

Velocity Software

General Information Guide



1 - What is Velocity Software

Velocity Software was conceived over 35 years ago to measure z/VM system performance. Since then, it has been innovating the market by consistently creating new products and reporting on new performance data in conjunction with new system releases (for z/VM, Linux, etc), products (such as OpenShift), and applications (such as Docker).

Velocity Software has over 100+ years of performance knowledge. This knowledge has been used to continually improve the offered products and create new products to solve performance issues and contribute to overall system and personnel efficiency.

The following sections will show:

- **Page - 4** - Addressing the complexities of performance management
- **Page - 7** - Velocity Software products features and benefits
- **Page - 16** - Modernizing the Mainframe
- **Page - 20** - Screen, Report and Graph categories with examples
- **Page - 38** - A useful performance tool checklist

2 - Velocity Software and Performance Management

Business Requirements of Performance Management

Performance management is not just performance tuning. It includes collecting data to meet other business requirements. Velocity Software addresses many business requirements that often displace multiple products, reducing system overhead and financial costs.

With the blending of traditional z/VM workloads and Linux on z/VM workloads, and the expansion of new applications/technologies, the world has changed. Hardware, software, network connections, system, subsystem and application efficiency and precision are paramount. Velocity Software is a key component to assist enterprises to perform at their best.

There are four business requirements of performance management that Velocity Software products address.

- **Performance Management:** The major performance management requirements are for catching performance problems before they occur (correct system configuration) and solving them when they do occur.
- **Capacity Planning:** This requires complete and accurate data. It is also important to have the ability to integrate the data into comprehensive enterprise capacity planning methodologies.
- **Chargeback / Accounting:** When instituted, it is important to have complete and accurate data categorized by enterprise requirements.
- **Operations:** As mainframe and distributed worlds are coming together in larger and more complex environments, the need for efficient, easy-to-understand operations tools is increasing.

Velocity Software Leads in Performance Management

Velocity Software offers quality products:

- The Velocity Software products are always maintained at the current levels of z/VM.
- The Velocity Software products are efficient - using minimal CPU, storage and DASD.
- The Velocity Software products are easily installed and maintained across enterprise systems.
- Velocity Software is always looking at what is new and upcoming and integrates those features into their products.
- Velocity Software uses the CP Monitor records and is constantly verifying the accuracy of the data collected and reported. Velocity Software was the first to discover the initial issues when running Linux under z/VM!
- All collected data can be easily grouped by users, classes, accounting information or servers to aid in problem determination, accounting and chargeback.
- Whether an operator, a systems programmer or upper level management, Velocity Software products are intuitive, easy to use and provide tools for personnel at all levels.

Velocity Software offers education:

- Velocity Software attends industry conferences to support customers.
- Velocity Software holds seminars to provide education to their customers.
- Velocity Software offers a Tuning Guide and case studies on the website to provide easy to follow performance tuning tips.

Velocity Software honors their customers:

- Customer performance issues are given agile attention via opened problem tickets.

- Customers can request enhancements. These are always given full consideration.
- Both technical and management level customers of Velocity Software products will find the necessary data and tools to make their jobs easier.

Velocity Software is innovative:

- The Velocity Software zPRO product was created to build and maintain private clouds.
- The Velocity Software zPRO product was enhanced to assist system programmers, operators and users to utilize browser-based tools for ease-of-use and efficiency.
- New applications and products such as Docker, Kubernetes and Grafana are constantly being integrated into features of the Velocity Software products.

3 - Velocity Software Products

Product offering summary

The ZVPS suite:

- **zMON** - a highly efficient performance monitor.
- **zWRITE** - a highly efficient real-time monitor data collector.
- **zMAP** - a report writer and performance database creator.
- **zTCP** - provides data acquisition of network data.
- **zVWS** - an extremely fast webserver running on native z/VM.
- **zVIEW** - a browser-based interface for viewing collected data.
- **zPORTAL** - a browser interface to support ZVPS products.
- **zOPERATOR** - a CMS console and operations manager.
- **zALERT** - provides operational support by monitoring and alerting.

Additional products:

- **zPRO** - Velocity Software's modernization product.
- **zTUNE** - a service offering that provides performance tuning.
- **zVRM** - a new resource management product.
- **zOSMON** - data and reports for monitoring z/OS systems.
- **VSEMON** - data and reports for monitoring VSE systems.

Product Features and Benefits

zMON/zWRITE

The zMON and zWRITE components of zVPS work with CP monitor records to gather a huge variety of performance data for use in real time displays and to create a performance database. This data is used for performance monitoring and reporting. Features and benefits include:

- **Hierarchical design** that groups displays by overview, system, user, subsystem and application data. This provides an easy-to-follow top down view of performance data.
- **Complete z/VM system reporting** including statistical summaries such as minimum, maximum and standard deviations applied to hardware, software, user, application and network data which is critical for performance management.
- **Correct calculations of performance data** especially when Simultaneous Multi-Threading (SMT) is active. This is critical to understand current resource usage and capacity planning.
- **Evaluation and displays of system resources by server, by user, accounting code and by class**, which allows for grouping related servers and users together for an easy ‘view at a glance’ summary of the system(s). This also makes detail information easier to find. For example, grouping all SUSE Linux systems together would make it easy to see if that group was utilizing the highest amount of CPU at a particular moment (summary) but also which of the SUSE systems is at the top of that list (detail).
- **Alerts** can be set for a plethora of conditions and alert methodologies. Colors on displays can also be used to indicate problems or potential problems.
- **Data is easily extracted** for specific performance evaluation and for use in other applications, such as MICS or MXG.
- **Data is gathered from the CP monitor** - The CP monitor data is the most stable and supported way to pull performance data.
- **Product upgrade time and costs are eliminated** when moving from one z/VM release to another.
- **There are large resource savings** based on how zMON and zWRITE collect data. Savings show in lower main storage, CPU and DASD requirements than other competitive products when performing identical functions.

zMAP

The zMAP component of zVPS is the report generator. It takes the CP monitor data from the files created by zMON and zWRITE to create daily reports that are easily customizable and ad hoc reports that can be used for capacity planning, accounting and charge back. Features and benefits include:

- **Many of the same features as zMON** - a hierarchical design; complete z/VM system reporting; correct calculations; grouping of information by server, user, accounting code and class; resource savings; and ease of upgrades between releases of z/VM.
- **The zMAP reports are nearly identical to zMON displays.** The differences are based on the timing of the data (some data is only used real time and other data needs to have gathered data so it is only in a report). This allows users of both to easily vacillate between the two with no additional training.
- **Service Level Agreement (SLA) reporting** provides the ability to measure user response time objectives against actual response times. This data is reported in both percentages and averages.
- **An extract facility (ESAEXTR)** which allows specific data to be gathered from history files. This data can be used to immediately calculate particular trends in the system or for exporting to other applications such as ESAPLOT or Excel.
- **MICS and MXG** ready data can also be created for use by other departments in the enterprise.
- **A chart creation facility (RUNCHART)** that is run nightly to create daily, weekly, monthly and rolling 30-day graphs that can be viewed via zVIEW. Many combinations of variables and time frames can also be used to create unique graphs.

zTCP

The zTCP component of zVPS is the facility that collects network data. Data is collected from most any system that is connected to z/VM that can run Simple Network Management Protocol or SNMP. Features and benefits include:

- **z/VM network data is measured and reported.** This gives performance data for the z/VM system.
- **Any network accessible system data is measured and reported.** Linux data, z/OS data, Unix data, anything that can connect to z/VM and run SNMP can have data shown in zMON and zMAP. This allows for all attainable performance management data to be accessed in one set of displays/reports/graphs.
- **Velocity Software has its own Management Information Bases or MIBs.** This allows for better data gathering (such as parent/child relationships in Linux).
- **zTCP is secure.** It collects data via SNMPv3 which is the version of SNMP that provides security.

- **All data from zTCP moves seamlessly into zMON.** This allows for one single data stream for zMON to process which enhances low resource utilization.
- **Alerts can be utilized.** zTCP provides a facility for outbound SNMPv3 alerts to integrate into other management consoles.

zVWS

The zVWS component of zVPS is the facility that is an extremely fast and efficient z/VM based web server. This is a very powerful tool. Features and benefits include:

- **The zVWS web server is z/VM based.** This facilitates the viewing of performance data from a browser and is easily maintainable on z/VM itself.
- **zVWS provides the ability to code additional browser based applications.** The web based applications that can be created are endless. CGI's are supported for assembler, REXX and PL/I. This also allows modern interfaces to be added to legacy applications.
- **zVWS provides security.** SSL (DES and TripleDES) are supported. Security can also be done by userid/password.

zVIEW

The zVIEW component of zVPS is the facility that is used to view performance data. This includes zMON real time data, zMAP report data, graphs, logs and alerts. Features and benefits include:

- **zVIEW is the browser interface that displays Velocity Software collected data.** Any performance data that has been gathered can be displayed on a 3270 screen or via zVIEW on a browser. Less z/VM experienced people and management tend to appreciate the ease of use of a browser interface to view performance data and graphical representations of the data. 'Point and click' is very easily learned.
- **zVIEW screens show a large amount of data.** Besides showing real time and reporting data, zVIEW has standard and customizable graphs, shows log data, alerts, zTUNE data (if installed), capacity reports and 'one pain of glass' views of the enterprise.
- **zVIEW is very customizable.** It can be arranged to show particular users and particular data. Management can have a customized screen that only shows high level information and a server administrator

can have a view of more data, but only for that server. It also has the ability to show certain time periods - like the last few hours that a particular problem was occurring.

