

## IBM z/VM

### Virtualization for extreme scalability, security, efficiency on IBM Z® and IBM LinuxONE™

---

#### Highlights

- Foundation for on-premise cloud on IBM Z and IBM LinuxONE
  - Host many virtual servers running different operating systems
  - Virtualize and share resources with very high levels of utilization
  - Continuous Delivery model enabling fast benefit
- 

IBM® z/VM® provides high levels of extreme security, scalability, and efficiency, providing a robust foundation for on-premise cloud computing. z/VM virtualization technology is designed to run hundreds to thousands of Linux servers on a single IBM Z or LinuxONE server with the highest degrees of efficiency, elasticity, and security.

Its ability to support numerous machine images and solution's architectures provides a highly flexible production and test environment for IBM Z and LinuxONE operating systems. That simplifies release migration, provides a test environment whenever needed, and deploys/consolidates several systems onto one physical server.

IBM z/VM V7.1 and IBM z/VM V6.4 will provide support for all IBM z15™ (z15) and IBM LinuxONE III servers.

#### IBM z/VM V7.1

The z/VM Continuous Delivery model delivers z/VM capabilities for faster adoption and benefit. New z/VM capabilities will be delivered in the service stream of the current release as Small Programming Enhancements<sup>1</sup>, thus providing the flexibility of selecting and deploying new capabilities immediately, along with moving from one release to another release in the new two-year cadence.

Details about new z/VM functions and subscribing to be alerted when they become available are found at:

[www.vm.ibm.com/newfunction](http://www.vm.ibm.com/newfunction)

More capabilities are delivered with z/VM V7.1<sup>2</sup>, including the enhancements planned for Continuous Delivery in 2019:

*SSI function integrated in z/VM base:* Continuous operation is extended with z/VM V7.1, since it includes the Single System Image feature as part of the base at no additional cost, which offers Live Guest Relocation. This helps continuous deployment of new capabilities as the new delivery model is adopted by enabling continuous availability of Linux applications while implementing new functions provided for z/VM.

*Dump Scalability:* This capability of the z/VM Dump process helps to reduce the time required to create, process, and transmit data from snap and Hard Abend dumps. The increased efficiency can save time, resources, and makes the deployment of z/VM configurations with large amounts of memory more feasible.

*ESM Authorization and Auditing of SMAPI Requests:* Two functions are delivered in support of an external security manager (ESM). First, some commands can use the current dynamic command protection setting of the LINK command when validating the required LINK authorizations, and second, programs can use the ESM for all Systems Management API (SMAPI) authorization decisions at the same granularity used with the SMAPI existing authorization mechanism.

*Elliptic Curve Support:* The z/VM TLS/SSL server is enhanced to improve security through the enablement of Elliptic Curve Cryptography cipher suites, providing a faster, more secure mechanism for asymmetric encryption than is provided by standard RSA or DSS algorithms.

*RSCS Query System Service:* This functionality provides a means to query the service level of each part that is included in RSCS. A new command parameter returns the highest level PTF that is applied to each part within the running RSCS server, eliminating ambiguity on what service is installed and active.

*VSwitch Priority Queuing:* When VSwitch Priority Queuing is enabled, z/VM establishes multiple OSA QDIO Output queues and transmits data to the external network at different priorities based on guest NIC importance.

*EAV Paging:* Sufficient paging capacity for z/VM partitions with large memory sizes can be defined on Extended Address Volume (EAV) devices. The paging space can be located anywhere on an EAV, thus reducing the number of paging devices required and the associated burden of managing a larger number of smaller paging devices. As systems continue to grow, this functionality helps when the need for paging space increases.

*80 Logical processor support:* 80 logical processors are supported on z15, z14™ (z14), LinuxONE III, LinuxONE Emperor™ II, and LinuxONE Rockhopper™ II, increasing the support from 64 logical processors per LPAR. This allows running more workload on a single z/VM instance, which is especially important when multithreading is enabled. Defining more logical processors for workloads on each LPAR may mean that fewer LPARs are required to support the same workloads.

