

PERFORMANCE NEWS

Issue 3 – March 2014

z/VM, Linux and Cloud on System z Updates



What's new at Velocity Software?

BY BARTON ROBINSON

Meeting the challenges of managing performance in an environment where everything changes on a regular basis will always keep our Velocity Software developers happy and busy. In the last year, z/VM 6.3 was announced and became available – with many changes to storage management, CPU management, and scheduling. We at Velocity Software had no choice but to upgrade our “old” z9 to a nice shiny BC12 – just so that we can do performance experiments in a controlled environment for z/VM 6.3 as well as SSI. The z9 had the processing power to deliver a large cloud environment and support development for many more years – but z/VM 6.x won't run there. So we now have a nice new toy to kick around with lots of strange configurations to play with and test performance. See the article on Expanded Storage to answer the first question I'm always asked about 6.3.

Newer versions of Red Hat and SUSE are always coming out with new metrics and capability. As the Linux operating system matures, it provides new metrics internally. We do our best to expose new information and evaluate how that can be used to assist your performance management needs. One of the new sets of metrics are for storage management at the process level. With large Oracle address spaces,

benefits in the past have been reported switching the Linux server to “large page support”, greatly reducing the size and quantity of page tables to support the workload. One of the new metrics shows the storage in use by page tables by process – now we have a good methodology for measuring the impact and projecting the impact of large page support.



Speaking of Oracle, Oracle 12C has just been made available – and Velocity Software is now producing Oracle metrics. These metrics are now fully integrated into our database and other system metrics. Our challenge is to take the Oracle metrics and the other new Linux metrics that are now collected every minute and understand what kinds of trends are interesting to expose, and what alerts should be included for operations to detect database issues. This is a project that after many years is finally off the ground – and of great interest to many of our customers. See the IBM Redpiece just published for Oracle 12C.

zVPS 4.2 will be available for Beta in April.

BARTON ROBINSON IS THE CTO OF VELOCITY SOFTWARE.

Join our Social Media Channels!

 facebook.com/pages/Velocity-Software/356098274460840

 twitter.com/VelocitySoftw

 xing.com/companies/velocitysoftwaregmbh

 linkedin.com/company/velocity-software-inc

Oracle Performance Support

BY BARTON ROBINSON

Velocity Software is pushing our technology in new directions. With a solid z/VM and Linux performance management suite, the next step is to provide a view into the application subsystems. The three application subsystems targeted in the next release of zVPS, 4.2, are Oracle, MQ and Java.

The Oracle support requires having an Oracle userid, an update to the Velocity Software SNMP MIB running on the Linux server, and zVPS 4.2. ZTCP 4.2 will detect the Oracle MIB if available and collect it, and zWRITE will process the data, including it in the performance database (history data), and zMAP will produce the reports. The initial data is focused on a subsystem level perspective rather than an SQL statement analysis or transaction analysis. Tying the Oracle metrics to the Linux system metrics at a system level will allow for reasonable settings for SGA sizes, and Linux server (virtual machine) sizing in ways that few people can accomplish today.

With this support, Oracle data will be captured every minute for detailed performance analysis, as well as supporting operational alerts. With the data in the database, long term analysis can be accomplished, assisting capacity planning and Oracle DBAs in trend analysis.

The Oracle 12C Redbook <http://www.redbooks.ibm.com/redpieces/abstracts/sg248159.html> recently published as a Redpaper included analysis using the new Velocity Software Oracle instrumentation. The initial release of this instrumentation allows more in depth storage analysis and I/O analysis.

The Oracle databases supported by this instrumentation include 10G, 11G and 12C on both system X and system Z, as shown in the following running on the Velocity Software demonstration site. This shows the version of Oracle running, and on which server. Our server names are encoded with “x” for distributed server and “v” for “VMWare servers”.

Report: ESAORAC

Oracle Database Configuration Report

Monitor initialized: 02/12/14 at 00:00:00 on

```
-----  
Node/          <-----Database Description----->  
Date  
Time          DatabaseName Instance      Version  
-----  
00:15:00  
oracle        orcl          orcl          10.2.0.4.0  
redhat6x      db01          db01          11.2.0.2.0  
sles11v2      db01          db01          12.1.0.1.0  
s11s2ora      db01          db01          11.2.0.2.0  
              db02          db02          11.2.0.2.0
```

The beta program for this support starts in April of 2014, as part of the zVPS 4.2 beta program.

Expanded Storage, is it really gone in z/VM 6.3?

