

VELOCITY

S O F T W A R E

Introduction to Performance Managing using zVPS for Linux on z

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

- Barton@VelocitySoftware.com
- [HTTPS://VelocitySoftware.com](https://VelocitySoftware.com)

- **Performance Management Overview**
- **Systems Management Features**
- **zVPS Objectives (and buzzwords)**
- **Single pane of glass**
- **End to End Performance Management**
- **zVPS (old and new)**
 - Data Collection
 - PDB
 - Technology
- **Applications**
- **Using zVPS: Case Studies**

- **SHARED resource environment,**
 - z/VM Performance critical (Bank runs their ATMs on Linux on Z)
 - Any server or application can impact other servers or applications
- **Linux is not z/OS**
 - No workload manager, workload prioritization is manual
- **This is not distributed Environment**
 - We do not have cycles to waste
 - We DO have capacity planning, chargeback requirements
 - We DO expect to run at very high utilization!
 - We do need to educate the users
- **Tools are needed specific to the environment**
 - “end to end”

Performance **Management** is a process

- Performance Analysis
- Operational Alerts
- Capacity Planning
- Accounting/Charge back

Our Product Objectives:

- Data Accuracy, product longevity, scalability, extensible
- Minimize complexity
- Ease of use, support
- Modernization

Performance Management Business Requirements

- Performance Analysis
- Operational Alerts
- Capacity Planning
- Accounting/Charge back

Correct data (Virtual Linux CPU data wrong)

SMT CPU data difficult to understand

Capture ratios (is the data valid?)

Instrumentation can NOT be the performance problem

- Why Performance Analysis: Service Level Mgmt
 - Diagnose problems real time
 - Manage Shared resource environment
 - Any application may impact other applications
- Infrastructure Requirements
 - Analyze all z/VM Subsystems in detail, real time
 - (DASD, Cache, Storage, Paging, Processor, TCPIP)
 - Analyze Linux
 - (applications, processes, processor, storage, swap)
 - Historical view of same data important
 - Why are things worse today than yesterday?
 - Did adding new workload affect overall throughput?
 - Know who/what is using resource and how to re-allocate

- Why Capacity Planning: Future Service Levels
 - How many more servers can you support with existing z15?
 - What is capacity requirements for an application?
 - **Avoid crises *in advance***
 - Consolidation Planning – Projecting requirements of the next 100 or 1000 servers
 - Impact of SMT?
- Infrastructure Requirements
 - Performance database (long term)
 - z/VM **AND** Linux data
 - Resource requirements by Server, **Application**, User
 - z/VM and z/Linux data must be usable by existing planners
 - **Interface to MICS, MXG, TUAM, TDS, IUE (BMC), Splunk**

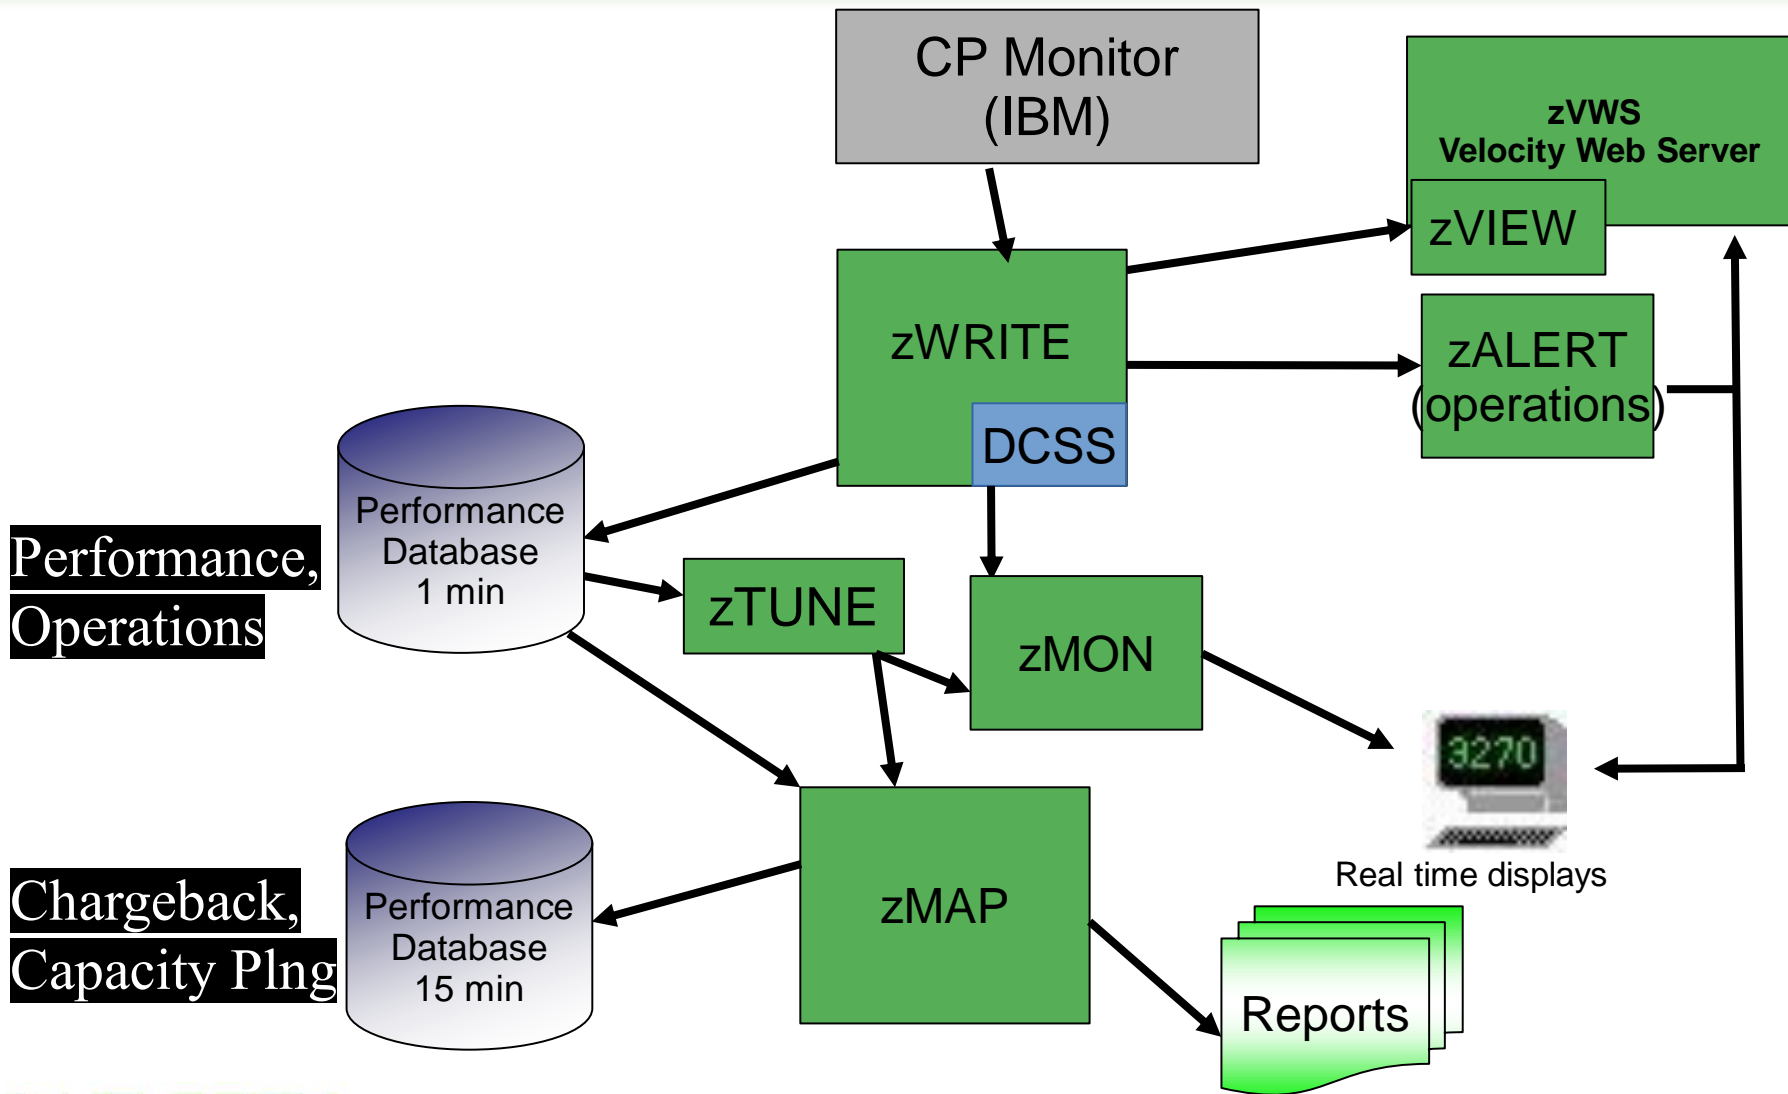
- Why Chargeback?
 - How much does an application cost IT to operate?
 - Distributed chargeback model is by server
 - Shared chargeback model is by resource utilized
 - Convincing customers to move applications to “z”
 - Encourages efficient/effective resource use
 - Align IT to your business model
- Infrastructure Requirements
 - Identify Resource by server
 - Identify Resource by Linux Application
 - **High capture ratio**
 - Every site does it differently, so flexible data is key

- **Operational Requirements**
 - Operations will manage 100's (1000's) of servers
 - Requires active performance management
 - Alerts for processes in loops, disks 90% full, missing processes
 - One test server in a loop impacts all other servers
 - Fast problem detection
 - Requires active performance management
 - **Requires AUTOMATION! (zALERT!)**
- **Infrastructure Requirements**
 - Interface to SNMP management console (NETCOOL, HPOpenView)
 - User tailored alerts
 - Web based alerts

zVPS does “End to End” Performance Management

- **Management wants**
 - “single pane of glass” - One tool that does all
- **Complete performance management includes:**
 - z/VM System Level: CEC, LPAR data, ALL SubSystems
 - Linux – Storage, CPU, file system, network
 - Process – applications, performance data
 - VSE – partitions, CPU, I/O
 - z/OS
- **Network analysis**
- **Application subsystem analysis**
 - Java, WAS, Oracle, MQ, DB2, postgres, gpfs
- **Outside “z” server analysis**
 - Linux on “x”, VMWare, KVM
 - Microsoft servers
 - VPN, gateways, utilities

- **Standard Interfaces mean less work**
 - Agentless
- **CP Monitor – z/VM**
 - LPAR data, CPU data
 - Disk, storage, paging data
 - Virtual machine data
 - Seeks data
- **SNMP – Standard (requires zTCP collector)**
 - Network data, microsoft servers, many appliances
 - VSE data
 - Linux “UCD” mib – high level ram, CPU
- **SNMP – Velocity Software mib**
 - Process details, applications, Java, WAS, Oracle



**Performance,
Operations**

**Chargeback,
Capacity Plng**

Screen: ESAMAIN Velocity Software - VSIVM4 ESAMON 4.090 01/18 16:32-17:04
1 of 3 System Overview LIMIT 500 2096 44B42

Time	<---Users--->			Transact.		CPUs	<Processor>		Cap- ture Ratio	<---Storage (MB)->		
	<-avg number-> On	Actv	In Q	per Avg. Sec.	Time		Utilization Total	Virt.		Fixed User	Active Resid.	Stor Load
17:04:00	137	67	17.0	28.5	0.14	1	13.7	12.3	100	60	2608	0.6
17:02:00	137	68	16.0	29.3	0.14	1	19.4	17.2	100	61	2611	0.6
17:01:00	137	71	20.0	28.9	0.17	1	55.5	48.4	100	60	2609	0.6
16:59:00	137	69	18.0	27.6	0.16	1	19.1	16.9	100	60	2598	0.6
16:58:00	137	67	17.0	28.4	0.12	1	16.2	14.7	100	60	2596	0.6
16:57:00	137	70	22.0	27.8	0.14	1	15.2	13.6	100	61	2597	0.6
16:56:00	137	62	22.0	27.9	0.14	1	64.5	62.9	100	60	2600	0.6
16:55:00	137	64	18.0	29.4	0.12	1	17.6	15.8	100	60	2594	0.6
16:54:00	137	62	20.0	28.8	0.13	1	16.3	14.9	100	61	2589	0.6
16:53:00	137	62	19.0	27.8	0.14	1	15.5	13.9	100	61	2592	0.6
16:52:00	137	68	20.0	27.8	0.13	1	18.0	16.3	100	60	2592	0.6
16:51:00	137	65	21.0	28.6	0.13	1	15.2	13.7	100	60	2594	0.6
16:50:00	137	62	17.0	28.2	0.15	1	16.8	15.3	100	61	2597	0.6
16:49:00	137	65	17.0	28.2	0.13	1	14.9	13.4	100	60	2597	0.6
16:48:00	137	62	18.0	28.2	0.12	1	16.2	14.8	100	61	2600	0.6
16:47:00	137	69	19.0	28.4	0.13	1	15.2	13.7	100	61	2598	0.6
16:46:00	137	63	20.0	27.1	0.14	1	63.9	62.2	100	60	2599	0.6
16:45:00	137	65	21.0	27.9	0.14	1	17.0	15.4	100	60	2599	0.6
16:44:00	137	65	25.0	28.6	0.13	1	14.9	13.6	100	60	2605	0.6
16:43:00	137	67	25.0	29.3	0.13	1	14.7	12.9	100	60	2603	0.6
16:42:00	137	70	22.0	28.8	0.14	1	17.3	15.6	100	59	2597	0.6
16:41:00	137	66	23.0	27.9	0.14	1	15.6	14.2	100	61	2611	0.6
16:40:00	136	63	25.0	27.8	0.15	1	16.0	14.7	100	59	2611	0.6
16:39:00	136	64	23.0	28.2	0.13	1	14.6	13.2	100	60	2611	0.6
16:38:00	136	62	21.0	27.8	0.14	1	16.1	14.7	100	61	2609	0.6
16:37:00	136	67	20.0	28.1	0.13	1	15.0	13.6	100	60	2609	0.6
16:36:00	136	65	21.0	27.5	0.15	1	63.4	62.0	100	61	2607	0.6
16:35:00	136	63	22.0	27.5	0.15	1	15.4	14.0	100	60	2605	0.6
16:34:00	136	64	20.0	27.9	0.12	1	16.1	14.7	100	61	2604	0.6
16:33:00	136	64	20.0	28.4	0.15	1	14.9	13.5	100	60	2609	0.6

PF1=Help PF2=Menu PF3=Quit PF4=Select PF5=Plot PF6=TOC PA1=CP
PF7=Backward PF8=Forward PF9=Sort PF10=Parms PF11=More PF12=Exit PA2=Copy
====>

```

Screen: ESATOC   Velocity Software - VSIVM4   ESAMON 4.090 01/18 17:07-17:08
1 of 1  Screen Table Of Contents   2096 44B42

Screen      Description
-----
Management Summary
ESAMAIN    System Overview
ESAHDR     System Configuration

System Management Summary
ESAMGMT    System Management
ESAMSLA    Management Service Level Analysis
ESAMTOP    Top Users Management Report

Performance Summary
ESASUM     System Load Summary
ESASUMCH   Channel Path Summary
ESASUMIO   Input/Output Summary
ESASUMPR   Processor Summary
ESASUMPS   Paging And Spooling Summary
ESASUMSM   Service Machine Summary
ESASUMSR   Scheduler Parameter Summary
ESASUMST   Storage Summary
ESASUMTR   Transaction Analysis Summary
ESASUMMD   Minidisk Cache Summary

Service Level Activity
ESAUSLA    User Service Level Analysis
ESAEXACT   Transaction Analysis

Transaction Activity
ESARATE    Transaction Rates And Response Times
ESASYSR    Transaction Rates And Response Times
ESAELAS    Transaction Classification
ESAEXCP    Transaction Exception Log

User Activity
ESAUSR1    User Log Activity
ESASRV1    Server Log Activity (Special)
ESAUSRC    User Configuration Analysis
ESASRVC    Server Configuration Analysis (Special)

PF2=View   PF3=Quit   PF7=Backward  PF8=Forward   PF12=Exit
====>
  
```

```
Screen: ESAUSP2 Velocity Software - VSIVM4 ESAMON 4.090 01/18 17:09-17:10
1 of 3 User Percent Utilization CLASS * 2096 44B42
```

Time	UserID /Class	<Processor>		<Resident->		<-----Main Storage----->		
		Total	Virt	Total	Actv	Lock -ed	<-WSSize--> Total	Actv
17:10:00	System:	15.32	14.23	667K	665K	5448	675K	665K
	REDHAT	4.58	4.53	281K	281K	1997	284K	284K
	TEST	3.56	2.98	161K	161K	844	161K	160K
	*TheUsrs	3.12	3.02	57661	57645	290	59127	57322
	SUSE	1.63	1.57	109K	109K	839	109K	108K
	ORACLE	0.96	0.96	50503	50503	66	50437	50437
	Velocity	0.93	0.90	4552	3444	28	7385	3401
	KeyUser	0.36	0.15	2973	2973	1379	1898	1573
	Servers	0.17	0.13	943	520	5	1874	495

Hit PF2 to zoom on SUSE class, get:

```
Screen: ESAUSP2 Velocity Software - VSIVM4 ESAMON 4.090 01/18 17:11-17:12
1 of 3 User Percent Utilization CLASS SUSE USER * 2096 44B42
```