- **zVIEW has a wide scope.** It can gather data from other systems/ LPARS that are being monitored to put all pertinent data on one screen (examples in the next section).
- **zVIEW is secure.** It can be organized so that each user can only see the data they are allowed to see. This is especially helpful for businesses hosting multiple customers.

zALERT

The zALERT component is an included feature of zVPS that has been added to support alerting. Features and benefits include:

- **zALERT provides full function alert capabilities.** Any performance metric provided by zVPS for z/VM, Linux, z/OS, VSE, network and applications can be used as input for an alert.
- **zALERT is easily customizable.** It is designed to integrate into any enterprise management console such as an operator console. Alert messages can be seen on the zALERT display on zMON; sent to CMS users; sent as an email; sent as a text or prompt an SNMP trap to be sent to a management console.
- **An alert can also be sent to or received from other LPARS.** This adds to the powerful displays on zVIEW.

zOPERATOR

The zOPERATOR component is an included feature of zVPS that has been added to support console operations. Features and benefits include:

- **zOPERATOR provides a scrollable and searchable console.** Messages can be from z/VM or configured to receive messages from Linux or other applications.
- **zOPERATOR is easily customizable.** Colors and highlighting can be used to make important messages stand out. Messages can also be 'held' to keep a critical message from scrolling off the screen.
- **As with zALERT, can handle automated actions.** Messages, emails, texts, etc can be created/executed based on the criteria coded.

zPORTAL

The zPORTAL component is an included feature of zVPS that has been added to support a web-based front end for configuration and display of zVPS components. Features and benefits include:

- **zPORTAL provides the ability to maintain zVPS via a front end.** It is easy to ‘point and click’ to see the status of the different zVPS products and analyze logs without having to log on to a z/VM system. The current versions of each product are also easily displayed across nodes.
- **zPORTAL supports unattended installs.** This provides the ability to upgrade an existing product, update product keys or send a product package to an RSCS connected system. This allows work across an SSI cluster to be done from one logon.
- **zPORTAL can display the status of zTCP defined notes.** Nodes can be viewed, stopped, started or restarted. A software list can be generated and displayed for each active node.
- **Web server logs can be displayed via zPORTAL.** The logs from zVWS can be analyzed along with showing hit counts for different criteria such as Path, File, IP, etc.
- **z/VM help files can be displayed on zPORTAL.** This allows an easy way to get z/VM information without having to log directly onto a z/VM system.

zPRO

The zPRO feature of the Velocity Software products provides a solution for implementing private clouds as well as providing an easy-to-use web-based front end for system programmers managing z/VM environments. zPRO’s full capabilities are explained in the next section on **Page - 16**.

zTUNE

The zTUNE feature of the Velocity Software products provides a mechanism to allow installations that do not have z/VM performance expertise access to basic performance solutions and over 35 years of performance knowledge. Features and benefits include:

- **When installed, zTUNE reports performance suggestions.** Real time data and daily reports are created that show common performance problems with possible solutions or analysis tips. For example; User xxxxx has a high diagnose rate - check ESADIAG and ESAUSRD to evaluate.
- **zTUNE has access to over 100 known performance issues.** This allows for fast problem resolution.
- **When zTUNE is acquired, Velocity Software offers accelerated performance assistance.** Velocity Software's performance experts will quickly analyze and resolve performance problems.

zVRM

zVRM is a NEW feature of the Velocity Software products that provides a resource manager that does automatic system tuning. This is especially helpful when the specifications for a Linux system are done from a more x86 perspective. Often systems are built too large or too small to performance efficiently when running on z/VM. Features and benefits include:

- **z/VRM is a powerful tool for managing Linux server configurations.** This resource manager will monitor Linux systems and work with z/VM to manipulate resources based on workload.
- **z/VRM monitors virtual CPUs.** When a system has virtual CPUs not in use, these will seamlessly and automatically be removed from the Linux system and given to z/VM for processing power elsewhere. If a system is in need of more cycles, virtual CPUs are added to the Linux system. This allows better sharing of resources between systems.
- **z/VRM monitors storage.** z/VRM uses Collaborative Memory Management (CMM) to adjust storage as needed for the current workload.
- **z/VRM is easily customized and maintained.** All settings are done via zPRO and are easily added and updated.
- **z/VRM saves money.** Since CPU processing power and storage are expensive commodities, being able to take an incorrectly configured system and have it seamlessly perform with the correct resources is of paramount importance.
- <https://velocitysoftware.com/products/zvrm.html> has more details on this powerful product!

zOSMON

The zOSMON feature of the Velocity Software products provides an interface to z/OS system SMF data. Features and benefits include:

- **zOSMON can monitor many z/OS systems.** Across an enterprise, there may be many LPARs and z/OS systems. All of these can be monitored with one connected zVPS subsystem.
- **zOSMON is precise.** It adds the ability to do real time performance monitoring with a one minute interval granularity - with negligible performance impact.
- **zOSMON works with zVIEW and zALERT.** Any z/OS system data that is collected can be seen on zVIEW and used as input to zALERT. This means performance data can easily be viewed along with z/VM, Linux and other application data, and incorporated into enterprise alerts.
- **zOSMON also reports on z/OS subsystems.** Subsystems such as CICS and DB2 are included in performance data gathering and reporting.
- **zOSMON saves money.** The overhead for performance monitor data gathering on z/OS is minimal. The processing of z/OS data is done on the less expensive z/VM processors instead of the more expensive z/OS MIPS.

zVSEMON

The zVSEMON feature of the Velocity Software products provides a performance interface to VSE/VSEn system performance data. Features and benefits include:

- **zVSEMON uses SNMPv3 to gather performance data.** System level data, partition level data and even job level data is collected for real time viewing and processed into daily reports.
- **zVSEMON works with TCP/IP data.** Velocity Software has worked with CSI International and Barnard Software to gather and process VSE network data.
- **zVSEMON also reports on the CICS subsystem.** CICS DMF records are sent to zVPS and are processed for viewing and reports.
- **zVSEMON works with z/VSE and VSEn.** Velocity Software works independently of which VSE provisioner is used.

Summary

It is easy to see, this is a powerful suite of products that has robust abilities and keeps up with the ever changing environment that is the mainframe.

4 - Modernizing the Mainframe

Velocity Software strives to keep the z/VM platform alive and successful. To this end, the z/PRO product was released and is continually enhanced. z/PRO covers some of the biggest hurdles for today's complex and often "lean" staffed environment:

- Provides a fast and easily-installed private Platform as a Service (PaaS) environment on z/VM - without the complexity of such products as SMAPI or Java.
- Provides a means to delegate functions to users that reduces the time highly skilled/highly paid personnel are required.
- Provides a web-based, intuitive tool collection for systems programmers and administrators to simplify tasks and enhance productivity, especially for new and/or inexperienced personnel.
- zPRO provides security to manage user and resource privileges.
- Providing a cloud-based facility for mainframe training.
- Provides an extremely easy way to maintain the Velocity Software products.

z/PRO Features and Benefits

General Benefits

Management Benefits include:

- z/PRO tracks system resources
- Managers can view these resources by individuals or groups
- Managers can view/audit system usage and provisioning
- z/PRO creates a mechanism for a more responsive end-user environment

System Administrator/Programmer Benefits include:

- z/PRO simplifies System professionals tasks
- Systems personnel can easily determine if end-users are authorized
- Systems personnel can easily determine if the required resources are available

- Systems personnel can easily create new “golden copy” systems to clone
- Systems personnel can easily start/stop servers as needed
- Systems personnel can easily adjust end-user server characteristics
- Systems personnel can easily migrate servers between SSI clusters

Operations Benefits include:

- z/PRO provides a means for the Operation staff to manipulate servers for scheduled outages, maintenance, load balancing, etc. It provides a simple browser-based, all-in-one-place tool that adds to personnel efficiency.

End-User Benefits include:

- z/PRO gives the end-users immediate access to provision servers, request resources, start/stop servers, etc, all without the intervention of those that might not have the time to action their needs.

Cloud Provisioning

z/PRO was used to create a brand new cloud environment on newly installed, bare metal hardware during a computer center migration. z/PRO allowed the exercise to take only two days.

z/PRO is used to provision and clone guest systems. This feature saves the system programmer/administrator valuable time by allowing authorized users to provision their own test systems based on their authority level. This feature is highly customizable for centralized or distributed cloud management, “gold image” creation, system characteristics (such as memory size, etc) and other capabilities (such as provisioned system termination). Currently available systems:

- Linux servers - SLES, RedHat, Ubuntu, etc
- VSE guests
- CMS users
- 2nd Level z/VM guests
- Coming soon - z/OS guests!

