

# Migrating z/VM Customers to V6.4

Bill Bitner (bitnerb@us.ibm.com)

z/VM Development Lab Client Focus and Care

*Last Updated May 1, 2017*



## Migrating Customers to z/VM 6.4

Everyone knows someone who is still hanging on and older release of z/VM. In many cases, it isn't because they don't have the hardware to run the latest z/VM release. They just haven't found a reason that is good enough to go ahead and migrate.

This document is meant to help open the discussion to reasons why migration might be of value to a customer. While the focus is on those with z/VM 5.4, it can be helpful to other releases as well.

The document lists questions to ask a client or to ask yourself if you're the client, and then maps those questions to value statements of newer releases.

If you need additional help in understanding the solutions or problems, let me know. I'm here to help.

Questions to ask your client	z/VM 6.4 Feature/Benefit/Value	Old z/VM Release where z/VM 6.4 advantage applies			
		V5.4	V6.1	V6.2	V6.3
Do you have systems that are at the limit for real memory?	Support for 2 TB of real memory <ul style="list-style-type: none"> <li>• Efficient usage of larger real memory sizes</li> <li>• Allows vertical scaling for larger workloads.</li> <li>• Helps reduce logical partition (LPAR) sprawl, by consolidating more virtual machines into a single LPAR.*</li> <li>• Reduces administrative expenses through managing a smaller number of large-capacity z/VM host servers.</li> </ul>				
Are you running with a higher Virtual to Real Memory ratio than optimal because of the limit for real memory?		X	X	X	X
Are you managing multiple z/VM logical partitions as you needed to split the workload due to memory constraints?					
Have you SET REORDER OFF for large virtual machines in order to get acceptable performance?	Memory Management changes <ul style="list-style-type: none"> <li>• improved the page steal selection process where the old reorder processing was replaced with new algorithms</li> <li>• Improved efficiency of page selection for virtual machines up to 1 TB in size*</li> </ul>	X	X	X	
Do you have a large amount of page space defined on smaller DASD (ECKD) because of limit of one concurrent page I/O?	HyperPAV Support for System Volumes <ul style="list-style-type: none"> <li>• Page, Spool, and resident volumes</li> <li>• Allows for greater paging bandwidth</li> <li>• Support for user volumes already existed.</li> </ul>	X	X	X	X

\* Dependent on workload characteristics.

Questions to ask your client	z/VM 6.4 Feature/Benefit/Value	Old z/VM Release where z/VM 6.4 advantage applies			
		V5.4	V6.1	V6.2	V6.3
Do you find yourself using a lot of paging space for an environment where you do not over commit memory to any large degree?	Ability to avoid duplication of page data <ul style="list-style-type: none"> <li>New KEEEPSLOT option allows you to avoid keeping a copy of a user page on disk after it is page faulted into real memory (this is a trade-off, but now you can make the choice where appropriate).</li> </ul>	X	X	X	X
Do you want more efficient paging I/O that exploits processor and storage server enhancements?	High Performance FICON (zHPF) Support <ul style="list-style-type: none"> <li>z/VM system I/O, paging, now exploits zHPF.</li> </ul>	X	X	X	X
Do you wish to use IBM FlashSystem Storage for z/VM system volumes (spooling, paging, EDEVs)?	New support for FlashSystems that removes the requirement to use a Spectrum Virtualize (aka SAN Volume Controller, SVC) device	X	X	X	X

\* Dependent on workload characteristics.

