

Introduction to Velocity Software Performance Management (zLOV)

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

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

Performance Management is a process with four key components:

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

zVPS Users:

- Centralized performance management
- Data analyst / Linux admin
- Enterprise capacity planning
- Accounting
- Centralized Operations
- Dashboards

Our Product Design Objectives:

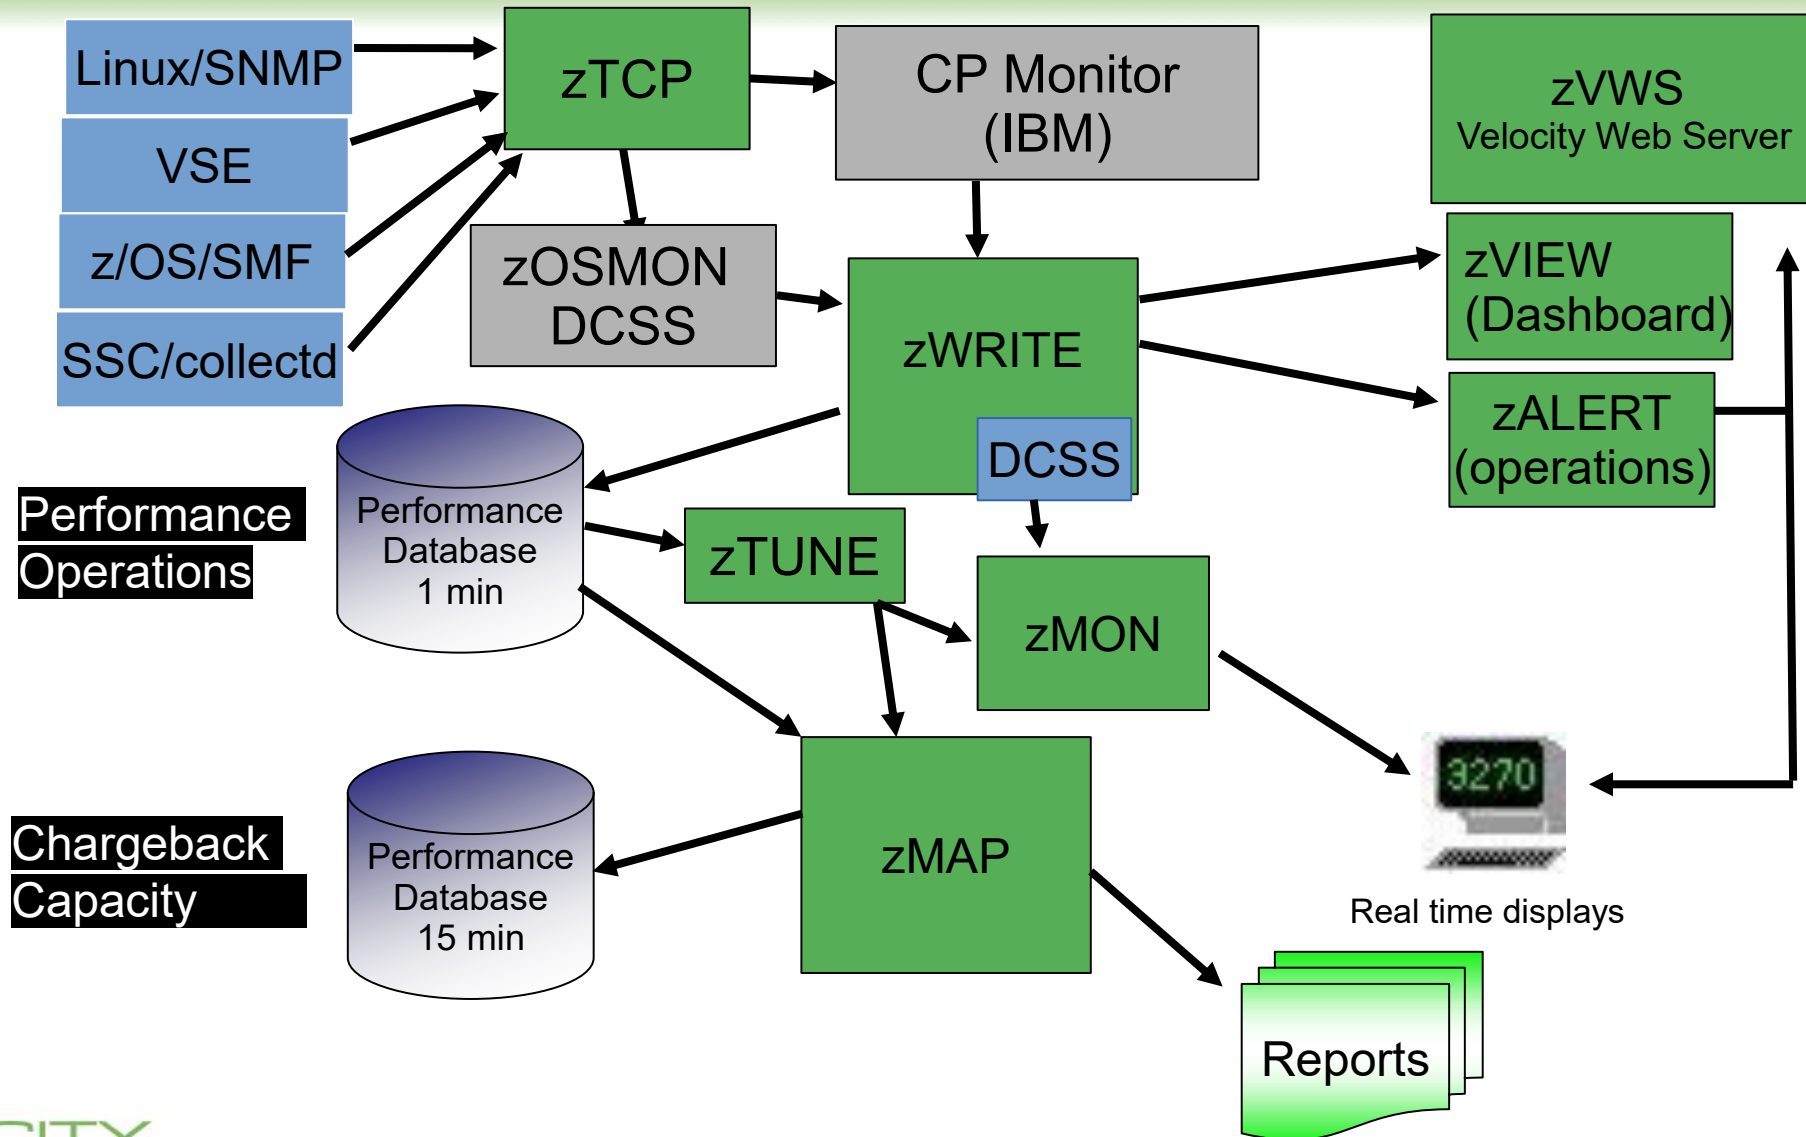
- Data Accuracy, product longevity, scalability, extensibility
- Minimize complexity - Keep it Simple (and elegant)
- Flexible (easily add / change function)
- INTUITIVE (navigation, report understanding)
- Ease of use, easy to support (Start with good architecture)
- Modernization (browser based, cross enterprise)
- Low resource consumption for zVPS (1% of one engine)
- Low resource consumption for data collection
- Provide data to dashboards (Grafana, Splunk, etc.)

