THE CLIPPER GROUP
Navigator™

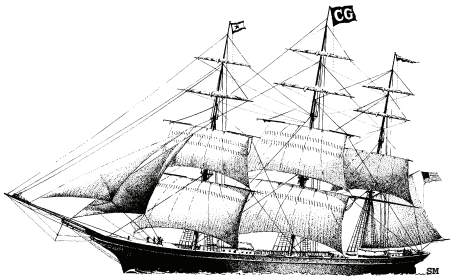

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2007030R February 28, 2007

Oh, The Things You Can Do … with z/VM 5.3!

Analyst: Anne MacFarland

Management Summary

With virtual machines, as with most things, it is not merely a question of "more" – an "and" is also involved.[1] The value of virtual machines is twofold. First, there is the envelope. When multiple virtual machines are running on a server, the envelopes protect the server and the guest environments from bad behavior (greed, corruption, illness, or death) on the part of one of the guests. Second, there is the improved resource utilization gained from the consolidation of multiple applications onto a single server. Its not just bragging rights that most IT folks are looking for, but hard-dollar savings in terms of power and cooling, as well as in hardware and software and management costs. **Not all virtual machines are alike**. The span of virtual machines runs from hardware partitions, through the software development realm in the form of *Java Virtual Machines*, to various software containers that both encompass and protect. In this latter category, ***VMware* is the best known, thanks to its early targeting of x86-based servers, but *z/VM*, which can trace its progenitors back to the 1960s, is the most evolved.** The growing popularity of mission-critical Linux workloads makes an articulation of z/VM's differentiating capabilities, and particularly the new capabilities of Release 5.3, timely.

- z/VM supports more virtual environments or guests per processor (*System z* Central Processor (CP)) or *Integrated Facility for Linux (IFL)* processor. It supports many hundreds, even thousands, not just a dozen, or two. z/VM can be deployed across more processors than other approaches to server virtualization. These features, taken together, mean you can consolidate more where you can, and support larger single processes, when you have a need for them. With Release 5.3, a virtual machine can utilize more kinds of processors.
- Because z/VM's resource controller, or hypervisor, does not just allocate memory but actively brokers reallocation of unused memory and other resources, more oversubscription of resources is supported safely, with a significant probability of lowering the TCO. The new release, Release 5.3, has some key new features in this area.
- If you add a z/VM instance of z/OS to your Linux environment, there are considerable tag-along enterprise-class benefits that the mainframe delivers.
- Because of the mega-consolidation that z/VM's granularity and complete resource virtualization support, and the management services inherent in z/VM, the carrying costs for large virtual architectures on a small footprint - even if it is a somewhat pricy footprint - can give a better TCO over time than virtualization on commodity hardware.

z/VM's latest upgrade is a good time to look at the new enhancements and at how they benefit both Linux and z/OS environments. For more details, please read on.

[1] For background on server virtualization and virtual machines, see the two-part series of **Clipper Notes:** (a) *Server Virtualization Made Real*, dated February 27, 2007, and available at http://www.clipper.com/research/TCG2007028.pdf, and (b) *Virtual Machines – Guest Environments for IT Infrastructure – A Tutorial*, dated February 25, 2007, and available at http://www.clipper.com/research/TCG2007029.pdf.

IN THIS ISSUE

Why z/VM is Different

z/VM is a software shell with a resource broker, or hypervisor, that coordinates the resource use of many operating systems. It is also part of IBM's *System z* mainframe architecture. z/VM provides support, not just for processor cycle and memory virtualization, but also for virtualization of networks and I/O. It works with mainframe capabilities, and leverages them for the benefit of its guests. This approach has the following benefits.

Efficiency

With z/VM, the virtualization is supported by z/VM, by *PR/SM*[2] and by the hardware at the instruction level. Linux[3] on System z runs unmodified as a guest environment, as can z/OS and other mainframe operating systems such as z/VSE and z/TPF, and also some special purpose applications. Each is all containerized and operates as if it owned the whole machine. Other virtual machines work by emulation, trap and translation, and other - less direct - methods, tracking commands and translating them to satisfy the underlying operating system, or to modify the operating system.

z/VM handles pooling, clustering for fail-over, resource sharing, and accounting, so that the use of one of its virtual environments does not add new needs for administrative management. Chargeback can be done at the level of a virtual machine, without the need for additional software to do so.

