 15172 32768 32768 0 15.6  
LINUX529  490861 429K 51239 32768 32768 0 15.6  
LINUX538  490861 465K 14472 32768 32768 0 15.6
```

Available storage

- Not used for cache
- Backed by z/VM
- Because that is what they do on x

Linux Storage from z/VM Virtualized (FreeFall): ESAUSPG (B)

Report: ESAUSPG User Storage Analysis
Monitor initialized: 04/25/22 at 14:00:00 on 8561 seri

<Storage Occupancy (MegaBytes)>				<-Page I/O-->		
UserID	<---Main Storage--->			Paged	Page	Page
/Class	Total	>2gb	<2GB	Out	Writes	Reads
14:45:00	2811K	2809K	1698	1241	1548K	102635
User Class Analysis						
Servers	65	40	25	59	17587	347
Velocity	36	36	0	25	10402	665
CAtech	15	8	7	10	2403	122
TheUsrs	2811K	2809K	1668	1147	1517K	101501
Top User Analysis						
LINUX3F8	497540	497540	0	69	94026	955
LINUX3FD	497558	497558	0	52	81951	22318
LINUX538	427501	426021	1480	73	97357	183
LINUX3FA	361117	360972	144	73	97493	33
LINUX411	45291	45291	0	50	89325	15935
LINUX51E	24505	24505	0	73	97428	61
LINUX3F5	497536	497536	0	70	96818	5829
LINUX3B2	249007	249007	0	66	99084	15752
LINUX400	24715	24715	0	56	83907	2888
LINUX529	60820	60820	0	73	97236	10
LINUX408	11406	11406	0	54	81802	2281
LINUX40A	62016	62016	0	59	89685	15470
LINUX41B	25663	25663	0	54	81802	2281
LINUX40D	10183	10183	0	54	81802	2281
LINUX404	9460	9460	0	54	81802	2281

User pages

- Are NOT just virtual
- All are Backed by z/VM
- Linux touches all pages over time
- 5 servers each have more than 10%

Linux Storage from z/VM Virtualized: ESAUSPG (B)

Report: ESAUSPG User Storage Analysis
 Monitor initialized: 04/25/22 at 14:00:00 on 8561 serial 1111

UserID /Class	<Storage Occupancy (MegaBytes)>			<-Page I/O-->		Pages	
	<---Main Storage---> Total	>2gb	<2GB	Paged Out	Page Writes	Page Reads	Moved <2GB
14:30:00	2806K	2805K	234	0	0	0	0
User Class Analysis							
Servers	124	75	48	0	0	0	0
Velocity	61	61	0	0	0	0	0
CAtech	25	19	7	0	0	0	0
TheUsrs	2805K	2805K	181	0	0	0	0
Top User Analysis							
LINUX3FD	497610	497610	0	0	0	0	0
LINUX3F8	497609	497609	0	0	0	0	0
LINUX538	421606	421505	101	0	0	0	0
LINUX3FA	360765	360759	6	0	0	0	0
LINUX411	45343	45343	0	0	0	0	0
LINUX3B2	249073	249073	0	0	0	0	0
LINUX529	60893	60893	0	0	0	0	0
LINUX400	24775	24775	0	0	0	0	0

User pages – Previous interval

- No Paging
- Passed edge of cliff....
- SET AGELIST settings?

cp query agelist

Target size = 37888K (37M) 2.0% **(5%)** of
 pageable storage

In use = 20480K

Pending writes = 18452K

Early writes = Yes

Sizing = Variable

Keep slot = Yes

Ready; T=0.01/0.01 17:37:49

Paging: setting best practices? (B)

cp query agelist **(DEFAULT)**

Target size = 37888K (37M) 2.0% of pageable storage **(5% option)**

In use = 20480K

Pending writes = 18452K

Early writes = Yes **(Write out unused pages early)**

Sizing = Variable

Keep slot = Yes

q agelist **(installation "B")**

Target size = 74562560K (72815M) **2.0%** of pageable storage **(Needed10%)**

In use = 8616968K

Pending writes = 8616928K

Early writes = No **(So everything waits until we write out pages)**

Sizing = Variable

Keep slot = No **(saves disk, but degrades performance of writing)**

Storage Process ramp up Analysis: ESALNXP (B)

Report: **ESALNXP** LINUX HOST **Process Statistics** Report Cit
Monitor initialized: 04/25/22 at 10:00:00 on 8561 serial 011108 Fir