zVPS Components

- zMAP – Performance reporting, long term performance data base
- zMON – Real time interface, short term performance data base
- zVWS – Native webserver
- zTCP – Data collector (SNMP, collectd, SMF, DMF)

No charge components with graphical interface

- ESAEXTR – Build your own reports, analyzer
- zVIEW – Performance management dashboard
- zOPERATOR
 - Fully integrated operations console (replaces Ops Mgr)
- zALERT
 - Performance alerts and notifications (integrated)
 - Alerts to SNMP management console (NETCOOL, HPOpenView)
 - Email alerts, Cell phone text alerts



Performance Operations

Chargeback Capacity

Why zTUNE? (KI Based, Knowledge based, AI: Actual Intelligence)

- Mainframes/LinuxOne environments are Complex
- Many things get overlooked even by experts
- Experts become experts by seeing many many performance problems
- What causes problems?
- When same problem multiple times, create a rule to look for it
- Checks for configuration “best practices”
- Inexpensive insurance to have the best skills when there is a problem

zTUNE Components

- zTUNE RULES – 100+ performance items that get checked
- ESATUNE report produced by zMAP, display on zVIEW
- Performance assistance on demand from Velocity Software experts
- Upload data for analysis at any time
- Ptrack ‘zTUNE’ sev1 alerts Velocity management phones....

Focus more now on simplifying problem resolution

User reports that applications complained about Linux on Z WAS performance:

Report: **ESATUNE** Tuning Recommendation Report

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

zVWS: Native generalized z/VM Webserver

- CMS Based
- VERY SIMPLE AND ELEGANT ARCHITECTURE
- Written in Assembler, because that is just fast
- Generalized server – completely eliminates need for SMAPI

VelocitySoftware.com (all runs on z/VM natively)

- VelocitySoftware.com, VelocitySoftware.de, etc.
- VelocitySoftware.net
- Linuxvm.org, MVMUA.org
- VMWorkshop.org (great conference for z/VM)
- GGWSC.ORG

Velocity Applications provided by Velocity Software

- zVIEW (Performance data presentation “dashboards”)
- zPORTAL (GUI interface to managing zVPS)
- zPRO (No SMAPI, no java, No Linux server requirements, no complexities)