System Programmer Tool Highlights

z/PRO information. There are multiple reports and configuration files that can be viewed and/or updated via z/PRO about z/PRO.

z/PRO handles z/VM and CMS management. Simple but time consuming tasks such as changing passwords, adding links or minidisks, updating virtual CPU counts and more can be done on z/PRO.

z/PRO shows resource usage for users/servers. Server status, server resource usage, server performance and IP addresses can all be viewed and actioned.

z/VM Administration Management tasks on z/PRO. Many different often time consuming tasks such as EDEV management, RSCS link management, storage server management and more are all in one place to increase system programmer efficiency.

Creating and expanding Shared File Systems. This is often a very time consuming task that takes a lot of reading of manuals to do. z/PRO makes it easy.

z/PRO Scheduler. z/PRO can manage the scheduling of events such as file backups, report generation, data collections, console actions and more.

z/PRO Backup and Restore. Daily, weekly and monthly backups can be created and scheduled. z/PRO can also be used to restore any resource that has been backed up by z/PRO. This is critical for certain z/VM resources.

z/PRO zSPOOL. Console files, whether open or closed, can be viewed and actioned with z/PRO's zSPOOL product. This is a major time saver as open spool files usually had to be closed to be viewed.

z/PRO zDIRECT. This allows systems programmers an easy way to create and update the z/VM system directory. Directory maintenance tasks are done via a browser, saving time and simplifying the process for the less experienced.

Security Management can be done on z/PRO. RACF commands and reports are simplified on z/PRO.

z/PRO has a Linux command interface. This is an easy way to communicate with the Linux systems.

Virtual Network Management. z/PRO simplifies the management of guest any LAN and VSWITCH.

Manage Kubernetes and Clusters. z/PRO can easily manage and create new Kubernetes clusters.

z/PRO can now access the HMC. z/PRO can be used to maintain the HMC configuration.

There are many more aspects of z/PRO that can be explored on the Velocity Software education site - www.velocitysoftware.com/educate

5 - Velocity Software Reports

Real Time Data (zMON) Screens and Historical Data (zMAP) Reports:

All of the Velocity Software reported data is ordered intuitively by function in a top down/basic to detailed hierarchy to lesson complexity and aid in fast problem resolution. There is also a handy index of screens and a table of contents for reports. Both real time data and reports can be viewed on a 3270 or via a browser by using zVIEW. It is then easy to use both real time and historical information to chase any performance issue. zVIEW also has the ability to do historical data from designated time frames to for better problem determination.

The **Tuning Guide** in the customer area highlights many of these reports and also has tips and tricks to allow even inexperienced system administrators to solve performance problems, or better yet, prevent them from happening.

Historical data is produced by zMAP and is broken out into the same categories as real time data. The majority of reports match the real time screens except when real time data is no longer relevant (such as current user data) or if a report is created with totals that aren't relevant in real time.

Collected data is categorized below:

System data - these show an overall view of the system

- Overall system view
- System configuration - Example shown below
- Operator consoles
- zTUNE data (if installed)

Service Level Analysis - these show how the system is responding

- Service level information
- Transaction analysis (Example shown below)

User data - these are a compilation of user/server data

- User/Server configuration data

- System resource use by user/server (Example shown below)
- Top user analysis

Shared File System - These show an overview of all of the shared file servers.

CPU - these show CPU activity - critical for good system performance

- System CPU consumption by LPAR (Example shown below)
- System CPU consumption by virtual/real CPU
- SMT information (if active)
- Diagnose commands
- Cryptography information
- CPU cache management

Main Storage - these show storage subsystem (memory) activity

- Overall storage information
- Shared address spaces data
- Virtual disk usage
- NSS/DCSS information
- Memory extensions
- System execution space data

Paging and Spooling - these show paging and spool activity

- Paging information
- Block paging information
- Paging/Spooling device activity

Input/Output Subsystem - these show the health of the DASD

- DASD device configuration/activity
- Channel configuration/activity
- Emulated device activity (EDEV, SCSI)
- FCP device activity
- DASD cache activity
- Minidisk cache activity

Network - these show information about the network

- Node configuration
- Information from all layers of the TCP/IP network
- VSWITCH activity
- NIC activity

Linux - these show information about Linux guests - three MIBs are represented throughout: HOST, UCD and an expanded Velocity Software MIB (which gives the parent/child data that is critical for process analysis)

- Defined Linux nodes/status
- Linux system process/application information (Example shown below)
- Linux resource usage - CPU, memory, I/O activity
- Linux file system activity
- Linux network activity

Linux Application - These show information for many applications that run on Linux

- GPFS - General Parallel File System
- Java Virtual Machine
- Kubernetes (Example shown below)
- MongoDB (Example shown below)
- Oracle

SSC/Docker - These show information about containers

- SSC - Secure Service Container from IBM
- Docker container configuration
- Docker container resource use - CPU, storage, paging (Example shown below)
- Docker file system information

zOS - These show information about z/OS systems collected from SMF data

- CEC information
- LPAR configuration and activity
- CPU utilization (Example shown below)
- Job and Job step activity/resource use
- Page dataset usage

VSE - These show information about VSE systems

- Configuration
- System activity (Example shown below)
- Partition activity

CICS - These show information about CICS systems

- System activity
- Transaction activity
- DSA information
- TCB activity

DB2 - These show information about DB2 systems

- Buffer analysis

- EDM pool information
- CPU usage
- Latch counters
- Storage usage
- Application and Transaction information
- DDF - Distributed Data Facility information

zALERT - These show alert messages from any system that is set up with zALERT

zOPERATOR - This shows the zOPERATOR console with customizable highlighting

Real Time Graphs:

Often a picture is worth a thousand words. zVIEW has default graphs built to create pictorial representations of important data. Customized graphs can easily be created for specific users (such as management) who only want particular information in a concise picture. Default graphs are:

- **System** - CPU, CPU Pool, LPAR and Spool utilization information (Example shown below)
- **User** - User and guest information
- **Linux** - CPU, File system, memory, Open Shift and swap information (Example shown below)
- **Linux Applications** - Java, Kubernetes, MongoDB and Oracle information
- **Storage** - Main, resident, user, VDISK storage information
- **I/O** - CHPID, DASD and MDC information
- **Paging** - Paging information
- **Network** - FCP, I/O, QDIO and VSWITCH information
- **z/OS** - CPU, I/O, ICF, LPAR and USS information
- **VSE** - CPU and CSI information
- **CICS** - DSA, task, response time and transaction information

Enterprise Main View:

Once the screens, reports and/or graphs have been reviewed, a condensed, customizable overall system view is provided by z/VIEW. This can be adjusted to show any screen/graph. Put all of your most important information on the front page of z/VIEW!

System Examples:

Below are some examples from zVIEW with explanations of the data:

ESAHDR

This is an example of the ESAHDR screen that shows the system configuration. The details in this screen/report show pertinent configuration information such as the processor model, if IFL's are being used, how many CPUs are on the system, SMT status, current DCSS size and more. Occasionally performance problems have been solved just by looking at this information!

```

Key expiration for VELOCITY      24/11/02
Components licensed:      ZOS TUN PRO

Monitor period started:      23/12/26 13:49:00      ZWRITE rel: 5150

z/VM version:      Version 07.3.0 SLU 2301      System: VSIWM4
LPAR:      VSIWM4

Time of last termination:      23/09/07 05:57:07
Abend code of last termination:
Time of last IPL:      23/09/07 05:57:37
Time zone adjustment from GMT:      05:00:00 West
System operator:      OPERATOR
Processor model:      z15-T02 8562-A02
Multithreading status(SMT):      Enabled
  Core Thread count:      2
  Enabled count:      2
System sequence code:      0000000000040F78

CPU capability factor:      8563
CPU(GP) capability factor:      8563
CPU cycles/ns (GP):      366
CPU cycles/ns (IFL):      4500

Operating on IFL processor(s)
Channel Path Measurement Facility(CPMF) Extended installed

Totals by Processor Type:
<-----CPU-----> <-Shared Processor Busy>
Type Count Ded Shared Total Assigned Ovhd Mgmt
-----
CP      2  0    2  87.8    87.5  0.2  0.2
IFL     4  0    4  90.5    88.7  1.5  1.8
ZIIP    1  0    1  17.4    17.3  0.0  0.0

Number of logical partitions defined:      13

Main storage installed (GB):      27.00
Main storage generated (GB):      27.00

Horizontal/Vertical Scheduling      GP CPUs      IFL CPUs
Unparking set to Medium
Excess use      .      Medium
Ceiling prediction CPU utilization:
  Confidence      .      95%
  Algorithm      . Std adjusted
Ceiling prediction of T/V ratio:
  Confidence      .      90%
  Algorithm      . Standard
Max parked CPUs      80
CPU Fudge Factor (SRM CPUPAD)      .      1.000
Horizontal capacity      .      1.000
CPU cap      .      .

Number of users monitored:      403
Number of DASD monitored:      408
Number of non-DASD monitored:      752

Monitor DCSS usage (pages)      Total      Used (%)
Sample configuration      500      177  35
Event configuration      50      .   0
Sample data      19076      921  5
Event data      4950      wrapped  0

EMON DCSS      16383      2954  18