Questions to ask your client	z/VM 6.4 Feature/Benefit/Value	Old z/VM Release where z/VM 6.4 advantage applies			
		V5.4	V6.1	V6.2	V6.3
Do you want to exploit processor features such as SIMD, Transactional Execution Facility, and EDAT 1 (1 MB pages)?	Allows guest exploitation of key z Systems processor facilities <ul style="list-style-type: none"> <li>• SIMD – Single instruction Multiple Data is the newest vector facility.</li> <li>• Transactional Execution Facility – allows more efficient serialization in compilers</li> <li>• EDAT 1 – allows for 1MB guest pages (backed by 4KB pages in z/VM).</li> </ul>	X	X	X	X
Do you have Linux guests that are tight on Memory, especially running databases?					
Are you running with a larger number of IFL cores? Greater than 16? 24? The limit of 32 (64 on a z13)	HiperDispatch <ul style="list-style-type: none"> <li>• Strengthens the affinity between where work is dispatched and where the data used by that work exists in processor cache</li> <li>• Improves throughput and processor efficiency</li> <li>• Allows greater elasticity in logical processor usage</li> </ul> Processor Scalability (via Service March 13, 2015) <ul style="list-style-type: none"> <li>• Improves the n-way curves on all supported processors*</li> <li>• Increases the IFL Cores supported to 64 on the z13 (Limit is 32 Cores on z13 with SMT enabled)</li> </ul>	X	X	X	
Do you limit the size of your logical partitions, in terms of number of IFLs, due to a drop off in performance?					
Do you change the number of online IFLs to your z/VM LPARs on a regular basis (VARY OFF/ON) to address processor resource needs and availability over time?					
Do you limit adding additional IFL cores because the capacity gain is too low compared to the software expense increase?					

Questions to ask your client	z/VM 6.4 Feature/Benefit/Value	Old z/VM Release where z/VM 6.4 advantage applies			
		V5.4	V6.1	V6.2	V6.3
Are you looking for increased capacity in a z/VM footprint?	Simultaneous Multithreading (SMT) <ul style="list-style-type: none"> <li>Allows z/VM to use two threads on IFL Cores to dispatch virtual CPUs with SMT support available on z13, z13s, or LinuxONE server.*</li> </ul>	X	X	X	
Do you want to explore the benefit of SMT, but are restricted due to the need to IPL to make SMT-related changes?	Dynamic SMT allows z/VM to dynamically set the number of threads per core for a z/VM system that has SMT enabled.	n/a	n/a	n/a	X
Do you have virtual machines that get “stuck” in the eligible-list?	Improvements in the z/VM Scheduler and Dispatcher <ul style="list-style-type: none"> <li>Remove the complex and confusing eligible-list construct</li> <li>Greatly improve the “fairness” and “predictability” of the z/VM scheduler</li> </ul>	X	X	X	X
Do you have trouble computing normalized Share values when tuning your z/VM system?					
Do you have a system where virtual machines do not need all of their entitled Share and want better distribution of the surplus?					

\* Dependent on workload characteristics.

Questions to ask your client	z/VM 6.4 Feature/Benefit/Value	Old z/VM Release where z/VM 6.4 advantage applies			
		V5.4	V6.1	V6.2	V6.3
Are you looking for ways to lower software license costs of IBM passport advantage products on Linux guests?	<p>CPU Pooling and IBM License Metric Tool (ILMT)</p> <ul style="list-style-type: none"> <li>• Available in z/VM service June 2014 and ILMT in August 2014</li> <li>• CPU pooling allows you to assign a limit or cap on the processor resources that can be used by a set of virtual machines in aggregate.</li> <li>• CPU pooling is similar to limit shares, but applies across a multiple virtual machines.</li> <li>• Support also allows ILMT to gather information on the CPU pool limits for IBM licensing purposes of Passport Advantage products. This is not 'usage' based pricing, but 'capacity' based.</li> </ul>	X	X	X	
Do you want a way to limit a group of virtual machines in aggregate without having to micromanage each virtual machine with SHARE settings?					
Are you running multiple LPARs today to limit the number of logical IFLs as part of managing software license costs?					
Are you avoiding the consolidation of additional workloads or applications onto an existing Linux on z/VM environment for fear of increasing software costs?					
Do you struggle with planned outages for critical workloads on your z/VM systems?	<p>Live Guest Relocation (LGR)</p> <ul style="list-style-type: none"> <li>• Introduced with z/VM 6.2, LGR is part of the optional VMSSI feature.</li> <li>• Allows most Linux guests to be relocated from one z/VM member of an SSI Cluster to another member without stopping the applications running in the virtual machine.</li> <li>• Helps clients avoid planned outages for virtual servers when performing maintenance</li> </ul>	X	X		