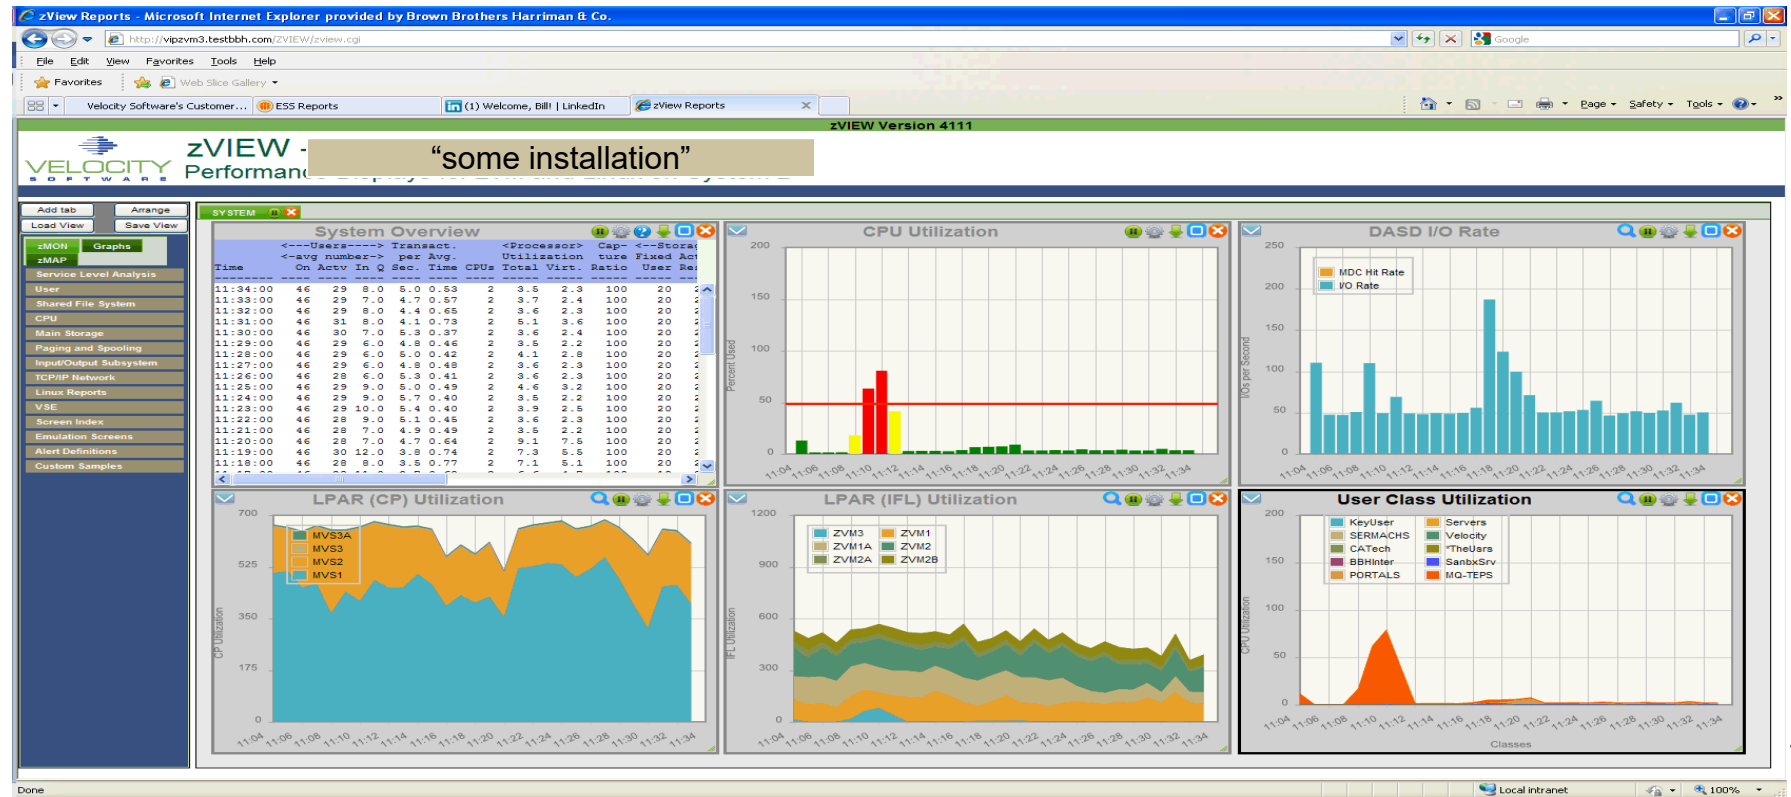
All platforms provided, one technology, one navigation architecture

- z/VM (CP monitor)
- Networks (SNMP)
- Linux (“z” and “x”) (SNMP)
- Linux for z (**ZLMD**)
- VSE (VSEMON – no charge, SNMP, DMF)
- z/OS (zOSMON: SMF record input)
- SSC (IBM Secure Software Container – collectd)
- Microsoft (SNMP – no charge)
- VMWare/ ESX (SNMP – no charge)

Many applications

- Oracle (SNMP)
- JVM (SNMP)
- GPFS (SNMP)
- Docker, Kubernetes (OpenShift, Rancher)
- MongoDB Enterprise

- End users define their dashboard(s)
- Many dashboards are provided (VSE, z/OS, Linux)
- Not just for Systems Programmers. (Applications, Operations)
- Menu driven



Single pane of glass to meet newer requirements

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

ADP Associate Portal | VLB6 - Enterprise View | vib6.mf.adp.com:1024/ZVIEW/

Today is Tuesday 23 Apr 2019

Enterprise Performance Summary - Automate “some installation”