*Dynamic Crypto*<sup>3</sup>: Dynamic Crypto support enables dynamic changes to the Adjunct Processor (AP) Cryptographic environment on a z/VM system, allowing the addition or removal of crypto hardware to be less disruptive to the system and its guests.

*Fast Minidisk Erase*<sup>3</sup>: This capability provides a means to remove data from ECKD minidisks in a faster, more efficient manner when a user ID or a minidisk is deleted.

*Dynamic Memory Downgrade*<sup>3</sup>: With this capability, z/VM memory management will allow the dynamic removal of real memory from a running z/VM system. Dynamic memory downgrade will require a z15, z14, LinuxONE III, Emperor II, or Rockhopper II.

## **z/VM Support<sup>4</sup> for IBM z15 and IBM LinuxONE III**

z/VM V7.1 and z/VM V6.4 will support z15 and LinuxONE III.

Both versions will enable guest exploitation for:

- Synchronous execution support for on-chip data compression
- Enhanced Vector and Vector packed decimal processing
- Crypto Express7S adapter and cryptographic enhancements
- FICON® Express16SA, OSA adapters
- Dynamic I/O enhancements, providing support for the configuration management of OSA-Express7S OSD CHPIDs, Crypto Express7S, FICON Express16SA FC and FCP CHPIDs, RoCE Express2, and Coupling Express3 LR adapters

In addition, z/VM V7.1 will support IBM System Recovery Boost, by providing a temporary utilization of the sub-capacity general purpose processors at full capacity during z/VM system initialization and workload bring-up, workload quiesce and system shutdown, and during system abend processing. System Recovery Boost will return the system to doing normal work faster, after any kind of planned or unplanned disruption. Support is primarily targeted to the z/VSE guest environment.

## **z/VM V7.1 and z/VM V6.4 capabilities**

z/VM V7.1 and z/VM V6.4 provide support for all z15, z14, IBM z13™ (z13), IBM z13s™ (z13s), IBM zEnterprise® EC12, IBM zEnterprise BC12 and all LinuxONE servers, and the Linux distributions from Canonical, Red Hat and SUSE, and the IBM Z operating systems IBM z/OS®, IBM z/VSE®, and IBM z/TPF.

Sub-Capacity pricing is available with z/VM V7.1 and z/VM V6.4. It allows for software pricing at less than full machine capacity and can provide more flexibility and improved cost of computing when managing the volatility and growth of new workloads. For more information read: [Sub-Capacity for z/VM](#).

## Efficiency and Scalability

*Encrypted paging support<sup>4</sup>*: Ciphering will occur as data moves between active memory and a paging volume owned by CP. z/VM provides support for encrypted paging, in support of the philosophy of encrypting all data in flight and at rest. Included in the support is the ability to dynamically control whether a running z/VM system is encrypting this data.

*Guest exploitation support for the Instruction Execution Protection Facility (IEPF)<sup>4</sup>*: The IEPF provides functionality to improve the security of programs running on IBM Z and LinuxONE by allowing virtual memory elements to be identified as containing only data. If an attempt is made to fetch an instruction from an address in such an element or if an address in such an element is the target of an execute-type instruction, a protection exception will occur. z/VM provides support for guest exploitation of the IEPF.

*Guest exploitation support for Pause-Less Garbage Collection<sup>4</sup>*: The guarded storage facility (GSF) is designed to improve the performance of garbage-collection processing by various languages, in particular Java™. z/VM provides support for guest exploitation of the GSF.

*Real memory and guest virtual support*: The maximum amount of real memory that z/VM exploits is 2 TB. The maximum supported virtual memory for a single guest is 1 TB. When configured with 2 TB of real storage and keeping the same over-commitment ratio for virtual-to-real memory, this can double the amount of virtual memory that can be efficiently used compared to older z/VM releases.