Time	UserID /Class	<Processor>		<Resident->		<-----Main Storage----->		
		Total	Virt	Total	Actv	Lock -ed	<-WSSize--> Total	Actv
17:12:00	SLES11X	0.39	0.39	24223	24223	247	23976	23976
	SLES11	0.32	0.32	12404	12404	181	12199	12199
	SUSELNX2	0.25	0.23	3648	3648	0	3628	3628
	SLES9X	0.21	0.21	14632	14632	35	14597	14597
	SLES10	0.20	0.20	28935	28935	299	28636	28636
	SLES9	0.20	0.20	12722	12722	177	12545	12545
	SLES8	0.06	0.03	11251	11251	0	11201	11201
	SLES8X	0	0	0	0	0	890	0
	SUSELNX1	0	0	0	0	0	219	0

zMAP listings and zMON displays SAME names

Report: ESATOC **Table Of Contents**
Monitor initialized: 12/23/14 at 13:55:
Monitor period: 660 seconds (

```

-----
Report      Title                                                                 Page (s)
-----
ESAHDR      z/VM Monitor Analysis ..... 3-      8
ESATUNE     Tuning Recommendation Report ..... 9-     21
            Performance Summary
ESASSUM     Subsystem Activity ..... 22
ESASUM      System Summary ..... 23-     28
            Transaction Activity
ESAUCLA     User Service Level Analysis ..... 29-    31
ESAXACT     Transaction Delay Analysis ..... 32-    40
ESARATE     Transaction Rates And Response Times ..... 41-    46
ESACLAS     Transaction Classification ..... 47
            User Activity
ESASRVC     Server Configuration ..... 48
ESASRV1     Server Log Activity ..... 49
ESAUsrc     User Configuration ..... 50
ESAUsrc1    User Log Activity ..... 51
  
```


History data format – long term

- All history in “daily” files, yyyyymmdd

ESAEXTR extracts data from history

- User designed reports, CSV files

Command Formats

- ESAMAP yyyyymmdd
- ESAMAP yyyyymm*
- ESAMAP (WEEK 51
- ESAMAP (MONTH 12
- Same for ESAEXTR

- **Performance database language:**
 - ESAEXTR – ZMAP feature
 - HISTORY KEYWORDS – describes variable names (~4000 metrics)
- **ESAEXTR Functions**
 - ESAEXTR filetype* (PARM ucdsys CSV
 - Filetype is history type, as in 201706* for “june, 2017”
 - CSV produces CSV format vs column aligned
- **ESAEXTR Statements**
 - EXTRACT:
 - x = 'NODE'
 - y = 'UCDSYS.REALSIZE'
 - y = 'UCDSYS.BUFFER'
 - y = 'UCDSYS.CACHE'
 - y = 'UCDSYS.REALSIZE-UCDSYS.REALAVAIL-UCDSYS.BUFFER-UCDSYS.CACHE' ; anonymous
 - y = 'UCDSYS.SWAPSIZE-UCDSYS.SWAPAVAIL'
 - y = 'UCDSYS.CMM'
 - **criteria = NODE = LNxD01*** ; wild card
 - TITLE = 'Linux Storage Analysis'
 - TITLE = '*Label Available Buffer Cache Anonymous'
- **or: ESAMAP filetype* (reportusr = 'LNxD01')**

- Requirement to go beyond z/VM and Linux metrics
- z/VPS provides over 4,000 unique metrics
 - z/VM System, storage, paging, dasd metrics (3,000)
 - z/VM Virtual machine metrics (~400)
 - Network metrics (~100)
 - Linux System metrics (~250 VSI, 80 HST, 80 UCD)
 - Linux Process metrics (~40)
- **Application subsystem users:**
 - **Oracle (70 metrics), supports (10G, 11G, 12C)**
 - **Websphere (30 metrics)**

Daily/Weekly/Monthly CSV files created, zview exposed....

```
; Chart Syntax: type days strt stop extr parm
; Where type is DAILY/WEEK/MONTH
; "days" is the previous "n" days
; TimesT is called with these parms
```

```
;      type    days  strttime stop    extractname  parm
CHART  DAILY    1    00:00  24:00  CPULPAR
CHART  DAILY    1    00:00  24:00  USERCPU

CHART  MONTHLY  31   00:00  24:00  CPULPAR
CHART  MONTHLY  31   07:00  17:00  USERCPU
CHART  MONTHLY  31   *      *      USERWAIT  SUSELNX1

CHART  WEEKLY   7    07:00  17:00  USERCPU
CHART  WEEKLY   7    00:00  24:00  CPULPAR
```

```
; Charts will be kept up to 12 months,52 weeks, 31 days
; format of chartcnt is "chartcnt mm ww dd"
CHARTCNT 2 2 2
```

```
ESAMAP      FILELIST A0  V 1
Filename Filetype Fm
TOTAL      CSVDC348 A1
TOTAL      CSVDU348 A1
TOTAL      CSVWC049 A1
TOTAL      CSVWU049 A1
TOTAL      CSVDC347 A1
TOTAL      CSVDU347 A1
TOTAL      CSVDC346 A1
TOTAL      CSVDU346 A1
```

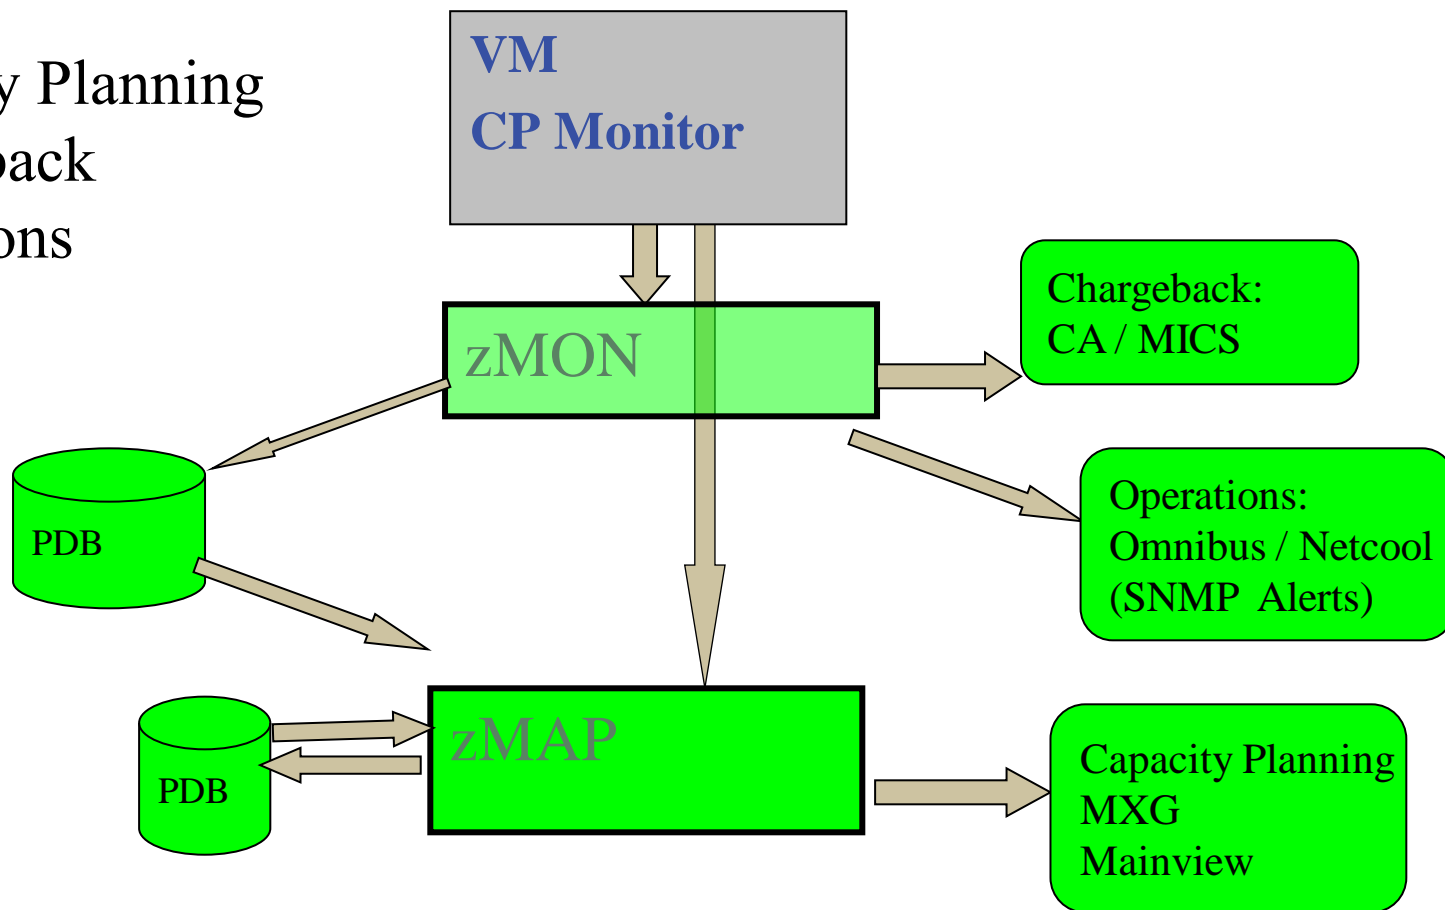
```
;LPAR Utilization over time
"STARTTIME","STOPTIME","LPARNAME","CPUUTIL"
"2010/12/13","02:00:00","Totals:","79.64"
"2010/12/13","02:00:00","VSIVM4","34.72"
"2010/12/13","02:00:00","VSIVM1","3.87"
"2010/12/13","02:00:00","VSIVM2","0.41"
```

- **Monitor data very large**
 - Process it real time, extract information, discard the data
 - All classification and collection done before discard
- **Classification functions**
 - User classes – defined by installation, some predefined
 - “Key users”, class 0 is special
 - DASD and non-DASD Control Units
- **USERCLASS statements**
 - userclass = 20
 - class_size = 140
 - nclasses = 1

 - Userclass(03,000) = 'ZVPS '
 - Userclass(03,007) = 'ZMAP'
 - Userclass(03,008) = 'ZMON'
 - Userclass(03,009) = 'ZWRITE'
 - Userclass(03,010) = 'ZSERVE'
 - Userclass(03,011) = 'ZTCP'
 - Userclass(03,016) = 'ZTCP'
 - Userclass(03,017) = 'ZWEB*'

Add “Enterprise” Support: (No silos)

Capacity Planning
Chargeback
Operations



- **Linux (and networks) adds requirement**
 - **Correct data (bad assumption)**
 - **Complete data (very difficult)**
 - **Low cost data (extremely difficult)**
- **Support requirements:**
 - SLES 7,8,9,10, 11, 12 (Installations still have 7 and 8)
 - RHEL 3,4,5, 6,7
 - UBUNTU, KVM
 - Other platforms (VSE, VMWare, SUN, P, **MicroSoft, ESX**)
- **Must support:**
 - Performance tuning (one minute granularity)
 - Capacity planning (15 minute granularity)
 - Operational alerts (one minute granularity)
 - Chargeback/Accounting (15 minute granularity)

- **Operational cost of agents**

- Does your agent use 2%? 5%? 95%? of a processor per Linux server?
- Does this matter on distributed servers where agents were created?
- Will local data collection fill up your file system?
- Does turning off performance monitoring solve the performance problem?
- Do you only turn on your agent when you have a problem???
- Diagnostics vs Performance Management?
- **Customer quote: an agent that costs 1% of a processor will cost me 10 IFLs**
- **(standard snmp host mib, about 1%, VSI mib .1%)**

- **Agents must provide correct data**

- Is your data correct? Or wrong by order of magnitude?
- Prior to SLES10/RHEL5, all “Virtual” agents provide wrong data
- **Why collect bad data?**

•Operational cost of running agents

- 2% per server costs 1 IFL per 50 servers,
- **Velocity targets less than .1% (point one percent)** of ONE processor with one minute data collection per Linux server
- Agents developed for INTEL such as Splunk are expensive

•Data Accuracy not easy

- Virtualized CPU (SMT) accounting must be normalized

•Capture ratios

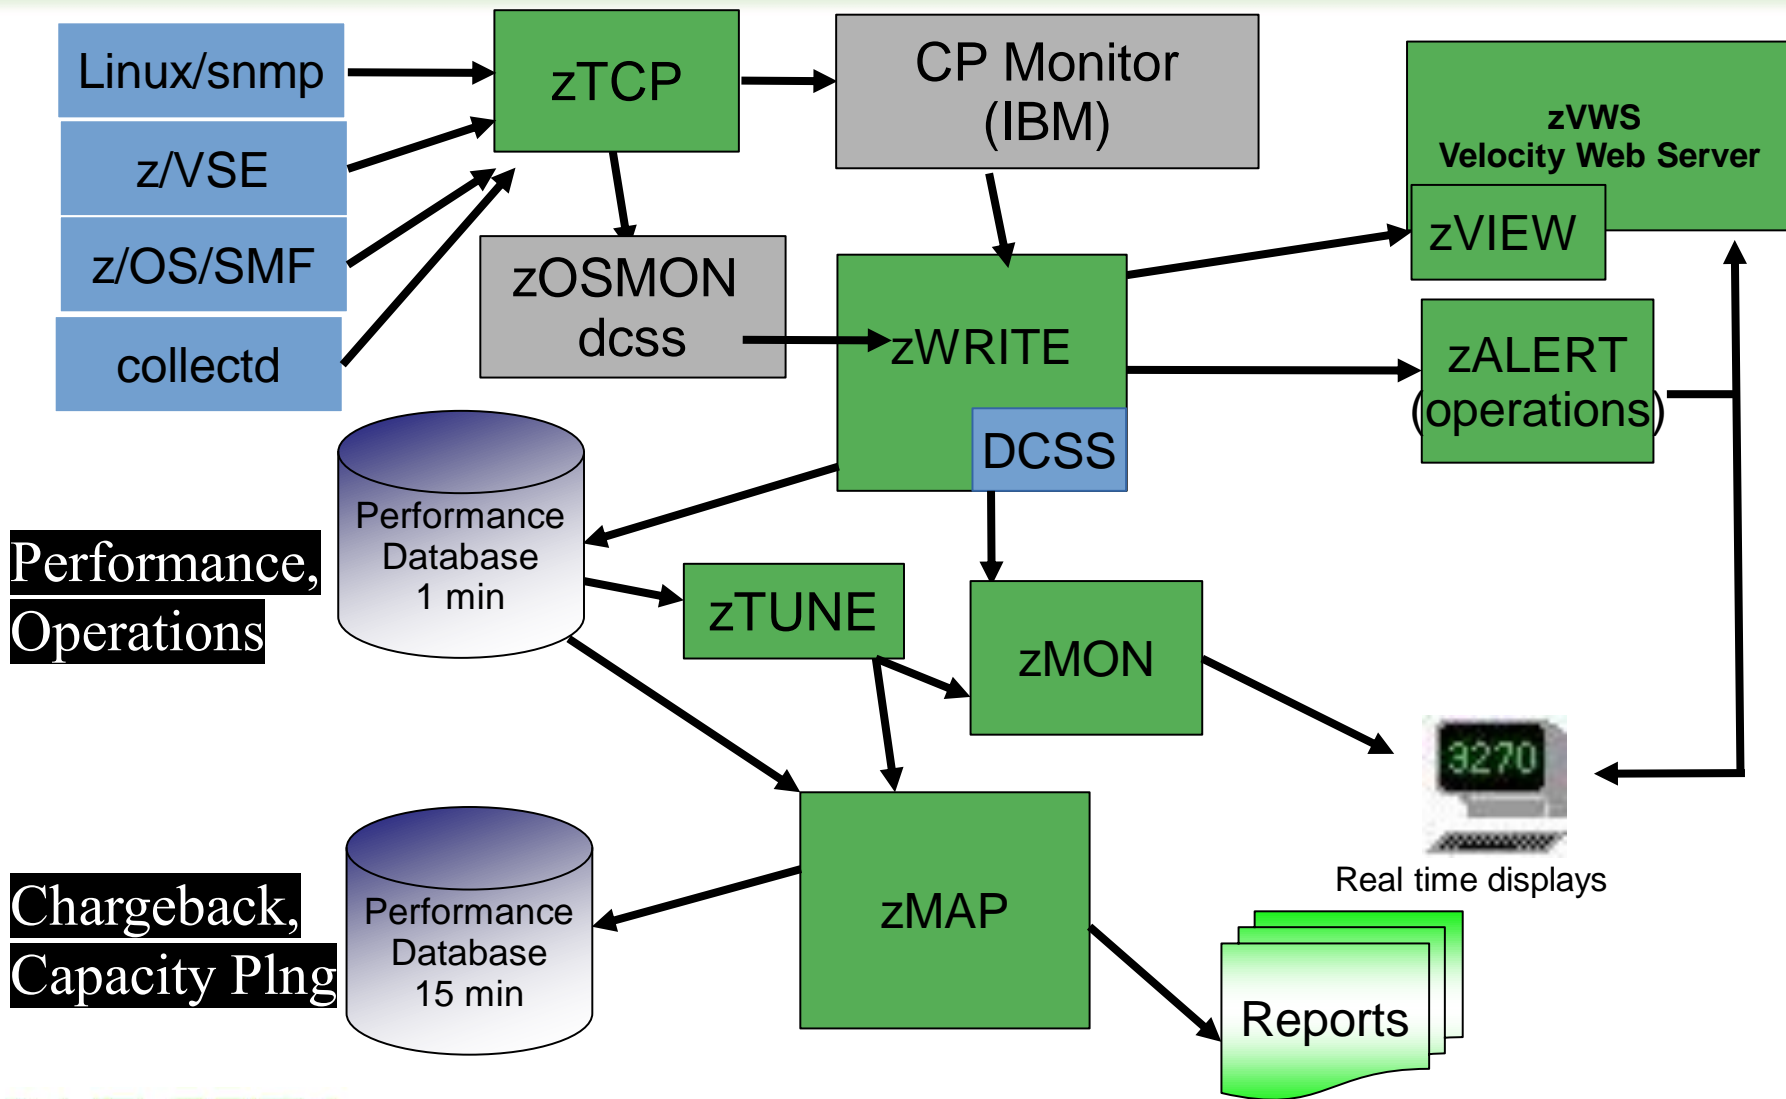
- Data must be complete,
- Capture ratio normally at 100% to the Linux process level

•Skills

- Skills are lacking in managing highly virtualized environments
- Access to skills critical when there are performance problems.