Finder | Search JASS Inventory

DC1			
V1P1	08:48	IFL Total (48) @ 736.41%	Expand
V1P2	08:48	IFL Total (48) @ 1144.00%	Expand
V1P3	08:48	IFL Total (48) @ 878.00%	Expand
V1P4	08:48	IFL Total (48) @ 1083.58%	Expand
V1N1	08:48	IFL Total (18) @ 912.66%	Expand
V1N2	08:48	IFL Total (24) @ 837.05%	Expand
P105	08:48	IFL Total (40) @ 873.80%	Expand
P106	08:48	IFL Total (40) @ 871.02%	Expand
P107	08:48	IFL Total (40) @ 1016.40%	Expand
P108	08:48	IFL Total (20) @ 894.27%	Expand
P109	08:48	IFL Total (24) @ 863.91%	Expand
P110	08:48	IFL Total (12) @ 872.85%	Expand
P113	08:48	IFL Total (24) @ 858.13%	Expand
P114	08:48	IFL Total (24) @ 876.48%	Expand
DC2			
V2P1	08:48	IFL Total (48) @ 736.37%	Expand
V2P2	08:48	IFL Total (48) @ 846.36%	Expand
V2P3	08:48	IFL Total (48) @ 813.07%	Expand
V2P4	08:48	IFL Total (48) @ 899.11%	Expand
V2P5	08:48	IFL Total (40) @ 892.33%	Expand
V2P6	08:48	IFL Total (40) @ 854.40%	Expand
P207	08:48	IFL Total (56) @ 822.13%	Expand
P208	08:48	IFL Total (64) @ 1065.63%	Expand
P209	08:48	IFL Total (56) @ 1875.48%	Expand
P210	08:48	IFL Total (64) @ 1329.40%	Expand
P211	08:48	IFL Total (44) @ 822.53%	Expand
P212	08:48	IFL Total (44) @ 895.54%	Expand
P213	08:47	IFL Total (40) @ 1173.87%	Expand
P214	08:48	IFL Total (56) @ 1365.42%	Expand
P215	08:48	IFL Total (56) @ 1460.97%	Expand
P216	08:48	IFL Total (40) @ 1382.33%	Expand
P217	08:48	IFL Total (40) @ 874.89%	Expand
P218	08:48	IFL Total (40) @ 108.81%	Expand
P219	08:48	IFL Total (48) @ 856.11%	Expand
P220	08:47	IFL Total (44) @ 899.74%	Expand
C203			
C203	08:48	IFL Total (32) @ 462.11%	Expand
C204	08:48	IFL Total (32) @ 885.28%	Expand
C205	08:48	IFL Total (20) @ 85.26%	Expand
C206	08:47	IFL Total (20) @ 885.34%	Expand
C207			
C207	08:48	IFL Total (24) @ 848.58%	Expand
C208	08:48	IFL Total (24) @ 792.81%	Expand
V2N1			
V2N1	08:48	IFL Total (20) @ 895.03%	Expand
V2N2	08:48	IFL Total (20) @ 1034.47%	Expand
V2N3			
V2N3	08:48	IFL Total (20) @ 498.21%	Expand
V2C1			
V2C1	08:48	IFL Total (24) @ 874.38%	Expand
V2C2			
V2C2	08:48	IFL Total (24) @ 823.02%	Expand
CDL			
VLB1	08:48	IFL Total (52) @ 2848.84%	Expand
VLB2	08:48	IFL Total (36) @ 2868.80%	Expand
VLB3	08:48	IFL Total (40) @ 8273.80%	Expand
VLB4	08:48	IFL Total (38) @ 8291.45%	Expand
VLB5	08:48	IFL Total (48) @ 646.22%	Expand
VLB6	08:48	IFL Total (28) @ 2287.44%	Expand
VLB8	08:48	IFL Total (24) @ 8823.21%	Expand
ZS01	08:48	IFL Total (16) @ 813.72%	Expand
ZS02	08:48	IFL Total (16) @ 9.82%	Expand
VLBX	08:48	IFL Total (3) @ 89.90%	Expand
HIL1	08:48	IFL Total (64) @ 85.85%	Expand
HIL2	08:48	IFL Total (60) @ 892.92%	Expand

Tailorable, expandable, zoomable:

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 (z/VM-Guests)				Linux Nodes (z/VM-Guests)			
LINUX9 (9)			3.93%	RH5X161			0.43%	000000-64			99.22%
suselnx3 (9)			2.57%	RH5Z161			0.37%				
REDHAT (2)			2.30%					LES11T			2.29%
								Linux Nodes (Distributed Servers)			
								PENSUSE 7.68%			

Demo System V4

Demo	13/12/02	18:29	IFL Total (1) 17.77%
Linux Nodes (z/VM-Guests)			
roblx1			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%
sles1lx			0.20%
sles1lx3			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%
vpuz			2.35%

Second level

Tims Test System

Tml2	13/11/27	13:09	IFL Total (1) 0.10%
Linux Nodes (z/VM-Guests)			
			1.85%
			1.50%
			0.85%
redhat56			0.57%

End users define their dashboard(s)

- Linux administrator dashboard provided, everything in one click
- Secure, no need for logon to Linux (no SSH, top)
- Fast and efficient, no restriction on numbers of viewers

Wednesday 7 Nov 2018 00:46 zVIEW Version 4310

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

Menu

mylinux

ESALNXP - Linux Process Con...

Node	Process Name	ID	P
ZSXL0006	systemd	1	
ZSXL0006	kthreadd	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	cio	23	
ZSXL0006	watchdogd	24	
ZSXL0006	kworker/0:1	26	
ZSXL0006	cmmthread	27	
ZSXL0006	kauditd	28	
ZSXL0006	kswapd0	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	qeth_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	CPU	Users	Percents	nice	prty	Storage Metrics (MB)
00:46:00	lxdb2001	*Totals*	0	0	0	0.6	0.1	0.1	0.1	0.3	0	4549 322 4557 0 1391 4.8 3.8 1
00:46:00	lxdb2001	init	1	1	1	0.0	0.0	0	0	0	0	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	0.0	-10	10	29.7 13.4 37.1 0 17.3 0.1 0.0 1
00:46:00	lxdb2001	cron	2223	1	2223	0.1	0.0	0	0.0	0	0	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	50.9 13.9 51.0 0 3.5 0.2 0.1 4
00:46:00	lxdb2001	db2sysc	2833	2831	2833	0.0	0.0	0	0	0	0	877 91.6 877 0 262 0.1 0.1
00:46:00	lxora12	*Totals*	0	0	0	1.2	0.3	0.9	0.0	0	0	3970 724 4157 115 1845 6.6 7.4 4
00:46:00	lxora12	amozxma0	1503	1	1503	0.0	0	0.0	0	0	0	250 10.1 314 0.9 66.3 0.1 0.4 2

ESAHST2 - LINUX HOST Storage Analysis Report - DEMO

Time	Node/Group	Index	<-Utilization-> MegaByte	Pct	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