*HyperPAV technology exploitation*: z/VM exploits the ability for an IBM DS8000® device to execute concurrent I/O requests to an ECKD paging volume. In HyperPAV mode, if the base volume is busy, z/VM selects a free alias device from a pool, binds the alias to the base device, and starts the I/O. When the I/O completes, the alias device is returned to the pool to be used for another I/O to the same logical subsystem (LSS). The primary benefit of exploiting HyperPAV is to improve paging throughput during periods of high-volume disk I/O, which will increase the efficiency of z/VM memory management for memory over-committed workloads. HyperPAV paging also enables the effective use of fewer and larger CPOWNEED volumes.

With HyperPAV paging taking advantage of DS8000® features the bandwidth for paging increases to allow managing dozens of paging volumes rather than 100s and allows for more efficient memory management of over-committed workloads.

HyperPAV is exploited also for the SYSRES volume and for volumes containing checkpoint and warm start data, volumes used for spooling and the z/VM user directory, and those associated with minidisk pools, as defined by a guest's use of MAPMDISK IDENTIFY.

*Support for the Enhanced-DAT facility:* A larger page size decreases the amount of guest memory needed for dynamic address translation (DAT) tables and also decreases the hardware overhead required to perform address translation. z/VM provides support for the Enhanced-DAT facility, which allows a guest to exploit 1 MB pages. In all cases, guest memory is mapped into 4 KB pages at the host level. With z/VM's large page support for its guests, Linux on Z, z/VSE, and z/OS virtual machines can benefit from reduced memory footprints and address translation times, which in turn can decrease overhead and improved throughput.

*Guest Transaction Execution (TX) support:* z/VM supports guest exploitation of the TX facility on supported machines. The TX facility allows a program to issue multiple instructions that appear to operate atomically, offering an alternative to a costlier mutual-exclusion mechanism such as software locks. This support can be exploited to improve the efficiency and scalability of multithreaded software such as Java or guest operating systems.

*Dump Processing Enhancements<sup>4</sup>:* As systems become larger, dumps also become larger, take more time to create, and are more difficult to store and transmit. With Dump Processing Enhancements, the amount of time it takes for z/VM to write a Hard Abend or snap dump to 3390 DASD may be reduced. The improvements were achieved via changes to the I/O channel program.

*Processor Scalability Efficiency Improvements<sup>4</sup>:* z/VM manages internal spinlocks more efficiently and thereby reduce system overhead. This functionality helps to improve performance and throughput for large n-way configurations and thereby to improve overall system capacity by allowing additional work to be performed. These improvements are greatest for workloads experiencing significant Scheduler Lock contention. Larger n-way configurations will tend to see greater benefit.

*Extended Address Volume (EAV) Minidisk Support<sup>4</sup>:* Enhanced EAV support for 3390-A DASD devices allows non-full pack minidisks to reside anywhere on the volume, including beyond the current restriction of the 64K cylinder boundary, and up to the one terabyte limit currently supported.

*Virtual Switch Enhanced Load Balancing<sup>4</sup>:* z/VM supports exclusive and shared Multi-VSwitch Link Aggregation configurations to improve load balancing and leverage both horizontal and vertical growth in single and cross virtual switch networking configurations. With this improvement, a VSwitch can better utilize the capacity of the OSA devices used for link aggregation whether as a single VSwitch growing vertically or a Multi-VSwitch growing horizontally.

## System Ease of Use

*Control Program environment variables:* This capability allows automation procedures to adapt more easily to changes in operating environments to help simplify the control and testing of a system setup. For example, an operator can indicate at IPL time that the system is running in a disaster recovery or test environment, which in turn enables automation routines to modify the devices used and alter the choice and sequence in which virtual machines are activated, as well as perform other environment-dependent functions.

*Query Shutdown command:* The new QUERY SHUTDOWN command enables a z/VM system programmer or a guest virtual machine to determine whether a system shutdown is in progress and to obtain additional information about the shutdown. This can help automate an orderly termination of the z/VM system and its virtual servers. This function can be of particular value to virtual machines that coordinate the shutdown of others. These coordinating virtual machines can receive the signal that the system is shutting down, issue the new QUERY command to get additional information, and take actions appropriate for an orderly shutdown.