```

Figure 2-1 ESAHDR Screen example

In the ESAHDR report at the end of the day, the same static information is shown, but it also then contains useful statistics such as the total CPU utilization for the day and what the top ten users were that contributed to that percentage. For total CPU utilization, it will be a percentage of the full capacity (ie it shows 96.6 of 200% meaning there are two CPUs for 200% and 96.6% was the utilization for the day)

```

CPU utilization:          96.6 of 200%
.
Top users and user classes by CPU consumption:
  UserID      <-Relative->      <---Absolute Percent CPU--->
  /Class      <-Pct CPU-->      <Out of 100%> <Out of 200%>
              Used   Cum      Util   Cum      Util   Cum
  -----
1. SLES15      41.2   41      9.9   10      39.8   40
2. MONGOO1     11.5   53      2.8   13      11.1   51
3. SLES12       8.2   61      2.0   15       7.9   59
4. RHOSCP1     7.6   68      1.8   17       7.4   66
5. SSNODE1     4.4   73      1.1   18       4.3   70
6. SSNODE3     3.9   77      0.9   19       3.8   74
7. SSNODE2     3.9   81      0.9   19       3.7   78
8. S11S2ORA    3.5   84      0.9   20       3.4   81
9. System      3.3   87      0.8   21       3.2   84
10. RHOSWK1    1.6   89      0.4   22       1.6   86

```

Figure 2-2 ESAHDR Report - CPU Utilization example

ESAXACT

The ESAXACT report is one of the most helpful screens in the Velocity Software arsenal. It shows what resources are causing users/servers to wait. This can show CPU issues, DASS issues, storage issues, paging issues, resource thrashing issues and possible configuration issues (such as a SHARE setting that is too low) and more. On zVIEW, if classes are defined, the data will show a summary for the class. Clicking on the class, shows the members of that class individually. Reports will also show the top users for each 15 minute increment and summary totals for the day.

Time	UserID /Class	<-Samples->		<---Percent non-dormant----->										non-dormant----->		Times					
		Total	Pct	In Q	Run	Sim	CPU	SIO	Pag	SVM	I/O	Pag	Ldg	Lim	Pct	E-	T-	Tst	D-	I/O	Throt
14:05:00	System:	7020	31.8	2.4	0	1.5	0	0	0.4	0	0	0	0	0	0	0	0.4	0	9.4	0	41.4
14:05:00	OpenShft	120	100.0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	100	0	0	0.1
14:05:00	ORACLE	180	99.4	0.6	0	2.8	0	0	0	0	0	0	0	0	0	0	0	97	0	0	1.1
14:05:00	RANCHER	360	83.3	14	0	3.0	0	0	0	0	0	0	0	0	0	0	0	83	0	0	31.0
14:05:00	RANCHS1	120	100.0	20	0	2.5	0	0	0	0	0	0	0	0	0	0	0	78	0	0	15.0
14:05:00	RANCHA1	120	100.0	13	0	3.3	0	0	0	0	0	0	0	0	0	0	0	84	0	0	12.3
14:05:00	RANCHA2	120	50.0	6.7	0	3.3	0	0	0	0	0	0	0	0	0	0	0	90	0	0	3.7
14:05:00	REDHAT	540	52.4	0	0	0	0	0	0	0	0	0	0	0	0	0	0	100	0	0	0.3
14:05:00	Servers	1740	1.7	0	0	0	0	0	13	0	0	0	0	0	0	13	0	87	0	18	0.2
14:05:00	SUSE	480	74.2	1.1	0	0	0	0	0	0	0	0	0	0	0	0	0	99	0	0	4.1
14:05:00	TheUsrs	1740	47.4	0.6	0	2.1	0	0	0.2	0	0	0	0	0	0	0.2	0	97	0	6.6	3.5
14:05:00	TEST	60	100.0	0	0	5.0	0	0	0	0	0	0	0	0	0	0	0	95	0	0	0.4
14:05:00	UBUNTU	60	100.0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	100	0	0	0.1
14:05:00	Velocity	480	3.8	0	0	0	0	0	11	0	0	0	0	0	0	11	0	89	0	0	0.8
14:05:00	Web	1260	0.1	0	0	0	0	0	100	0	0	0	0	0	0	0	100	0	19	0	0.0

Figure 2-3 ESAXACT Screen example

ESAUSR2

The ESAUSR2 data shows virtual machine resource use. This screen allows the top users of resources to be seen at a glance. CPU, storage, paging and spooling resources are all shown. This information quickly shows a user/server/class resource usage overview that leads to the next step of problem determination. The report again shows summary data for each 15 minute interval in a day, plus daily totals.

ESAUSR2 - User Resource Utilization - VM4																				
Time	UserID /Class	<---CPU time-->		<-Main Storage (pages)-->				<-Paging (pages)-->			<Spooling(pages)-->			Q'd	Total	<VMDBLK Rebalances>				
		Total	Virt	Resid	Lock	Paged	Out	Read	Write	Alloc	Read	Write	Spl			CPU	sec	NL1	NL2	NL3
17:17:00	System:	36.51	35.54	1.0	6585K	6580K	8986	5000	8058K	30	0	1028K	0	48	0	4863899	0	0	0	0
17:17:00	RANCHER	18.73	18.35	1.0	2293K	2293K	161	0	1052K	0	0	57	0	0	0	1361345	0	0	0	0
17:17:00	ORACLE	12.23	12.08	1.0	267K	267K	726	0	438K	2	0	23391	0	0	0	1603333	0	0	0	0
17:17:00	SUSE	2.23	2.15	1.0	1341K	1341K	1524	0	1519K	4	0	86	0	0	0	769109	0	0	0	0
17:17:00	SLES12	1.55	1.53	1.0	925K	925K	189	0	566K	0	0	17	0	0	0	366201	0	0	0	0
17:17:00	SLES15	0.30	0.29	1.1	92658	92658	122	0	150K	0	0	19	0	0	0	316582	0	0	0	0
17:17:00	SLES11X4	0.16	0.11	1.4	39211	39211	682	0	74223	1	0	5	0	0	0	20967.1	0	0	0	0
17:17:00	LXDB2001	0.12	0.12	1.0	46296	46296	52	0	166K	3	0	4	0	0	0	38232.3	0	0	0	0
17:17:00	SLES8X	0.06	0.05	1.1	127K	127K	191	0	245K	0	0	28	0	0	0	14941.9	0	0	0	0
17:17:00	LXSUGAR	0.04	0.04	1.0	73755	73755	99	0	235K	0	0	8	0	0	0	11141.1	0	0	0	0
17:17:00	SLES11X2	0.00	0.00	1.1	37532	37532	189	0	83264	0	0	5	0	0	0	1043.8	0	0	0	0
17:17:00	TheUsrs	2.03	1.85	1.1	1656K	1655K	847	0	2907K	22	0	36015	0	0	0	758079	0	0	0	0
17:17:00	Velocity	0.55	0.53	1.0	6180	5675	6	5000	13601	0	0	431K	0	24	0	63067.5	0	0	0	0
17:17:00	TEST	0.21	0.16	1.3	30495	30495	24	0	64910	0	0	0	0	0	0	56802.3	0	0	0	0
17:17:00	Servers	0.20	0.10	2.1	22552	21137	1412	0	96331	0	0	470K	0	2	0	47091.2	0	0	0	0
17:17:00	REDHAT	0.17	0.16	1.0	451K	451K	4181	0	639K	1	0	164	0	0	0	153708	0	0	0	0
17:17:00	OpenShft	0.08	0.08	1.1	479K	479K	54	0	1238K	0	0	0	0	0	0	41199.4	0	0	0	0
17:17:00	UBUNTU	0.04	0.04	1.0	31886	31886	51	0	59194	1	0	7	0	0	0	9663.3	0	0	0	0
17:17:00	Web	0.04	0.04	1.1	8142	5529	0	0	31579	0	0	67226	0	22	0	523.2	0	0	0	0

Figure 2-4 ESAUSR2 Screen example

ESALPAR

The LPAR configuration and usage data gives another high level picture of overall system performance information. How many CPUs are on the system? What is each processor's weight and polarization? How much steal time is there? All of this is important to understand where processing power is going - and if it is efficient. Logical processor data is reported for each LPAR on the CEC. CPU utilization information is reported for the LPAR that is running zMON/zMAP.