ESAHUCD2 - LINUX UCD Memory Analysis Report - DEMO

Time	Node/Group	<Real Storage (MB)>	Total Avail	Used	<-SWAP Storage (MB)>	Total Avail	Used	MIN	Avail	CMM	Buffer	Cache							
00:46:00	ZPRO	1400	3	1403	2105	0	2075	2075	10	6	107	5	4320	0	0	353	6	1070	0

ESAHUCD4 - LINUX UCD System Statistics Report - DEMO

Time	Node/Group	<Processor Pct Util>	Idle	<-Swaps>>> Rates (per sec)	<-Disk IO>>> Switch	Intrpt	<-Load A
00:46:00	ZPRO	2.7 1.2 1.4	0 1188	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	180.5 733.9	0.33 0	

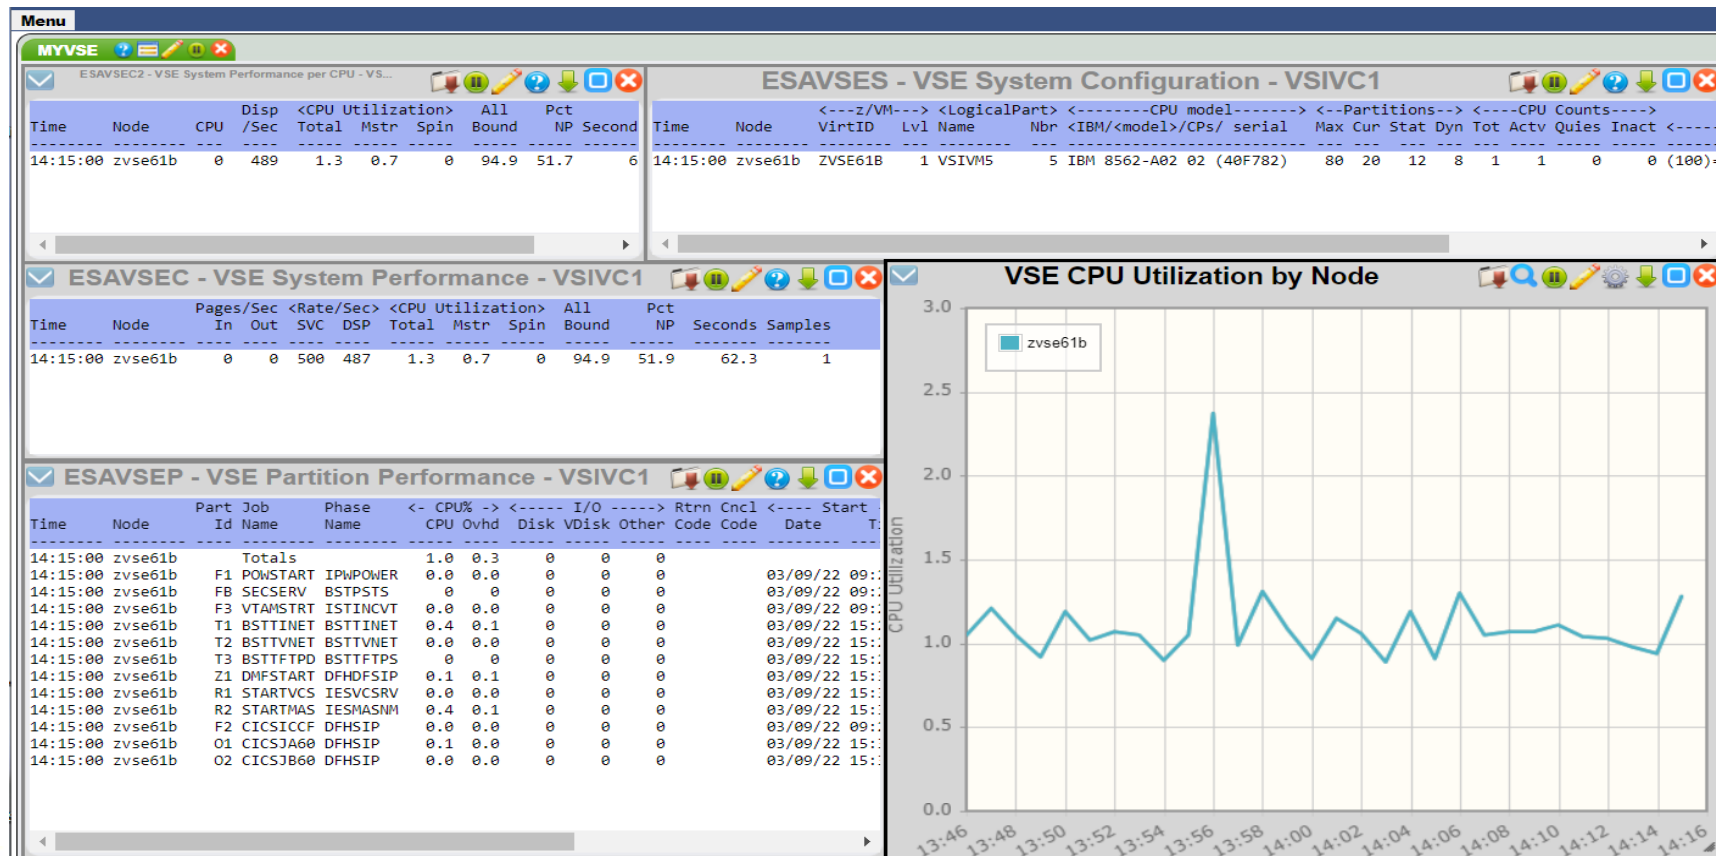
ESAHST4 - LINUX HOST System Statistics Report - DEMO

Time	Server	Num	<Processes>	Stgs=	<---Local---	Date	Time	System	<-System Initiali
			Users	Current	Max (MB)			Uptime	Dev