Questions to ask your client	z/VM 6.4 Feature/Benefit/Value	Old z/VM Release where z/VM 6.4 advantage applies			
		V5.4	V6.1	V6.2	V6.3
Are you currently running an SSI Cluster with z/VM 6.2 or 6.3, and would like to gain benefits of z/VM 6.4 but are afraid of needing a cluster-wide outage?	Upgrade in place <ul style="list-style-type: none"> <li>• New installation process that allows step-wise approach to bring an SSI cluster forward to a new release</li> <li>• Can continue to push out the new release individually on separate members of the SSI cluster</li> </ul>	n/a	n/a	X	X
Is full support for your z/VM system important to you?	Fully supported <ul style="list-style-type: none"> <li>• z/VM 6.1 reached End of Service on April 30, 2013</li> <li>• z/VM 6.2 will reach End of Service on June 30, 2017</li> <li>• z/VM 5.4 &amp; 6.3 will reach End of Service on December 31, 2017</li> <li>• Note while some older releases are still in service, enhancements will not be moved back to them. This may include support for newer processors</li> </ul>	X	X	X	X



Questions to ask your client	z/VM 6.4 Feature/Benefit/Value	Old z/VM Release where z/VM 6.4 advantage applies			
		V5.4	V6.1	V6.2	V6.3
Do you want greater availability for your virtual network environment?	z/VM VSwitch support for LACP Groups with shared OSAs via service to z/VM 6.3 when used with IBM z13 <ul style="list-style-type: none"> <li>• PTFs available June 24, 2015</li> <li>• Link aggregation allows VSwitches to aggregate bandwidth of multiple OSAs and offer more effective failover handling</li> </ul>				
Have you been interested in using Link Aggregation (LACP) with z/VM VSwitches, but cannot afford to dedicate OSA ports to your z/VM LPARs?		<b>X</b>	<b>X</b>	<b>X</b>	

## References for z/VM 6.4

- z/VM Home Page for z/VM 6.4
  - <http://www.vm.ibm.com/zvm640/>
  
- z/VM 6.4: A Customer Driven Release
  - Live Virtual Class November 11, 2016
  - <http://www.vm.ibm.com/education/lvc/zvmlvc.html>
  
- Security for z/VM 6.4 – News and How To's
  - Live Virtual Class March 15, 2017
  - <http://www.vm.ibm.com/education/lvc/zvmlvc.html>
  
- z/VM 6.4 Data Sheet
  - <https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=POD03122USEN&>

# Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

BladeCenter*	FICON*	OMEGAMON*	RACF*	System z9*	zSecure
DB2*	GDPS*	Performance Toolkit for VM	Storwize*	System z10*	z/VM*
DS6000*	HiperSockets	Power*	System Storage*	Tivoli*	z Systems*
DS8000*	HyperSwap	PowerVM	System x*	zEnterprise*	
ECKD	IBM z13*	PR/SM	System z*	z/OS*	

\* Registered trademarks of IBM Corporation

## The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.  
 Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.  
 Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.  
 IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.  
 ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.  
 Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.  
 Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and  
 Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.  
 Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.  
 OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the [OpenStack website](#).  
 TEALEAF is a registered trademark of Tealeaf, an IBM Company.  
 Windows Server and the Windows logo are trademarks of the Microsoft group of countries.  
 Worklight is a trademark or registered trademark of Worklight, an IBM Company.  
 UNIX is a registered trademark of The Open Group in the United States and other countries.

\* Other product and service names might be trademarks of IBM or other companies.

### Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g., zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the "Authorized Use Table for IBM Machines" provided at [www.ibm.com/systems/support/machine\\_warranties/machine\\_code/aut.html](http://www.ibm.com/systems/support/machine_warranties/machine_code/aut.html) ("AUT"). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.