•Performance Management Education (on demand)

- **Performance Data infrastructure existed (zMON/zMAP)**
 - PDB already existed for performance analysis and Capacity Planning
 - Data presentation tools existed
- **Data source needed for Linux and Network:**
 - **Low overhead (want to monitor 100 / 1000 servers under z/VM)**
 - **Agents developed for Intel and Distributed servers did not care about overhead**
 - Open Source (fast development time), instrumentation **MUST** be part of the platform
 - Standard interface
- **SNMP: Standard interface for network and host data**
 - Provided by TCPIP Vendor
 - **Not proprietary agent! – Can't charge for it....**
 - Used to collect network, host data from NT, SUN, HP
 - NETSNMP available for Linux - Meets all requirements
 - (Distributed with RHEL 3,4,5 SLES 7,8,9,10,11)
 - **Platform independent (Intel, P-series, Microsoft, Linux)**
 - **.03% of ONE IFL (z10,sles9) per server, ONE MINUTE COLLECTION**



**Performance,
Operations**

**Chargeback,
Capacity Pnlg**

Longevity requires consistency and standards

- Correct data implies standard data
- Data sources must be consistent, low overhead, integrated
- zVPS uses standard sources (mostly....)

z/VM: CP Monitor (IBM) Exclusively

Networks: snmp mib-ii (standard, open source)

Linux: netsnmp (standard with Linux, “z” and “x”)

- Standard Netsnmp is 1% “agentless” agent (ucd mib, host mib)
- Velocity Software snmp mib (“z” and “x”) replaces most metrics for .1%
- ALL Distributions (suse, redhat, ubuntu), all releases (z & x)

VSE: IBM mib, Velocity Software mib, CICS (DMF) (2021)

- (“<http://VelocitySoftware.com/vsecics.html>”)
- BSI/CSI TCPIP from the vendors

z/OS: SMF records (IBM/logstream) (70/30/75/113,CICS, DB2, etc)

MibII: Transport layer (ESATCP1), IP layer(ESATCP2), hardware layer (ESATCP4), icmp (ESATCP3)

- Transport layer data shows connections, TCP rates, UDP rates

```
Report: ESATCP1          TCPIP Transport Layer Data Report
-----
```

Date/	<-----TCP Connections----->					<-TCP Communications / sec				
Time/	Current	<Opens/Second>		<Closes/Sec>	<----Segments Transmitted-					
Node	Connects	Active	Passive	Fails	Resets	Input	Outpt	ReTran	InError	

00:15:00										
Node Groups										
KeyUser	1.1	0.0	0.0	0.0	0	0.04	0.06	0.07	0.00	
*TheUsrs	21.9	1.3	1.6	0.2	0	48.74	48.75	0.00	0	
VsLPARs	5.3	0.1	0.5	0.1	0.3	8.02	11.95	0.29	0.08	
*** Nodes *****										
oracle	16.9	0.5	1.0	0	0	24.51	24.52	0.00	0	
RH5X161	0	0	0	0	0	0	0	0	0	
S11R20RA	5.0	0.8	0.6	0.2	0	24.25	24.26	0	0	
TCPIP	0	0.0	0.0	0.0	0	0.02	0.03	0.03	0.00	
TCPIP2	0	0	0	0	0	0.01	0.01	0	0.00	
TCPIP2	1.1	0.0	0.0	0.0	0	0.02	0.03	0.03	0.00	
VSIVM1	2.0	0.0	0.0	0.0	0	0.49	0.59	0.10	0.01	
VSIVM2	1.1	0.0	0.0	0.0	0	0.02	0.03	0.03	0.00	
VSIVM4	2.2	0.0	0.4	0.0	0.3	7.51	11.33	0.15	0.07	

Analyzing “distributed” Disks (snmp data)

HOST MIB data:
Provides disk data
Percent full
Supports WinNT, Unix
Alerts by disk full

Standard data!!!

Report: ESA**HST2** LINUX HOST Storage Analysis Report
Monitor initialized: 02/05/07 at 10:41:41 on 2084 serial 55BAF

```

-----
NODE/          <-Utilization->          <-----Storage----->
Time/          <MegaByte>  Pct          Alloc
Date          Index  Size  Used  Full  Errors  Units  Description
-----
10:43:00
acme
          1    495  14.2  2.9          0    1024  Memory Buffers
          2    495   487 98.4          0    1024  Real Memory
          3   2031  12.8  0.6          0    1024  Swap Space
          4   2310   775 33.6          0    4096  /
          6   2310  1293 56.0          0    4096  /usr

dominoz1
          1   2002  38.5  1.9          0    1024  Memory Buffers
          2   2002  1994 100          0    1024  Real Memory
          3   2031  97.4  4.8          0    1024  Swap Space
          4   2310  1556 67.4          0    4096  /
          6   2310  1398 60.5          0    4096  /usr
          7   984K  238K 24.2          0    4096  /notesdata

ebiz1
          1    997   9.0  0.9          0    1024  Memory Buffers
          2    997   992 99.5          0    1024  Real Memory
          3   2031   514 25.3          0    1024  Swap Space
          4   2310  1607 69.6          0    4096  /
          6   2310  1451 62.8          0    4096  /usr
          7   101K   10K 10.3          0    4096  /notesdata

```

Distributed Systems Process data

Windows NT

Screen: **ESAHST1 NT Data**

ESAMON V3.2 07/30 14:56-14:57

1 of 1 LINUX HOST Software Analysis Report

NODE * LIMIT 500

Time	Node	Name	ID	Type	Status	Total	Intrval	CPU Pct	Storage (K) Current
14:57:00	ENTWDB	NetTime.	2648	4	1	4259	0.68	1.12	1320
		NetTime.	2452	4	1	982	0.57	0.94	1040
		sqlagent	2408	4	1	100	0.03	0.05	3724
		snmp.exe	2268	4	1	73	0.07	0.12	3888
		taskmgr.	2224	4	1	21076	0.28	0.46	2524
		sqlservr	2136	4	1	50038	9.53	15.72	511624
		NetTime.	1808	4	1	10481	1.47	2.42	1092
		sqlmangr	1660	4	1	189	0.01	0.02	3664
		DLLHOST.	1648	4	1	102	0.02	0.03	4684
		liccheck	1352	4	1	1272	0.04	0.07	1584
		DLLHOST.	1284	4	1	2158	0.09	0.15	6660
		inetinfo	1208	4	1	3063	0.10	0.16	9708
		WinVNC.e	1160	4	1	20742	0.56	0.92	3536
		explorer	788	4	1	2252	0.14	0.23	5336
		SERVICES	272	4	1	6892	1.50	2.47	7480
		msdtc.ex	164	4	1	71	0.02	0.03	5108

“Grouped” Linux user cpu by process name

Report: ESAHSTA LINUX HOST Application Report
 Monitor initialized: 21/01/11 at 07:03:00 on

```

-----
Node/      Process/      <Application Status Counts> <-----Processor----->
Date      Application      Run-   Res   Load  <---Utilization--->
Time      name            Total Actv ning   Wait   -ed   Percent seconds Avg
-----
07:04:00
***Node Groups***
TheUsers *Totals*      840.0  138  11.0   829    0    88.0   52.7  0.1
      automoun      1.0   1.0    0    1.0    0    0.0    0.0  0.0
      events/0      1.0   1.0    0    1.0    0    0.0    0.0  0.0
      httpd        277.0  106   1.0   276    0   86.0   51.5  0.3
      java          2.0   2.0    0    2.0    0    0.0    0.0  0.0
      ksoftirq      3.0   1.0    0    3.0    0    0.0    0.0  0.0
      rotatelo     72.0  14.0    0   72.0    0    1.0    0.6  0.0
      sendmail      6.0   3.0    0    6.0    0    0.0    0.0  0.0
      sidd          1.0   1.0    0    1.0    0    0.2    0.1  0.2
      snmpd        9.0   9.0   9.0     0    0    0.7    0.4  0.1
  
```

Note .1% per snmp

Report: ESA**UCD2** LINUX UCD Memory Analysis Report Linux Test
Monitor initialized: 02/05/07 at 10:41:41 on 2084 serial 55BAF First recor

```
-----Storage Sizes (in MegaBytes)-----
<---Real Storage--> <-----SWAP Storage-----> Total <-----Storage in Use--
Node/      Total Avail Used Total Avail Used MIN Avail Shared Buffer Cache
Date
-----
```

Node/	Total	Avail	Used	Total	Avail	Used	MIN	Avail	Shared	Buffer	Cache
10:43:00											
acme	494.7	7.7	487.0	2031	2018	12.8	15.6	2026	0	14.2	39.1
dominoz1	2002.1	8.0	1994	2031	1934	97.4	15.6	1942	0	38.6	1417
ebiz1	997.1	5.7	991.4	2031	1517	513.7	15.6	1523	0	8.9	635.8
ebiz2	997.1	13.0	984.2	2031	1878	152.8	15.6	1891	0	26.9	607.8
ibmds1	2002.1	11.6	1990	2031	2029	2.0	15.6	2041	0	84.0	1484
ebizdev2	997.1	6.8	990.4	2031	1980	51.3	15.6	1986	0	63.3	530.9
ebizdev1	997.1	8.0	989.2	2031	1754	277.3	15.6	1762	0	43.8	521.2
ibmedge1	1007.3	497.1	510.2	2031	2031	0	15.6	2528	0	174.9	165.4
ibmds3	8031.8	81.5	7950	2031	2031	0	15.6	2112	0	320.3	6494
ibmedge2	1007.3	492.7	514.6	2031	2031	0	15.6	2524	0	175.3	167.4
ibmred2	997.1	4.5	992.6	2031	2026	4.6	15.6	2031	0	98.4	586.4
ibmred1	997.1	9.7	987.4	2031	2026	4.6	15.6	2036	0	98.7	578.5
tdirdb2	4012.0	31.9	3980	2031	1613	418.1	15.6	1645	0	250.1	3017
tdirtam	4012.0	1294	2718	2031	2031	0	15.6	3325	0	235.1	2106
tdirtds	4012.0	1061	2951	2031	2031	0	15.6	3092	0	324.8	2259
tdirtim	4012.0	1007	3005	2031	2031	0	15.6	3038	0	239.7	1981
tdsds-a1	997.1	124.0	873.1	2031	2031	0	15.6	2155	0	87.1	569.0
ibmds2	8031.8	78.0	7954	2031	2031	0.4	15.6	2109	0	251.7	6546

Linux data shows
Real storage
Swap storage
“cache”

Swapping is “good”

If not swapping,
reduce vm size
Dynamic add stg
if needed

- z/VM new releases supported day 1 (Note stg size)**

Report: ESASTR1

Velocity Software Corporate

```

-----
          Users <-----Pages-----Over
          Loggd System <Available> System  User  <-AddSpace> VDISK Commit  Capt-
Time          On Storage<2gb  >2gb  ExSpc Resdnt  System User  Rsdnt Ratio  Ratio
-----
15:29:00      69 138412K   84 1258K 10296   134M 1641K    0    0 1.915   1.000
15:30:02      69 138412K  109 1649K 10194   134M 1699K    0    0 1.915   1.000
15:31:00      69 138412K  178 1540K 10059   134M 1732K    0    7 1.915   1.000
15:32:00      67 138412K  838 2974K  9869 89040K 1764K    0   21 1.915   0.686
15:33:03      66 138412K 200K   46M  9717 58977K 1064K    0   26 1.915   0.776
15:34:08      66 138412K 390K   77M  9277 23615K  707K    0   22 1.915   0.743
15:36:01      65 138412K 486K  136M  8669 204895 81839    0   10 1.915   0.995
15:37:00      65 138412K 486K  136M  8426 205722 78891    0   10 1.915   0.996
15:38:00      62 138412K 486K  137M  8333 206309 14582    0    0 1.915   1.000
*****Summary*****
Average:      66 138412K 250K   65M  9392 57454K  912K    0   16 1.915   0.903
  
```

Benefit of using standard (snmp) interface?

- z/VM new releases supported day 1 (Note stg size)

Report: **ESAUCD2**

LINUX UCD Memory Analysis Report

Veloc

```

-----
Node/      <-----Storage Sizes (in MegaBytes)-----
Time/      <--Real Storage--> <-----SWAP Storage----> Total <----Storage i
Date       Total  Avail Used  Total Avail Used  MIN  Avail CMM  Buffer
-----
15:29:00
ZLNXT030   994.8  407.7  587.1  256.1  256.1      0  15.6  663.8      0  21.8
ZLNXT006   494.7  388.8  105.9  511.5  511.5      0  15.6  900.3      0  19.3
ZLNXT017  3008.7  2612  396.9  1279  1279      0  15.6  3891      0  29.5
ZLNXT002  2001.3  902.4  1099  512.0  512.0      0  15.6  1414      0  53.9
ZLNXT007  201192  96151  103K  1023  1023      0  15.6  97174      0  18.1
ZLNXT009  201192  165K  32356  640.1  640.1      0  15.6  166K      0  19.9
ZLNXT013  201192  171K  26563   7.9   7.9      0  15.6  171K      0  22.2
ZLNXT010  201192  181K  15917  1535  1535      0  15.6  182K      0  24.3
ZLNXT011  201192  194K   2280  128.9  128.9      0  15.6  194K      0  28.4

```

• High cpu capture ratio

Report: ESALNXV LINUX Virtual Processor Analysis Report

Node/ Name	VM ServerID	<Linux Pct CPU>			<Process Data>			Capture Ratio	Prorate Factor
		Total	Syst	User	Total	Syst	User		
10:03:00									
NEALE1	LNEALE1	100.0	11.4	88.6	100.2	11.5	88.7	1.002	1.000

Report: ESALNXP LINUX HOST Process Statistics Report

node/ Name	<-Process Ident->			Nice	<-----CPU Percents----->					
	ID	PPID	GRP	Valu	Tot	sys	user	sys	usr	
10:03:00										
NEALE1	0	0	0	0	100	0.43	3.35	11.0	85.4	
kswapd0	100	1	1	0	0.12	0.12	0	0	0	
snmpd	1013	1	1012	-10	0.13	0.03	0.10	0	0	
sh	3653	3652	30124	0	52.7	0	0	9.37	43.3	
gmake	9751	9750	30124	0	43.4	0.02	0.02	1.37	42.0	
sh	10129	9751	30124	0	0.02	0.02	0	0	0	
sh	10130	10129	30124	0	0.63	0.03	0.23	0.28	0.08	
cc1	10307	10306	30124	0	3.12	0.18	2.93	0	0	
rpmbuild	30124	16382	30124	0	0.07	0.03	0.03	0	0	
sh	30125	30124	30124	0	0.02	0	0.02	0	0	
gmake	30126	30125	30124	0	0.02	0	0.02	0	0	

Report: ESALNXC LINUX Process Conf

Node/ Name	<-Process Ident->			<-----Pr Path
	ID	PPID	GRP	
NEALE1				
init	1	0	0	init [3]
migratio	2	1	0	migratio
ksoftirq	3	1	0	ksoftirq
events/0	4	1	0	events/0
khelper	5	4	0	khelper
kblockd/	6	4	0	kblockd/
cio	41	4	0	cio
cio_noti	42	4	0	cio_noti
kslowcrw	43	4	0	kslowcrw
apldata	96	4	0	apldata
aio/0	101	4	0	aio/0
pdflush	5266	4	0	pdflush
pdflush	26647	4	0	pdflush
kswapd0	100	1	1	kswapd0
kmcheck	158	1	1	kmcheck
syslogd	976	1	976	/sbin/sy
klogd	979	1	979	/sbin/kl
snmpd	1013	1	1012	snmpd
portmap	1030	1	1030	/sbin/po
rpciod	1034	1	1	rpciod
lockd	1035	1	1	lockd
sshd	1072	1	1072	/usr/sbi
sshd	16272	1072	16272	sshd: bu
sshd	16288	1072	16288	sshd: bu
sshd	16290	16288	16288	sshd: bu
bash	16291	16290	16291	bash
python	16312	16291	16291	python
do-bui	16313	16312	16291	/bin/sh
bb_do	16382	16313	16291	/usr/bin
rpmb	16415	16382	16415	rpmbuild
rpmb	30124	16382	30124	rpmbuild

- **Standard mib**

- MIB II: 1.3.6.1.2.1 (Network)
- HOST 1.3.6.1.2.1.25 (process, file system device, memory)