LPAR

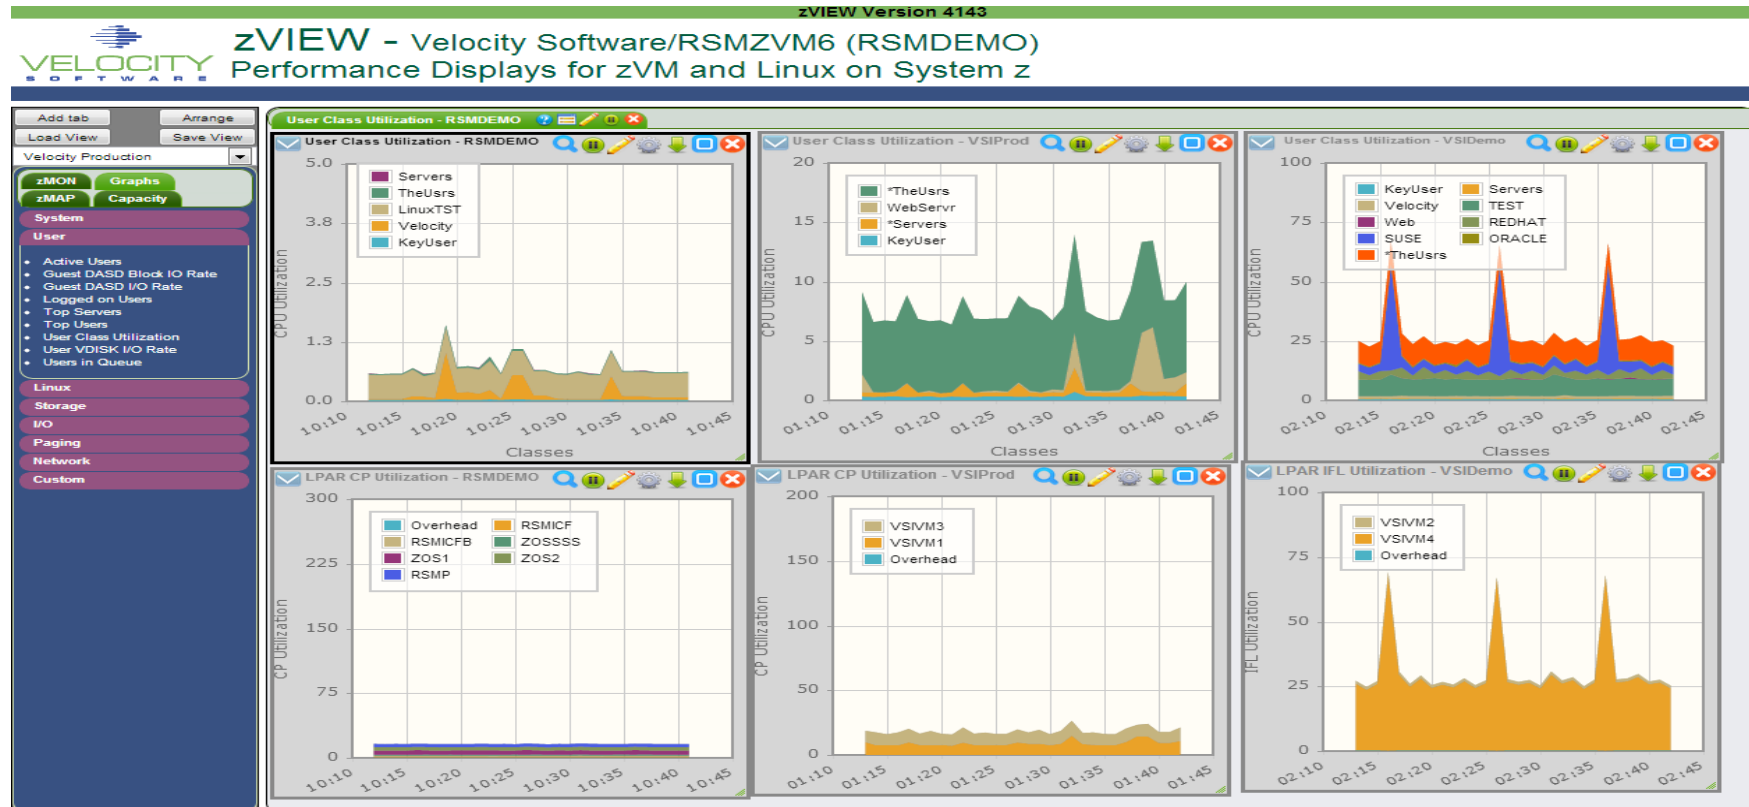
End users define their dashboard(s) – VSE at one click

- Secure, no need for logon to the server
- Fast and efficient, system partitions, jobs, pretty chart at one click