```
-----  
node/      <Process Ident> Nice PRTY <-----CPU Percents-----> <STORAGE (MB)  
Name       ID      PPID Valu Valu Tot  sys user syst usrt  Size RSS  
-----  
LINUX538   0       0     0     0  2.82 0.57 1.56 0.08 0.61  13K  487  
LINUX538   0       0     0     0 12.4 0.67 7.85 0.40 3.44  14K  533  
LINUX538   0       0     0     0 61.8 5.48 52.7 0.16 3.45  46K  31K  
LINUX538   0       0     0     0 21.8 2.24 19.5 0.01 0.02  94K  80K  
LINUX538   0       0     0     0 53.4 5.06 47.6 0.08 0.64 126K 112K  
LINUX538   0       0     0     0 28.0 2.81 24.9 0.04 0.28  96K  81K  
LINUX538   0       0     0     0 17.2 1.63 15.5 0.01 0.02 149K 134K  
LINUX538   0       0     0     0 54.7 5.39 49.3 0.01 0.00 173K 159K  
LINUX538   0       0     0     0 62.1 4.18 57.2 0.07 0.58 204K 189K  
LINUX538   0       0     0     0 69.5 4.50 65.0 0.00 0.00 207K 193K  
LINUX538   0       0     0     0 50.9 3.93 46.8 0.03 0.15 183K 169K  
LINUX538   0       0     0     0 27.2 2.10 24.4 0.08 0.63 208K 193K  
14:45:00  
LINUX538   0       0     0     0 23.6 1.76 21.2 0.08 0.55 208K 193K  
  chef-cli  2808    1     0    20 0.61  0     0 0.07 0.55  200  51  
  BESClie  2887    1     0    20 1.65 0.21 1.44  0     0  669  34  
  mongod-g 45390  1  0  20 20.6 1.16 19.4  0  0 195K 193K  
LINUX538  0  0  0  0  102 4.96  97 0.01 0.00 208K 194K  
  BESClie  2887    1     0    20 1.66 0.19 1.47  0  0  669  34  
  mongod-g 45390  1  0  20  99 4.45 94.8  0  0 195K 193K ←--
```

One large serve ramped up (one application)

- No storage available, system responded

z/VM 6.3 Invalid but Resident Storage CLIFF Analysis (B)

Report: ESAUSTR User Storage Analysis
 Monitor period: 3600 seconds (1:00:00) Last record:

UserID /Class	Size	Alloc	Resident	UFO Activ	Storage (MB)				Resident		Mbyte	slots	Made	
					IBR		AgeList		Unreferd		used	IBR		
					TOT	<2gb	>2gb	<2gb	>2gb	<2gb	>2gb			
TheUsrs	4231K	2529K	2529K	2529K	0	0	0	0	0	0	0	0	0	
TheUsrs	4231K	2613K	2613K	2613K	0	0	0	0	0	0	0	0	0	
TheUsrs	4231K	2654K	2654K	2654K	0	0	0	0	0	0	0	0	0	
TheUsrs	4231K	2712K	2712K	2712K	0	0	0	0	0	0	0	0	0	
TheUsrs	4231K	2745K	2745K	2745K	3.5	0.5	3.0	21.4	413	0	0	0	33.3	←2% line
TheUsrs	4231K	2681K	2681K	2681K	0.9	0.1	0.8	5.4	103	0	0	0	31.3	
TheUsrs	4231K	2765K	2765K	2763K	10.9	1.5	9.4	64.3	1533	0	0	0	0.8	
TheUsrs	4231K	2774K	2774K	2772K	10.9	1.5	9.4	64.3	1532	0	0	0	1.8	
TheUsrs	4231K	2778K	2778K	2776K	10.6	1.5	9.1	64.3	1540	0	0	0	0	
TheUsrs	4231K	2783K	2783K	2781K	10.6	1.5	9.1	64.3	1536	0	0	0	0.0	
TheUsrs	4231K	2775K	2775K	2773K	10.8	1.5	9.3	64.3	1535	0	0	0	2.4	
TheUsrs	4231K	2791K	2791K	2789K	10.6	1.5	9.1	64.3	1532	0	0	0	0	
TheUsrs	4231K	2805K	2805K	2804K	10.6	1.5	9.1	64.3	1530	0	0	0	0	
TheUsrs	4231K	2812K	2811K	2808K	12.9	1.5	11.4	31.5	2023	0.0	37.0	1147	137	←Fall over cliff
TheUsrs	4231K	2833K	2812K	2810K	36.9	1.5	35.4	0.1	2319	0.0	737	20588	632	
TheUsrs	4231K	2810K	2805K	2803K	17.8	1.5	16.3	40.0	1851	0.0	193	5434	720	
TheUsrs	4231													
TheUsrs	4231													
TheUsrs	4231													
TheUsrs	4231													
TheUsrs	4231													
TheUsrs	4231													
TheUsrs	4231K	2904K	2813K	2808K	1788	1.5	1786	0.0	2333	0	3695	91340	35.2	

Invalid but Resident (IBR as of z/VM 6.3)

- Are correct servers losing pages? (NO)
- Same server, LINUX538 is “root cause”

Virtual machine Storage sizes very high

- Application consumes a lot and builds fast

z/VM configuration parameters

- suited for SMALL
- Change AGELIST

Virtual machine VCPU count higher than available?

- Determine real requirements?

z/VM With Hiperdispatch parks engines

- 20 Physical LPAR VCPU
- SMT gives 40 threads
- Parking drops

Virtual processors

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

Infrastructure costs

- Minimize – shared resource architecture

Linux CPU Details: ESALNXS (B)

Report: ESALNXS LINUX VSI System Analysis Report
Monitor initialized: 04/25/22 at 14:00:00 on 8561 seri

```
-----  
Node/      <---Load Numbers--> CPU <Processor Pct Util>  
Time       Users Procs MaxProc NBR Total  Syst  User  Idle  
-----  
LINUX3F8   0    439      0 Tot 651.4 34.7  606 2490  
           1    30.2   3.3 21.2 64.1  
           2    24.6   1.6 22.7 73.3  
  