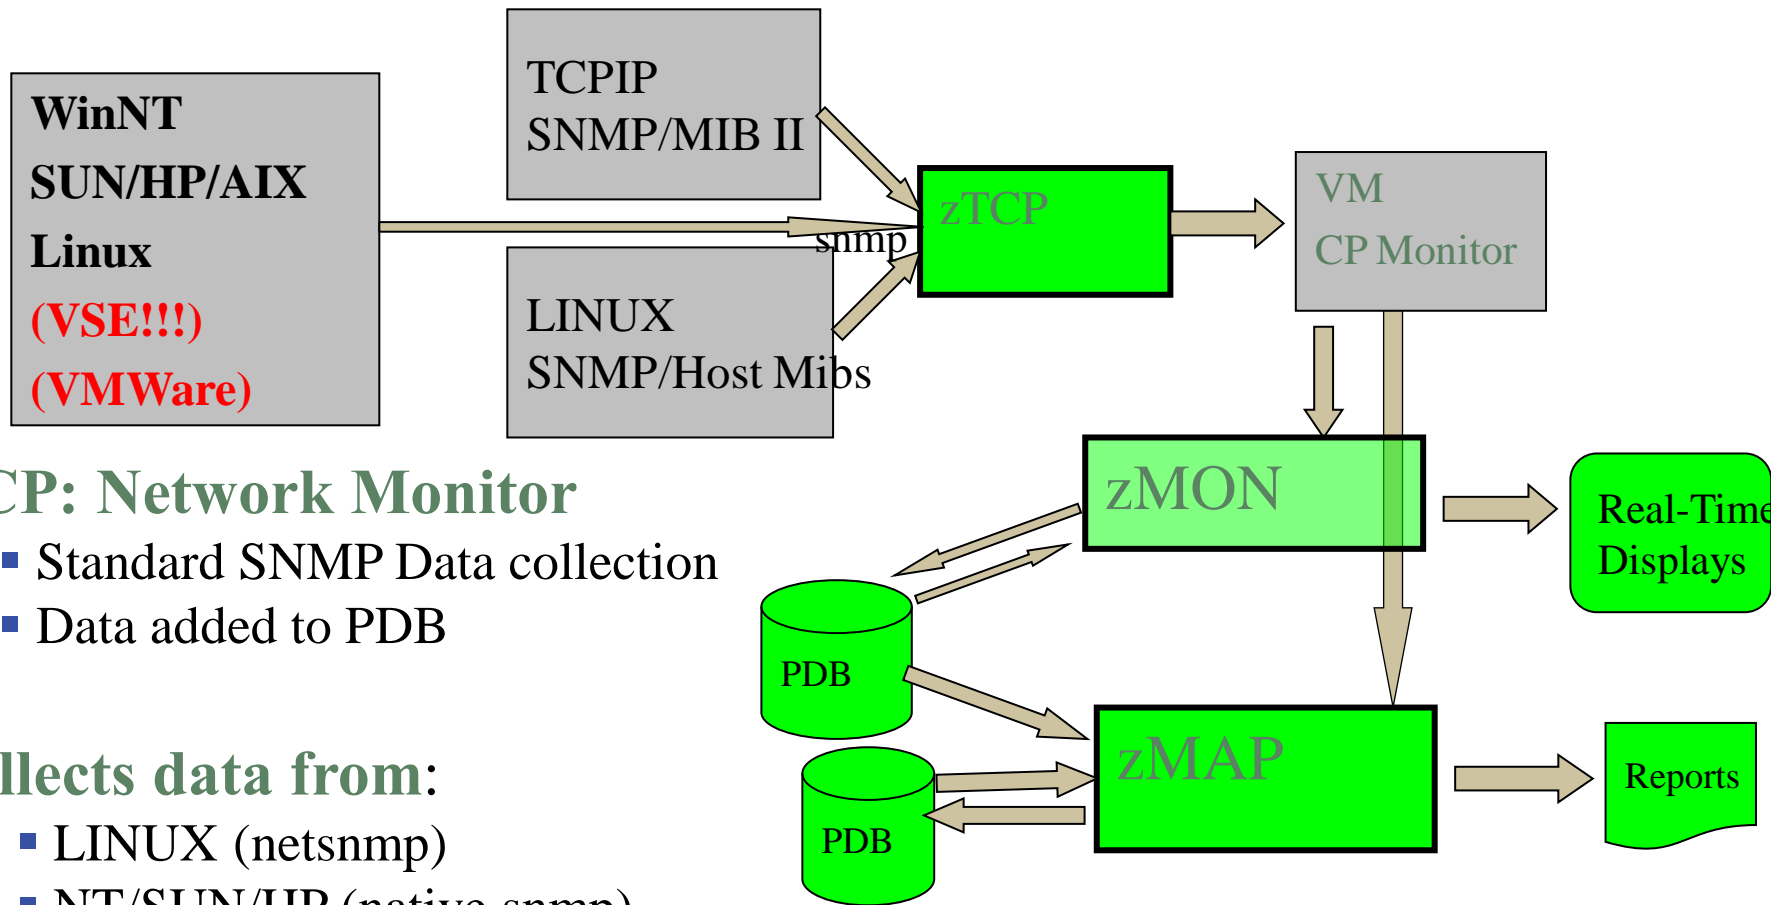
- **Private mibs:**

- Private: 1.3.6.1.4
- **ucd-snmp 1.3.6.1.4.1.2021.**
- **Velocity 1.3.6.1.4.1.F971**
- **VeloJava 1.3.6.1.4.1.F971.100.**
- **VeloOracle 1.3.6.1.4.1.F971.11**
- **VeloVSE 1.3.6.1.4.1.F971.10.1**
- **C21**

- **IBM 1.3.6.1.4.1.2**
- **IBMVSE 1.3.6.1.4.1.2.6.81FD**

- **Why Velocity mib?**

- Performance management



zTCP: Network Monitor

- Standard SNMP Data collection
- Data added to PDB

Collects data from:

- LINUX (netsnmp)
- NT/SUN/HP (native snmp)
- Printers/Routers....
- **THOUSANDS OF SERVERS?**

Support enterprise operations console

Issue with SNMP alerts

- How many control points? (one per server?)
- How many configuration files? (one plus per server?)

ZVPS SNMP Alert Architecture

- Centralized alert
- One point of control (ZALERT)

SNMP alerts sent to any SNMP operations console

- Create “SNMP TRAPDEST” file
- * format is ip address, and community name
67.100.74.25 velocity

Sending SNMP alerts by other functions:

```
/* authorized user can send alerts */  
parse arg msg  
'CP SMSG ZTCP ALERT' msg
```

- **Valid and Correct?**

- **Process data from “virtualized” Linux was wrong**

- Compare VERY accurate VM performance data to Linux data, easy to see
- All process accounting based on timer ticks

- Sample wrong by factor of 10-100 prior to SLES10

- Known issue since 2001
- [HTTP://velocitysoftware.com/present/CaseAFS](http://velocitysoftware.com/present/CaseAFS)
- Mostly corrected in SLES10, RHEL5 (now underreports) by “steal timer”

- **Leads to solving performance problems?**

- z/VM owns the shared resources
- **“Native” tools will not detect many problems**
- **Native “tools” are mostly for after the fact diagnostics**

Analyzing Linux CPU by process

Velocity MIB data:
Provides process data
Parent/Child relationship

Note ALL application
processes are owned by
“24445”.

```
Report: ESALNXP          LINUX HOST Process Statistics Report
Monitor initialized: 02/05/07 at 10:41:41 on 2084 serial 5
-----
```

node/ Name	<-Process ID	Ident->	Nice	<-----CPU Percents----->					
	ID	PPID	GRP	Valu	Tot	sys	user	syst	usr
10:43:00									
dominoz1	0	0	0	0	9.9	3.20	6.69	0	0
ksoftirq	5	1	0	19	0.03	0.03	0	0	0
ksoftirq	7	1	0	19	0.05	0.05	0	0	0
kswapd0	134	1	1	0	0.05	0.05	0	0	0
kjournal	1140	1	1	0	0.08	0.08	0	0	0
snmpd	1775	1	1774	-10	0.27	0.16	0.11	0	0
scontrol	24521	24445	24414	0	0.03	0	0.03	0	0
server	24539	24521	24414	0	1.46	0.41	1.06	0	0
logasio	24553	24539	24414	0	0.14	0.11	0.03	0	0
event	28636	24539	24414	0	0.16	0.03	0.14	0	0
replica	28663	24539	24414	0	1.76	0.27	1.49	0	0
update	28665	24539	24414	0	5.36	1.92	3.44	0	0
amgr	28667	24539	24414	0	0.03	0	0.03	0	0
adminp	28670	24539	24414	0	0.19	0.08	0.11	0	0
sched	28676	24539	24414	0	0.03	0	0.03	0	0
rnrmgr	28686	24539	24414	0	0.03	0	0.03	0	0
clrepl	28920	24539	24414	0	0.22	0	0.22	0	0

Analyzing Linux CPU by Application

Velocity MIB data:

- Provides process data
- Parent/Child relationship
- Allows combining into “applications”
- Note the “bash/24445” “application”

Define alerts based on application



```

Report: ESALNXA          LINUX HOST Application Report
Monitor initialized: 02/05/07 at 10:41:41 on 2084 ser
-----
Node/      Process/      ID    <---Processor Percent--->
Date       Application
Time       name          Total sys  user syst usrt
-----
10:43:00
dominoz1  *Totals*      0     9.9  3.2  6.7   0   0
          bash      24445  9.4  2.8  6.6   0   0
          kernel    1     0.2  0.2   0     0   0
          snmpd     1775  0.3  0.2  0.1   0   0
    
```

Velocity MIB data:

Provides process data

Parent/Child relationship

And reporting by Linux userid

Allows alerts by userid

Report: ESALNXU **LINUX USER** Analysis Report

Monitor initialized: 02/05/07 at 10:41:41

```

-----
Node/                               <---Processor Percent--->
Date   <-----User and Group Identity----->   <Process><Children>
Time   Userid      GroupID      usrpid  grppid  Total  sys  user  syst  usrt
-----
10:43:00
dominoz1 bin          root          1        0        0        0        0        0        0
         daemon      daemon        2        2        0        0        0        0        0
         lp           lp           4        7        0        0        0        0        0
         notes       notes        1001     1001     9.4     2.8     6.6     0        0
         root        root         0        0        0.5     0.4     0.1     0        0
  
```

- Many installations lack z/VM and Linux on z/VM tuning skills
- Velocity Software's objective is to ensure our customer performance problems are resolved – quickly.
- zTUNE includes configuration guidance, health checks when ever installation requests, and assistance in all areas of Linux on z/VM and z/VM performance
- **no more “performance was unexplainably bad so we abandoned the project”**

- **Focus more now on simplifying problem resolution**
- **User reports that applications complained about zLinux / WAS performance:**

```
Report: ESATUNE          Tuning Recommendation Report
Monitor initialized:      on 2084 serial 9ABED
-----
```

```
The following changes are suggestions by Velocity Software
to enhance performance of this system.
```

```
However, Velocity Software takes no responsibility -
all tuning is the responsibility of the installations.
Please call 650-964-8867 if you have any questions about
these values, or suggestions on report enhancements.
```

```
USR2 User LINUX160 is paging excessively (75.0 per second)
This user can be protected using SET RESERVED
```

```
SPL5 Spool utilization is 100% full.
Perform Spool file analysis and purge large
spool files, or force users currently writing
excessively to spool.
```

```
*****zTUNE Evaluation *****
```

```
XAC1 User total PROCESSOR WAIT excessive at 33 percent.
Current reporting threshold set to 20.
This is percent of inqueue time waiting for
specific (PROCESSOR)resources to become available.
```

```
LPR3 LPAR share is too low, causing USER CPU Wait
VM LPAR allocated share: 0.94 percent of total
VM LPAR used 389 percent of allocated share
```

- **Performance Management**

- VMWare (node grouping)
- Mixed mode capacity planning (IFL vs CP)
- Linux DISK reporting (ESAUCDD)
- Granularity in virtual cpu reporting (ESAUSCP, ESALNXS)
- VSE Initial support (ESAVSES, ESAVSEC)
- LGR support
- zVIEW V2
- zALERT

- **Operational support**

- Global Installer
- Portal

- **Other**

- z/VM 5.4 support, up to 96 CPUs per LPAR, 200 physical CPUs / CEC
- Framework for future products and enhancements
- **Customer enhancement requests - MANY**

- **Performance Management**
 - **Application support** (JAVA, Oracle)
 - VSE partition support, job support (ESAVSEP)
 - Linux process metrics for RAM, I/O, Swap (ESALNXP, ESALNXI)
 - Linux system metrics for ram (ESALNXR)
 - **MFC Support** for z114, z196, EC12/BC12, Z13/Z13S
 - CP Pooling support
 - LINMON support
 - APPLE sever support (decimal process ID up to 99,999)
- **Operational support**
 - Recognize lpar, vmid for linux servers, LGR support
 - Peer to peer support
 - SNMP V3
- **Other**
 - z/VM 6.2, z/VM 6.3 Support

- **Performance Management**
 - Java Thread support (ESAJVMT)
 - HiperPav Support (ESAHPP)
 - SMT Support (ESASMT, ESAUSR5)
 - Diagnose support (ESADIAG)
 - OSA Support (ESAOSA)
- **Operational support**
 - Move linux nodes to correct LPAR (requires VSI mib) (ESALNXV)
 - DNS Support for zTCP
- **Other**
 - z/VM 6.4 support
 - **Many “small” Customer enhancement requests**

- **Performance Management**
 - User Diagnose support (ESAUSRD)
- **Operational support**
 - Enterprise server inventory
 - UBUNTU Support with snmpv3
 - Class C subnet node discovery
- **Other**
 - Z14 Support (model numbers, MFC)
 - Specter apar / status recognition

- **Added DNS Names capability**

```
community = 'velocity'   TCPIP='TCPIP'   nodegrp = 'VSILPARs'
```

```
dnsport = 53
```

```
dnsIPADDR = '64.105.172.26'
```

```
node = 'VSIVM1' domain='vsivm1.VelocitySoftware.com'
```

```
node = 'VSIVM2' domain='vsivm2.VelocitySoftware.com'
```

```
node = 'VSIVM3' domain='vsivm3.VelocitySoftware.com'
```

```
node = 'VSIVM4' domain='demo.VelocitySoftware.com'
```

- **Added SSI Support – monitor where operating**

```
TCPIP='TCPIP'   peerport = 1998
```

```
peeraddr = '67.218.99.132' peerport = 1998   ;vsivm2
```

```
peeraddr = '67.218.99.134' peerport = 1998   ;vsivm4
```

```
peeraddr = '67.218.99.135' peerport = 1998   ;vsivm5
```

```
Ssiflag = '1'b   ;Following can be LGR'd, ztcp will move
```

```
Node = 'lnxssi1' domain='prod.mylinux.mycompany.com'
```

msg ztcp query peers

Ready; T=0.01/0.01 18:14:46

	,TCPIP	184.105.60.11	1998,	0:00:00	.
VSIVM2	,TCPIP	184.105.60.12	1998,	18:14:00	Y. 5131
VSIVM4	,TCPIP	184.105.60.14	1998,	18:14:00	Y. 5130
VSIVM1	,TCPIP2	192.168.5.41	1998,	18:14:00	N. 5126
VSIVM2	,TCPIP2	192.168.5.42	1998,	18:14:00	Y. 5131
VSIVM4	,TCPIP2	192.168.5.44	1998,	18:14:00	Y. 5130
VSIVM5	,TCPIP2	192.168.5.45	1998,	18:14:00	N. 5130
	,TCPIP2	192.168.5.46	1998,	0:00:00	N.
VSIVM3	,TCPIP2	192.168.5.43	1998,	18:14:00	N. 5123

End Display

**Objective is for an enterprise communications model
“peers” will transfer data collection to local ztcp**

Concept is to centralize server inventory

- Zwrite parameter: `zpro='1'b` ; creates dcss area
- Local zTCP will then be “manager”
- Multiple managers ok
- zTCP peers communicate

Node “manager” keeps track of all nodes in enterprise

- `“Smsg ztcp query zpro”` shows where nodes run
- (Can be very long list)

Tailorable, expandable, zoomable

- (See "demo.VelocitySoftware.com")

Today is Monday 2 Dec 2013 zVIEW Version 4159

VELOCITY SOFTWARE zVIEW
Enterprise View - Velocity Software - VSIVM4 (DEMO)

First level

VSIVM1		VSIVM2		VSIVM3(old)	
VM1	13/12/02 18:29 CP Total (2) 6.63%	VM2	13/12/02 18:29 IFL Total (1) 0.91%	VM3	13/12/02 21:29 024B42-0 99.22%
Linux Nodes (Distributed Servers)		Linux Nodes (zVM-Guests)		Linux Nodes (zVM-Guests)	
LINUX9 (9)	3.93%	RH5X161	0.43%	000000-64	99.22%
suselnx3 (9)	2.57%	RH5Z161	0.37%		
REDHAT (2)	2.30%				

Demo System V4	
Demo	13/12/02 18:29 IFL Total (1) 17.77%
Linux Nodes (zVM-Guests)	
rob1x1	2.83%
redhat6	1.18%
oracle	0.82%
redhat56	0.47%
redhat5x	0.43%
lxsugar (2)	0.41%
redhat64	0.31%
sles8 (2)	0.31%
sles10	0.29%
redhat5	0.27%
redhat3	0.25%
redhat6x	0.24%
suselnx2	0.22%
sles11 (2)	0.22%
sles11x	0.20%
sles11x3	0.19%
sles9x	0.18%
scsil0s	0.17%
sles10x4	0.17%
sles9	0.16%
Linux Nodes (Distributed Servers)	
linux93 (2)	100.00%
opensuse (2)	8.97%
JIRA (2)	5.88%
vpnbrz	5.50%
vpnbrc	4.76%
mail (9)	3.42%
vpnz	2.35%

Second level

Tims Test System	
Tim1_2	13/11/27 13:09 IFL Total (1) 0.10%
Linux Nodes (zVM-Guests)	
	1.85%
	1.50%
	0.85%
	0.57%

Single pane of glass (did I say “scalable”?, tailorable?)