Expanded storage was originally a high speed paging device to supplement the need for real storage, limited in size by older architectures. If it wasn't for Linux, the need for expanded storage might have gone away years ago.

Linux itself did not exacerbate storage requirements – though original releases of Linux would poll every 10ms (100 times per second), and this caused early storage problems for those running Linux on “z”. This problem was corrected by the “timer patch”, so that in virtual environments, an idle Linux server would accurately be considered idle.

It's normal and expected to overcommit storage in a z/VM environment. Over-commitment of storage causes paging – not a problem, as z/VM has always had extremely efficient paging and storage management. Paging was a manageable problem – up to a point,

with systems having been measured at paging 1,000s of pages per second to disk – and providing great application response times.

What broke the z/VM storage management was Linux application subsystems such as WAS or DB2 that started polling – even after we learned this was a bad idea in virtual environments. DB2 was caught once polling at 200,000 times per second – blamed at the time on “high availability”. This polling is probably not noticeable on dedicated “x” environments, but in storage overcommitted virtual environments, this greatly changes the paging subsystem requirements.

The definition of an idle server is one that did not use CPU for 300ms. This 300ms delay originates many years ago when teletypes were still in use (true story), as the virtual machines would incur extra overhead if they dropped from queue between characters being typed. The definition of an idle server is then used to determine which server to “steal” storage from to make room for other active workloads.

In a traditional virtual environment, in an overcommitted environment, unused pages should be paged out to make room for active pages. Idle servers should be paged out first, as they should have no workload needing storage to operate. The growing pains z/VM installations have suffered is that when supporting maybe one hundred 2 or 4 GB servers, it became impossible to define an idle server, but it was obvious something had to be paged out. It became a difficult problem to determine which server and which pages should be paged out to make room for active workloads. The risk was always paging the wrong pages out to disk meant very slow response times for applications requiring those pages.

To overcome this problem, most large installations would take 20% of their LPAR storage as a short term buffer to ensure that if the wrong page was stolen, it didn’t go to disk for a period of at least 30 seconds. If the workload was 100 servers at 2GB each, then there would be about 200GB of storage perceived to be in use. Of that 200GB, a large percentage really was idle. Many large installations tested the theory of using the sometimes recommended 4GB or 8GB maximum of Expanded

Storage vs the Velocity Software recommendation of 20%, resulting in those large installations using the 20%, even if that meant Expanded Storage was sized at sometimes 60 or 80GB of storage on large LPARs.

Now with large changes in z/VM 6.3, what is the requirement for Expanded Storage? With Expanded Storage, 20% of the pages would be put in a “test mode”, not immediately accessible to the application, but almost immediately made available when a page fault occurs. Page faults for pages in expanded storage were resolved very quickly because of a synchronous transfer between expanded storage and real storage.

In z/VM 6.3, pages are made “Invalid but Resident” (IBR) – almost the same concept as Expanded Storage. After a time, IBR pages are put on the AGE LIST. In the following report from zVPS 4.2, from a terabyte LPAR, the virtual machine sizes total 345 million pages, of which 337 million have been touched, and 263 million pages are resident in real storage. Note that the AgeList is 5.3 million pages – this would be similar to having 20 GB of expanded storage, not 20% but a considerable buffer against stealing the wrong pages. These 5.3 million pages could be page faulted by the application, but immediately made available by updating pointers in the page tables – much faster even than moving pages in from expanded storage. There is more information available, and some tuning knobs to be evaluated – but the conclusion is that Expanded Storage really will be something of the past...

```
Report: ESAUSTR      User Storage Analysis
-----
          <-----Virtual Server Storage (Pages)----->
UserID   Size  Alloc Resi-  UFO   <--IBR--> <AgeList>
/Class           dent Activ TOT   >2gb <2gb >2gb
-----  -----  -----  -----  -----  -----  -----
15:09:10 345M  337M  263M  258M  23K  20K    0 5.3M
***User Class Analysis***
Servers  57856 12501   320  40.0  120 72.0    0 160
Linux    4194K 147K 17793  3277  128 128    0 14K
TheUsers 340M  337M  263M  258M  22K  20K    0 5.3M
```

BARTON ROBINSON IS THE CTO OF VELOCITY SOFTWARE.

Mainframe Hard!

BY MAGGIE NECOY

It's always great to see some of you wearing different Velocity Software shirts at conferences. It always amazes me that some of the shirts we see out there are from way back, when Velocity Software was first founded.

Velocity Software has a long tradition in giving away nice shirts at conferences, but I think the latest one will be one of the most popular ones.