           23   17.2   0.8 16.2 81.4  
           24   13.7   0.7 12.9 85.1  
           25   16.5   0.8 15.5 82.1  
           26   21.0   0.7 20.1 77.6  
           27   16.0   0.7 15.2 82.7  
           28   14.1   0.6 13.4 84.8  
           29   16.6   0.9 15.5 81.9  
           30   11.6   0.5 11.1 87.5  
           31   12.0   0.4 11.5 87.1
```

Linux balances across all CPUs

- Spin locks are expensive
- More use of locks with more CPUs
- What if vcpu holds lock, but paged out?
- Is there concurrency with 32 VCPU?

Virtual Machine Diagnose Counts: ESAUSRD (B)

Report: ESAUSRD Virtual Machine Diagnose Analysis(720)
 Monitor initialized: 04/25/22 at 12:00:00 on 8561 seri: 04/

UserID /ClassID	Total rate	<-----diag counts / second								
		000	004	008	00C	010	014	024	044	09C
LINUX3FD	1274	.	.	0	98.2	1176
LINUX3F8	304	.	.	0	31.3	312
LINUX3F8	1197	.	.	0	36.7	1161
LINUX3F8	744	.	.	0.0	41.0	703
LINUX3F8	2052	.	.	0	164	1888
LINUX3F8	1696	.	.	0	112	1584
LINUX3F8	2669	.	.	0	397	2273
LINUX3F8	953	.	.	0	224	729
LINUX3F8	1843	.	.	0	224	1619
LINUX3F8	203	.	.	0	12.9	190
LINUX3F8	320	.	.	0	20.6	299
14:45:00	3830	0.5	0.0	0.5	0.6	0.0	0	0.1	394	2978
LINUX3F8	1557	.	.	0	209	1348
LINUX3F8	453	.	.	0	58.3	395
LINUX3F8	633	.	.	0	75.2	558
LINUX3F8	242	.	.	0	15.6	227
LINUX3F8	269	.	.	0	18.6	250
LINUX3F8	2360	.	.	0	19.8	2340
LINUX3F8	266	.	.	0	15.6	250
LINUX3F8	784
LINUX3F8	784
LINUX3F8	287

Spin locks normal when many vcpu

- Spin locks rate did not change
- What changed? Paging started
- Spin lock uses diag 44, 9c to give up control.
- If owning process paged out, long delay

z/VM CPU Details: ESACPUA (B)

Report: ESACPUU CPU Utilization Report Cit
 Monitor initialized: 04/25/22 at 14:00:00 on 8561 serial 011108 Fir

Time	<----Load---->			CPU CPU	CPU Type	<-----CPU (percentages)----->					
	<-Users-> Actv	In	Q /sec			Total util	Emul time	User ovrhd	Sys ovrhd	Idle time	Steal time
14:45:00	45	188	2.3	0	IFL	42.8	39.9	1.2	1.8	55.9	1.3
				1	IFL	43.9	41.5	1.2	1.3	54.7	1.3
				2	IFL	45.1	42.7	1.1	1.3	53.7	1.2
				3	IFL	44.0	41.6	1.1	1.3	54.7	1.3
..				14	IFL	45.8	43.3	1.1	1.4	52.9	1.3
				15	IFL	44.3	41.9	1.2	1.3	54.3	1.3
				16	IFL	45.1	42.6	1.2	1.4	53.6	1.3
				17	IFL	44.4	41.9	1.2	1.3	54.2	1.3
				18	IFL	26.7	25.4	0.6	0.7	23.4	49.9
				19	IFL	25.8	24.5	0.6	0.7	24.3	49.9
				20	IFL	14.9	13.8	0.5	0.6	22.1	62.9
				21	IFL	14.7	13.6	0.5	0.5	22.4	62.9
				22	IFL	10.3	9.7	0.3	0.3	12.6	77.1
				23	IFL	10.3	9.7	0.3	0.3	12.6	77.1
				24	IFL	7.3	7.0	0.1	0.2	7.2	85.4
				25	IFL	7.3	7.0	0.1	0.2	7.3	85.4
				26	IFL	4.2	4.1	0.1	0.1	5.0	90.7
				27	IFL	4.2	4.1	0.1	0.1	5.1	90.7
				28	IFL	0.0	0.0	0.0	0.0	0.2	99.8
				29	IFL	0.0	0.0	0.0	0.0	0.2	99.8
				30	IFL	0.0	0	0	0.0	0	100.0
				31	IFL	0.0	0	0	0.0	0	100.0
				32	IFL	0.0	0	0	0.0	0	100.0
				33	IFL	0.0	0	0	0.0	0	100.0
				34	IFL	0.0	0	0	0.0	0	100.0
				35	IFL	0.0	0	0	0.0	0	100.0
				36	IFL	0.0	0	0	0.0	0	100.0
				37	IFL	0.0	0	0	0.0	0	100.0
				38	IFL	0.0	0	0	0.0	0	100.0
				39	IFL	0.0	0	0	0.0	0	100.0
System:						943.1	891.3	23.9	27.9	1125	1932

VM pushes work to High Entitled
 Linux balances across all CPUs

- Spin locks are expensive
- More use of locks with more CPUs
- What if vcpu holds lock but paged out?
- Is there concurrency with 32 VCPU?

Linux CPU Details: ESALNXS (B)

Report: ESALNXS LINUX VSI System Analysis Report
Monitor initialized: 04/25/22 at 15:00:00 on 8561 seri

Node/ Time	<---Load Numbers-->			CPU	<Processor		Pct Util>	
	Users	Procs	MaxProc	NBR	Total	Syst	User	Idle
LINUX538	0	426	0	Tot	129.2	7.7	120	3054
				1	4.4	0.3	4.0	95.0
				2	9.3	0.5	8.8	89.8
				3	8.3	0.5	7.8	90.9
				4	8.3	0.5	7.7	90.8
				5	6.3	0.4	5.9	92.9
				6	5.3	0.3	5.0	93.9
				7	3.0	0.2	2.8	96.3
				8	6.4	0.3	6.1	92.9
				9	3.1	0.2	2.9	96.4
				10	4.6	0.3	4.2	94.7
				11	3.7	0.2	3.5	95.8
				12	5.5	0.3	5.1	93.7
				13	2.3	0.1	2.1	97.3
				14	3.7	0.2	3.4	95.7
				15	2.8	0.2	2.6	96.8
				16	3.3	0.2	3.1	96.2
				17	3.8	0.2	3.5	95.7
				18	4.2	0.2	4.0	95.5
				19	4.1	0.3	3.8	95.2
				20	3.9	0.3	3.6	95.7
				21	3.1	0.1	2.9	96.7
				22	3.8	0.2	3.5	95.6
				23	4.2	0.2	3.9	95.3
				24	2.6	0.2	2.4	97.0
				25	3.1	0.2	2.9	96.7
				26	3.0	0.2	2.7	96.6
				27	2.2	0.1	2.1	97.6
				28	2.5	0.2	2.3	97.1
				29	2.4	0.1	2.2	97.3
				30	2.2	0.2	1.9	97.5
				31	1.9	0.1	1.8	97.7

Linux balances across all CPUs

- To belabor the point, WHY 32 CPUs
- One major process
- Dispatched on cpus sequentially
- How does this help?

z/VM CPU Parking Details: ESAOPER (B)

