

PERFORMANCE NEWS

Issue 4 – October 2014

z/VM, Linux and Cloud on System Z Updates



What's new at Velocity Software?

BY BARTON ROBINSON



The best part about my job at Velocity Software is still the performance research. Looking at data, and understanding the value of that data to a business is why I get up early and go to work. Current research in Oracle (last newsletter), processor speeds and z/VM 6.3 is more than enough to keep me busy. I hope the research on processor speeds is of interest to you, I'm looking forward to see how the data is utilized.

Gigahertz is Gigahertz! Just how fast is a mainframe? MIPS, Meaningless Indicator of Performance, Bogo (self-explanatory), or just use Gigahertz? Over the years, presentations from the Washington Systems Center have shown that one processor can have several significantly different "MIPS" ratings – all based on different workload on the same processor with the same cycle time. When Linux first arrived on the mainframe scene, the first question from those with capacity planning experience was for a Megahertz/MIP comparison. In early analysis of Linux workloads, comparing workload on the mainframe to that workload on an Intel, my number was 4, as in 4 megahertz resulted in approximately 1 MIP. Listening to an IBM presentation, the presenter referred to this number as "Barton's number", IBM legal would not allow IBMers to state any such number, but the presenter could quote "someone from the industry". A few years later, machines got larger with more IFLs, further benchmarks on large machines suggested that "Barton's number" was closer to 5.

How was this number utilized? If a workload on a distributed platform was using (as measured) 20%

of 10 2GHz processors, then it is using 4 GHz. How many IFLs are needed? Depends on how fast the IFL is. Today, we know the EC12 processor is 5.5 GHz, so one IFL replaces those 102 GHz processors for that workload. That's very simple capacity planning – but a very important part of the plan.

The value to knowing this ratio was more significant when IBM only provided MIPS ratings for new boxes. Now IBM provides the cycle time of the processors and only rarely in the Linux world do we hear the term "MIPS". For capacity planners wanting to know how to forecast a given workload's IFL requirement when consolidating workloads to a mainframe environment, this number is of great value – the total cost of a consolidation project depended on knowing the number of required IFLs. Early in the life of "Linux on System/z" many consolidation projects were stopped due to incorrect capacity plans. One plan suggested a 3 IFL requirement for a workload, but the reality was 11 IFLs when an honest Gigahertz to Gigahertz analysis was performed. When reality set in, the cost was too high and the project aborted after significant effort and costs by the customer and business partners. Another

more conservative plan at another installation asked for 25 IFLs, but only half were really required. Then they overpaid and had “extra capacity”. Valid capacity planning needs proper input.

Now for the reality: if you can't measure it, I'm just not interested. With z/VM 5.4 and current releases, the hardware reports the number of instructions processed and the number of cycles consumed by the LPAR. Given the ratio, we now have the true number. On the Velocity Software DEMO LPAR, running 2 IFLs, and at z/VM 5.4, the ratio running a Linux workload is about 2.7 cycles per instruction, so the 4.2 GHz processor could be measured at 1,500 MIPS, very impressive, especially remembering our first Velocity Software system running 5 MIPS. For our VSE workload on the GP processors (.29 GHz), the cycles per instruction was 2.5 at some times, and 3.5 other times. The workload made a difference!

Another very large LPAR on a very large machine running z/VM 6.3, the average cycles per instruction runs at about 3.4, with some processors at 2.8, and all running between 2.8 and 3.8. That makes the 5.5 GHz processors about 1,600 MIPS for this Linux and Oracle workload. For z10, the only data analyzed so far, and not a fair analysis had a ratio of 8 to 1. Which if accurate, shows that though the z10 was 4.4 GHz, the EC12 could be 3 times faster in terms of real work. For comparison, a 5 IFL 6.2 LPAR was also measured at 3.2 cycles per instruction that might suggest that vertical scheduling in z/VM 6.3 has little benefit – but your mileage will vary, and now there is a way for YOU to measure the difference.