Our „Mainframer“-shirt is a limited edition T-shirt, that was designed by internationally noted graphic designer Götz Gramlich and hand-printed by Fabian Lauer in Heidelberg, Germany. Two young entrepreneurs, to whom the word „mainframe“ wasn't more than a big box when we first started working on this design.



We think they did an excellent job adopting the look and feel of old school Hard Rock bands and integrating the Mainframe on the shirt. We are happy to be giving away these iconic collector's items to customers and prospects at our conferences. Check out the 2014 events on page 8 to find out where you can find us and pick up your T-shirt along with „Mainframer“ stickers and bags.



MAGGIE NECOY IS MANAGING DIRECTOR OF GLOBAL SALES AND MARKETING AT VELOCITY SOFTWARE.

VM WORKSHOP 2014

LEARN z/VM AND LINUX ON SYSTEM z FROM THE EXPERTS



The VM Workshop is a grassroots non-profit organization comprised of customers, ISVs, and business partners with an interest in the use and growth of the z/VM platform and Linux on System z. The original VM Workshops were held during the 70's, 80's and 90's at universities around the U.S. and Canada.

This newly formed re-birth of the popular VM Workshops of the past is being re-launched to bring back the low-price, high-quality focus that was so successful in the early days of VM.

All funds paid to the VM Workshop are used on the events and attendees. The goal is to provide a forum for z/VM education with topics that are relevant to customers supporting z/VM and Linux on System z in today's IT environment.

Threads will include VM and VM maintenance, Networks and VM, Tools for VM, labs covering introductions to VM and Linux on VM, and more.

Inexpensive dorm rooms and nearby hotel accommodations are available to make attending the VM Workshop one of the most cost-effective conferences available.

Whether you are new to z/VM or Linux on System z or an experienced master, don't miss this great opportunity to take your skills to the next level for z/VM and Linux on System z!

WHEN:
June 26-28, 2014

WHERE:
North Carolina A&T
State University,
Greensboro, North Carolina

WORKSHOP FEE:
\$100.00 per person

Learn More &
Register Today at
VMWorkshop.org

Installer 4200

BY TIM KESSLER

Velocity's installer program, VSIMAIN, has a new version that is now available. Significant changes have been made to improve the speed and ease with which Velocity products can be installed.

The following are the major enhancements to the VSIMAIN installer program:

- Automatic installation (new installations) – For new installations, the installer will now ask if you want an Automatic install. If you are using the recommended disk configuration of minidisks for ZMAP 191 and 192 and ZWRITE 191, the installer will now automatically create the SFS directories, run the SFS authorization for those directories, create the VM directory statements and create RACF authorization files.
- Automatic installation (existing installations) – PF11 Install All has been added to automatically install all products that have been downloaded and not installed. In addition to the product install, if the product requires a post installation exit to be run or uses a DCSS, these will also be automatically run or defined.
- SSI installation – SSI installations combined with Automatic installation can be completed with little intervention. On the installation of the first member of an SSI cluster, the VM directory entries and RACF definitions are created for all members in the cluster. Once the VM directory has been updated, the RACF definitions run on each member of the cluster, and the zVPS configuration completed, VSIMAIN can install all the base zVPS products on the second and following members of the cluster without any intervention. zVPS Configuration changes made on the first member of the cluster are automatically copied to the other members.
- VM directory updates – For DIRMAINT users, DMALL EXEC is created that can be run to create all the required VM userids. For VMDIRECT or VMSECURE users, VMSALL EXEC will do the same. If a directory manager product is not in use, all the required directory

statements will be created to add to your directory file. For all cases, if in an SSI environment, SUBCONFIGs are also created.

- FTP support on Check for Updates – In addition to the HTTP support on the Check for Updates function (PF9 on main VSIMAIN screen), FTP support has now been added. This allows to easily see which updates are needed for your system and provides downloads directly to your VM system. In a SSI environment, the download needs to be done only once per cluster. This facility also allows product key update, uploads to your customer area and the ability to delete members in your customer area (FTP only). FTP and HTTP proxy servers and FTP firewalls are also supported.
- Restart server support – On the VSIMAIN Installation and Configuration screen, there is a new option to restart the product's server machine(s). This can be used after installing a new release of a product to allow the servers to pick up the new version. Note this may require an update to the product to support this.
- Server version checking – VSIMAIN will now display the installed version in red if one or more of the product's server machines are not running the installed version of the product. To determine which server(s) need to be updated, place the cursor on the red installed version and press PF1. Note this may require a product update to support this.
- Product definition errors – If all the product definitions have not been completed, an error will appear in the Status column in red. Place the cursor on the error and press PF1 to obtain more information on the error.
- AUTOG, AUTOV and subpool support – On the Disk Configuration screen, DIRMAINT's AUTOG and AUTOV can now be specified along with the group or volume. VMSECURE and VMDIRECT subpools are also supported.