```
14:33:59 CPU Unpark from 18 to 20 CPUUtil= "8.06", Projected= "8.92"  
14:34:01 CPU Unpark from 20 to 24 CPUUtil= "9.5", Projected= "10.9"  
14:34:03 CPU Park from 24 to 22 CPUUtil= "9.16", Projected= "9.28"  
14:34:21 CPU Park from 22 to 20 CPUUtil= "6.95", Projected= "8.77"  
14:34:27 CPU Park from 20 to 18 CPUUtil= "6.69", Projected= "7.98"  
14:34:39 CPU Unpark from 18 to 22 CPUUtil= "8.10", Projected= "9.39"  
14:34:41 CPU Park from 22 to 18 CPUUtil= "7.63", Projected= "7.84"  
14:34:49 CPU Unpark from 18 to 20 CPUUtil= "5.63", Projected= "8.00"  
14:34:53 CPU Park from 20 to 18 CPUUtil= "6.22", Projected= "7.99"  
14:34:59 CPU Unpark from 18 to 20 CPUUtil= "7.36", Projected= "8.26"  
14:35:01 CPU Unpark from 20 to 24 CPUUtil= "8.69", Projected= "10.0"  
14:35:03 CPU Park from 24 to 20 CPUUtil= "7.44", Projected= "8.12"  
14:35:33 CPU Park from 20 to 18 CPUUtil= "6.18", Projected= "7.91"  
14:36:01 CPU Unpark from 18 to 20 CPUUtil= "6.92", Projected= "8.63"  
14:36:03 CPU Park from 20 to 18 CPUUtil= "5.78", Projected= "6.73"  
14:36:05 CPU Unpark from 18 to 20 CPUUtil= "6.98", Projected= "8.19"  
14:36:07 CPU Unpark from 20 to 22 CPUUtil= "8.14", Projected= "9.30"  
14:36:09 CPU Park from 22 to 20 CPUUtil= "8.42", Projected= "8.70"  
14:36:11 CPU Unpark from 20 to 22 CPUUtil= "8.74", Projected= "9.06"  
14:36:15 CPU Unpark from 22 to 24 CPUUtil= "9.7", Projected= "10.2"  
14:36:19 CPU Park from 24 to 22 CPUUtil= "7.42", Projected= "9.9"
```

Hipervisor makes some projection of requirements

- Every two seconds, monitor record provided
- 20 VCPU = 40 threads.
- If only 18 threads available, why 32 vcpu in linux servers?

Top down - capture ratio 100% at all levels

- CEC / LPAR
- LPAR / z/VM
- Virtual machine
- Linux process (capture ratio still 100%)

Report: ESALPARS Logical Partition Summary

Totals by Processor type:
<-----CPU-----> <-Shared Processor busy->
Type Count Ded shared Total Logical Ovhd Mgmt

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

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: ESALPARS

Report: ESALPARS Logical Partition Summary SOMEgroup -
 Monitor initialized: 04/25/22 at 14:00:00 on 8561 serial 111108 First record

Time	<-----Logical Partition----->						<-Assigned Shares----->					
	Name	Nbr	Virt CPUs	CPUs Type	<%Assigned> Total	Ovhd	<---LPAR--> Weight	<VCPU Pct> Pct	/SYS /CPU	<-Thread-> Idle	cnt	
14:45:00	Totals:	00	1	CP	0.3	0.0	10	100				
	Totals:	00	318	IFL	1962	25.9	805	100				
	VLXPMW08	08	20	IFL	673.2	7.3	60	7.5	0.37	40.2	373.3	2
	VLXPMW03	01	30	IFL	86.5	2.2	60	7.5	0.25	26.8	76.23	2
	VLXPMW04	02	20	IFL	154.2	4.0	60	7.5	0.37	40.2	133.3	2
	VLXPMW05	05	20	IFL	425.1	2.2	60	7.5	0.37	40.2	345.2	2
	VLXPMW06	06	20	IFL	74.5	1.6	60	7.5	0.37	40.2	64.12	2
	VLXPMW07	0B	20	IFL	151.6	1.8	60	7.5	0.37	40.2	128.4	2
	VLXTMW03	0E	4	IFL	36.2	0.6	10	1.2	0.31	33.5	29.94	2
	VLXTMW04	03	22	IFL	151.5	3.8	200	24.8	1.13	122	135.8	2
	VLXTMW05	0F	6	IFL	5.7	0.1	10	1.2	0.21	22.4	4.84	2
	VLXTMW06	04	8	IFL	181.1	1.6	10	1.2	0.16	16.8	102.9	2
	VLXTMW07	07	7	IFL	9.8	0.2	10	1.2	0.18	0.18	0	1

Totals by Processor type:

<-----CPU----->		<-Shared Processor busy->					
Type	Count	Ded	shared	Total	Logical	Ovhd	Mgmt
CP	1	0	1	0.4	0.3	0.0	0.1
IFL	108	0	108	1986.1	1935.7	25.9	24.5

Look for “Shared processors”

- IFLs shared (all)
- CPU is low utilized

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 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:

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
-----
08/29/18
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 ***
TCP/IP    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

Linux Process Load: 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 0669C7 First record analyzed: 08/29/18 20:00:35

node/ Name	<-Process Ident->			Nice	PRTY	<-----CPU Percents----->					<-----Storage Metrics (MB)----->									
	ID	PPID	GRP			Valu	Valu	Tot	sys	user	syst	usrt	Size	RSS	Peak	Swap	Data	Stk	EXEC	Lib
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_v															0.1	261	18	0	0.44	

Look for processes within Linux, in percent of cpu

- By relevant server (L233P)
- Correct? Relevant? Cron? Init?

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 01/21/
Monitor initialized: 04/15/11 at 10:00:00 on 2097 serial 72655 First record analyzed: 04/15/11 10:00:00

Time	Users <-----				Pages-----				>							
	Loggd On	System Storage	Fixed Store	Non-Pgble	Free Stor	Frame Table	<Available> <2gb	>2gb	System ExSpc	User Resdnt	NSS/DCSS Resident	<-AddSpace> System User	VDISK Rsdnt	<MDC> Rsdnt	Diag	
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 large percent

Storage Utilization: ESASTR1

Report: ESASTR1 Main Storage Analysis Velocity Software Corporate ZMAP 4.4.0 08/31/18 Pg 4
 Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669C7 First record analyzed: 08/29/18 20:00:35

Users <-----MegaBytes-----> Over																	
Time	Loggd On	System Storage	Fixed Store	Non-Pgble	Free Stor	Frame Table	<Available> <2gb >2gb	System ExSpc	User Resdnt	NSS/DCSS Resident	<-AddSpace> System User	VDISK Rsdnt	<MDC> Rsdnt	Diag 98	Commit Ratio	Capt-Ratio	
08/29/18																	
20:01:35	48	147456	0	163	4	1152	2 255	43	143720	27	1589	0	395	0	18	1.387	0.998
20:02:35	48	147456	0	163	5	1152	2 256	43	143718	27	1589	0	396	0	18	1.387	0.998
20:03:35	48	147456	0	163	4	1152	2 256	43	143718	27	1589	0	396	0	18	1.387	0.998

Report: ESASTR1 Main Storage Analysis Velocity Software Corporate ZMAP 4.4.0 08/31/18
 Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669C7 First record analyzed: 08/29/18 20:00:35

Users <-----MegaBytes----->																
Time	Loggd On	System Storage	Fixed Store	Non-Pgble	Free Stor	Frame Table	<Available> <2gb >2gb	System ExSpc	User Resdnt	NSS/DCSS Resident	<-AddSpace> System User	VDISK Rsdnt	<MDC> Rsdnt			
08/29/18																
20:01:35	48	147456	0	163	4	1152	2 255	43	143720	27	1589	0	395	0		
20:02:35	48	147456	0	163	5	1152	2 256	43	143718	27	1589	0	396	0		
20:03:35	48	147456	0	163	4	1152	2 256	43	143718	27	1589	0	396	0		
20:04:35	48	147456	0	163	4	1152	2 259	43	143715	27	1589	0	396	0		
20:05:35	48	147456	0	163	4	1152	2 256	43	143718	27	1589	0	396	0		
20:06:35	49	147456	0	163	4	1152	2 256	43	143721	27	1589	0	393	0		
20:07													1589	0	393	0
20:08													1589	0	392	0
20:09													1589	0	392	0
20:10													1589	0	392	0
20:11													1589	0	392	0

Total storage analysis (in pages)

- MDC? SET MDC MAX/MIN (ZVPS uses MDC)
- VDISK use? Which server?
- User resident should be large percent

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***  
LN XUWA01 2889K 2798K 90725  65398 258675 10999    112    0    5390 13806      0    0  
LN XUWA03 3848K 3762K 85186  63975  8378  21875    277    0 221201  6714      0    223173  
LN XUWA02 685385 648345 37040 296256 84613 36427    2443    0 22943  1983      0    0  
LN XQWA01 1246K 1218K 28190 541178 51075 35529    2727    0 14094  2787      0    1428  
LN XDWA02 713091 672702 40388 56215 148406 16314    649    0    451  1828      0    0  
LN XDWA04 1152K 1120K 31859 592756 96720 13708   63725 63261  1189    942      0    0  
LN XDWA03 330601 324021  6581  4194  39207  3926    5601  5345    120    734      0    8  
LN XTWA04 883228 860363 22865 90734 129722  7768     31    0    182    66      0    1889  
LN XUWA15 693689 664995 28694 53516 137150 10556    1382    0    553    457      0    0
```