- Data from “Many” multiple LPARs(50) / geographies(3)

“some installation”

The screenshot displays the 'Enterprise Performance Summary - Automatic Data Processing, Inc(VLB6)' interface. It features a grid of system performance data organized by Data Center (DC1, DC2) and CDL (CDL). Each system entry includes a name (e.g., V1P1, V1N1, P107), a time stamp (08:48), 'IPL Total' count, and a percentage value. Each entry has an 'Expand' button. A callout box labeled '“some installation”' points to a specific entry in the grid.

System Name	Time	IPL Total	Percentage	Action
V1P1	08:48	Total (48)	276.4%	Expand
V1N1	08:48	Total (18)	317.16%	Expand
P107	08:48	Total (40)	3016.40%	Expand
P113	08:48	Total (24)	353.13%	Expand
V2P1	08:48	Total (48)	296.5%	Expand
V2P5	08:48	Total (40)	397.3%	Expand
P209	08:48	Total (56)	1372.48%	Expand
P213	08:47	Total (40)	1173.87%	Expand
P217	08:48	Total (40)	375.83%	Expand
C203	08:48	Total (32)	462.1%	Expand
C207	08:48	Total (24)	649.58%	Expand
V2N3	08:48	Total (20)	499.91%	Expand
VLB1	08:48	Total (52)	2948.84%	Expand
VLB5	08:48	Total (48)	646.2%	Expand
ZS02	08:48	Total (16)	9.82%	Expand

Very fast access

- VSE
- z/OS
- Linux

Cloud 1

VSIVC1		14:08	IFL Total (4) ⊕	4.21%
zOS Systems				
V25A		4.60%	<div style="width: 4.60%; height: 10px; background-color: #008000;"></div>	
V25A		0.02%	<div style="width: 0.02%; height: 10px; background-color: #008000;"></div>	
VSE Systems				
zvse61c		1.17%	<div style="width: 1.17%; height: 10px; background-color: #008000;"></div>	
zvse61b		1.07%	<div style="width: 1.07%; height: 10px; background-color: #008000;"></div>	
zvse62c		0.89%	<div style="width: 0.89%; height: 10px; background-color: #008000;"></div>	
zvse62b (2)		0.70%	<div style="width: 0.70%; height: 10px; background-color: #008000;"></div>	
Top 15 Linux Nodes(z/VM-Guests)				
⊕ MONG505A (1)		0.47%	<div style="width: 0.47%; height: 10px; background-color: #008000;"></div>	
⊕ VSIEXTRN (1)		0.31%	<div style="width: 0.31%; height: 10px; background-color: #008000;"></div>	
⊕ RHKSNFS1 (1)		0.22%	<div style="width: 0.22%; height: 10px; background-color: #008000;"></div>	
⊕ JSVEXTRN (1)		0.15%	<div style="width: 0.15%; height: 10px; background-color: #008000;"></div>	
⊕ JSVSVR13 (1)		0.08%	<div style="width: 0.08%; height: 10px; background-color: #008000;"></div>	
⊕ S15PSTG1 (1)		0.06%	<div style="width: 0.06%; height: 10px; background-color: #008000;"></div>	
⊕ SLFSRV10 (1)		0.05%	<div style="width: 0.05%; height: 10px; background-color: #008000;"></div>	
⊕ JSVSVR10 (1)		0.04%	<div style="width: 0.04%; height: 10px; background-color: #008000;"></div>	
⊕ JSVWRK01 (1)		0.04%	<div style="width: 0.04%; height: 10px; background-color: #008000;"></div>	
⊕ CBSVR010 (1)		0.03%	<div style="width: 0.03%; height: 10px; background-color: #008000;"></div>	
⊕ JSVSVR12 (2)		0.03%	<div style="width: 0.03%; height: 10px; background-color: #008000;"></div>	
⊕ RS327001 (1)		0.03%	<div style="width: 0.03%; height: 10px; background-color: #008000;"></div>	
⊕ DSVSVR01 (1)		0.02%	<div style="width: 0.02%; height: 10px; background-color: #008000;"></div>	
⊕ GOLDDVM71 (1)		0.02%	<div style="width: 0.02%; height: 10px; background-color: #008000;"></div>	
⊕ JSVSVR20 (1)		0.02%	<div style="width: 0.02%; height: 10px; background-color: #008000;"></div>	
Remaining 1 servers		0.02%	<div style="width: 0.02%; height: 10px; background-color: #008000;"></div>	
Top 5 Users				
ZALERT		0.66%	<div style="width: 0.66%; height: 10px; background-color: #008000;"></div>	
ZVWS		0.49%	<div style="width: 0.49%; height: 10px; background-color: #008000;"></div>	

Close

End users define their environment(s)

- Linux administrators get most everything in one click
- Secure, no need for logon
- Fast and efficient

Wednesday 7 Nov 2013 00:46 zVIEW Version 4310

VELOCITY SOFTWARE
zVIEW - Velocity Software - VSIVM4 (DEMO)
Performance Displays for zVM and Linux on System z

mylinux

Node	Process Name	ID	P
ZSXL0006	systemd	1	
ZSXL0006	kthread	2	
ZSXL0006	kworker/0:0	3	
ZSXL0006	kworker/0:0H	4	
ZSXL0006	mm_percpu_wq	6	
ZSXL0006	ksoftirqd/0	7	
ZSXL0006	rcu_sched	8	
ZSXL0006	rcu_bh	9	
ZSXL0006	migration/0	10	
ZSXL0006	cpuhp/0	11	
ZSXL0006	kdevtmpfs	12	
ZSXL0006	netns	13	
ZSXL0006	khungtaskd	14	
ZSXL0006	oom_reaper	15	
ZSXL0006	writeback	16	
ZSXL0006	kcompactd	17	
ZSXL0006	ksmd	18	
ZSXL0006	crypto	19	
ZSXL0006	kintegrityd	20	
ZSXL0006	kblockd	21	
ZSXL0006	md	22	
ZSXL0006	clio	23	
ZSXL0006	watchdogd	24	
ZSXL0006	kworker/0:1	26	
ZSXL0006	cmthread	27	
ZSXL0006	kauditd	28	
ZSXL0006	kswapd	29	
ZSXL0006	ecryptfs-kthrea	30	
ZSXL0006	kthrotld	72	
ZSXL0006	khvcd	73	
ZSXL0006	kmcheck	74	
ZSXL0006	ipvs_addrconf	75	
ZSXL0006	kworker/0:1H	148	
ZSXL0006	kworker/u128:3	150	
ZSXL0006	jbd2/dasda1-8	172	
ZSXL0006	ext4-rsv-conver	173	
ZSXL0006	vfio-ccw	301	
ZSXL0006	geth_wq	322	
ZSXL0006	kworker/u128:0	7826	
ZSXL0006	systemd-journal	212	
ZSXL0006	systemd-udev	235	
ZSXL0006	systemd-timesyn	272	
ZSXL0006	cron	353	

ESALNXP - VSI Linux Percent Usage by Process - DEMO

Time	Node	Name	ID	PPID	GRP	Tot	sys	user	syst	usr	nice	prty	Size	RSS	Peak	Swap	Data	Stk	EXEC
00:46:00	lxdb2001	*Totals*	0	0	0	0.6	0.1	0.1	0.1	0.3	0	0	4549	322	4557	0	1391	4.8	3.8
00:46:00	lxdb2001	init	1	1	1	0.0	0.0	0	0	0	0	20	2.4	0.9	2.4	0	0.2	0.1	0.0
00:46:00	lxdb2001	snmpd	2200	1	2199	0.1	0.1	0.1	0.1	0	-10	10	29.7	13.4	37.1	0	17.3	0.1	0.0
00:46:00	lxdb2001	cron	2223	1	2223	0.1	0	0	0.0	0.0	0	20	2.6	0.9	2.7	0	0.2	0.1	0.0
00:46:00	lxdb2001	db2fmc	2245	1	2245	0.4	0	0	0.1	0.3	0	20	50.9	13.9	51.0	0	3.5	0.2	0.1
00:46:00	lxora12	*Totals*	2833	2831	2833	0.0	0.0	0	0	0	0	20	377	91.6	377	0	262	0.1	0.1
00:46:00	lxora12	amozxm	1503	1	1503	0.0	0	0	0.0	0.0	0	20	250	10.1	314	0.9	66.3	0.1	0.4

ESAHST2 - LINUX HOST Storage Analysis Report - DEMO

Time	Node/Group	Index	Size	Used	Full	Err	Alloc	Units	R/W	Boot	Storage	Description
00:46:00	ZPRO	0	196K	109K	55.7	0	1K				Totals	
00:46:00	VPNS	0	5376	5376	100	0	1K				Totals	

ESAUCD2 - LINUX UCD Memory Analysis Report - DEMO

Time	Node/Group	Total	Avail	Used	Total	Avail	Used	MIN	Avail	CMM	Buffer	Cache
00:46:00	ZPRO	4600	3	1423	3166	0	3075	3076	40	6	107	5
00:46:00	VPNS	4600	3	1423	3166	0	3075	3076	40	6	107	5

ESAUCD4 - LINUX UCD System Statistics Report - DEMO

Time	Node/Group	Processor	Total	Syst	User	Nice	Pct	In	Out	In	Out	Rate	Intrpt	Load	Min	S
00:46:00	ZPRO		2.7	1.2	1.4	0	1188	0	0	0	0	56.7	2080.5	1023.7	0.49	0
00:46:00	VPNS		10.1	4.2	5.9	0	389	0	0	0	0	180.5	733.9	0.33	0	0

ESAHST4 - LINUX HOST System Statistics Report - DEMO

Time	Server	Num	Users	Current	Max	StgSz	Local	Date	Time	System	Uptime	Dev	Parameter	Initiali
00:46:00	ZPRO													
00:46:00	VPNS													

LPAR

IFL Utilization

The screenshot shows the ZMON interface. On the left is a sidebar menu with categories: zMON, Graphs, zMAP, System, Service Level Analysis, and User. The 'User' category is selected, showing a list of users including ESAUSR1 through ESAUSPG. On the right is a window titled 'ESAUSPG' displaying a 'User Storage Analysis' table. The table has columns for Time, UserID / Class, Total, >2GB, <2GB, Xstor, DASH, Xstor, Disk, and Migration. The row for 'REDHAT5X' is highlighted in green, and an arrow points to it from the 'User' menu.

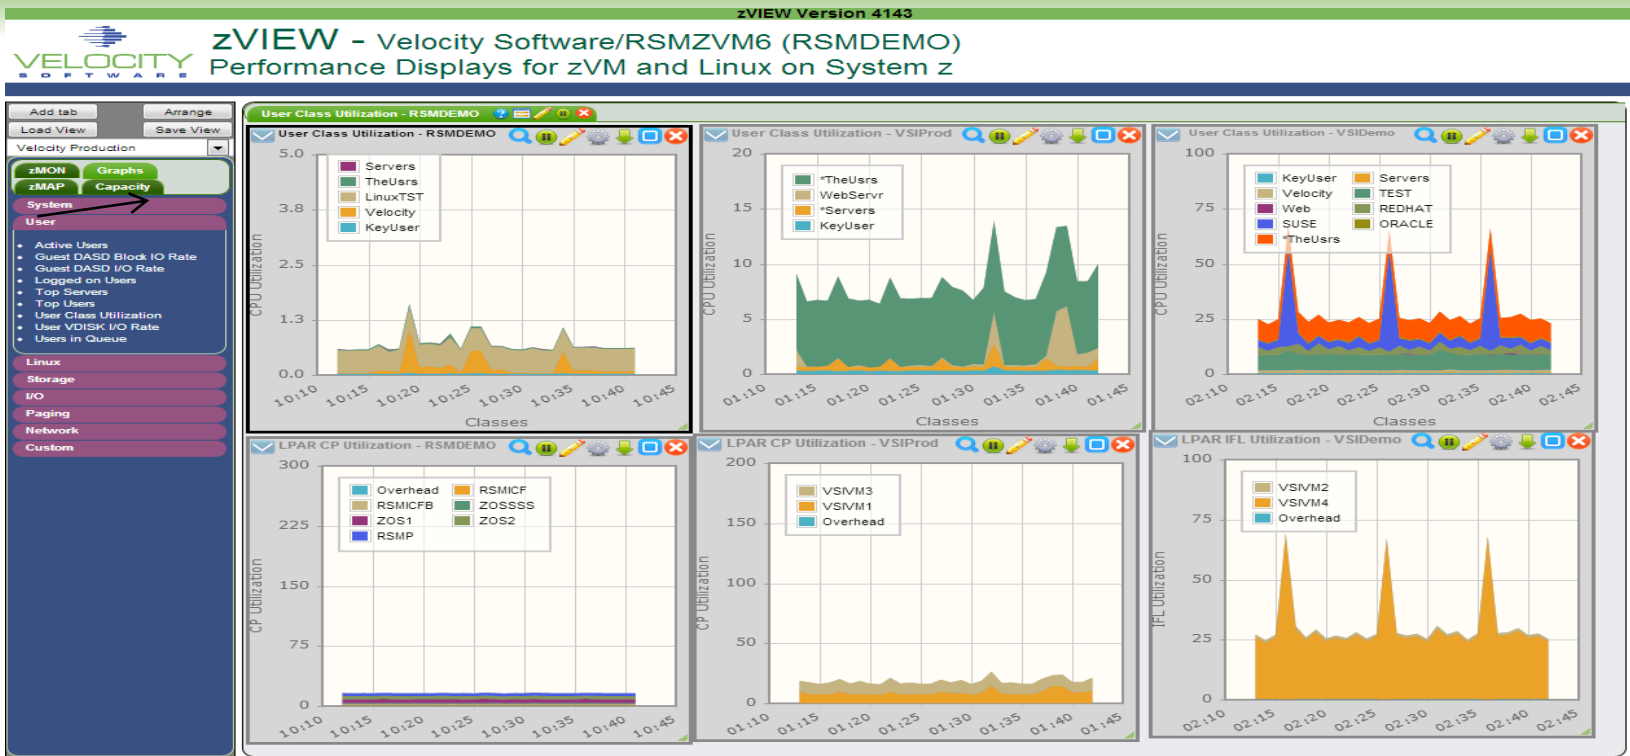
Time	UserID /Class	Total	>2GB	<2GB	Xstor	DASH	Xstor	Disk	Migration
17:10:00	System:	664879	197480	467399	747999	2609K	28	0	
17:10:00	*TheUsrs	41674	12525	29149	22170	199418	0	0	
17:10:00	KeyUser	3430	1901	1529	349	8276	0	0	
17:10:00	ORACLE	34842	11904	22938	6711	188759	0	0	
17:10:00	REDHAT	258455	78708	179747	536580	592529	0	0	
17:10:00	REDHAT5X	87333	33358	53975	485474	31158	0	0	
17:10:00	REDHAT5	46665	12525	34140	5737	108832	0	0	
17:10:00	REDHAT6	19821	5939	13882	23266	105537	0	0	
17:10:00	Servers	1210	810	400	1978	30403	0	0	
17:10:00	SUSE	176464	48152	128312	51280	864768	0	0	
17:10:00	TEST	142604	39842	102762	125496	693779	0	0	
17:10:00	Velocity	4105	2268	1837	1591	23659	0	0	
17:10:00	Web	2095	1370	725	1844	7372	28	0	

Click on “user” to see user screens

Click on “redhat” class to see “redhat users”


Oracle data available 4.2

Multiple System View (3 LPARs)



Oracle data from multiple lpar's visible on one tab

zMAP Capacity/Trend Graphs



- **Alerts**
 - User tailorable
 - 3270 based, web based, and / or SNMP
 - Alerts can be set on any variable or calculated variable
- **Linux alert examples:**
 - Disk full
 - Missing processes (requires complete data)
 - **Looping processes (requires correct data)**
- **z/VM alert examples**
 - Page/spool space full (avoid abends)
 - Looping servers
 - DASD service times
- **Network alert examples**
 - Transport errors
 - ICMP rates
 - Bandwidth thresholds

3270 Style Alerts (50+ sample alerts provided)

```

Screen: LINALERT          Velocity Software          25 Mar 2015 06:42:29
----- Exceptions Analysis Alerts -----

Type Description
LNDX / area on oracle is 79.51% full
LNDX /opt area on oracle is 82.24% full
LNDX /home area on oracle is 59.02% full
LNDX / area on RH5X161 is 32.54% full
LNDX / area on S11R20RA is 81.56% full
LNDX /boot area on S11R20RA is 2
LNDX /opt area on S11R20RA is 95
LNDX /mnt/oracle area on S11R20R
LNSU Swap utilization for Linux
LNSU Swap utilization for Linux
    
```

Today is Wednesday 25 Mar 2015 zVIEW Version 4174
zVIEW - Velocity Software - VSIVM4 (DEMO)
Performance Displays for zVM and Linux on System z

LINALERT - Exceptions Analysis Alerts - 15/03/25 at 06:47 - DEMO

Code	Alert Description
LNDX	CPU utilization on Linux node BlakeMC is 13.86%
LNDX	/ area on lxsugar is 90.74% full
LNDX	/usr area on lxsugar is 57.59% full
LNDX	/ area on opensuse is 39.71% full
LNDX	/home area on opensuse is 53.23% full
LNDX	/iso/sles11s area on opensuse is 100.00% full
LNDX	/iso/s11sp2- area on opensuse is 100.00% full
LNDX	/iso/s11sp2- area on opensuse is 100.00% full
LNDX	/iso/s11sp3- area on opensuse is 100.00% full
LNDX	/iso/s11sdk- area on opensuse is 100.00% full
LNDX	/iso/s10sp2 area on opensuse is 100.00% full
LNDX	/iso/r64 area on opensuse is 100.00% full
LNDX	/iso/r62 area on opensuse is 100.00% full
LNDX	/iso/s10v1 area on opensuse is 100.00% full
LNDX	/iso/r7 area on opensuse is 100.00% full
LNDX	/iso/sles11s area on opensuse is 100.00% full
LNDX	/iso/s12-1 area on opensuse is 100.00% full
LNDX	/iso/s12-2 area on opensuse is 100.00% full
LNDX	/iso/s12sdk1 area on opensuse is 100.00% full
LNDX	/iso/s12sdk2 area on opensuse is 100.00% full
LNDX	/ area on oracle is 79.51% full
LNDX	/opt area on oracle is 82.24% full
LNDX	/home area on oracle is 59.02% full
LNDX	/ area on redhat5 is 52.26% full
LNDX	/ area on redhat5x is 32.54% full
LNDX	/ area on redhat56 is 95.80% full
LNDX	/mnt area on redhat56 is 53.23% full
LNDX	/ area on redhat6 is 30.00% full
LNDX	/ area on redhat6x is 94.92% full
LNDX	/dev/shm area on redhat6x is 51.42% full
LNDX	/ area on redhat64 is 36.09% full
LNDX	/boot area on rhel7v is 23.79% full
LNDX	/ area on roblnx2 is 78.74% full

Or Browser based
Click Thru
or SMS, email...