Flexibility/Granularity

With z/VM on System z, additional virtual machines can be added on the fly – in minutes – without taking the system down. This is not possible with other container-based approaches to server virtualization.

z/VM has multiple dimensions of scope that dwarf other approaches. Hundreds – even thousands - of Linux server images can run on a single instance of z/VM. You can run many copies of z/VM on a single mainframe[4] or across several mainframes. z/VM virtual machines can be nested within other z/VM virtual machines. Since little management headache is added by the use of z/VM, you can envision setting up a platform for interrelated applications, with appropriate clustering and redundancy that would take advantage of the mainframe's very large internal communications bandwidth.

Sharing

Over the decades, z/VM has honed a number of data-in-memory techniques, such as virtual networking, execute-in-place file systems, in-memory disk cache, etc., to allow co-located virtual machines to work together synergistically. These techniques are part of what enables such a large number of virtual machines to be supported on a single System z engine. In turn, this large number of co-located virtual machines means that you can parallelize applications to accelerate throughput when needed.

But there is more. With z/VM, guests can share executable processes. An administrator can build, not just a set of an isolated environments, but a multi-dimensional ecosystem, sharing where appropriate and shielding where that is needed, that can virtualize entire data center systems. **If you are tired of the hardware needed to support identical processes in a classical share-nothing scale-out physical environment, z/VM is an alternative you may have overlooked.**

Synergies with the System z Platform

Finally, as part of an integrated platform, z/VM offers some inherent advantages. Development of hardware, software, and environments like z/VM are all made with a comprehension of, and consideration for, the mainframe platform as a whole. Product support is, obviously, a one-stop experience. z/VM, together with System z Logical Partitions (LPARs), (see Exhibit 1, below), give the administrator a way to manage, isolate, share, and failover a number of processes using well-honed, highly-evolved automation. Co-residency with z/OS (something not available on Linux-only mainframes, of course) lets Linux workloads partake of some of z/OS's high-performance and quality-of-service features.

As part of a System z environment, z/VM also benefits from the mainframe's internal communications network, inherent security, and sharing features.

Exhibit 1 - Comparing z/VM and LPARs

System z hardware is and always has been architected for concurrent running of multiple processes. The classic partitions to support the multiple processes are called LPARs, and they have a heritage of dynamic flexibility of configuration and management that cannot be matched. z/VM and LPAR share a similar architecture, and together, their virtualization-on-virtualization provides a very large scale, not just of server consolidation, but also of system consolidation. The cooperation and coordination of z/VM with System z architecture has been productive and synergistic for years. z/VM 5.3, in a balanced way, just ups the ante across the board.

[2] PR/SM Stands for Processor Resource/Systems Manager., and is part of the firmware support for logical partitions (LPARs).

[3] The latest Linux kernel has support for a guest mode, which makes this even more efficient.

[4] Since Linux has an execute-in-place file system, you can load Linux onto mainframe storage, link it to z/VM, and get even more scalability.

Look below to see some of the new features that are being added in the latest release — z/VM 5.3, that will become generally available in June.

New in z/VM 5.3

More Scale

More virtual processors per z/VM instance are supported. A single copy of z/VM can now address up to 32 real System z processors (up from 24), which now may include the full array of CPs, *IFLs*[5], *zIIP*[6] engines, and *zAAP*[7] engines. (See details under "Specialty Processors", to the right.)

With z/VM Release 5.3 software, a z/VM system can exploit far more than the previous limit of 128 GBs of memory[8]. How much more depends of the nature of the hosted applications and their memory use, but the thrashing of memory contention does not occur until well over 200 GB of virtual memory is actively in use.

This "memory oversubscription" is supported in part by some z/VM memory management enhancements. A new *z/VM Collaborative Memory Management Assist* harnesses the ability of the z9 mainframe to coordinate memory paging to the level of individual pages.[9] This allows the system to identify memory that can be paged out without contacting the Linux guests. This new capability is paired with a new enhancement to z/VM called *Cooperative Memory*. Here, the *Virtual Machines Resource Manager (VMRM)* notifies Linux Guest Applications directly when memory contention starts to become a performance problem.