Total storage analysis (in pages, or now default “megabyte” option)

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

Virtual Machine Storage : ESAUSPG

Report: ESAUSPG User Storage Analysis Velocity Software Corporate ZMAP 4.4.0 08
 Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669C7 First record analyzed: 08/29/18 20:00:35

UserID /Class	<-Storage Occupancy in MegaBytes->			<--Main Storage page Read/Write-->					Pages		<Address Spaces->			
	<---Main Storage---> Total	>2gb	<2GB	<--Paging--> Xstor	DASD	<-Page Writes to:--> Xsto	Disk	Migr	<Page Reads:> Xstor	Disk	Moved <2GB	Page Faults	<MegaB Resident> VirtDisk	AddSpce
08/29/18														
20:01:35	143720	141862	1857	0	88145	0	26584	0	0	27283	0	41610	395	0
***Key User Analysis ***														
TCPIP	6	6	0	0	12	0	0	0	0	0	0	252	0	0
User Class Analysis														
Servers	2	2	0	0	57	0	0	0	0	0	0	348	0	0
Velocity	25	25	1	0	35	0	0	0	0	0	0	3191	0	0
*Prod	142570	140729	1841	0	85284	0	23749	0	0	3261	0	14196	395	0
*Util	197	195	3	0	1027	0	97	0	0	4	0	846	0	0
TheUsrs	919	906	12	0	1730	0	2738	0	0	24018	0	22777	0	0
Top User Analysis														
L24BP	25578	25247	331	0	12451	0	2798	0	0	1	0	96	1	0
L233P	14602	14413	189	0	8964	0	1097	0	0	383	0	436	0	0
L200P	8491	8382	110	0	7773	0	179	0	0	948	0	2527	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?

Virtual Machine Storage : ESAUSPG

```

<-Storage Occupancy in MegaBytes-> <--Main Storage page Read/Write--> Pages Page <Address Spaces-
UserID <---Main Storage---> <--Paging---> <-Page Writes to:--> <Page Reads:> Moved Faults <MegaB Resident>
/Class Total >2gb <2GB Xstor DASD Xsto Disk Migr Xstor Disk <2GB VirtDisk AddSpce
-----
*****User Summary*****
L24BP 25552 25222 330 0 12473 0 11364 0 0 1563 0 1247 1 0
L233P 14568 14380 188 0 8977 0 5845 0 0 1028 0 8267 0 0
L20BP 12266 12110 157 0 0 0 0 0 0 0 0 34054 5 0
L203P 5493 5422 71 0 3468 0 2606 0 0 4670 0 13644 11 0
L23BP 5476 5405 71 0 3485 0 1347 0 0 7885 0 19404 38 0
L200P 8504 8394 110 0 7773 0 1904 0 0 6203 0 17463 0 0
L239P 22283 21995 288 0 5013 0 13806 0 0 4079 0 23939 0 0
L244P 9876 9749 128 0 9853 0 10761 0 0 1526 0 2976 0 0
L20DP 6888 6799 89 0 4310 0 12507 0 0 753 0 7263 0 0
L208P 5687 5614 74 0 2807 0 1094 0 0 1421 0 9855 0 0
L215P 7609 7511 98 0 3799 0 2218 0 0 435 0 14899 11 0
L224P 7199 7105 94 0 9272 0 10286 0 0 10981 0 2699 0 0
L24FP 1988 1962 26 0 518 0 3 0 0 160 0 7482 324 0
  
```

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

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

VDISK for Swap: ESAVDSK

Report: ESAVDSK VDISK Analysis Report Velocity Software Corporate

Owner	Space Name	-----<--Size--->	<AddSpce>	Priv	VIO	<--pages----->							
		AddSpc	VDSK	Cre-	Del-	or	rate	User	Resi-	Lock-	Sto-	DASD	
		Pages	Blks	ates	etes	Shrd	/sec	Links	dent	ed	len	Read	
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	
LN XUWA03	VDISK\$LN XUWA03\$0206\$051E	250K	2002K	0	0	Shrd	0.65	1	14K	0	1.6	6.7	
LN XUWA03	VDISK\$LN XUWA03\$0207\$051F	375K	3002K	0	0	Shrd	0.29	1	4980	0	0.4	0.7	
LN XUWA03	VDISK\$LN XUWA03\$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													
LN XQWA01	VDISK\$LN XQWA01\$0206\$0530	64256	512K	0	0	Shrd	0	1	46.9	0	0.1	0	
LN XQWA01	VDISK\$LN XQWA01\$0207\$0531	64256	512K	0	0	Shrd	0	1	1381	0	0.3	0	
LN XTWA04	VDISK\$LN XTWA04\$0206\$051C	131K	1049K	0	0	Shrd	0	1	3984	0	11.7	0	
LN XUWA03	VDISK\$LN XUWA03\$0206\$051E	250K	2002K	0	0	Shrd	10.1	1	46K	0	12.9	58.4	
LN XUWA03	VDISK\$LN XUWA03\$0207\$051F	375K	3002K	0	0	Shrd	16.2	1	88K	0	6.1	19.7	
LN XUWA03	VDISK\$LN XUWA03\$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!!!**

```
Report: ESAVDSK          VDISK Analysis Report          Velocity Softwar
Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669C7  First record ana
-----
Maximum VDISK:          Blocks  (MB)
System storage:         --No Limit--
Storage per user:       --No Limit--

                                <--Size--> <AddSpce> Priv  VIO          <--pages-->
                                AddSpc VDSK Cre- Del-  or   rate   User Resi- Lock-
                                Pages  Blks ates etes Shrd /sec  Links dent  ed
-----
08/29/18
20:01:35
L23BP      VDISK$L23BP$$$$0160$00BB  128K 1024K    0    0 Priv 0.32    1 9903    0
L24FP      VDISK$L24FP$$$$0160$0036  128K 1024K    0    0 Priv 5.57    1  83K    0
-----
System Totals:                2334K   37M    0    0    . 16.0    38 101K    0
-----
```