Time	CEC	<-Logical->	LPAR	<-----Logical Processor----->							<-----CPU (percentages)----->					<Multi-thread>			
	Phys CPUs	<Partition> No	Pool Name	<----CPU---->	<%Assigned>	Weight	Cap	Abs	Wt	Total	Emul	User	Sys	Idle	Stl	Idle	Time	cp1/cp2	
				Type Cnt ID	Total Ovrhd	/Polar	ped	Cap	Cmp	Util	Time	Ovrhd	Ovrhd	Time	Pct	Time			
14:52:00	7	Totals:		CP 7 Tot	41.8 0.2	225													41.62
14:52:00	7	Totals:		IFL 11 Tot	86.5 1.5	740													65.65
14:52:00	7	01	VSIVM1	IFL 1 Tot	1.0 0.0	25	No	No	No										0.95
14:52:00	7	02	VSIVM2	IFL 1 Tot	3.2 0.0	20	No	No	No										3.21
14:52:00	7	03	VSIVM3	IFL 1 Tot	0.5 0.0	100	No	No	No										0.47
14:52:00	7	04	VSIVM4	IFL 3 Tot	68.1 1.2	200	No	No	No	80.8	74.9	2.8	3.1	447.7	71.5	50.57			
14:52:00	7	04	VSIVM4	IFL 0	30.7 0.6	200	VMe	No	No	37.3	34.6	1.3	1.5	160.4	2.22	21.88	0	1	
14:52:00	7	04	VSIVM4	IFL 1	25.5 0.4	200	VMe	No	No	30.3	28.2	1.1	1.1	168.1	1.51	19.05	2	3	
14:52:00	7	04	VSIVM4	IFL 2	11.9 0.3	200	VLo	No	No	13.1	12.1	0.5	0.5	119.2	67.8	9.65	4	5	
14:52:00	7	05	VSIVM5	CP 2 Tot	16.7 0.1	75	No	No	No										16.61
14:52:00	7	05	VSIVM5	IFL 1 Tot	0.2 0.0	150	No	No	No										0.18
14:52:00	7	05	VSIVM5	ZIIP 1 Tot	0.3 0.0	75	No	No	No										0.29
14:52:00	7	07	VSIVC1	IFL 1 Tot	11.6 0.2	20	No	No	No										8.66
14:52:00	7	08	VSIVC2	IFL 1 Tot	1.3 0.0	100	No	No	No										1.10
14:52:00	7	09	VSIVC3	IFL 1 Tot	0.4 0.0	50	No	No	No										0.33
14:52:00	7	0A	VSIVC4	CP 1 Tot	0.5 0.0	50	No	No	No										0.51
14:52:00	7	0A	VSIVC4	IFL 1 Tot	0.2 0.0	75	No	No	No										0.18
14:52:00	7	0A	VSIVC4	ZIIP 1 Tot	0.0 0.0	75	No	No	No										0.02
14:52:00	7	0E	ZOSLP1	CP 2 Tot	6.3 0.0	50	No	No	No										6.28
14:52:00	7	0E	ZOSLP1	ZIIP 1 Tot	25.1 0.0	100	No	No	No										0.39
14:52:00	7	0F	ZOSLP2	CP 2 Tot	18.3 0.0	50	No	No	No										18.22
14:52:00	7	0F	ZOSLP2	ZIIP 1 Tot	0.2 0.0	50	No	No	No										0.14

Figure 2-5 ESALPAR Screen example

Linux Examples:

Let's jump now to some more specific information, such as data for Linux.

ESALNXP

Often if a Linux server is having a performance issue, it is some process that is either grabbing all the resources or a process that may have started at the same time over multiple servers (such as a CRON job). The ESALNXP report shows how each server's processes behaved. It is easy to see real time and as a summary if there is a mis-behaving process. Also, having the parent/child relationship in the Velocity Software MIB allows even better analysis of the servers and their running processes.

Report: ESALNXP		LINUX HOST Process Statistics Report										Velocity Software Corporate			ZMAP 5.1.5 12/23/23		Pg 3769					
Monitor initialized: 12/22/23 at 00:00:00 on 8562 serial 040F78										First record analyzed: 12/22/23 00:00:00												
node/ Name	<Process Ident> ID	Nice PPID	PRTY Valu	<-----CPU Percents-----> Valu Tot	sys	user	sys	usr	<-----Storage Metrics (MB)-----> Size	RSS	Peak	Swap	Data	Stk	EXEC	Lib	Lck	PTbl	<-Faults/Second-> min	maj	mint	majt
*****Summary*****																						
Average:																						
lxdb2001	0	0	0	0.39	0.05	0.06	0.07	0.22	4609	318	112K	0	35K	115	90.6	34K	0	65.7	22	0	897	0
snmpd	2213	1	-10	10.06	0.04	0.02	0	0	25	9	890	0	304	3.2	0.8	281	0	0.89	17	0	0	0
db2fmc	2258	1	0	20.00	0.00	0.00	0.07	0.22	51	14	1224	0	83.0	5.6	1.7	1K	0	2.02	4	0	888	0
db2sysc	2813	2811	0	20.03	0.00	0.03	0	0	941	91	24K	0	7831	3.2	2.5	5K	0	14.2	0	0	0	0
lxdb2002	0	0	0	0.65	0.05	0.06	0.13	0.42	5269	451	127K	0	35K	137	165	44K	0	88.1	27	0	40K	0
snmpd	2213	1	-10	10.06	0.03	0.02	0	0	25	9	889	0	299	3.2	0.8	281	0	0.98	19	0	0	0
db2fmc	2292	1	0	20.55	0.00	0.00	0.13	0.42	51	14	1224	0	83.0	5.6	1.7	1K	0	2.02	8	0	40K	0
db2sysc	2847	2845	0	20.03	0.00	0.03	0	0	877	91	21K	0	6295	3.2	2.5	5K	0	14.3	0	0	0	0
lxdb2003	0	0	0	0.07	0.04	0.03	0.00	0.00	1572	276	44K	0	26K	102	317	5K	5K	38.2	20	0	4	0
snmpd	1360	1	-10	10.06	0.04	0.02	0	0	53	45	1364	0	939	3.1	0.7	304	0	3.19	20	0	0	0
lxora12	0	0	0	1.00	0.30	0.56	0.04	0.10	28K	3732	680K	2960	46K	265	189K	17K	0	645	103	0	474	0
init	1	0	0	20.14	0.00	0.00	0.04	0.10	2	0	58.5	2.34	4.41	3.2	0.9	48	0	0.23	0	0	461	0
kswapd0	20	2	0	20.00	0.00	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
java	1715	1	0	20.10	0.01	0.08	0	0	856	290	22K	2033	19K	3.3	1.2	439	0	27.8	0	0	0	0
java	1926	1	0	20.08	0.01	0.07	0	0	1029	334	28K	784	23K	3.3	1.2	439	0	28.4	0	0	0	0
ora_vktm	3515	1	0	-2.17	0.11	0.06	0	0	785	3	19K	0	71.3	3.2	5731	407	0	10.6	0	0	0	0
ora_dia0	3529	1	0	20.03	0.00	0.03	0	0	787	13	19K	0	71.3	3.2	5731	407	0	16.1	0	0	0	0
ora_ckpt	3535	1	0	20.02	0.01	0.01	0	0	785	30	19K	0	71.3	3.2	5731	407	0	18.3	0	0	0	0
ora_mmon	3543	1	0	20.04	0.00	0.03	0	0	802	350	19K	0	100	12.2	5731	407	0	32.3	0	0	0	0
ora_mml1	3545	1	0	20.02	0.00	0.02	0	0	787	196	19K	0.09	100	6.7	5731	407	0	28.9	1	0	0	0
ora_cjq0	3579	1	0	20.02	0.00	0.02	0	0	796	378	19K	0	71.3	12.1	5731	407	0	31.5	0	0	0	0
snmpd	4657	1	-10	10.26	0.12	0.14	0	0	30	9	872	32.3	420	8.2	0.8	281	0	1.22	77	0	0	0
ora_vkrm	48847	1	0	20.00	0.00	0.00	0	0	785	12	3140	0	11.9	0.5	955	68	0	2.07	0	0	0	0

Figure 2-6 ESALNXP Report example

ESAUCD2

A very important aspect of running Linux on z/VM is storage. Often, guest Linux servers are given too much memory, which can cause performance degradation and even server death. Several screens/reports are dedicated to Linux storage analysis. Again, if classes are defined, data is organized by class with the ability to drill down into the class to see the users/servers.