With Release 5.3, z/VM will be beefing up its virtual network capabilities in a number of useful ways. Of course, the high-speed hardware hypervisor within System z lets the hosted applications communicate with each other much faster than if they had to communicate over an external IP network. Within the z/VM, the communication is even faster. To improve the resilience of external communications, z/VM plans to support the IEEE 802.3ad Link Aggregation standard.[10] Link aggregation lets shared HBAs[11] be linked into a non-blocking mesh, so that, if an external switch fails, System z can re-route around the problem. The removal of another potential source of interruption enhances the consolidation of more workloads on System z via z/VM.

More System z Addressable Capabilities

Specialty Processors

With Release 5.3, System z specialty processors[12] can be emulated in z/VMs. They can be accommodated in two ways.

- In *virtualization support*, z/VM sends targeted workloads to physical specialty engines (zIIPs and zAAPs) in the mainframe server. (Obviously, this cannot be a Linux-only box.) Parts of a CP can be shared across several virtual machines, and parts of a virtual machine can be spread across multiple processors.
- With *simulation support*, z/VM creates zIIPs and zAAPs on System z standard CPs. This enhancement is a great boon to developers working on z/OS, because they can compare different *what-ifs* across a complete (though virtual) System z environment.

Back End Enhancements

- In addition, z/VM virtual machines can now take advantage of z/Architecture's MIDAW (Modified Indirect Address Word) capability. This is a tool for improving database I/O operations across multiple data sources. This is a timely enhancement, because enterprises of any size have found themselves with fistfuls of data sources and growing needs for customer relationship management and business analytics that can span them all.
- With z/VM 5.3, IBM has enhanced SCSI disk support for Linux guests.
- z/VM 5.3 can support IBM's 3592 Model E05 tape drive encryption when guest systems that have no inherent support for encryption It accesses an outboard encryption manager through an out-of-band connection.. It can also encrypt data of DDR and SPXTape back-up utilities. Alternatively, an added z/OS guest can also run the *Encryption Key Manager* application.
- *Hyper Parallel Access Volume (PAV)* support, requested by existing customers, has also been added. This is another I/O enhancement.

---

[5] Special, less-costly processors for executing Linux workloads.

[6] zIIP Processors are dedicated to database information integration. For more on zIIPs, see **The Clipper Group Navigator** dated January 24, 2006, entitled *System z adds zIIP to ally with DB2 on z/OS to better Serve the Onslaught of Business Data,* and available at http://www.clipper.com/research/TCG2006006.pdf.

[7] zAAP is an application assist processor to offload JAVA processing. For more, see **The Clipper Group Navigator** entitled *zSeries Zips through Java with zAAP*, dated April 7, 2004, and available at http://www.clipper.com/research/TCG2004030.pdf.

[8] If you think this is irrationally large, think of the business value of in-memory databases in key business processes.

[9] It uses System z's Host Page Management Assist.

[10] When System z supports *OSA Express 2* link aggregation support, sometime this year.

[11] Host Bus Adapters.

[12] Specialty processors offered to date include zAAP (to offload Java processing) and zIIP (to offload data serving); these processors are cheaper than standard engines and their workloads do not incur z/OS license charges.

Management and Security

- An enhanced API lets customers use the external management of *IBM Director 5.20*[13], IBM's multi-platform hardware manager, to manage their z/VM virtual environments. With its enhanced hooks into z/VM, administrators who do not have mainframe skills can manage their z/VM virtual environments with familiar controls and look/feel. IBM Director now includes two priced options targeted at z/VM.
  - ***z/VM Center*** is a way for z/OS systems to run low-priority mainframe utility workloads easily and more cheaply on Linux in a z/VM.
  - ***Software Distribution Premium Edition*** is a package that lets an administrator create software distribution bundles of commercial applications, custom software, or a little of both, for an IBM Director-managed system. For an enterprise with multiple mainframes and a significant investment in custom software, this can ease the distribution of application clusters.
- RACF support for z/VM environments is now available as a separately packaged feature called *RACF Security Server for z/VM*.
- RSCS (Remote Spooling Communications Subsystem) is now a one-time-charge feature option on z/VM, allowing for licensing on IFLs. RSCS is another way for Linux environments to connect to other systems.
- New z/VM LDAP support gives a single point of control for authentication of resource use by both physical (in LPAR) and virtual (in z/VM) resources. z/VM has added more native authentication, access controls, and SSL communication.
- A new guest ASCII console support gives a way to manage guest environments, if other means fail.
- With z/VM 5.3, IBM will pursue EAL 4 certification[14], featuring security of the hypervisor as well as of Linux and of the LPARs.[15] This is of interest to those in industries where regulations require such certification.