Virtual Disk Analysis

- Which server is using vdisk?
- Are there multiple vdisks, and **PRIORITIZED!!!**

Report: ESAUCD2 LINUX UCD Memory Analysis Report Velocity Software

```

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
LNXDPOB02 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 large buffer, CMM?

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?
 - High available (greater than 5%)
 - High buffer (greater than 20mb)
- Issues? Swap
- Swap used vs available....

Top down:

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

Paging rules change in 6.3

Paging Subsystem: ESAPSDV

Report: ESAPSDV Page And Spool Device Activity Veloc
 Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669C7 First

-----Paging/Spooling----->												
Dev	<-----Slots----->					<per sec>		</Sec><Device->		%Alloc		
No.	Serial	Avail	Used	%Use	Max	Read	Writ	Queue	+RSCH	Serv Time	Resp Time	Select
08/29/18												
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

Report: ESADSD2 DASD Performance Analysis Velocity Software Corporate
 Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669C7 First record analyzed: 08/2

```

-----DASD Response times (ms)-----
Dev          Device <--SSCH--> <%DevBusy> <SSCH/sec-> <-----DASD Response times (ms)----->
No. Serial  Type  Total  ERP  Avg  Peak  avg  peak  Resp  <--Service times--> <--Queueing-->
-----

```

08/29/18
 20:01:35

Top DASD by Device busy

Dev	Serial	Type	Total	ERP	Avg	Peak	avg	peak	Resp	Serv	Pend	Disc	Conn	DASD	Cntl	THR
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”

Paging Subsystem: ESAPSDV (B)

Report: ESAPSDV Page And Spool Device Activity
 Monitor initialized: 04/25/22 at 00:00:00 on 8561 serial 111108

```

-----
Dev          <----Paging/Spooling----->      </Sec><Device->
  No. Serial Avail Used %Use  Max Read Writ Queue +RSCH Time Time
  -----
  
```

04/25/22
 00:15:00

*****Summary*****

Average:

Page Devices

5D16	V51P0A	12M	0	0	0	0	0	0	0.0	0	0
5D17	V51P0B	12M	0	0	0	0	0	0	0.0	0	0
5D18	V51P0C	12M	0	0	0	0	0	0	0.0	0	0
5E95	V51P00	5897K	0	0	0	0	0	0	0.0	0	0
5EA1	V51P01	5897K	0	0	0	0	0	0	0.0	0	0
5C03	V51P07	12M	0	0	0	0	0	0	0.0	0	0
5D14	V51P08	12M	0	0	0	0	0	0	0.0	0	0
5D15	V51P09	12M	0	0	0	0	0	0	0.0	0	0

Total Page 130M 0 0

Spool Devices

5E97	V51D00	5897K	6M	100	0	0	0	0	0.0	0	0
5A46	V51D01	12M	10M	88	0	0	0	0	0.0	0	0
5A47	V51D02	12M	0	0	0	0	0	0	0.0	0	0
5A4D	V51D08	12M	0	0	0	0	0	0	0.0	0	0
5F9B	V51S01	5897K	113K	2	114K	0.0	0.1	0	0.8	0.1	0.1

Total Spl 106M 16M 15 16M 0.0 0.1 0 1.5 0.0 0.0

Paging Configuration:

- How many page devices (13)
- Equal sizes?
- How full? (70% target?)
- Rates reasonable? Device type dependent

Spooling Configuration (10 spool devices)

Paging Subsystem: ESAPSDV (B)

Report: ESAPSDV Page And Spool Device Activity SOMEg
Monitor initialized: 04/25/22 at 14:00:00 on 8561 serial 111108 First

-----Paging/Spooling----->												
Dev	<-----Slots----->					<per sec>		</Sec><Device->				%Alloc
No.	Serial	Avail	Used	%Use	Max	Read	Writ	Queue	+RSCH	Serv Time	Resp Time	Select
15:00:00												
Page Devices												
5D16	V51P0A	12M	442K	4	698K	351	968	17.3	23.0	0.0	0.6	100.0
5D17	V51P0B	12M	444K	4	699K	352	970	32.3	22.9	0.1	1.1	100.0
5D18	V51P0C	12M	442K	4	700K	355	974	4.3	22.7	0.0	0.0	100.0
5E95	V51P00	5897K	440K	7	694K	358	971	4.3	23.7	0.0	0.0	100.0
5EA1	V51P01	5897K	442K	7	696K	357	971	27.1	25.0	0.0	0.9	100.0
5F90	V51P02	5897K	443K	8	700K	350	966	24.5	23.4	0.0	0.4	100.0
5F91	V51P03	5897K	441K	7	697K	350	966	60.3	23.6	0.1	0.4	100.0
5C00	V51P04	12M	442K	4	698K	349	965	24.9	23.6	0.0	0.4	100.0
5C01	V51P05	12M	444K	4	701K	351	970	0	24.2	0.0	0.0	100.0
5C02	V51P06	12M	443K	4	700K	355	972	0	22.8	0.1	0.1	100.0
5C03	V51P07	12M	440K	4	696K	348	963	34.1	23.5	0.1	6.6	100.0
5D14	V51P08	12M	441K	4	697K	349	963	33.8	23.3	0.1	0.1	100.0
5D15	V51P09	12M	442K	4	699K	352	970	35.9	22.9	0.1	0.1	100.0
Total Page		130M	6M	4	9M	4577	13K					

Paging Configuration:

- How many page devices (13)
- Equal sizes?
- How full? (70% target?)
- High queuing, 20 I/O per second (AVG)
- Response time “looks ok”? (avg per page)

Paging Subsystem: ESADSD2 (B)

Report: ESADSD2 DASD Performance Analysis SOMEgroup -
 Monitor initialized: 04/25/22 at 14:00:00 on 8561 serial 111108 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
-----
```