This analysis is part of zVPS 4.2, available soon, and can be seen on our DEMO system today, look in the zMAP reports section for ESAMFC. It should be

possible to build a profile showing different processors and different releases of z/VM (with hopefully enough data to evaluate horizontal vs. vertical scheduling).

The original “Barton's number” of 4 was close, and even today with a more accurate understanding of the real speed of the processors, the one significant figure was a very good estimator, and provided results that were a lot better than some of the disastrous projects that failed.

So next question, how to make your processor faster? On current processors there is a data cache and an instruction cache. Each instruction executed must have the data and instruction in the Level 1 cache. If the data or instruction is not in the Level 1 cache, then the data must be written to the Level 1 cache prior to instruction execution. Data coming from Level 2 cache is fast, Level 3 cache not as fast, and if the data must come from Level 4, or another book or memory, takes even more cycles.

The data provided shows that for the large LPAR with a 3.4 ratio, lost cycles to Level 1 cache writes was 3.4 (1.3 for instructions, 2.1 for data) – so a good assembler loop should be able to run at 5,500 MIPS, or to convert that to business perspective, a machine running my assembler will really run faster than your java, and not because it's assembler, but because the Level 1 caching characteristics will be better.

As more installations are installing significantly more IFLs, there is an ROI on understanding performance data and more importantly what to do with that data. Saving 5% of an IFL isn't much, but saving 5% of 100 IFLs should make somebody happy.

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

zPRO V2 Cloud Enablement

BY TONY NOTO

zPRO V2 Overview

In recent years, there has been a major shift in the IT industry. Many of the applications that were deployed on, or migrated to, a server farm, are coming back to “mother mainframe”. As “Cloud” becomes more prevalent, more demands are being placed on the administration of z/VM. This is true not merely as a virtualization platform, but as the basis for “Cloud on the mainframe”.

Following the migration back to the mainframe, is the talent pool normally tasked with maintaining a Linux server. For those of us who have been speaking “mainframe” for a few decades (or more ... sigh), we can only shake our heads and look to the horizon. We look to that day when we will cross over that line, move to Florida, wear long brown socks with sandals and drive a golf cart around the assisted living data center. However, until that day arrives, it is our job to maintain our system and teach GenX how to step up to the plate in the “Big Blue League”.

So how do we perform our normal jobs and delegate “doable” tasks to our newly arrived protégés, with a minimum amount of pain?

Enter zPRO V2, “the” web enabled UI tool, dare I say it... extension of YOU. You, my friend, are a z/VM system programming wizard. You know what a CP directory is. You know where your SYSTEM CONFIG lives. You know how to generate CMS, GCS, blah blah blah ... in your sleep, assuming, of course, that you have a 3270 emulator immediately handy. But, what about GenX. All (s)he has is SSH, vi (ewww), grep, and a browser.

zPRO V2 allows you to “can” your skills so that others may perform, some, if not many, of the tasks with a much smaller knowledge base (for the time being). It enables you to view and manage your enterprise through a “single pane of glass”.

You, the genius that you are, can build a workflow template. This template, once put in place, can be presented to GenX, as a point-and-click front-end. Nothing has to be known about CP directory, ESM (External Security Manager), SFS, et.al. All of that knowledge can be imbedded into a workflow template, as we’ll see in just a bit.

zPRO V2 is designed, to:

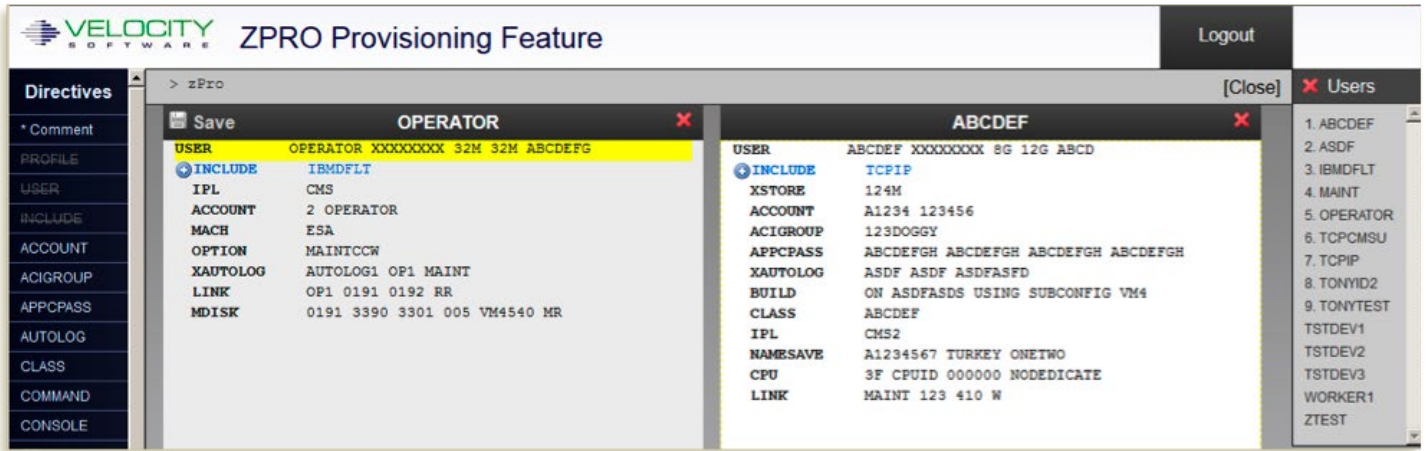
- Expedite the completion of the SYSPROG tasks
- Offload some of the provisioning responsibilities to other areas that may not have the technical skills
- Minimize repetitive tasks by incorporating zPRO scripting
- Guide inexperienced techies to a higher level of proficiency by presenting end-result of “canned” requests.

zPRO V2 as the SYSPROG

As the sysprog, and given a 3270 emulator, you can make your z/VM system sing and sit up and beg. These days, that is just tedious work. And, depending on your environment, it could be required frequently, or just every now and then. So, just how can zPRO V2 help you address your “SYSPROG” workload?

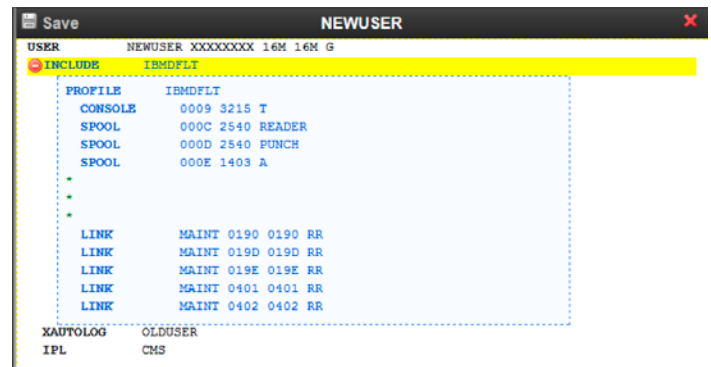
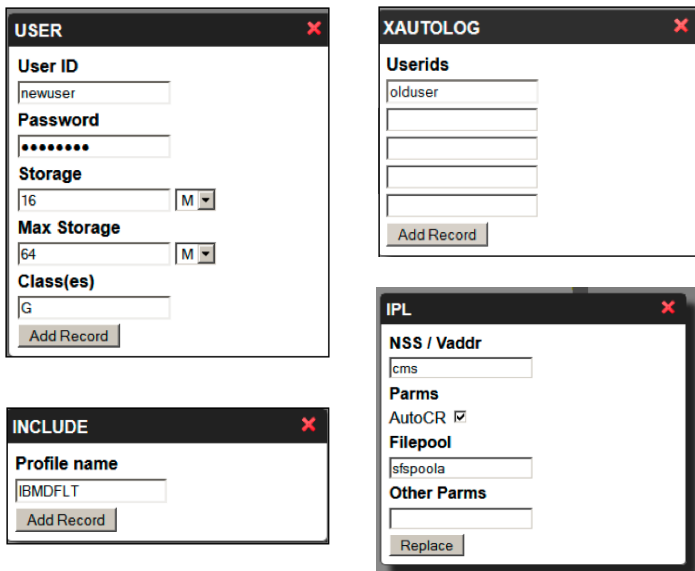
zPRO V2 speaks your language.