Several requests to extend ESAOPER screen

zOPERATOR:

- Optional **no-charge** application component of zMON
- Uses standard zMON 3270 screen driver, existing function
- Scrollable 3270 console
- Messages can be searched by text, date or time
- All messages logged in daily files
- Actions can be set based on messages received
- Can be viewed from DIALed terminal or **zVIEW (WEB!)**

Operator Function browser based

- Click Thru for problem analysis – LPCP example

Today is Wednesday 25 Mar 2015 zVIEW Version 4174

VELOCITY SOFTWARE zVIEW - Velocity Software - VSIVM4 (DEMO)
Performance Displays for zVM and Linux on System z

ZOPER - zOPERATOR Console - DEMO

```

05:57:10 ZALERT  LPCP LPAR VSIVM5 IS AT 99%
05:57:10 ZALERT  LNPR CPU UTIL FOR PROCESS java-6500 ON BlakeMC IS 14%
05:58:10 ZALERT  LPCP LPAR VSIVM5 IS AT 99%
05:58:10 ZALERT  LNPR CPU UTIL FOR PROCESS java-6500 ON BlakeMC IS 14%
05:59:10 ZALERT  LPCP LPAR VSIVM5 IS AT 100%
05:59:10 ZALERT  LNPR CPU UTIL FOR PROCESS java-6500 ON BlakeMC IS 14%
05:59:45 DXTSERVE Inactivity detected but VM system monitor is still active.
06:00:00 OPERATOR HCPMXE6224I Event recording is pending because there are no users connected to *MONITOR for this type of data.
06:00:00 OPERATOR HCPMXE6224I Sample recording is pending because there are no users connected to *MONITOR for this type of data.
06:00:10 ZALERT  LPCP LPAR VSIVM5 IS AT 100%
06:00:10 ZALERT  LNPR CPU UTIL FOR PROCESS java-6500 ON BlakeMC IS 14%
06:01:10 ZALERT  LPCP LPAR VSIVM5 IS AT 99%
06:01:10 ZALERT  LNPR CPU UTIL FOR PROCESS java-6500 ON BlakeMC IS 14%
06:02:10 ZALERT  LPCP LPAR VSIVM5 IS AT
06:02:10 ZALERT  LNPR CPU UTIL FOR PROCI
06:03:10 ZALERT  LPCP LPAR VSIVM5 IS AT
    
```

LPAR Utilization - DEMO

Time	VSIVM5 IFL (CPU Utilization)	VSIVM5 CP (CPU Utilization)
05:33	0	0
05:34	0	0
05:35	0	0
05:36	100	100
05:37	100	100
05:38	100	100
05:39	100	100
05:40	100	100
05:41	100	100
05:42	100	100
05:43	100	100
05:44	100	100
05:45	100	100
05:46	100	100
05:47	100	100
05:48	1.88	100
05:49	100	100
05:50	100	100
05:51	100	100
05:52	100	100
05:53	100	100
05:54	100	100
05:55	100	100
05:56	100	100
05:57	100	100
05:58	100	100
05:59	100	100
06:00	100	100
06:01	100	100
06:02	100	100
06:03	100	100

Single pane of glass – all LPARs console

The screenshot displays a multi-pane interface with a menu bar at the top left. The main area is a grid of 12 windows, each showing system logs for a different LPAR. The windows are titled 'ZOPER - zOPERATOR Console - VSIVC4', 'ZOPER - zOPERATOR Console - CLOUD1', and 'ZOPER - zOPERATOR Console - CUST...'. Each window contains a list of log entries with timestamps and details such as user names, system components, and actions performed. For example, one window shows '16:48:24 OPERATOR AUTO LOGON *** ZWSSL03 USE' and '16:54:06 ZVWS ZADMIN VSIL0G0100I 184.105.60.17'. The logs are organized in a consistent format across all panes, providing a unified view of system activity across multiple environments.

Console Management View

Performance | zVWS administration | zTCP administration | zVPS administration

Available zVPS Console Logs

zPRO Available Log Files

Select one or more logfiles that you wish to view or download

Download View Upload Reset

DXTZMAP			
<input type="checkbox"/> 14 Jun 2013 (4)	<input type="checkbox"/> 03 Jun 2013 (4)	<input type="checkbox"/> 23 May 2013 (4)	<input type="checkbox"/> 12 May 2013 (4)
<input type="checkbox"/> 13 Jun 2013 (4)	<input type="checkbox"/> 02 Jun 2013 (4)	<input type="checkbox"/> 22 May 2013 (4)	<input type="checkbox"/> 11 May 2013 (4)
<input type="checkbox"/> 12 Jun 2013 (4)	<input type="checkbox"/> 01 Jun 2013 (4)	<input type="checkbox"/> 21 May 2013 (648)	<input type="checkbox"/> 10 May 2013 (4)
<input type="checkbox"/> 11 Jun 2013 (4)	<input type="checkbox"/> 31 May 2013 (4)	<input type="checkbox"/> 20 May 2013 (6)	<input type="checkbox"/> 09 May 2013 (39)
<input type="checkbox"/> 10 Jun 2013 (4)	<input type="checkbox"/> 30 May 2013 (4)	<input type="checkbox"/> 19 May 2013 (6)	<input type="checkbox"/> 08 May 2013 (4)
<input type="checkbox"/> 09 Jun 2013 (4)	<input type="checkbox"/> 29 May 2013 (4)	<input type="checkbox"/> 18 May 2013 (6)	<input type="checkbox"/> 07 May 2013 (4)
<input type="checkbox"/> 08 Jun 2013 (4)	<input type="checkbox"/> 28 May 2013 (4)	<input type="checkbox"/> 17 May 2013 (6)	<input type="checkbox"/> 06 May 2013 (4)
<input type="checkbox"/> 07 Jun 2013 (4)	<input type="checkbox"/> 27 May 2013 (4)	<input type="checkbox"/> 16 May 2013 (6)	<input type="checkbox"/> 05 May 2013 (4)
<input type="checkbox"/> 06 Jun 2013 (4)	<input type="checkbox"/> 26 May 2013 (4)	<input type="checkbox"/> 15 May 2013 (4)	
<input type="checkbox"/> 05 Jun 2013 (4)	<input type="checkbox"/> 25 May 2013 (4)	<input type="checkbox"/> 14 May 2013 (4)	
<input type="checkbox"/> 04 Jun 2013 (4)	<input type="checkbox"/> 24 May 2013 (4)	<input type="checkbox"/> 13 May 2013 (4)	
INSTALL			
<input type="checkbox"/> 14 Jun 2013 (263)	<input type="checkbox"/> 02 Jun 2013 (553)	<input type="checkbox"/> 20 May 2013 (78)	<input type="checkbox"/> 09 May 2013 (6)
<input type="checkbox"/> 13 Jun 2013 (16)	<input type="checkbox"/> 31 May 2013 (12)	<input type="checkbox"/> 17 May 2013 (153)	<input type="checkbox"/> 08 May 2013 (257)
<input type="checkbox"/> 12 Jun 2013 (38)	<input type="checkbox"/> 30 May 2013 (6)	<input type="checkbox"/> 16 May 2013 (887)	<input type="checkbox"/> 06 May 2013 (5)
<input type="checkbox"/> 10 Jun 2013 (8)	<input type="checkbox"/> 29 May 2013 (317)	<input type="checkbox"/> 15 May 2013 (494)	<input type="checkbox"/> 05 May 2013 (155)
<input type="checkbox"/> 05 Jun 2013 (6)	<input type="checkbox"/> 28 May 2013 (6)	<input type="checkbox"/> 14 May 2013 (48)	
<input type="checkbox"/> 04 Jun 2013 (6)	<input type="checkbox"/> 23 May 2013 (63)	<input type="checkbox"/> 13 May 2013 (434)	
<input type="checkbox"/> 03 Jun 2013 (1050)	<input type="checkbox"/> 22 May 2013 (20)	<input type="checkbox"/> 10 May 2013 (14)	
SFPURGER			
<input type="checkbox"/> 14 May 2013 (8)	<input type="checkbox"/> 11 May 2013 (8)	<input type="checkbox"/> 08 May 2013 (8)	<input type="checkbox"/> 05 May 2013 (8)
<input type="checkbox"/> 13 May 2013 (8)	<input type="checkbox"/> 10 May 2013 (8)	<input type="checkbox"/> 07 May 2013 (8)	
<input type="checkbox"/> 12 May 2013 (8)	<input type="checkbox"/> 09 May 2013 (8)	<input type="checkbox"/> 06 May 2013 (8)	
ZALERT			
<input type="checkbox"/> 13 Jun 2013 (40)	<input type="checkbox"/> 03 Jun 2013 (48)	<input type="checkbox"/> 24 May 2013 (50)	<input type="checkbox"/> 14 May 2013 (52)
<input type="checkbox"/> 12 Jun 2013 (48)	<input type="checkbox"/> 02 Jun 2013 (38)	<input type="checkbox"/> 23 May 2013 (42)	<input type="checkbox"/> 13 May 2013 (70)
<input type="checkbox"/> 11 Jun 2013 (42)	<input type="checkbox"/> 01 Jun 2013 (52)	<input type="checkbox"/> 22 May 2013 (63)	<input type="checkbox"/> 12 May 2013 (42)

VSE 4.3 adds SNMP Interface plus some mibs:

- IBMVSE “vse mib” – system data

Report: ESAVSES VSE System Configuration Report

```
-----
NODE      <---z/VM---> <LogicalPart> <-----CPU model----->
/Time    VirtID   Lvl  Name          Nbr <IBM/<model>/CPs/ serial
-----
```

06:26:00

```
vse2     ZVSE      1  VSIVM3        0  IBM 2096-A02 02 (14B4202)
-----
```

06:27:00

```
vse2     ZVSE      1  VSIVM3        0  IBM 2096-A02 02 (14B4202)
-----
```

```
-----
NODE      <---z/VM---> <--Partitions--> <-----CPU Counts----->
/Time    VirtID   Lvl Max Cur Stat Dyn Tot Actv Quies Inact
-----
```

06:26:00

```
vse2     ZVSE      1  120  20  12   8   2   2   0   0
-----
```

06:27:00

```
vse2     ZVSE      1  120  20  12   8   2   2   0   0
-----
```

VSE 4.3 adds SNMP Interface plus some mibs:

- IBMVSE “vse mib” adds CPU data for system, and by virtual cpu

```

Report: ESAVSEC          VSE System Performance Report          VSIVM3
-----
NODE          Pages/Sec <Rate/Sec> <CPU Utilization>    All      Pct      Seconds
/Time        In   Out   SVC  DSP  Total  Mstr  Spin  Bound  NP      OfData
-----
06:26:00
vse2          0     0   196  428   83.6   4.0   0.0     0     4.8     64.6
CPU- 0                270   40.5   2.7     0     0     6.6     64.6
CPU- 1                160   43.7   1.4   0.0     0     3.2     64.6
-----
06:27:00
vse2          0     0   295  597   82.5   4.2   0.0     0     5.1     56.0
CPU- 0                359   36.6   3.2   0.0     0     8.8     56.0
CPU- 1                238   45.3   1.0     0     0     2.2     56.0

```

Velocity Software proof of concept for “Plug in”

- SNMP Support is “extensible”

What do customers want?

- TCPIP?
- VSAM?
- CICS?
- DB2
- High Capture ratio?

Report: ESAVSEP VSE Partition Performance

NODE /Time	Part ID	Job Name	Phase Name	<-CPU CPU	Time-> Overhd
06:26:00					
vse2		Totals		52.0	1.5
	FB	SECSERV	BSTPSTS	0	0
	F7	TCPIP00	IPNET	0.3	0.0
	F6	TCPIP01	IPNET	0.0	0.0
	F3	VTAMSTRT	ISTINCVT	0.0	0.0
	F2	CICSICCF	DFHSIP	0.6	0.0
	F1	POWSTART	IPWPOWER	0.0	0.0
	R2	STARTMAS	IESMASNM	0.6	0.0
	R3	STRTMAS1	IESMASNM	0	0
	S1	STGPLAY5	STGPLAY	6.6	0.4
	S2	STGPLAY2	STGPLAY	0.6	0.1
	S3	STGPLAY4	STGPLAY	11.6	0.3
	S4	STGPLAY1	STGPLAY	17.3	0.3
	R1	STARTVCS	IESVCSRV	0.0	0.0
	S5	STGPLAY3	STGPLAY	14.3	0.3

zVPS VSE agent collects, sends DMF records
z/OS, zVSE(n) all supported

Screen: **ZOSCIX1** Velocity Software - VSIVM4 ESAMON
2 of 3 CICS Analysis SYSID

Time	SYSID	<--CICS Program-->		<---Transactions-->		
		APPLID	JobName	Total	Resp	CPU
15:27:00	V24A	C24ASTND	C24ASTND	1952	0.006	0.002
15:26:49	V63B	CICSNB01	CICSNB01	1520	0.121	0.013
	V63C	CICSNC01	CICSNC01	2040	0.307	0.013
15:26:00	V24A	C24ASTND	C24ASTND	1940	0.007	0.002
15:25:52	V63C	CICSNC01	CICSNC01	2000	0.370	0.014
15:25:49	V63B	CICSNB01	CICSNB01	1520	0.117	0.013
15:25:00	V24A	C24ASTND	C24ASTND	1940	0.006	0.002
	V63B	CICSNB01	CICSNB01	1480	0.137	0.013
	V63C	CICSNC01	CICSNC01	1920	0.384	0.013
15:24:00	V24A	C24ASTND	C24ASTND	1937	0.006	0.002
	V63B	CICSNB01	CICSNB01	1520	0.147	0.014
	V63C	CICSNC01	CICSNC01	2040	0.370	0.013

CICS performance analysis. 5 new screens

Screen: **ZOSCIX1** Velocity Software - VSIVM4 ESAMON 5.140 06/12 15:20-15
 3 of 3 CICS Analysis SYSID * JOB *

Time	SYSID	CICS Program		Task Statistics					Peak		Intv Sec
		APPLID	JobName	Total	MXT	Actv	Qued	PctM	Actv	Qued	
15:28:00	V24A	C24ASTND	C24ASTND	1953	75	1	0	1.3	15	0	60
15:27:50	V63C	CICSNC01	CICSNC01	2105	25	2	0	8.0	25	9	60
15:27:49	V63B	CICSNB01	CICSNB01	1497	25	2	0	8.0	21	0	60
15:27:00	V24A	C24ASTND	C24ASTND	1934	75	1	0	1.3	19	0	60
15:26:49	V63B	CICSNB01	CICSNB01	1534	25	3	0	12.0	22	0	60
	V63C	CICSNC01	CICSNC01	1967	25	2	0	8.0	25	9	60
15:26:00	V24A	C24ASTND	C24ASTND	1932	75	1	0	1.3	13	0	60
15:25:52	V63C	CICSNC01	CICSNC01	1976	25	2	0	8.0	25	10	60
15:25:49	V63B	CICSNB01	CICSNB01	1526	25	2	0	8.0	22	0	60
15:25:00	V24A	C24ASTND	C24ASTND	1932	75	1	0	1.3	17	0	60
	V63B	CICSNB01	CICSNB01	1484	25	2	0	8.0	22	0	60
	V63C	CICSNC01	CICSNC01	1989	25	2	0	8.0	25	10	60

F1=Help

PF3=Quit

PF4=Select

PF5=Plot

PF6=Reset

Screen: **ZOSCIX2** Velocity Software - VSIVM4
1 of 2 CICS Transaction Analysis

ESAMON 5.140 06/12 15:28-15:
SYSID * APPLID * JOB *

Time	SYSID	APPLID	<-Transactions-> Group	Count	<-Response Time->			<---Dispatch---->		
					Total Resp	Susp Time	Disp Time	CPU Time	PC Load	zIIP CPU
15:29:00	V24A	C24AAOR1	Totals	0	0	0	0	0	0	0
			Inflight	16	174.2	174.2	0.000	0.000	0	0
		C24ASTND	Totals	1940	0.006	0.003	0.004	0.002	0	0.000
			Inflight	12	60.04	59.93	0.114	0.046	0	0
		C24ATOR1	Totals	0	0	0	0	0	0	0
			Inflight	8	60.03	60.03	0.001	0.001	0	0
15:28:49	V63B	CICSNB01	Totals	1520	0.128	0.106	0.022	0.013	0	0
			Inflight	0	0	0	0	0	0	0
	V63C	CICSNC01	Totals	2120	0.339	0.313	0.026	0.013	0	0
			Inflight	0	0	0	0	0	0	0

ESAUSCP – VCPU Analysis, how many cpus needed?