15:00:00
 Top DASD by Device busy

No.	Serial	Type	Total	ERP	Avg	Peak	avg	peak	Resp	Serv	Pend	Disc	Conn
5D17	V51P0B	3390-9	20635	0	5.6	9.7	22.9	39.5	4.5	2.4	0.1	0.6	1.8
5D15	V51P09	3390-9	20630	0	5.5	9.6	22.9	39.0	4.5	2.4	0.1	0.6	1.7
5C02	V51P06	3390-9	20499	0	5.2	8.9	22.8	38.6	5.7	2.3	0.1	0.6	1.6
5D18	V51P0C	3390-9	20447	0	5.0	8.6	22.7	37.6	4.7	2.2	0.1	0.6	1.5
5D16	V51P0A	3390-9	20734	0	5.0	8.3	23.0	39.3	3.7	2.2	0.1	0.5	1.5
5D14	V51P08	3390-9	21009	0	5.0	8.6	23.3	42.7	5.1	2.1	0.1	0.5	1.5
5C01	V51P05	3390-9	21829	0	4.9	8.2	24.3	43.5	4.4	2.0	0.1	0.5	1.4
5C03	V51P07	3390-9	21150	0	4.9	8.5	23.5	41.5	4.5	2.1	0.1	0.5	1.4
5C00	V51P04	3390-9	21269	0	4.7	7.9	23.6	42.2	3.0	2.0	0.1	0.5	1.4

DASD Rates, Performance ESADSD2

- average service/response time
- Pend, disconnect low -> Else dasd cache
- Connect low -> Else faster channels
- Disconnect, get solid state
- Response = service, else queueing
- Peak busy for device (1 minute peak)
- Service time of 2.4 is high by today's standards

Paging Subsystem: ESADSD2 (B)

Dev No.	Serial	Device Type	<--SSCH-->		<%DevBusy>		<SSCH/sec->		<-PAV/HiperPAV>		
			Total	ERP	Avg	Peak	avg	peak	Q-d	SSCH	Alias
Top DASD by Device busy											
5D17	V51P0B	3390-9	20635	0	5.6	9.7	22.9	39.5	0	0	0
5D15	V51P09	3390-9	20630	0	5.5	9.6	22.9	39.0	0	0	0
5C02	V51P06	3390-9	20499	0	5.2	8.9	22.8	38.6	0	0	0
5D18	V51P0C	3390-9	20447	0	5.0	8.6	22.7	37.6	0	0	0
5D16	V51P0A	3390-9	20734	0	5.0	8.3	23.0	39.3	0	0	0
5D14	V51P08	3390-9	21009	0	5.0	8.6	23.3	42.7	0	0	0
5C01	V51P05	3390-9	21829	0	4.9	8.2	24.3	43.5	0	0	0
5C03	V51P07	3390-9	21150	0	4.9	8.5	23.5	41.5	0	0	0

Paging I/O Response times (ESADSD2)

- Service time of 2.4 is high by today's standards
- Disconnect time high
- **Why is there a queue? (hiperpav?) – NOT ON!**

Paging Analysis: ESABLKP (B)

Report: ESABLKP Block Paging Analysis SOMEgroup - VLXPM
Monitor initialized: 04/25/22 at 14:00:00 on 8561 s First record anal

Time	<----Load----> <-Users-> Tran Actv In Q /sec	Serv Time (ms)	<-Block-> <-Reads-> /sec Size	<-Blocks <-Steal-> /sec Size	Block Fault /sec	<--Block Exce <Single Read> User System
14:15:00	45 190 2.3	.	0 0	0 0	0	0 0.0
14:30:00	45 195 2.3	.	0 0	0 0	0	0 0.0
14:45:00	45 188 2.3	0.0	4.4 18.7	64.5 26.9	4.4	0.1 0.2
15:00:00	45 189 2.3	1.2	123.8 18.9	391.3 27.8	123.8	73.3 7.7
15:15:00	46 179 2.4	3.0	148.2 18.7	434.4 27.5	148.2	207.3 18.5
15:30:00	45 194 2.3	2.6	99.3 18.3	111.7 30.0	99.3	173.4 11.2
15:45:00	47 193 2.4	3.3	128.7 17.6	329.3 28.7	128.7	185.4 41.2
16:00:00	48 191 2.2	4.1	59.9 19.1	53.6 31.3	59.9	78.3 7.5

Block Paging Analysis

- Block page read – optimal 10 pages
- Service time is from scheduler, time to “read page”
- Pages stolen, unreferenced – Storage stress
- Single page read – goes up with 6.3
- As space fills, block size drops.

Top down:

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

DASD Configuration: ESADSD1

Report: ESADSD1 DASD Configuration Velocity Software Corporate

Dev No.	Sys ID	Serial	Device Type	SHR	<CHPIDS OnLn>				MDisk Links	<----Extent---->		<--MDC St		
					01	02	03	04		Type	Start	Size	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

DASD Configuration: ESADSD1

Report: ESADSD1 DASD Configuration Velocity Software Corporate ZMAP 4.4.0 08/31
Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669C7 First record analyzed: 08/29/18 20:00:35
Monitor period: 3600 seconds (1:00:00) Last record: 08/29/18 21:00:35

Dev No.	Sys ID	Serial	Device Type	SHR	<CHPIDS	OnLn	<-Cntrl	Unit->	UserID	MDisk	Cyl/Blk	<----Extent----	<MDC	Status>	STG				
					01	02	03	04	OBR/CU Model	(if ded)	Links	Count	Type	Start	Size	Elig	Now	Shr	ID
2A20	17A9	LV1P00	3390-9	NO	58	53	40	55	3C/00 2107		0	10017	Page	1	10K	Yes	On	No	2A00
					42	39	4C	43											
2A21	17AA	LV1P05	3390-9	NO	58	53	40	55	3C/00 2107		0	10017	Page	1	10K	Yes	On	No	2A00
					42	39	4C	43											

DASD Configuration

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

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

Dev No.	Device Serial Type	<--SSCH-->		<%DevBusy>		<SSCH/sec->		<-----DASD Response tim <--Service times-->				
		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				
Dev	Device	<--SSCH-->	<%DevBusy>	<SSCH/sec->	<-----DASD Response tim									
No.	Serial	Type	Total	ERP	Avg	Peak	avg	peak	Resp	Serv	Pend	Disc	Conn	
<--Service times-->														
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)