Single pane of glass

- Data from multiple LPARs / geographies
- Menu driven end user designed view



Dynamic Charts

- Data extracted from database dynamically to create graph, example last month




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 24.71% full
LNDX /opt area on S11R20RA is 95.24% full
LNDX /mnt/oracle area on S11R20RA is 30.60% full
LNSU Swap utilization for Linux node is 30.60% full
LNSU Swap utilization for Linux node is 30.60% full
  
```

Or Browser based:
Click through
or SMS, email...



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	/ 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 24.71% full
LNDX	/opt area on S11R20RA is 95.24% full
LNDX	/mnt/oracle area on S11R20RA is 30.60% full
LNSU	Swap utilization for Linux node is 30.60% full
LNSU	Swap utilization for Linux node is 30.60% full
LNDX	/ area on Ixsguar is 90.74% full
LNDX	/usr area on Ixsguar 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 30.60% 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

Single pane of glass – all LPARs console

The screenshot displays a grid of terminal windows, each representing a different LPAR (Logical Partition) in a z/OS environment. The windows are titled with the LPAR name and the console type, such as 'ZOPER - zOPERATOR Console - VSIVC4', 'ZOPER - zOPERATOR Console - CLOUD1', and 'ZOPER - zOPERATOR Console - CUST...'. Each window shows a stream of system messages, including logon events, event manager executions, and system status reports. The messages are color-coded for readability, with green for normal operations and red for errors or warnings. The interface also includes a menu bar at the top and a toolbar with various icons for window management and system actions.

OpenShift data source

- Closed system, performance management not included
- Prometheus support very expensive
- Velocity Software provides containerized SNMP

SNMPd container from Velocity Software

- Very light weight (target overhead .1% one IFL per node)
- Simple and elegant
- One per node
- Standard SNMP data collection with zVPS
- Standard performance management provided
- Supports OpenShift, Rancher, Docker platforms
- Supports zCX!

By CPU Consumption by pod, by container

Report: **ESAK8S2** Kubernetes Resource Utilization Report

```

-----
NODE/          <---Container-->  <--Container CPU----->
Time/ PodName   <---Process ID-->  <-----CPU Percents----->
Date  ContainerName  ProcID ProcName   Tot  sys user syst usrt
-----
rhoscpl
  console-operator-59d
    console-operator      17395 console    0.62 0.12 0.50    0    0
  openshift-controller
    openshift-controller  29430 cluster-   0.38 0.08 0.30    0    0
  kube-controller-mana
    kube-controller-mana  12030 cluster-   1.15 0.15 1.00    0    0
  insights-operator-7f
    insights-operator     14827 insights   0.30 0.05 0.25    0    0
  node-exporter-chc49
    node-exporter         3207 node_exp   0.38 0.13 0.25    0    0
  kube-apiserver-rhosc
    kube-apiserver        3378 watch-te  11.1 1.04 10.1    0    0
    kube-apiserver-cert-  5683 cluster-   0.18 0.08 0.10    0    0
    kube-apiserver-check  6166 cluster-   0.47 0.08 0.38    0    0
  prometheus-k8s-1
    prometheus         1817314 promethe 42.1 0.95 41.1 0 0
    thanos-sidecar        1817435 thanos     0.15 0.02 0.13    0    0
    prometheus-proxy      1817486 oauth-pr   0.57 0.03 0.53    0    0
  prometheus-operator-
    prometheus-operator-  11347 promethe   0.17 0.02 0.15    0    0
  vsi-SNMPd-vk5vd
    vsi-SNMPd           1762314 SNMPd    0.53 0.20 0.33 0 0
  
```

CPU numbers provided by IBM with SMT broken

- ZIIP / zCX has same issue
- See Velocity Software's "capture ratio" presentation
- **VSI Prorated** based on HMC and MFC data

```
Report: ESAUSP5      User SMT CPU Consumption Analysis
Monitor initialized: 03/08/23 at 07:00:01 on 8562 serial 040F78
-----
      <-----CPU Percent Consumed   (Total)-----> <--TOTAL CPU-->
UserID  <Traditional> <MT-Equivalent> <IBM Prorate> <VSI Prorated>
/Class  Total  Virt  Total  Virtual  Total  Virtual  Total  Virtual
-----  -----  -----  -----  -----  -----  -----  -----  -----
07:02:00 414.9  408.0  322.7   317.3  239.7   235.8  208.2   204.7
***User Class Analysis***
OpenShif 355.0 350.3 276.0   272.3  204.9 202.2  178.1 175.7
***Top User Analysis***
RHOSCP1 142.4 140.8 110.1   108.9  82.93  82.01  71.43  70.65
RHOSCP3 125.2 123.8  97.38   96.34  72.35  71.60  62.80  62.14
RHOSCP2  86.79 85.04  68.00   66.64  49.31  48.30  43.55  42.67
```

The Velocity Software Resource Manager based on zVPS

Server modification “happens”

- More CPU and/or RAM needed and must be added
- Application resource requirements grow

Why Excessively large servers?


- That's the way they do it on Intel / VMWare
- Avoid future outage, hardware changes

zVRM, Velocity Resource Manager automates management

- CMM to reduce over sized memory when not needed
- CMM to return memory as workload increases
- Vary vCPU on/offline to meet demand
- Allows definitions of oversized servers to operate efficiently

Requires zPRO APIs, zVPS for data input

Sizing Linux by T-shirt Size




Sizing Linux by T-shirt Size

z/VM can give the system back “free memory”

z/VM provides great feedback mechanisms for memory and vCPU needs

z/VM manages all virtual to real processors assignments



Linux touches all memory assigned to it

Linux can grow and shrink memory and vCPU's

Linux spreads work out over all the vCPU's assigned to it



Often created a death spiral for Linux on IBM z Systems

Velocity Software's new Resource Manager (zVRM) dynamically optimizes Linux configurations for running on IBM z and LinuxONE systems with z/VM

zVRM uses Velocity Software's Knowledge-based Intelligence (KI) to accomplish this

Feedback loops make small incremental adjustments every minute

Reduced resources for your workloads allows you to run more with less providing greater value and quantifiable savings

Too many vCPUs hurts because:

- DIAG 9C overhead
- Linux balances across all vCPUs
- Pollutes Hardware cache
- vCPUs will wait for each other
- vCPUs will wait to be dispatched, but will do little work

Customer Critsit:

- Customer has excess processor capacity
- Bad performance, “CPU WAIT” (ESAXACT)
- Totals by processor type on box (25% utilization - ESALPARS)