- DB2 workload has very strange overhead....

Report: ESAUSCP **Virtual Machine** VCPU Analysis

UserID	<---CPU time-->				<---Percent						
	CPUvadd	<-Percent->		<-SHARE-->	CPU	<-Samples->					
	Cnt	TOT	Virt	Type	Value	TYPE	Total	In Q	Run	Sim	CPU
07:17:00	0	57.73	35.96	.	.	.	4307	1045	3.2	1.1	1.4
TSTDB2	2	22.88	2.85	ABS	4.0	IFL	118	117	6.0	7.7	1.7
CPU-00		2.21	2.16	ABS	0	IFL	59	59	5.1	0	1.7
CPU-01		20.68	0.69	ABS	0	IFL	59	58	6.9	16	1.7

Report: ESALNXS **LINUX VSI** System Analysis Report

Node/ Time	<---Load Numbers-->			CPU NBR	<Processor Pct Util>			NICE Time	<CPU Overhead%>			IO Wait
	Users	Procs	MaxProc		Total	Syst	User		Idle	Krnl	IRQ	
01/16/17												
TSTDB2	0	346		0 Tot	0	0	0	0	0	0	0	0
				1	0	0	0	0	0	0	0	0
				2	0	0	0	0	0	0	0	0

ESAUCD2 – The most useful storage report available

- Note, page tables are “anonymous / overhead”

```

Report: ESAUCD2          LINUX UCD Memory Analysis Report          Velocity Softwar
Monitor initialized: 05/13/14 at 00:00:00 on 2828 serial 414C7          First recordana
-----
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
-----
00:15:00
oracle    994.8  18.1  976.7  123.9  74.0  49.9  15.6  92.1  0  240.6  581.4  154.7
redhat5   499.2  17.9  481.3  4095  4095  0.0  15.6  4113  0  140.5  206.6  134.2
redhat5x  497.1  19.8  477.3  4095  4095  0.0  15.6  4114  0  150.0  170.6  156.7
redhat56  497.0  24.3  472.7  1051  1051  0.0  15.6  1075  0  170.1  174.6  128.0
redhat6   492.7   7.8  484.9  4095  4090  5.2  15.6  4098  0  167.9  182.6  134.4
redhat6x  994.8  10.7  984.1  495.8  404.0  91.9  15.6  414.7  0  29.7  785.4  169.0
rhel64v   996.4  70.0  926.4  2047  2047  0  15.6  2117  0  152.0  601.8  172.6
roblx2    241.7  11.1  230.6  0  0  0  15.6  11.1  0  44.2  107.6  78.8
sles10    493.0  19.8  473.2  4219  4219  0  15.6  4238  0  140.9  281.1  51.2
sles11    494.7  172.8  322.0  4087  4087  0  15.6  4260  0  139.3  122.7  59.9
sles11v2  2006.7  85.9  1921  1542  699.6  842.4  15.6  785.5  0  3.0  894.9  1023
sles11v3  868.8  91.2  777.6  2046  1759  287.2  15.6  1850  0  4.2  65.8  707.6
suselnx2  247.3  158.6  88.6  255.8  255.8  0  15.6  414.5  0  29.0  37.3  22.4
s11s2ora  996.5  23.7  972.8  743.8  598.2  145.5  15.6  621.9  0  41.2  777.9  153.7
  
```

New metrics

- RSS, Size - Same
- Locked: Locked memory size (mlock)
- Peak: peak RSS (high water mark)
- Data: size of data, stack
- **PTBL: page table entries (linux 2.6.10) - Use to evaluate LARGE PAGES**
- EXEC: size of executable (text)
- Lib: shared library code size
- **Swap: Swapped out**
- Stack: size of stack

Report: ESALNXP LINUX HOST Process Statistics Report Velocity Software Corporate ZMAP 4.2.0

node/ Name	<-Process Ident->			PRTY	<-----CPU Percents----->					<-----Storage Metrics (MB)----->										
	ID	PPID	GRP		Valu	Tot	sys	user	syst	usrt	Size	RSS	Peak	Swap	Data	Stk	EXEC	Lib	Lck	PTbl
00:15:00																				
oracle	0	0	0	0	1.87	0.11	1.05	0.16	0.55	7345	845	108K	0	1997	62.8	28K	6K	0	130	
init	1	1	0	16	0.60	0	0	0.12	0.48	1	0	12.5	0	2.17	1.2	8.9	0	0	0.12	
oracle	21131	1	21131	16	0.88	0.00	0.87	0	0	403	52	3585	0	18.4	1.4	965	139	0	5.98	
redhat6x	0	0	0	0	1.66	0.38	0.67	0.22	0.38	19K	1216	275K	462	15K	103	74K	18K	0	219	
init	1	1	1	20	0.59	0.00	0	0.21	0.38	3	1	46.6	0.53	3.11	1.3	2.2	38	0	0.21	
sles11v2	0	0	0	0	5.96	3.54	1.83	0.19	0.40	105K	4321	1.5M	6958	21K	517	347K	34K	0	1498	
init	1	1	1	20	0.58	0.00	0.00	0.19	0.38	11	0	135	1.27	2.34	1.7	0.5	25	0	0.51	
ora_vktm	5963	1	5963	-2	1.65	1.65	0	0	0	1137	2	17K	28.2	46.3	2.1	3546	285	0	7.03	
ora_vktm	10254	1	10254	-2	1.62	1.33	0.29	0	0	926	2	14K	27.8	46.3	2.1	3546	285	0	7.27	
slls2ora	0	0	0	0	1.86	0.42	0.68	0.26	0.50	16K	1063	238K	830	2353	141	70K	9K	0	207	
init	1	1	1	20	0.75	0	0	0.26	0.50	2	0	34.0	1.31	2.57	1.9	0.5	28	0	0.14	

Benchmark process analysis (2G SGA, oversized)

Report: ESALNXP Velocity Software Corporate ZMAP 4.2.0

node/ Name	<-Proc ID	<----- Size	----- RSS	----- Peak	----- Swap	Storage Metrics (MB)----->					
						Data	Stk	EXEC	Lib	Lck	PTbl
NO HUGE PAGES											
oracle	43146	2303	265	2249	0	3.07	0.1	181	13	0	0.96
oracle	43148	2310	81	2256	0	8.95	0.1	181	13	0	1.06
oracle	43152	2303	57	2249	0	3.07	0.1	181	13	0	0.69
oracle	43158	2308	141	2254	0	3.20	0.3	181	14	0	1.21
oracle	43160	2303	101	2249	0	3.07	0.1	181	13	0	0.84
HUGE PAGES											
oracle	51439	2304	18	2250	0	4.26	0.1	181	14	0	0.31
oracle	51451	2303	22	2250	0	3.07	0.1	181	14	0	0.32
oracle	51453	2314	23	2259	0	3.07	0.1	181	13	0	0.32
oracle	51455	2303	16	2249	0	3.07	0.1	181	13	0	0.31
oracle	51457	2310	23	2256	0	8.95	0.1	181	13	0	0.31
oracle	51459	2318	17	2263	0	3.07	0.1	181	13	0	0.32

Java/WebSphere Metrics (Management vs diagnostics)

Report: ESAJVM

Java Subsystem Analysis Report

Velocity Sof

```

-----
Node/          <JavaClass> Memory <-----Heap data----->
Date          <-----Application-----> <--Loaded-> pending <-----sizes----->
Time          Name                Type Curr  /Sec  Final  Init  Used  Commit Max
-----
13:06:00
S11R20RA WAS Server1                JVM   15287    0      0 52.4M  100M 107.5M  268M
           WAS Server2longerna    JVM   15312    0      0 52.4M  85.4M 103.3M  268M
  
```

Report: ESAJVM

Java Velocity Software Corporate

ESAMAP 4.2.0 06/19/13

```

-----
Node/          <---Non Heap Data----> <--Thread Count data->
Date          <-----Application- <-----sizes-----> Curr Daemon Peak start
Time          Name                Init Used  Commit Max Live Count thrds /sec
-----
13:06:00
S11R20RA WAS Server1                0  101M 184.7M    0 58.0   55.0  55.0    0
           WAS Server2longerna    0  101M 171.9M    0 58.0   55.0  55.0    0
  
```

- The Velocity Software mib extracts threads

Report: ESAJVMT Java Subsystem Analysis Report Velocity Sof
 Monitor initialized: 12/05/16 at 14:35:40 on 2828 serial 0314C7 First record

Node/ Date Time	<-----Thread ID-----> Name	nr	<--Blocks--> /Second	Time	<Thread >/Sec	Waits> Time	CPU (ms)
14:37:00	Totals: AppSrv01-server1	0	0.2	0	73.1	0	170.8
lxora12	CommunicatorServer	7	0	0	0	0	7.8
	Thread-11	17	0	0	0.0	0	1.1
	Deferred Alarm Manager	30	0	0	2.2	0	4.0
	Non-Deferred Alarm Manager	31	0	0	2.0	0	2.9
	Deferrable Alarm : 0	43	0	0	1.0	0	3.1
	LT=0:P=315710:O=0:port=9100	49	0	0	0	0	1.4
	LT=1:P=315710:O=0:port=9403	50	0	0	0	0	1.8
	ThreadService-0	90	0.0	0	1.6	0	19.2
	Deferrable Alarm : 1	99	0	0	1.0	0	3.9
	Deferrable Alarm : 2	135	0	0	1.0	0	3.8
	Thread-79	140	0	0	0.3	0	1.1
	ThreadService-1	148	0.0	0	1.8	0	16.7
	Deferrable Alarm : 3	149	0	0	1.0	0	3.4
	ThreadService-2	150	0.0	0	1.7	0	11.7
	ThreadService-3	151	0.0	0	1.6	0	14.5
	ThreadService-4	153	0.1	0	1.5	0	16.1
	ThreadService-5	154	0.0	0	1.5	0	25.8
	AIO Timer Thread 1	183	0	0	1.0	0	1.7
	WebContainer : 2	186	0	0	1.0	0	1.5
	WebContainer : 15	226	0	0	1.0	0	1.0
	WebContainer : 17	228	0	0	1.0	0	1.6

Oracle Database Configuration (management vs diagnostics)

ESAORAC: Oracle Configuration, SGA, PGA High Level information

Report: ESAORAC Oracle Database Configuration Report

```

-----
Node/      <-----Database Description-----> <-----Database----->
Date                                             <----Start----->
Time      DatabaseName      Instance      Version      Date          Time Status
-----
PAZXXT10  soedb                    soedb         12.1.0.1.0   2014/01/27   10:15 OPEN
redhat6x  db01                     db01          11.2.0.2.0   2013/12/19   14:42 OPEN
sles11v2  db01                     db01          12.1.0.1.0   2013/11/08   13:20 OPEN
  
```

```

Node/      <-----
Date                                             <-----Storage Overview (MB)----->
Time      Database      Max      Fixed      Free      Size      Max      MaxMan
-----
PAZXXT10  soedb                    1598      2.3      557K     557.1     293.7     1040.0
redhat6x  db01                    399.6     2.2     139K     139.3     164.8     529.0
sles11v2  db01                    334.4     2.2    32768    106.5     355.2     12950
  
```

Measuring Oracle – Linux Process Perspective

Report: **ESALNXA** LINUX HOST Application Report Velocity Software ZMAP 4.2.0

Node/ Date Time	Process/ Application name	<---Processor Percent--->					<Process->		<---Percent Process Status--->					
		Total	sys	user	sys	usr	Total	Actv	Run-	Sleep	Zom	Disk	Page	Stop

08:30:00														
PAZXXT10	*Totals*	6.6	2.0	2.6	0.7	1.3	149.0	24.5	0.7	99.3	0	0	0	0
	init	1.9	0.0	0.0	0.6	1.3	1.0	0.3	0	100	0	0	0	0
	ora_vktm	1.9	1.0	0.8	0	0	1.0	1.0	0	100	0	0	0	0

08:45:00														
PAZXXT10	*Totals*	55.9	7.5	46.1	0.8	1.6	164.9	42.5	1.9	94.7	0	3.4	0	0
	init	2.3	0.0	0	0.7	1.6	1.0	0.2	0	100	0	0	0	0
	ora_vktm	1.3	0.7	0.6	0	0	1.0	1.0	0	100	0	0	0	0
	oracle_1	19.8	2.9	16.8	0	0	12.0	12.0	15.0	48.3	0	36.7	0	0
	xterm	27.8	1.7	26.1	0	0	3.3	1.0	0	100	0	0	0	0

09:00:00														
PAZXXT10	*Totals*	69.4	11.1	56.9	0.5	0.9	181.6	57.7	1.8	95.1	0	3.1	0	0
	init	1.3	0.0	0.0	0.5	0.8	1.0	0.3	0	100	0	0	0	0
	ora_dbw0	2.2	1.5	0.7	0	0	1.0	1.0	6.7	0	0	93.3	0	0
	ora_lg00	0.7	0.4	0.2	0	0	1.0	1.0	0	46.7	0	53.3	0	0
	ora_vktm	1.2	0.7	0.5	0	0	1.0	1.0	0	100	0	0	0	0
	oracle_1	43.5	5.0	38.5	0	0	20.0	20.0	8.0	73.7	0	18.3	0	0
	xterm	15.7	1.6	14.2	0	0	5.0	1.3	0	100	0	0	0	0
	Xvnc	1.3	0.5	0.8	0	0	1.0	1.0	6.7	93.3	0	0	0	0

Measuring Oracle Database Storage

ESAORAG: General Storage Areas – SGA, no changes

Report: **ESAORAG** **SGA/PGA Analysis Report** Velocity Software

Monitor initializ/14 at 08:00:00 on 2094 serial 53E5D First record anal

Node/		<-----Shared Global Area (SGA) in Megabytes----->										
Date	<--Data	Max	Fixed	Redo	Buffer	<-----Pool sizes----->						
Time	Name	Size	Size	Buffr	Cache	Free	Shrd	Large	Java	Stream	ShrIO	

08:30:00												
PAZXXT10	soedb	1598	2.3	6.9	655.4	557K	295	32.8	16.4	32.8	49.2	

08:45:00												
PAZXXT10	soedb	1598	2.3	6.9	658.6	557K	295	32.8	16.4	29.5	49.2	

09:00:00												
PAZXXT10	soedb	1598	2.3	6.9	671.7	557K	295	32.8	16.4	16.4	49.2	

ESAORAG: General Storage Areas – PGA Grows with workload

Report: ESAORAG

```

-----
Node/          <--PGA Data (in Megabytes)-->
Date          <--Data  <-Target-> InUse Alloc Free
Time         Name      Parm Auto      -able
-----
01/28/14
08:30:00
PAZXXT10 soedb          557  451  56.8  84.1  17.5
-----
08:45:00
PAZXXT10 soedb          557  426  84.2 128.9  30.4
-----
09:00:00
PAZXXT10 soedb          557  404 109.4 170.3  43.6

```

Measuring Oracle Workloads

ESAORAS: User Committs/Rollbacks, Session CPU, Recursive CPU

Report: ESAORAS Oracle Subsystem Analysis Report

```

-----
Node/      <---Database----> <-User Activity-> <--CPU---->
Date                               <Rate per second>  Sess  Re-
Time      Name          Instance Calls  Comm Rollbk  -ion  Cur
-----
08:30:00
PAZXXT10 soedb          soedb      0.2    2.3    3.8    0.0    0
-----
08:45:00
PAZXXT10 soedb          soedb      0.2  241.2  73.1  22.0    0.0
-----
09:00:00
PAZXXT10 soedb          soedb      0.2  569.5  168.2  52.4    0.11

```

ESAORAS: Oracle Subsystem

Report: ESAORAS Velocity Software Corporate ZMAP 4.2.0 12/21/13

```

-----
Node/          <----Physical Reads Activty----> <-Physical Write Activity-->
Date          <-----Rate per second-----> <-----Rate per second----->
Time      Name  Rds Hits Direct  I/O  Bytes  Writs  CHits  Dirct  I/O  Bytes
-----
08:30:00
PAZXXT10 soedb  0.8  0.2  0.2  0.2      0   0.1  1415   0.3  0.2  0.1
-----
08:45:00
PAZXXT10 soedb  18.1  8.3 172.8 172.8      0  172.2  715K   42.8  42.3  0.5
-----
09:00:00
PAZXXT10 soedb  36.5 13.8 279.9 279.9      0  276.6  9733K  178.6 178.6  0.1
-----

```

Performance analysis vs “averages”

Report: ESAORAS Oracle Subsystem Analysis Report Velocity Software Corporate ZMAP 4.2.0
 Monitor initialized: 08/04/14 at 00:00:00 on 2828 serial 414C7 First record analyzed: 08/04/14 00:00:00

Node/Date/Time	<---Database---> Name	<-User Activity-> Instance	<--CPU--> <Rate per second> Calls	Sess Re- Comm	<-----Physical Reads Actvty---> Rollbk	<-Physical Write Actvty-> -ion	Cur	<-----Rate per second-----> Rds	Hits	Direct	I/O Bytes	Writs	CHits	Dirct	I/O	
00:01:00	roblx1	orcl	0.3	302.3	10.9	2.0	0	10.0	4.9	6.9	6.9	0	6.7	56361	1.0	1.0
00:02:00	roblx1	orcl	0.2	26.7	10.4	0.2	0	6.9	2.0	0	0	0	0	0	1.7	1.7
00:03:00	roblx1	orcl	0.2	40.3	9.5	0.3	0	5.7	0.8	0	0	0	0	0	1.5	1.5
00:04:00	roblx1	orcl	0.2	7.6	9.6	0.2	0	5.2	0.2	0	0	0	0	0	1.9	1.9
00:05:00	roblx1	orcl	0.2	23.3	9.3	0.4	0	4.9	0.3	0	0	0	0	0	0.8	0.8
00:06:00	roblx1	orcl	0.2	16.6	10.2	0.2	0	5.8	0.6	0	0	0	0	0	43.6	43.6
00:07:00	roblx1	orcl	0.2	37.3	10.3	0.2	0	5.1	0.0	0	0	0	0	0	1.3	1.3
00:08:00	roblx1	orcl	0.5	26.7	15.0	0.9	0	6.9	0.7	4221	122.2	4099	267.7	2499K	1.6	1.6
00:09:00	roblx1	orcl	0.3	25.5	16.1	1.3	0	36.9	41.5	6781	4202	2579	333.9	60.6M	9.7	9.7
00:10:00	roblx1	orcl	0.2	26.4	14.3	1.7	0.0	45.0	18.6	6489	645	38.6	248.9	3648K	1.9	1.9

Possible metrics for alerting:

- For server, by database:
- SGA size > x
- PGA allocated > x
- Users?
- I/O
- Cache hits?