V2: DASD Cache: ESADSD5

Report: ESADSD5 3990-3 Cache Analysis Velocity Software Corporate ZM
 Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669C7 First record analyzed: 08/29/18

```

-----
                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
-----
08/29/18
20:01:35
***Top DASD by Device busy***
2A8C  LV1P22  100  9.0  5.3 58.7  88.6  8.0  4.3 53.3    0    0    0    1.0  1.0  1.0  0.2  100    0
2A9D  LV1P16  100  8.7  4.6 52.1  95.4  8.3  4.2 49.8    0    0    0    0.4  0.4  0.4  0.1  100    0
2A89  LV1P19  100  8.1  4.1 50.1  91.0  7.4  3.3 45.1    0    0    0    0.7  0.7  0.7  0.1  100    0
***End Top DASD by Device busy***

-----
2200 CtlUnit  100 17.5 14.6 83.7  49.0  5.3  2.7 51.1  3.3  3.0 91.5  8.9  8.9  8.9  2.8  100    0
2479 L20804   100  4.9  3.3 67.5  92.2  2.6  1.1 44.2  2.0  1.8 91.5  0.4  0.4  0.4  0.0  100    0
-----
2400 CtlUnit  100 19.8 16.5 83.4  65.8  9.7  6.7 68.8  3.3  3.0 92.3  6.8  6.8  6.8  2.3  100    0
  
```

DASD Cache: ESADSD5

- 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

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

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-  
LNXPB02    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

Network Configuration: ESATCPI

Report: ESATCPI TCPIP Interface Configuration Report Velocity Software Corpo
Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669C7 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

- Ethernet adapter
- Loop back
- MTU check

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 0669C7      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 Output 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    0 34.08    0    0    0    0    0
```

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

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 Output Input Output Inpt Output
-----
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
LNXDPA02   - 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
LN XUWA01  - 1 21167   21167      0 57.81  57.8      0      0      0      0
           - 4  109K   122K      0 236.9  268.5      0      0      0      0
LN XDWA02  - 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

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

Date/ Time	Dev. Nmbr	owner	Virt DevN	QDIO Fmt	Number		<QDIO SIGA Instructions/Sec->		<-Throughput / sec->		<Buffers>		<--Bytes-->		
					Queues	In Out	<---Guest--->	<---CP----->	Read	Writ	Read	Writ	Read	Writ	
11:00:00	0000	Totals	0000	QDIO	0	0	0	0	0	693	0	1066	676	644K	422K
	F3D8	VSWCTRL2	F3D8	QDIO	1	1	0	0	0	573	0	895	535	527K	306K
	F3E0	VSWCTRL2	F3E0	QDIO	1	1	0	0	0	119	0	171	141	118K	117K
	F53E	LN XUWA02	7002	HPER	1	4	0	0	0	0.6	0	1	0	89	0
*****Summary*****															
Average:	0000	Totals	0000	QDIO	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
	F3D8	VSWCTRL2	F3D8	QDIO	1	1	0	0	0	530	0	891	491	529K	322K
	F3E0	VSWCTRL2	F3E0	QDIO	1	1	0	0	0	108	0	149	130	85716	119K
	F3F0	VSWCTRL1	F3F0	QDIO	1	1	0	0	0	0	0	0	0	0	0
	F515	LN XDPB02	7002	HPER	1	4	0	0	0	0	0	0	0	0	0
	F518	LN XDWA01	7002	HPER	1	4	0	0	0	0	0	0	0	0	0
	F53B	LN XUWA01	7002	HPER	1	4	0	0	0	0	0	0	0	0	0
	F53E	LN XUWA02	7002	HPER	1	4	0	0	0	0.6	0	1	0	92	0
	F542	LN XUWA03	7002	HPER	1	4	0	0	0	0	0	0	0	0	0
	F545	LN XUWA04	7002	HPER	1	4	0	0	0	0	0	0	0	0	0
	F548	LN XDMS2A	7002	HPER	1	4	0	0	0	0	0	0	0	0	0

QDIO activity

- Hipersockets
- Virtual switch

QDIO Data Rates: ESAQDIO

Report: ESAQDIO Queued I/O Report Velocity Software Corpor
Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669C7 First record analyzed: 0

Date/ Time	Dev. Nmbr	owner	Virt DevN	QDIO Fmt	Number		QDIO SIGA		Instructions/Sec->		<-Throughput / sec->		<--Bytes-->		
					Queues	In Out	<---Guest----	> <----CP----->	<Buffers>	<--Bytes-->					
								"s"		"s"	Sent	From	Sent	From	
08/29/18															
20:01:35	0000	Totals	0000	QDIO	0	0	0	0	0	815	0	3511	1711	101M	842K
	1790	DTCVSW1	0600	QDIO	1	1	0	0	0	130	0	195	117	150K	150K
	7584	L24BP	7584	FCP	1	1	0	0	0	0	0	4	4	51939	0
	7588	L24FP	7588	FCP	1	1	0	0	0	0	0	57	28	2212K	0
	758E	L10CU	758E	FCP	1	1	0	0	0	0	0	2	2	3519	0
	759D	L25FP	759D	FCP	1	1	0	0	0	0	0	3	3	78867	0
	75C4	L24BP	75C4	FCP	1	1	0	0	0	0	0	4	4	48033	0

QDIO activity

- Hipersockets
- Virtual switch

Guest Lan / Virtual Switch Data Rates: ESANIC / ESATCP4

Report: ESANIC Virtual NIC Activity Report Velocity Software Corporate
 Monitor initialized: 08/29/18 at 20:00:35 on 2827 serial 0669C7 First record analyzed: 08/29/18

Date/Time	Virt	NIC	Tranp	<---network lock requests--->		<-stack->		<-bytes->		<-----Pack						
Userid	Virtual	Devc	BASE	/type	<--Per Second-->	waits/sec	wait	</second>	</Second>	<-Rate-->						
	LanName	ADDR	ADDR	Ntwrk	send	rcv	send	rcv	time	rqst	dfrd	Sent	Rcvd	Sent	Rcvd	
08/29/18																
20:01:35																
L233P	VSW4	0340	0340	02/02	579.8	580	175	0	0	0.1	595	0	497K	474K	598	595
L233P	VSW4	0360	0360	02/02	98.4	98.4	60.2	0	0	0.0	82.8	0	24K	25K	100	82.8
L24BP	VSW5	0360	0360	02/02	12.8	12.8	8.2	0	0	0	10.5	0	6699	2580	12.9	10.5

Screen: **ESATCP4** Velocity Software - VSIWM4
 1 of 2 TCPIP Hardware Layer / Interfaces

Time	Node/ Group	Interface	<Total Octets>	
			<-Per second->	
			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

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 initialized: 06/15/16 at 00:00:00 on 2828 seria 314C7
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
```

There IS logic to doing performance analysis

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