Time	Node/ Group	<Real Storage (MB)>			<--SWAP Storage (MB)-->			Total	<-----Storage in Use (MB)----->					Error
		Total	Avail	Used	Total	Avail	Used		MIN	Avail	CMM	Buffer	Cache	
12:55:00	VPNs	1051.6	751.1	300.6	0	0	0	62.5	751.1	0.0	16.7	42.8	241.1	0
12:55:00	VMWARE	995.6	226.4	769.2	4096	3861	234.8	15.6	4088	0.0	0.0	338.6	430.7	50.0
12:55:00	UBUNTU	234.6	13.0	221.6	371.9	370.2	1.6	15.6	383.2	0	38.7	63.7	119.2	3.0
12:55:00	TheUsrs	23699	2701	20998	3911	3103	807.8	156.3	5805	505.9	1558.2	9369	10071	866.3
12:55:00	SUSE	23047	2463	20584	11764	10601	1163	93.8	13064	249.0	539.7	17295	2749	585.8
12:55:00	REDHAT	4699.7	866.2	3833.5	7804	7460	344.3	78.1	8326	632.0	166.6	1784	1883	115.4
12:55:00	redhat9	970.1	182.5	787.6	2048	1984	64.0	15.6	2167	0.0	0.6	483.3	303.7	22.6
12:55:00	redhat85	779.8	153.1	626.7	1536	1418	118.4	15.6	1571	0.0	0.1	454.9	171.7	12.9
12:55:00	redhat75	988.4	242.4	746.0	2048	1950	97.9	15.6	2192	424.0	5.3	65.0	675.7	1.3
12:55:00	redhat74	991.4	105.7	885.7	124.0	124.0	0	15.6	229.6	208.0	160.1	297.1	428.5	56.0
12:55:00	redhat01	970.1	182.5	787.6	2048	1984	64.0	15.6	2167	0.0	0.6	483.3	303.7	22.6
12:55:00	RANCHER	11921	3943	7977.9	6144	6041	103.0	46.9	9984	0	962.2	3662	3354	236.5
12:55:00	ORACLE	996.8	15.3	981.4	123.9	60.0	64.0	15.6	75.3	0.0	266.4	574.4	140.6	0

Linux Application Examples:

ESAK8S2

Kubernetes is one of the latest technologies for managing containers. Velocity Software has it covered with several screens/reports that show how the Kubernetes management servers are performing. Below shows a partial end of day summary for several managed containers.

```
Report: ESAK8S2      Kubernetes Resource Utilization Report      Velocity Software Corporate      ZMAP 5.1.5 12/27/23
Monitor initialized: 12/26/23 at 00:00:00 on 8562 serial 040F78      First record analyzed: 12/26/23 00:00:00
-----
```

NODE/ Time/ Date	ContainerName	ProcID	ProcName	Container CPU----->				Container Memory----->				Container----->					
				CPU Percents----->				Storage metrics(mb)>				Faults/Second->					
				Tot	sys	user	sys	usr	Size	RSS	Peak	Swap	ptbl	min	maj	mint	majt
	rancher	4157	tini	0.58	0.12	0.33	0.09	0.04	1146	205	17K	0	16.7	21	0	42	0
*****Summary*****																	
Average:																	
	ranchal																
	rancher-7c5dbf46fc-t																
	rancher	40762	tini	0.95	0.24	0.54	0.11	0.05	2049	315	49K	0	37.9	35	0	56	0
	s390x-hello-world-5d																
	s390x-hello	3958	sample	0.00	0.00	0.00	0	0	690	1	1379	0	0.19	0	0	0	0
	gitjob-5bd78d7cd9-xs																
	gitjob	31520	gitjob	0.08	0.02	0.06	0	0	730	32	18K	0	4.64	0	0	0	0
	fleet-agent-86bcc746																
	fleet-agent	31514	fleetage	0.18	0.04	0.14	0	0	752	45	18K	0	7.26	57	0	0	0
	rke2-canal-5n722																
	calico-node	2890	runsvdir	4.54	0.25	0.22	0.75	3.33	3807	210	91K	239	30.5	4	0	5K	0
	kube-flannel	2924	flanneld	0.33	0.02	0.03	0.20	0.08	735	18	18K	16.6	4.59	0	0	510	0
	kube-proxy-rancher-a																
	kube-proxy	1807	kube-pro	0.02	0.00	0.01	0.00	0.00	739	19	18K	15.9	5.16	0	0	4	0
	rke2-ingress-nginx-c																
	rke2-ingress-nginx-c	4328	catatoni	0.10	0.02	0.08	0	0	1838	37	45K	3038	22.5	0	0	0	0
	cert-manager-ddd4d6d																
	cert-manager	31796	controll	0.23	0.05	0.18	0	0	743	27	18K	0	5.30	0	0	0	0
	s390x-hello-world-5d																

Figure 2-7 ESAK8S2 Report example

ESAMNG2

Many businesses are using MongoDB to store their information. The Velocity Software screens/reports show many different metrics for MongoDB such as assertions, cursors, connections, errors, commits, journal entries and performance timings. Byte counts are shown below:

Report: ESAMNG2		MONGODB Processing Report				Velocity Software Corporate				ZMAP 5.2.0 08/09/25				Pg 814	
Monitor initialized: 08/08/25 at 00:00:00 on 8562 serial 040F78						First record analyzed: 08/08/25 00:00:00									
Node Name	Host Name	Threads	<Network: Bytes /s><Disk: Bytes /s>		<-Cache-----<---Pages----->>		<---Connections-->								
			In	Out	Read	Write	Used	Max	Read	Ckpt	Written	Dirty	Open	Avail	Total
08/08/25 00:15:00															
mongo8pr305590	4234	95	165888	10912	84K	299K	4069	7542	5.7G	502M	14	51186	213678	162034	1763
mongo8s1301295	3854	170	367242	38512	4.9M	24M	4236	7651	5.8G	502M	94	51106	787390	165840	2204
mongo8s2306372	3522	94	154368	9520	12K	22K	3874	7548	3.1G	502M	16	51184	90016	153233	1423
mongo804617255	14491	68	460115	37120	4.6M	19M	13K	12K	4.0G	502M	26	51174	609886	606283	6

Figure 2-8 ESAMNG2 Screen example

ESADOCK2

Docker is one of the more popular service platforms for running containers. Available metrics such as CPU, storage and paging are used to show container performance.

Time	Node	Container Index	Container Image	Container Name	<CPU Percent>		<-----Storage in "MB"----->						<PageRate>		Major Page Fault		
					User	System	Current Use	Max	Cache	RSS	<Anonymous>	<---File-->	</Second>	In		Out	
16:05:00	sles12	1b0ca9ad002a	httpd2	bouncer70	0	0	3.32	16.45	0.58	0.39	0.36	0.35	0.12	0.12	0	0	0
16:05:00	sles12	4d0978ac4b9f	httpd2	bouncer71	0	0	3.53	12.97	0.58	0.57	0.45	0.45	0.12	0.12	0	0	0
16:05:00	sles12	4891e692cf98	httpd2	bouncer72	0	0	3.53	13.03	0.58	0.58	0.45	0.45	0.12	0.12	0	0	0
16:05:00	sles12	9def13e9cd32	httpd2	bouncer73	0	0	3.49	13.04	0.60	0.50	0.51	0.31	0.12	0.16	0	0	0
16:05:00	sles12	2fa9a95b2a67	httpd2	bouncer74	0	0	3.54	13.01	0.58	0.58	0.45	0.45	0.12	0.14	0	0	0
16:05:00	sles12	4c32dd30b405	httpd2	bouncer75	0	0	3.53	13.03	0.58	0.56	0.53	0.35	0.11	0.15	0	0	0
16:05:00	sles12	c25d53117894	httpd2	bouncer76	0	0	3.50	13.07	0.58	0.55	0.44	0.44	0.12	0.12	0	0	0
16:05:00	sles12	0f635cf931b6	httpd2	bouncer77	0	0	3.46	12.99	0.57	0.50	0.41	0.41	0.12	0.12	0	0	0
16:05:00	sles12	2f784d370066	httpd2	bouncer78	0.0	0	3.43	14.94	0.58	0.47	0.40	0.40	0.12	0.12	0	0	0
16:05:00	sles12	d5a99d727da4	httpd2	bouncer79	0	0	3.48	12.98	0.59	0.51	0.51	0.33	0.12	0.14	0	0	0
16:05:00	sles12	956fca5916fa	httpd2	bouncer0	0	0	3.46	12.98	0.58	0.49	0.49	0.33	0.12	0.12	0	0	0
16:05:00	sles12	fd2deaab9e6f	httpd2	bouncer1	0	0	3.43	15.94	0.57	0.49	0.41	0.41	0.12	0.12	0	0	0
16:05:00	sles12	ef6b9ecd47c1	httpd2	bouncer2	0	0	3.53	12.98	0.55	0.60	0.45	0.45	0.12	0.12	0	0	0
16:05:00	sles12	fd20320d457	httpd2	bouncer3	0	0	3.50	13.00	0.57	0.55	0.44	0.44	0.12	0.12	0	0	0
16:05:00	sles12	b6976b2b1ab9	httpd2	bouncer4	0	0	3.53	12.94	0.57	0.58	0.45	0.45	0.12	0.12	0	0	0
16:05:00	sles12	4998fff62d47	httpd2	bouncer5	0	0.0	3.52	12.99	0.47	0.67	0.45	0.45	0.12	0.12	0	0	0
16:05:00	sles12	a494cdd53ed	httpd2	bouncer6	0	0	3.43	12.98	0.58	0.47	0.43	0.37	0.12	0.12	0	0	0
16:05:00	sles12	fab51fc16cc3	httpd2	bouncer7	0	0	3.49	12.96	0.58	0.50	0.45	0.38	0.12	0.12	0	0	0
16:05:00	sles12	83153772a262	httpd2	bouncer8	0	0	3.42	12.98	0.57	0.47	0.50	0.30	0.12	0.12	0	0	0
16:05:00	sles12	a679be3a94b7	httpd2	bouncer9	0	0	3.53	12.94	0.58	0.57	0.45	0.45	0.12	0.12	0	0	0

Figure 2-9 ESADOCK2 Screen example

z/OS Examples:

ZOSCPU

Run z/OS systems? Velocity Software has zOSMON. It collects SMF records that are then processed on z/VM, which is more efficient and cost effective. All of the Velocity Software products work with this information so it is integrated into screens/reports and alerts like any other user/server data. It has the added benefit of being able to display one minute granularity. Running jobs and job steps are also reported along with any CICS data when running. The help contains the SMF records id for each field.