Linux applications by Group

Report: ESALNXA LINUX HOST Application Report

Monitor initialized: 21/01/11 at 07:03:00 on

```
-----
Node/      Process/   ID    <---Processor Percent--->
Date       Application
Time       name              <Process><Children>
              Total sys  user syst usrt
-----
```

07:04:00

Node Groups

Node	Process	ID	Total	sys	user	syst	usrt
WASApps	*Totals*	0	90.8	9.4	78.6	0.5	2.2
	automoun	0	0.0	0	0.0	0	0
	httpd	0	5.4	1.8	3.6	0	0
	httpd1	0	44.4	4.0	40.5	0	0
	httpd18	0	8.8	0.2	8.6	0	0
	httpd19	0	2.8	0.2	2.6	0	0
	httpd2	0	2.5	0.2	2.3	0	0
	httpd3	0	4.1	0.7	1.3	0.3	1.8
	httpd4	0	6.0	0.9	5.1	0	0
	httpd5	0	1.1	0.1	1.1	0	0
	httpd6	0	2.7	0.2	2.5	0	0
	httpd7	0	6.4	0.3	6.1	0	0
	httpd9	0	4.3	0.0	4.2	0	0
	kernel	0	0.6	0.0	0	0.2	0.4
	snmpd	0	0.7	0.4	0.2	0	0

Linux Application Accounting

- Defining applications

```

appname = 'PIDFILE'   '   appstring = 'PidFile'
appname = 'MineCrft'  '   appstring = 'minecraft'
appname = 'HVC'       '   appstring = 'hvc'
appname = 'DESKTOP'   '   appstring = 'desktop'
appname = 'oracle'    '   appstring = 'oracle'
    
```

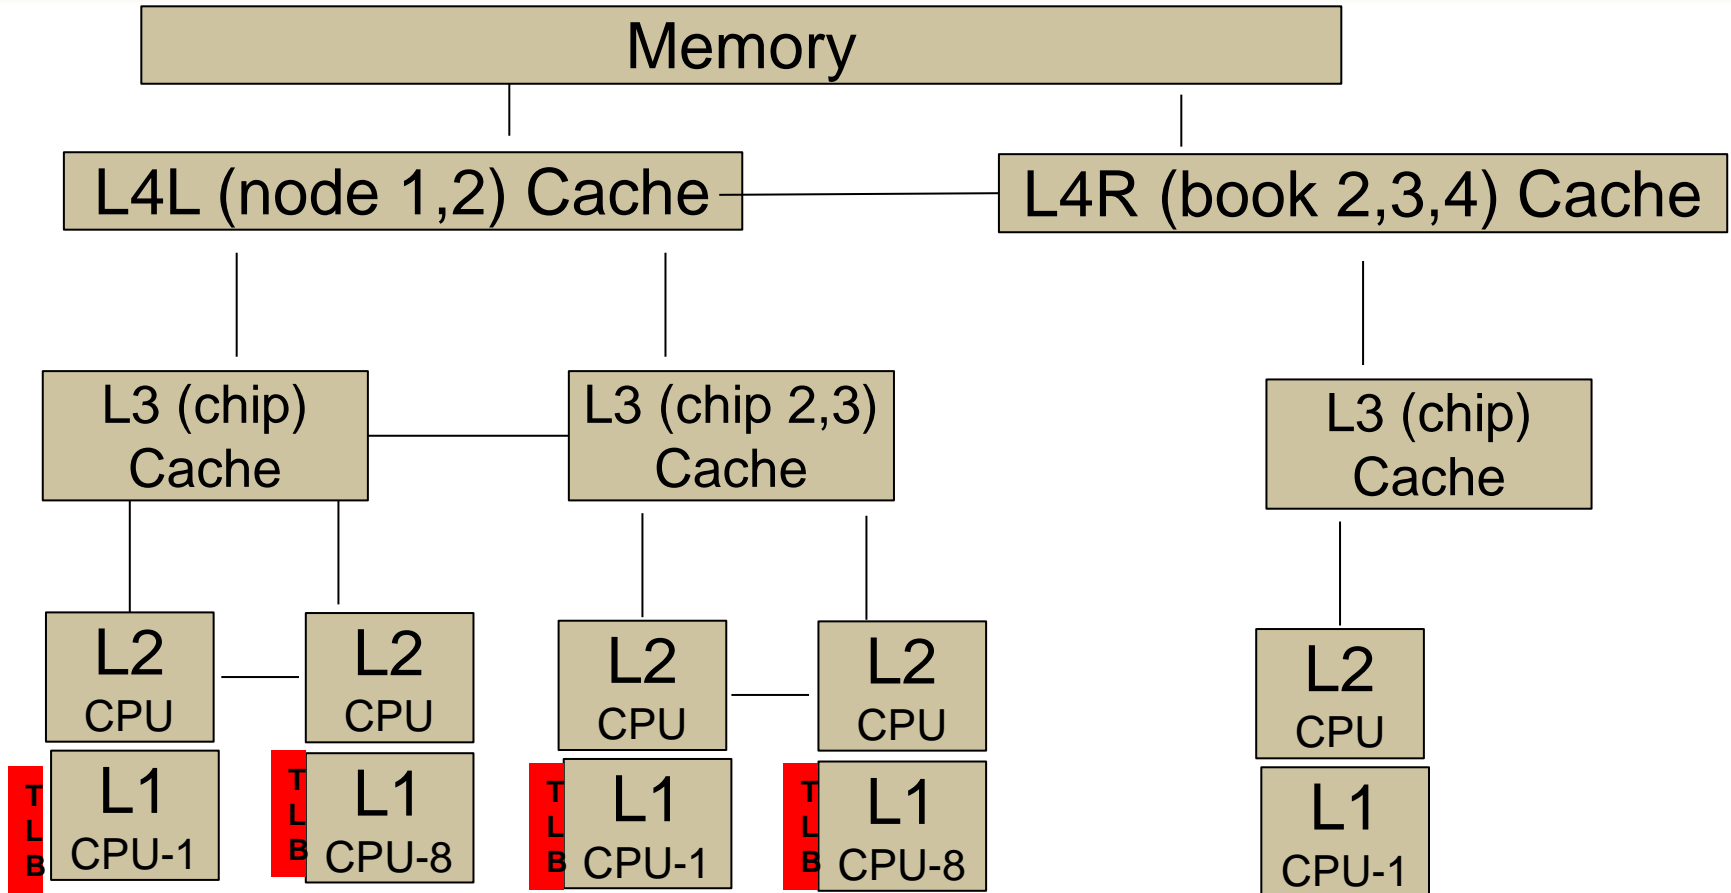
Screen: **ESALNXA** Velocity Software - VSIVM4
1 of 3 LINUX VSI Host Application Report

ESAMON 4.201 02/25
CLASS * NODE BLAKEMC

Time	Node/ Group	Process/ Application name	ID	<---Processor Percent--->				
				Total	sys	user	syst	usrt
21:20:00	BLAKEMC	*Totals*	0	15.6	0.1	15.5	0	0
		kernel	2097K	0.0	0.0	0	0	0
		snmpd	98934	0.1	0.1	0.0	0	0
		MineCrft	81176	15.5	0.0	15.5	0	0
21:19:00	BLAKEMC	*Totals*	0	14.5	0.1	14.4	0	0
		snmpd	98934	0.1	0.0	0.0	0	0
		MineCrft	81176	14.4	0.0	14.4	0	0
21:18:00	BLAKEMC	*Totals*	0	14.4	0.1	14.3	0	0
		snmpd	98934	0.1	0.0	0.0	0	0
		MineCrft	81176	14.3	0.0	14.3	0	0

- **Websphere argument string used for accounting**
- `wasadmin 27144 6846 0 Feb06 ? 00:43:13 /u01/was61/java/bin/java -
Declipse.security -Dwas.status.socket=34229 -Dosgi.install.area=/u01/was61 -
Dosgi.configuration.area=/u01/was61/profiles/appsrv/configuration -
Dosgi.framework.extensions=com.ibm.cds -
Xshareclasses:name=webspherev61_%g,groupAccess,nonFatal -Xscmx50M -
Xbootclasspath/p:/u01/was61/java/jre/lib/ext/ibmorb.jar:/u01/was61/java/jre/lib/ext/
ibmext.jar:/u01/J2EEProbe/DiagnosticsAgent/classes/IBM/1.5.0/instr.jre:/u01/J2EE
Probe/DiagnosticsAgent/classes/boot -classpath
/u01/was61/profiles/appsrv/properties:/u01/was61/properties:/u01/was61/lib/startu
p.jar:/u01/was61/lib/bootstrap.jar:/u01/was61/lib/j2ee.jar:/u01/was61/lib/Improxy.ja
r:/u01/was61/lib/urlprotocols.jar:/u01/was61/deploytool/itp/batchboot.jar:/u01/was
61/deploytool/itp/batch2.jar:/u01/was61/java/lib/tools.jar -
Dibm.websphere.internalClassAccessMode=allow -verbose:gc -Xms1024m -
Xmx1200m -
Dws.ext.dirs=/u01/was61/java/lib:/u01/was61/profiles/appsrv/classes:/u01/was61/c
lasses:/u01/was61/lib:/u01/was61/installedChannels:/u01/was61/lib/ext:/u01/was61
/web/help:/u01/was61/deploytool/itp/plugins/com.ibm.etools.ejbdeploy/runtime -
Dderby.system.home=/u01/was61/derby -Dcom.ibm.itp.location=/u01/was61/bin -
Djava.util.logging.configureByServer=true -
Duser.install.root=/u01/was61/profiles/appsrv -`

- **Node Classes (installation defined)**
 - By application (capacity planning)
 - By VMWare/ESX box
 - By department (chargeback)
- **Other grouping (automatic)**
 - Linux Process by user (ESALNXU)
 - Linux Process by process name (ESAHSTA)
 - Linux Process by application (ESALNXA)
 - Requires Parent/Child relationship
 - Linux Disk storage by NODE class
- **Define alerts (Operational support)**
 - Based on application
 - Based on node group
 - Based on linux user



Question, If 10,000 dispatch / second / cpu, impact?

- What is the CPU Measurement Facility (Basic)
- CPI: Cycles per Instruction

Report: ESAMFCA MainFrame Cache Hit Analysis
 Monitor initialized: 12/10/14 at 07:44:37 on 282

```

-----
                <CPU Busy> <-----Processor----->
                <percent>  Speed/<-Rate/Sec->  CPI
Time           CPU  Totl  User  Hertz  Cycles  Instr  Ratio
-----
07:48:35      0  20.8  18.4  5504M  1121M  193M  5.807
                1  21.6  19.6  5504M  1161M  221M  5.264
                2  24.4  22.5  5504M  1300M  319M  4.078
                3  22.4  19.7  5504M  1248M  265M  4.711
                4  19.6  17.6  5504M  1102M  194M  5.683
                5  20.4  18.6  5504M  1144M  225M  5.087
                6  23.9  22.0  5504M  1341M  341M  3.935
                7  17.6  15.4  5504M   949M  160M  5.927
                8  18.5  16.5  5504M  1005M  194M  5.195
                9  22.5  20.6  5504M  1259M  347M  3.629
-----
System:                212   191  5504M  10.8G  2457M  4.733
  
```

Why you should be interested – what is a MIP?

Report: ESAMFC MainFrame Cache Analysis Rep

Time	CPU	<CPU Busy>		<-----Processor----->			
		Totl	User	Speed/ Hertz	<-Rate/Sec-> Cycles	Instr	Ratio
14:05:32	0	92.9	64.6	5000M	4642M	1818M	2.554
	1	92.7	64.5	5000M	4630M	1817M	2.548
	2	93.0	64.7	5000M	4646M	1827M	2.544
	3	93.1	64.9	5000M	4654M	1831M	2.541
	4	92.9	64.8	5000M	4641M	1836M	2.528
	5	92.6	64.6	5000M	4630M	1826M	2.536
System:		557	388	5000M	25.9G	10.2G	2.542
14:06:02	0	67.7	50.9	5000M	3389M	2052M	1.652
	1	67.8	51.4	5000M	3389M	2111M	1.605
	2	69.0	52.4	5000M	3450M	2150M	1.605
	3	67.2	50.6	5000M	3359M	2018M	1.664
	4	60.8	44.5	5000M	3042M	1625M	1.872
	5	70.1	53.8	5000M	3506M	2325M	1.508
System:		403	304	5000M	18.8G	11.4G	1.640

**1830 mips
(at 100%)**

**2828 Mips
(at 100%)
Doing 10%
more work**

Why working sets are important,

Why we need large pages?

DAT Translation consumes 30% of the cycles for both threads

Report: ESAMFC MainFrame Cache Magnitudes Report ZMAP 4.2.4

Time	CPU	<CPU Busy> <percent>		<-----> Speed/ Hertz Ratio		<-Translation Lookaside buffer (TLB)- <cycles/Miss><Writs/Sec>				CPU Cycles	
		Totl	User	Hertz	Ratio	Instr	Data	Instr	Data	Cost	Lost
07:45:01	0	25.9	24.4	5000M	1.704	159	742	473K	244K	19.77	257M
	1	35.9	34.7	5000M	1.491	138	731	530K	249K	14.17	255M
	2	15.8	13.9	5000M	2.868	206	826	419K	245K	36.30	289M
	3	16.6	15.4	5000M	2.508	212	825	411K	247K	34.90	291M
	23	18.1	17.0	5000M	2.144	197	815	412K	229K	29.44	268M
	24	21.4	19.9	5000M	1.865	114	533	598K	302K	21.35	229M
	25	26.2	24.9	5000M	1.742	98	503	736K	346K	18.71	246M
	26	12.9	11.6	5000M	2.050	154	631	378K	214K	29.92	194M
	27	13.1	11.9	5000M	1.987	156	630	378K	217K	29.64	195M
System:		514	476	5000M	2.257	176	724	14M	7641K	30.69	7917M

- For z/VM, OSA MIB installs on a Linux Server
- Two sources, Shows configuration, totals, by LPAR
- Not sure which source is accurate or why discrepancy

Report: ESAOSA OSA System Configuration Report
Monitor initialized: 05/14/16 at 06:02:00 on

```

-----
Collector <-----OSA Configuration--> MacAddress
Node      Idx   Name  Nbr   Type Level Shrd Active
-----
06:03:00
OSA178    2    OSA1   0 1G Eth 6.00  Yes 6CAE8B483FD4
  
```

Report: ESAOSA OSA Velocity Software Corporate
Monitor initialized: 05/14 First record analyzed: 05/14

```

-----
Collector <-----OSA LPAR Bus CPHID KBytes/Sec Packets/sec
Node      Idx   Name  Nbr   NBR Util Util   IN   OUT   In   OUT
-----
06:03:00
OSA178    2    OSA1   0    Tot  0    0    7.0  8.2  30.1  23.2
          2    0    .    3    1
          4    0    .   17   17
          5    0    .    4    4
  
```

Management vs Diagnostics

- **Cost of management must be low**

Performance Management:

- **Performance Analysis**
- **Capacity Planning**
- **Operational alerting**
- **Chargeback capability**