When it comes to CP user directory maintenance, zPRO V2 will hold your hand and/or accelerate your process. For instance, when doing user directory maintenance, you may need to compare one user to another. With zPRO split screen, it’s just a click and done. (See graphs on the next page)



Or perhaps you are building a new user, let zPRO V2 guide you with just the dialogs that you need to expedite the process. As seen below, zPRO V2 not only inserts the records for you, but in the correct order. You can always review your work, even as you perform your tasks, for instance if you need to verify the contents of an “expandable” section such as an included profile.

Or, the expanded view would be:



Maybe, just maybe this is a bit too much work for such a mundane task. Enter the zPRO V2 modeling facility. By creating a directory model, your installation can assign the task of creating a user id without having to know or understand the underlying processes. For instance, using the zPRO V2 scripting language, to create the user described in the above example, a model similar to the following can be created:

```

*!+ DETERMINE IF USER ALREADY EXISTS
*!FID &USERID DIRECT
*!CP QUERY &USERID
*!FAILRC 0 45 | &USERID already exists
*!SFS ENROLL USER &USERID SFSP00LA: (STO 2 BLO 180
USER &USERID NOLOG 16M 16M G
INCLUDE IBMDFLT
IPL CMS PARM AUTOCR FILEPOOL SFSP00LA:
XAUTOLOG OLDUSER
*? Build class G user
*? USERID |Y| UID | User-id | 8
  
```

The resulting directory would look like this.



Once in place, the end user simply opens a browser window and selects which model is to be created. In this case, the end user would be presented with the following dialog box:



Once the end user clicks the “Create” button, zPRO V2 does the rest.

Of course this is a simplistic example and does not fully represent the capabilities of the modeling feature of zPRO V2, some of which are shown in the following example.



```
#!# DETERMINE IF USER ALREADY EXISTS
#!CP QUERY AUSERID
#!FAILRC 0 45 | AUSERID already exists
#!FID AUSERID DIRECT
#!CHS EXEC GETDEV DEV1 DEV2 DEV3 DEV4
USER AUSERID WORKIT AORAGE AORAGE 0
INCLUDE SLESPROF
MACHINE ESA
SHARE RELATIVE 100 ABSOLUTE 20% LIMITSOFT
DEDICATE 0201 ADEV1
DEDICATE 0600 ADEV2
DEDICATE 0601 ADEV3
DEDICATE 0602 ADEV4
DISK 0203 FB-512 Y-DISK 256000 U
DISK 0203 FB-512 Y-DISK 256000 U
*? Build SuSe SLES 11 R 10
*? USERID | Y | UID | USER ID | 8
*? STORAGE | Y | RADIO | MEMORY SIZE | 64M 128M 512M 1024M 2048M
```

zPRO V2 speaks security

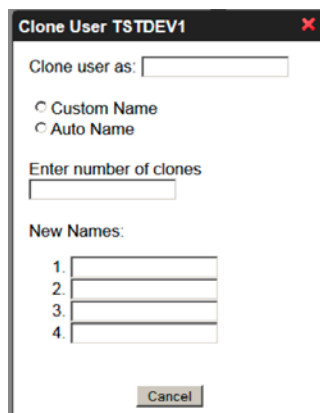
Security in zPRO V2 has been carefully considered. In addition to normal security features on-board the zVWS product, zPRO V2 implements additional layers:

- User identification and verification can be managed as a CP directory entry, an ESM entity, or an installation defined and manipulated process
- Configuration files have been put in place to define the scope of user authorities
- Exit points are available so that Installation designed/written exits can be put in place to facilitate the extension of security to address site specific concerns
- Velocity Software recommends the use of SSL connections to ensure secure transmission of sensitive data

zPRO V2 speaks “self service”

Should your installation have a need to implement a Platform As A Service (PAAS) environment, zPRO V2 can assist in that process as well. An end user can request that one, or multiple copies of a pre-existing virtual image, be created very dynamically.