Time	<---CPU-->			Samp Cnt	<---CPU Util-->			<-----Rates Per Second----->					Late Retrn
	SYSID	ID	Type		Total	Wait	Park	SRB	TCB	Wait	SIGP	I/O	
14:26:00	V24A	0	GP	1	14.8	85.2	0	804.0	811.3	316	141	11.7	8.4
14:26:00	V24A	2	GP	1	7.3	92.7	0	323.1	369.6	134	93.8	4.4	4.9
14:26:00	V24A	Tot	GP	2	22.1	178	0	1127	1181	450	235	16.1	13.2
14:26:00	V24A	4	ZIIP	1	0.1	99.9	0	0.0	172.8	55.6	9.9	4.4	0
14:26:00	V24A	5	ZIIP	1	0.0	100	0	0.2	98.8	10.0	1.3	0.0	0
14:26:00	V24A	Tot	ZIIP	2	0.2	200	0	0.2	271.6	65.6	11.2	4.5	0
14:26:00	V24C	0	GP	1	6.5	93.5	0	126.1	213.4	210	19.7	3.7	0
14:26:00	V24C	Tot	GP	1	6.5	93.5	0	126.1	213.4	210	19.7	3.7	0
14:26:00	V24C	1	ZIIP	1	0.3	99.7	0	0.0	206.9	19.7	4.7	0.1	0
14:26:00	V24C	Tot	ZIIP	1	0.3	99.7	0	0.0	206.9	19.7	4.7	0.1	0

Figure 2-10 ZOSCPU Screen example

VSE Examples:

ESAVSEJ

Velocity Software has a new product zVSEMON uses SNMP to collect data from VSE. Partition data, job data, network data and CICS data.

Time	Node	Part Id	Job Name	Phase Name	<- CPU% -> CPU Ovhd	<----- I/O -----> Disk VDisk Other	Rtrn Code	Cncl Code	<----- Start -----> Date Time	Stop Time	User Info	Prtly	Share	Job Num
14:54:00	vsen63c	F1	POWSTART	IPWPOWER	0.0 0.0	22.0 0 0			09/07/23 13:23:12			1	0	0
14:54:00	vsen63c	FB	SECSERV	BSTPSTS	0.0 0.0	0 0 0			09/07/23 13:23:12			2	0	0
14:54:00	vsen63c	F3	VTAMSTR	ISTINCVT	0.0 0.0	0 0 0			09/07/23 13:23:16	SYSA		3	0	63548
14:54:00	vsen63c	T1	TCPIP00	IPNET	0.2 0.0	0 0 0			10/30/23 08:02:04			4	0	44060
14:54:00	vsen63c	S1	SVSESRVR	SVSESRVR	0.0 0.0	0 0 0			10/30/23 08:02:22			5	0	44062
14:54:00	vsen63c	Z1	DMFSTART	DFHDFSIP	0.1 0.0	36.0 0 0			11/02/23 14:34:50			6	0	55103
14:54:00	vsen63c	R1	STARTVCS	IESVCSRV	0.0 0.0	0 0 0			10/30/23 08:02:31			7	0	44063
14:54:00	vsen63c	R2	STARTMAS	IESMASNM	0.1 0.0	0 0 0			10/30/23 08:02:35			7	0	44064
14:54:00	vsen63c	F2	CICSICCF	DFHSIP	0.0 0.0	0 0 0			09/07/23 13:23:20	SYSA		8	0	63549
14:54:00	vsen63c	O1	CICSINC01	DFHSIP	0.0 0.0	0 0 0			12/02/23 17:31:31			9	0	25313
14:54:00	zvse62b	F1	POWSTART	IPWPOWER	0.0 0.0	0 0 0			10/31/23 12:48:09			1	0	0
14:54:00	zvse62b	FB	SECSERV	BSTPSTS	0 0	0 0 0			10/31/23 12:48:07			2	0	0
14:54:00	zvse62b	F3	VTAMSTR	ISTINCVT	0.0 0.0	0 0 0			10/31/23 12:48:20	SYSA		3	0	24897
14:54:00	zvse62b	T1	BSTTINET	BSTTINET	0.2 0.1	0 0 0			10/31/23 12:48:25			4	0	24899
14:54:00	zvse62b	T2	BSTTVNET	BSTTVNET	0.0 0.0	0 0 0			10/31/23 12:48:57			4	0	24900
14:54:00	zvse62b	T3	BSTTFTPD	BSTTFTPS	0 0	0 0 0			10/31/23 12:48:56			4	0	24901
14:54:00	zvse62b	Z1	DMFSTART	DFHDFSIP	0.1 0.0	0 0 0			10/31/23 12:48:51			5	0	24902
14:54:00	zvse62b	R1	STARTVCS	IESVCSRV	0.0 0.0	0 0 0			10/31/23 12:49:03			6	0	24905
14:54:00	zvse62b	R2	STARTMAS	IESMASNM	0.1 0.1	0 0 0			10/31/23 12:49:02			6	0	24906
14:54:00	zvse62b	F2	CICSICCF	DFHSIP	0.0 0.0	0 0 0			10/31/23 12:48:48	SYSA		7	0	24903
14:54:00	zvse62b	O1	CICSJA95	DFHSIP	0.0 0.0	0 0 0			10/31/23 12:49:08			8	0	24904
14:54:00	zvse62c	F1	POWSTART	IPWPOWER	0.0 0.0	22.0 0 0			10/31/23 12:52:05			1	0	0
14:54:00	zvse62c	FB	SECSERV	BSTPSTS	0.0 0.0	0 0 0			10/31/23 12:52:04			2	0	0
14:54:00	zvse62c	F3	VTAMSTR	ISTINCVT	0.0 0.0	0 0 0			10/31/23 12:52:09	SYSA		3	0	18961
14:54:00	zvse62c	T1	TCPIP00	IPNET	0.3 0.1	0 0 0			10/31/23 12:52:12			4	0	18963
14:54:00	zvse62c	T2	SVSESRVR	SVSESRVR	0.0 0.0	0 0 0			10/31/23 12:52:27			4	0	18966
14:54:00	zvse62c	Z1	DMFSTART	DFHDFSIP	0.1 0.0	39.0 0 13.0			10/31/23 12:52:51			5	0	18965
14:54:00	zvse62c	R1	STARTVCS	IESVCSRV	0.0 0.0	0 0 0			10/31/23 12:52:48			6	0	18969
14:54:00	zvse62c	R2	STARTMAS	IESMASNM	0.1 0.0	0 0 0			10/31/23 12:52:48			6	0	18970
14:54:00	zvse62c	F2	CICSICCF	DFHSIP	0.0 0.0	0 0 0			10/31/23 12:52:33	SYSA		7	0	18967
14:54:00	zvse62c	O1	CICSJA69	DFHSIP	0.0 0.0	0 0 0			10/31/23 12:52:48			8	0	18968

Figure 2-11 ESAVSEJ Screen example

Graph Examples:

CPU Utilization

This graph will show the utilization for each of the shared IFL engines. Hovering over each column will show more specific information. Colors, high water marks, etc have defaults but are also highly customizable.




Figure 2-12 CPU Utilization graph example

Linux Swap Utilization

This graph will show the Linux swap utilization. It is easy to keep or remove the legend, change the systems shown, change the time frames or download the graph with the icons shown.




Figure 2-13 Linux Swap Utilization by Class example

Enterprise View Example:

Enterprise Main View

This shows the default Enterprise View. This can be customized by adding/deleting screens to the default view. In addition, different LPARs/systems can be added to give an even broader view of the enterprise. This is very beneficial for Operations staff.