- FBA blocks – The Disk Configuration screen now supports FBA or SCSI blocks for the minidisk definitions.
- Easier product installs – In addition to the Automatic product install, individual products that have multiple releases to be installed (zero version plus update) can now be installed together. Select the latest release you wish to install and the zero release will automatically be installed or select the zero release and VSIMAIN will ask if you wish to install the latest release too.
- Console Aggregation Facility – An INSTALL configuration option has been added to allow running the Console Aggregation Facility when ZVPS logs on. Options are: YES, NO or ASK. The Console Aggregation Facility will receive all the zVPS product server's console logs and save them by date and product. Another option, Max console age, will automatically prune the files based on the number of days entered.

With these new features, you will be able to easily and quickly install and update your Velocity products. By using Automatic install, an entire SSI complex can be installed in under two hours.

Early user Mike Giglio of HealthPlan Services says, "The installer brings zVPS into the Single System Image (SSI) environment. We now have the ability to do overall standardization across the clusters yet we still have individual controls on each system. When you need to apply an update, you download it once per cluster and then apply the update to the individual z/VM systems."

Manny Valdez of Lockheed Martin reports, "The installer provides easier and quicker installs, especially supporting the SSI architecture."

Download the new version of the installer here: <http://www.velocitysoftware.com/customer/Ftpv4.cgi>

TIM KESSLER IS SENIOR SOFTWARE DEVELOPER
AT VELOCITY SOFTWARE.

An Enterprising View: zVIEW sees it all!

BY RICH SMRCINA

zView is Velocity Software's web-based application for viewing and analyzing performance data for z/VM and Linux systems.

zView's feature allowing the user to see multiple VM systems, called the CECLIST has been expanded and revamped with a feature we call Enterprise View.

Configuring Enterprise View is similar to the old CECLIST. There's a file on the config disk called ZVIEW CECLIST. Each non-comment entry in the file is a VM system that will appear in Enterprise View.

```
SYSTEM = VM1 HTTP://WWW.VELOCITYSOFTWARE.  
COM/ZVIEW/ 'VSIVM1' 'First level'
```

The SYSTEM keyword identifies each VM system that will be displayed, and also becomes a link to the URL. The URL is a pointer to the zView application on

that system. When the system link is clicked, the zView application of the associated system is opened in a new tab or window. The quoted value after the URL is the text that will be displayed in the window for that system. The second quoted value is a group identifier.

The Enterprise View shows one window for each VM system defined (see next page). Each window contains the heading information defined for that system. Within the window is the zView link for that system, the current local time of that LPAR, the type and number of processors defined to that system (CP/IFL). If it's a mixed mode LPAR, each processor type and count is shown separately. Also, the CPU Utilization by processor type is displayed as a value and a graph.

Beneath the system information for the LPAR is a section that shows the top defined Linux systems and the CPU Utilization of each, as a value and graph.

Today is Monday 17 Feb 2014 zVIEW Version 4159

VELOCITY SOFTWARE
zVIEW Enterprise View - Velocity Software - VSVM4 (MYNAME)

First level

VSVM1 Expand VM1 14/02/17 06:48 IFL Total (1) 0.29% Linux Nodes (Distributed Servers) LINUX9 (9) 1.01% suselx3 (9) 0.93%	VSVM2 Expand VM2 14/02/17 06:48 IFL Total (1) 0.58% Linux Nodes (z/VM-Guests) oracle 0.55% RH5X161 0.17% Linux Nodes (Distributed Servers) sles11z2 (2) 2.16%	VSVM3(old) Expand VM3 14/02/17 09:49 0214C7-0 3.57% 000000-64 3.57% Linux Nodes (Distributed Servers) OPENSUSE 0.53% sles11x 0.12%
Demo System V4 Expand Demo 14/02/17 06:48 IFL Total (2) 13.87% Linux Nodes (z/VM-Guests) redhat6x 1.01% redhat6 0.49% oracle 0.47% Ixsugar (2) 0.16% redha6x 0.14% redha6 0.14% rob1z2 (2) 0.13% redha66 0.13% sles11x 0.12% sles11x3 0.09%	VSVM5 Expand VM5 14/02/17 06:48 CP Total (1) 2.21% IFL Total (2) 0.01%	Demo System V3.5(old) Expand DemoV3 14/02/17 06:48 0414C7-0 7.36% 0414C7-1 6.51% Total 13.87% Linux Nodes (z/VM-Guests) redhat6x 1.01% redhat6 0.49% oracle 0.47% Ixsugar 0.16%
Second level Tims Test System Expand ZSERVE not responding	RKS Level 2 Expand RKS2LV 14/02/17 06:48 IFL Total (1) 0.08%	DXT Level 2 Expand DXT2LV 14/02/17 09:48 IFL Total (1) 0.12% Linux Nodes (z/VM-Guests) redhat6 0.46% redha66 0.16% redha6x 0.14% redha64 0.13% redha6 0.12% rob1z2 (2) 0.12% sles10x4 0.11% Linux Nodes (Distributed Servers) opensuse (2) 0.71% Ixsugar (2) 0.73%