Again, through the use of user authorizations, and the model feature, this type of function can be delegated to a non-SYSPROG end user.



Along with the ability to provision resources, comes the ability to reclaim resources. This can be accomplished on demand and/or automatically through the use of assigned resource expirations.

zPRO V2 features

The intent is to make zPRO V2 the most robust z/VM based cloud enablement and system management tool available. That is my/our definition. Your definition may be something different. So to address this potential gap, Velocity Software is looking for a few good installations to be part of the steering committee as zPRO V2 works its way to market. If you are willing, and brave enough, to join the few, the proud, the steering committee, contact either Maggie Necoy (necoy@velocitysoftware.com) or Tony Noto (tony@velocitysoftware.com), or call our corporate office, located in Mountain View, CA.

zPRO V2 Play Dates

Product availability is scheduled for early 4Q2014. The initial release will be limited to the steering committee members and will be running on the Velocity Software system for demonstration. Two environments will be established. One environment will be a z/VM 5.4 environment and the other a z/VM 6.3 environment. All members, will be given access to both systems.

TONY IS SENIOR SOFTWARE ARCHITECT AT VELOCITY SOFTWARE.

Life is zeasy with zPRO.



zPRO. The modern and easy way to manage your System z:

- light weight cloud enablement for z/VM and Linux on System z
- modern browser based systems management interface
- simplifies cloning, provisioning and virtualizing servers
- web based system administration

www.velocitysoftware.com

The VM Workshop – Don't miss it next year!

BY RICK TROTH

The VM Workshop is both educational and entertaining in addition to growing with each consecutive year. For 2014, the Workshop was at NC A&T, where the facilities made for a fantastic conference. Some attendees stayed in the conveniently located honors dormitory. Others stayed in nearby hotels. All participants gathered at the main classroom building where there were up to four simultaneous sessions. There were over 100 present at the event this time.


As an added bonus, Velocity Software held our in-depth Performance Seminar two days prior to the VM Workshop. This popular training class was made available to both Velocity customers and all Workshop attendees this year. Those who were able to attend both events got a full week of education and enlightenment.

The Workshop recognized long timers and newcomers. Of note, Fran Hart, renowned support person with a career spanning several VM-oriented companies, was formally honored for her work supporting the platform and the customers. Also, several NC A&T students, under tutelage of Professor Cameron Seay, were recognized. Dr. Seay has given his students a real enthusiasm for VM and Linux and other mainframe environments.



Topics included ...

- CMS admin "bootcamp"
- SES, DIRMAINT, Xedit, and related topics
- how to compete with VMware for Linux workloads
- Linux Guest Migration
- several REXX language sessions
- Backup and Recovery as well as Business Continuity (D/R)
- CMS SFS and related administration
- secure web using SSL
- automated operations and monitoring (a zVPS specialty)
- product updates from IBM (for z/VM) and SuSE and RedHat
- real life customer stories
- hands-on labs



Most of the session materials are available online at the Workshop web site:

<http://www.vmworkshop.org/2014/presentations>

If you were there, great! If you did not make it, we hope you'll be there next time!

RICK IS SENIOR SOFTWARE ENGINEER AT VELOCITY SOFTWARE.



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 in 2014? 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:

6 – 10 October: Enterprise2014
The Venetian, Las Vegas, Nevada, USA

17 – 21 November: SUSECon 2014
Coronado Springs Resort, Lake Buena Vista, Florida, USA

Events EMEA:

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

Imprint

Responsible for all contents: Velocity Software
Design: Götz Gramlich, Kolja van Boekel
Contributing Editors: Tony Noto, Barton Robinson, Rick Troth
Project Management: Magnelia Necoy
To subscribe for Velocity's Performance News please email:
info@velocitysoftware.com