```
Totals by Processor type:
<-----CPU-----> <-Shared Processor busy->
Type Count Ded shared Total Logical Ovhd Mgmt
-----
IFL      70   0    70 1706.9  1628.8 30.1 48.0
```

Linux balances small tasks across vCPUs

Each vCPU waits and then does a small task


Having more vCPUs just means the number of delays increases

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
-----
17:01:00  0 430.6 425.0      .      .      . 3480 1682 18 0.2 21
LINUXM1   4 61.53 59.64   REL   400 IFL   240  240 15 1.3 36
  CPU-00   15.89 15.26   REL   100 IFL    60   60 17 1.7 35
  CPU-01   13.77 13.41   REL   100 IFL    60   60 8.3 3.3 38
  CPU-02   17.69 17.19   REL   100 IFL    60   60 15   0 32
  CPU-03   14.18 13.78   REL   100 IFL    60   60 18   0 38
LIMEDI1   6 49.66 48.93   REL   600 IFL   360  360 5.3  0 8.9
  CPU-00    8.85  8.67   REL   100 IFL    60   60 10   0 10
  CPU-01    8.25  8.13   REL   100 IFL    60   60 6.7   0 6.7
  CPU-02    7.24  7.11   REL   100 IFL    60   60 1.7   0 10
  CPU-03    8.47  8.36   REL   100 IFL    60   60 1.7   0 8.3
  CPU-04   10.13 10.04   REL   100 IFL    60   60 8.3   0 10
  CPU-05    6.73  6.61   REL   100 IFL    60   60 3.3   0 8.3

```



zVRM Test System – vCPU Added/Removed

ESAUCD4

ESAUCD4 - LINUX UCD System Statistics Report - VM4

Time	Node/ Group	<Processor Pct Util>				Idle Pct	VCPU Cnt	<-Swaps->		<-Disk IO->		Switch Rate	Intrpt Rate	<-Load Averages->			<-CPU Overhead%->			%IO- wait	Pro- rate
		Total	Syst	User	Nice			In	Out	In	Out			1Min	5Min	15Min	Krnl	IRQ	Inrpt		
13:06:00	redhat75	0.1	0.1	0.0	0	100	1	0	0	0	0	20.9	12.5	0.02	0.42	0.29	0	0	0.02	0	0.90
13:05:00	redhat75	0.1	0.1	0.0	0	99.6	1	3.3	0	35.5	0	24.9	14.5	0.06	0.51	0.31	0	0	0	0	0.98
13:04:00	redhat75	0.1	0.1	0.0	0	99.9	1	0.5	0	4.0	0	20.0	12.3	0.16	0.63	0.34	0	0	0	0	0.84
13:03:00	redhat75	0.1	0.1	0.1	0	131	1	0.1	0	0.4	3.2	27.2	18.4	0.45	0.78	0.36	0	0.01	0	0	0.77
13:02:00	redhat75	0.1	0.1	0.0	0	196	2	17.6	0	693.1	13.9	55.8	46.5	1.24	0.95	0.39	0	0	0.02	0.02	1.18
13:01:00	redhat75	125.6	42.9	82.7	0	0.1	2	23K	22K	203K	179K	6059.0	3215.0	3.37	1.17	0.42	0	1.20	0.08	0.15	0.79
13:00:00	redhat75	67.8	21.7	46.1	0	52.0	1	35K	36K	310K	287K	9129.7	4606.2	2.09	0	0.17	0	3.35	0.15	0.17	1.50
12:59:00	redhat75	0.1	0.1	0.0	0	99.8	1	0	0	0	0	19.9	12.6	0	0	0	0	0	0.02	0	0.99
12:58:00	redhat75	0.1	0.1	0.0	0	99.8	1	0.1	0	1.9	0	19.6	12.3	0	0	0	0	0	0	0	0.67
12:57:00	redhat75	0.1	0.1	0.0	0	99.9	1	0.0	0	0.3	2.5	22.4	13.3	0	0	0	0	0	0	0	0.97
12:56:00	redhat75	0.1	0.1	0.0	0	99.8	1	0	0	0	0	21.0	12.5	0	0	0	0	0.02	0	0	1.11

zVPS PERFORMANCE SUITE

THE WORLD'S LEADING PERFORMANCE MANAGEMENT SOLUTION FOR Z/VM WITH LINUX AND/OR VSE. NEWLY ADDED SUPPORT FOR Z/OS PLUS NO-CHARGE MONITORING FOR X86 LINUX AND WINDOWS.



zPRO CLOUD MANAGEMENT

SOLUTION FOR MODERNIZING THE Z/VM PLATFORM, ADDING ON-PREM CLOUD SUPPORT WITH WEB-BASED CONTROLS FOR Z/VM SYSTEM MANAGEMENT PLUS LPAR CLONING AND ANSIBLE PLAYBOOKS USING VELOCITY'S HIGHLY AVAILABLE SMAPI-FREE API'S.



zTUNE SUPPORT SUBSCRIPTION

VELOCITY SOFTWARE'S ELITE PERFORMANCE SUPPORT SERVICE THAT CREATES AUTOMATED DAILY REPORTS WITH RECOMMENDED ACTIONS



zVRM RESOURCE MANAGER

REAL-TIME MANAGEMENT FACILITY FOR Z/VM WITH LINUX THAT AUTOMATES SYSTEM SETTINGS USING KNOWLEDGE-BASED INTELLIGENCE TO TUNE WORKLOAD REQUIREMENTS AND OPTIMIZE OVERALL SYSTEM PERFORMANCE.





Helping To Make
zPenguins Fly
Since 2001