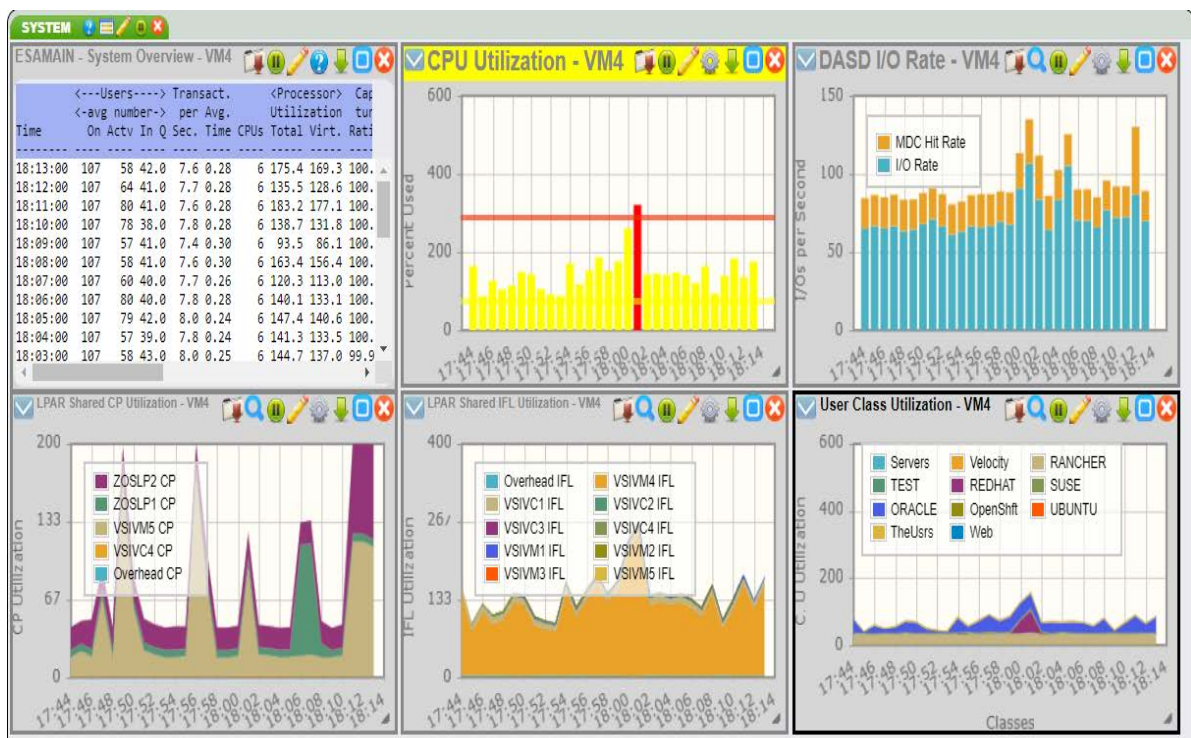


Figure 2-14 Enterprise Main View example 1

Here is another example of the Enterprise view.

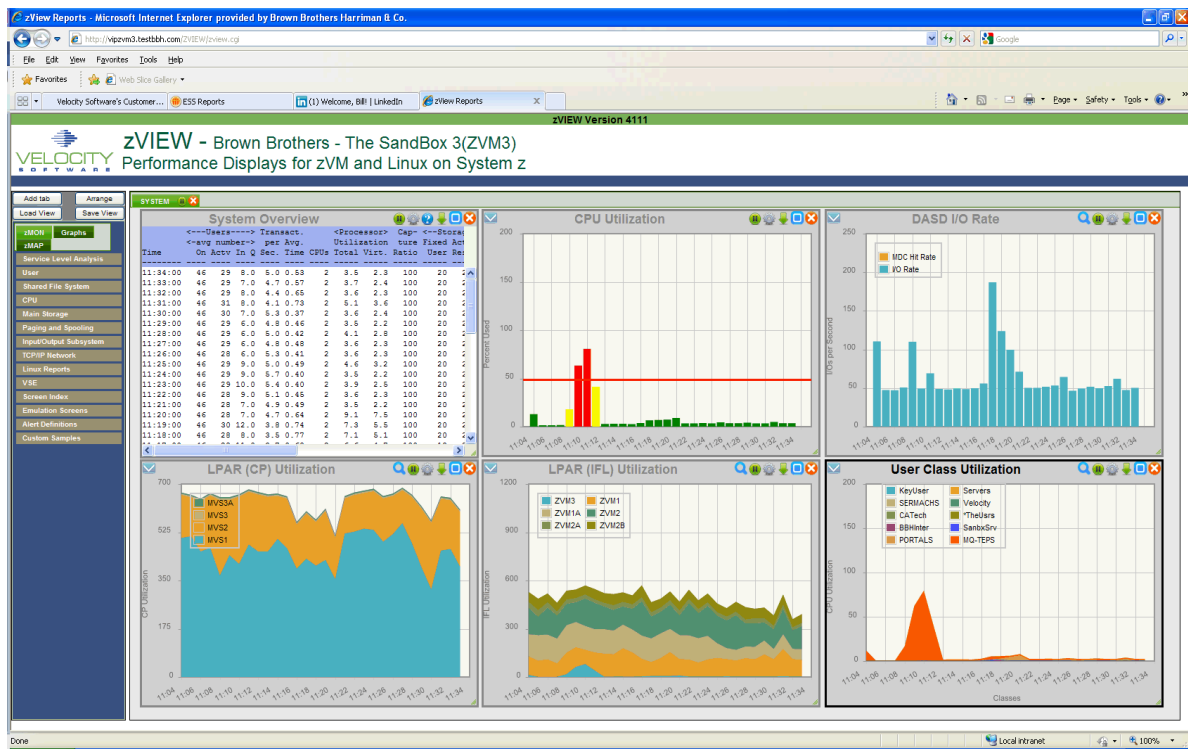


Figure 2-15 Enterprise Main View example 2

Summary

There are over 160 screens/reports available to view performance data. These are well organized into top-down/general-to-specific sections for intuitive viewing. For more information or a demo - go to demo.velocitysoftware.com.

6 - Performance Tool Checklist

When comparing products, there are several considerations that may not be obvious for those without experience on this platform. The following lists some of the considerations:

Ease of Use

- Are the screens and report reading intuitive?
- Are the products immediately usable, PF key driven, browser capable?
- Are education needs provided for, eliminated (or minimal)?
- Are reports structured in multiple levels, providing a management overview level which increases in detail to ultimately show the subsystem level?
- Does the product incorporate time saver screen-to-screen transition?
- Does it provide consistent data between both the real-time and historical reporting modules?

Ease of Maintainability

- Is it based on IBM's warranted CP Monitor?
- Are upgrades required for the product to continue to run when going from one release or service level to another?
- Is there a charge for a new release?
- Is the real-time monitor easily tailorable for specific user or installation needs?

Resource Reduction

- Can the product eliminate raw data collection, minimizing DASD requirements?
- Does the product have embedded capabilities to produce multiple kinds of output, minimizing CPU resources by processing data only one time?
- What kind of main storage requirement is needed?

- What is the personnel time required for conversion and release upgrades?
- Can the current release of the product run in an upgraded system, allowing for the product upgrade to be done later?
- What kind of education will be required of systems personnel and management?

Security

- Security is a concern in most environments. Is there an ability to restrict users to only the data for which they are authorized?

Access to Developers and VM Performance Experts

- Is there access to someone knowledgeable in VM performance?
- Is there extra costs for phone support?
- Is there support for performance analysis/questions?
- What kind of lead time is typical for a new feature to be introduced into the product?
- What is the typical lead time to problem resolution?

Reporting

- What is the completeness of reports that are available within the product - are YOUR important metrics represented?
- Does this product report by both user and by user class?
- Are Service Level Objectives measurable using percentiles (the more accurate method of measuring SLAs) presented?
- Can this product read both raw data and history files?
- Can it write to history archive files for future analysis, eliminating DASD and/or tape requirements?

Interface Capabilities

- What kind of interface capabilities does it have? Can it connect to what YOU need - such as MICS or splunk?
- Is a graphical presentation available to enhance investigation of existing problems?

- Is a rexx interface available to make changes to monitor panels and reports?

Price

- Consider the price, do features and functionality warrant the designated price?

Summary

This document has show the features and benefits of the Velocity Software products. This is why it is said Velocity Software is THE number one choice in performance software!

To read more about Velocity Software - <https://www.velocitysoftware.com/>

Velocity Software

General Information Guide

Velocity Software Contact Information:

Velocity Software, Inc.
zVPS Product Manager
PO Box 390640
Mountain View, CA 94039-0640

TEL: (650) 964-8867
FAX: (650) 964-9012

SUPPORT@VelocitySoftware.com

©Copyright 2025 Velocity Software, Inc. All Rights Reserved.

ESAMAP, ESAMON, ESAWRITE, ESAWEB, ESATCP, ESALPS, zALERT, zMON, zMAP, zOPERATOR, zOSMON, zPORTAL, zPRO, zTCP, zTUNE, zVIEW, zVPS, zVRM, zVWS and zWRITE are trademarks of Velocity Software, Inc.

All other products and/or company names are trademarks or registered trademarks of their respective owners.