All of the new z/VM features represent a balanced build-out of processing capacity, memory-management enhancements, networking and I/O capabilities, and the security and management to make larger virtual environments truly supportable.

For z/OS customers, z/VM 5.3 is useful more broadly. It gives new ways to build a logical-level *proximity* that can accelerate business processes. It gives reasons to bring more application workloads onto the mainframe.

For Linux users, z/VM 5.3 offers a great deal more. The non-disruptive creation of z/VM virtual environments has always had great appeal for gaming and other real-time business models based on viral allure, but now this kind of flexibility can be brought to the composites of allocations that support business processes. It offers an alternative to the limitations of open system boxes and their management, which gave us three-tier and other physically-defined deployment models. Consider the business ramifications of moving from fixed arrays of physical assets (think armies marching in formation) to logically aggregated, segregated, and secure cells of business process support.

Pricing

z/VM Version 5 has moved from a monthly charge to a flat, one-time fee. This means that if you migrate an existing z/VM to a new mainframe, you do not have to pay IBM again. IBM has also tweaked its pricing to encourage purchasing of new z/VM licenses. While there has always been a volume discount, in the past z/VMs were priced per enterprise per processor per location. Now they are priced per processor per enterprise. Global enterprises will get more value from volume pricing. And the price is uniform across geographies (not counting the vagaries of currency exchange).

Conclusion

The z/VM environment is like a very large group of people with rather complicated names who know each other very well that are prepared to do large amounts of serious work - very efficiently. They can support joint processes and each other with far less explanatory babble. Like such a group, z/VM has a highly-evolved *modus operandi*.

z/VM is so highly evolved for two reasons. For one, it has been around longer and has had an investment of tens of thousands of person-hours. In addition, it is a complete server environment that has grown in a balanced, sensible fashion. It is capable of hosting many servers, and has the resiliency to be a suitable mega-host. It offers the kind of hyper-data-center-consolidation that can give the aggregate costs of other kinds of infrastructures a run for their money. If you are using more infrastructure and liking it less, or if you would prefer fewer constraints on how you build support for a business process, consider z/VM.

[13] IBM Director 5.20 was released in January. Workload Management and the Dependency Service for z/OS, Linux, UNIX, AIX, HP/UX, i5/OS, Windows, and Solaris) are now available as separate products through *IBM Virtualization Engine*, separate from IBM Director 5.20.

[14] z/VM already has EAL 3+ certification.

[15] Coincidentally, the IBM Labs are developing, together with Xen, a secure shell for open systems hypervisors named *sHype*.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Anne MacFarland **is Director of Data Strategies and Information Solutions for The Clipper Group.** Ms. MacFarland specializes in strategic business solutions offered by enterprise systems, software, and storage vendors, in trends in enterprise systems and networks, and in explaining these trends and the underlying technologies in simple business terms. She joined The Clipper Group after a long career in library systems, business archives, consulting, research, and freelance writing. Ms. MacFarland earned a Bachelor of Arts degree from Cornell University, where she was a College Scholar, and a Masters of Library Science from Southern Connecticut State University.

- ***Reach Anne MacFarland via e-mail at Anne.MacFarland@clipper.com or at 781-235-0085 Ext. 128. (Please dial "128" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, **The Clipper Group Explorer**, **The Clipper Group Observer**, **The Clipper Group Captain's Log**, **The Clipper Group Voyager**, **Clipper Notes**, and *"clipper.com"* are trademarks of The Clipper Group, Inc., and the clipper ship drawings, *"Navigating Information Technology Horizons"*, and *"teraproductivity"* are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.