If distributed Linux nodes are defined to a system, they are shown as a group under the local Linux nodes. Distributed Linux systems are those that are not running on the LPAR.

Each processor type for each LPAR is also represented as a link. Clicking on the link will take you to a zView graph of that processor type in the CEC in which it is running. The expand button on the top right of the window expands the view of that LPAR in a larger window.

LPARs can be grouped on the display. The last quoted item in the CECLIST definition for each system is a group name. Our Enterprise View contains two groups, First Level and Second level. The order of display is defined by the GROUPORDER statements in the CECLIST file.

```
GROUPORDER 'First level'
GROUPORDER 'Second level'
```

The Enterprise View extends into zView itself. If a CECLIST file is present, zView has the ability to show information from any defined system, not only the local system. Information from multiple systems can be shown simultaneously on one zView display.

A list box shows each VM system from which zView items can be selected. The local system is highlighted.

To show information from multiple systems, select the system from the list, then select the ZMON display or graph. The information from the selected system will appear.

RICH SMRCINA IS SENIOR SOFTWARE DEVELOPER AT VELOCITY SOFTWARE.

The screenshot shows a menu with the following options: Add tab, Arrange, Load View, Save View, Color config, Demo System V4 First level (selected), VSIVM1 First level, VSIVM2 First level, VSIVM3 First level, Demo System V4 First level (highlighted), VSIVM5 First level, Demo System V3.5 First level, Tims Test System Second level, RKS Level 2 Second level, DXT Level 2 Second level, Shared File System, and CPU.

Imprint

Responsible for all contents: Velocity Software
 Design: Götz Gramlich, Kolja van Boekel
 Contributing Editors: Tim Kessler, Magnelia Necoy,
 Barton Robinson, Rich Smrcina, Cheri Stephens
 Project Management: Magnelia Necoy

To subscribe for Velocity's Performance News please email: info@velocitysoftware.com



Velocity Inc. – Headquarters
P.O. Box 390640
Mountain View, CA 94039-0640
Phone: 650/964-8867 | Fax: 650/964-9012

PRSR FIRST CLASS
US POSTAGE
PAID
SAN JOSE, CA
PERMIT #10

2014 System z Events and Conferences

Looking to improve your System z knowledge and skills? Here are some of the System z events we will be attending in the upcoming months. Check our next newsletters and social media sites for updates.

Events North America:

10 – 14 March: SHARE

Anaheim Marriott, Anaheim, CA, USA

2 – 3 April: zSeries Oracle Special Interest Group Conference 2014

IBM Institute for Electronic Government (IEG) Briefing Center, Washington DC, USA

24 – 25 June: Velocity Software Performance Workshop

Greensboro, North Carolina, USA

26 – 28 June: VM Workshop

North Carolina Agricultural and Technical University in Greensboro, North Carolina, USA

3 – 8 August: SHARE

David L. Lawrence Convention Center Pittsburgh, PA, USA

6 – 10 October: Enterprise 2014

Venetian, Las Vegas, Nevada, USA

17 – 21 November: SUSEcon

Coronado Springs Resort, Lake Buena Vista, Florida, USA

Events EMEA:

7 – 9 April: GSE Frühjahrstagung 2014

für z/VSE, z/VM und Linux on System z
Le Méridien Parkhotel, Frankfurt, Germany

12 – 16 May: IBM System z Technical University

Corinthia Grand Hotel Royal, Budapest, Hungary

20 – 22 October: European IBM / GSE Technical University for z/VSE and z/VM and Linux on System z

Dresden, Germany

4 – 5 November: GSE UK

Wittlebury Hall Hotel, Northamptonshire, UK