*SCSI enhancements for z/VM:* Improved Small Computer System Interface (SCSI) support for guest attachment of disk and other peripherals to IBM Z and LinuxONE servers:

- Enables ease of use with enhanced management for SCSI devices to provide information needed about device configuration characteristics.
- Enhances interoperability between the SCSI driver and SAN Volume Controller (SVC) and devices incorporating SVC technology such as the IBM Storwize® V7000 and IBM FlashSystem™ V840 and V9000.
- Allows a z/VM storage administrator to use FlashSystem storage as a z/VM-system-attached disk without the need for an intermediate SAN Volume Controller (SVC). Previously, while FlashSystem could be used by a Linux virtual machine without an SVC, to use it for z/VM system volumes or EDEVs for virtual machines, an external or internal SVC was required.
- Improves reliability when SCSI disk devices are attached to the z/VM hypervisor for system use, without the need to be attached behind an SVC.
- Supports issuing multiple I/O requests concurrently to EDEVs.<sup>4</sup>

SCSI management EDEVICE query commands are enhanced to improve the usability and problem diagnosis of EDEV-intensive environments and provide a clearer end-to-end view of the storage configuration. This simplifies the process of verifying that the storage configuration is consistent between z/VM and the disk storage subsystem.

The following updates are designed to further enhance the reliability of SCSI devices:

- The CP missing interrupt handler is disabled for EDEVICES, allowing the SCSI driver to manage its outstanding requests in a more appropriate manner.
- The SCSI driver is updated to provide additional path recovery.
- Debug facilities within the SCSI driver are enhanced, allowing IBM support teams to more quickly diagnose and debug issues in the field.
- Guidelines for multi-path configuration are provided for SVC and devices incorporating SVC technology to ensure path recovery is optimal.

*z/VM CMS Pipeline Update:* The integration of new CMS Pipelines functionality provides a much more inclusive set of tools for application developers. This upgrade eliminates the need to download Pipelines code, includes fixes not previously integrated into the z/VM product, broadens the ecosystem, enables innovation and includes additional functionality.

*DirMaint RACF connector with z/VM:* The DirMaint RACF® connector improves how z/VM security is handled in a managed environment. The connector allows appropriate security policy changes to be passed directly to RACF, allowing a z/VM environment managed by IBM Wave for z/VM or via an OpenStack environment to function properly with RACF installed on the system.

*z/VM RACF automate control of access list authority:* The ADDCREATOR and NOADDCREATOR options are added to the RACF SETROPTS command and determine whether the creator of a RACF profile is automatically added to its access control list. This enhancement removes the need for manual intervention in RACF resource configuration and eliminates a point of potential human error from security policy management.

*z/VM Performance Toolkit enhancements<sup>4</sup>:* z/VM Performance Toolkit exploits z/Architecture® and its expanded set of instructions. Consequently, the PERFSVM virtual machine must run on z/Architecture CMS (z/CMS). New and updated performance reports are provided within the Performance Toolkit Feature in support of HyperPAV Paging. These new reports include information that helps tune the z/VM HyperPAV Paging Subsystem.

*Network Security Enhancement:* The NICDEF Security Controls introduces Directory Network Authorization capability with which each virtual NIC can be configured and authorized entirely within the user directory. This eliminates the need to issue SET VSWITCH and COUPLE commands to complete the network configuration.

*Encryption of TCPNJE connections:* RSCS TCPNJE traffic can be encrypted by directing the flow through an SSL server. The secure TCP/IP protocols that were previously implemented to support VMCF clients and servers are extended for IUCV clients and servers.

*z/VM Cloud Connector:* The z/VM Cloud Connector is a development toolkit that manages the z/VM host and its virtual machines. Upper layer cloud management solutions can consume these RESTful APIs directly to manage z/VM. For additional information: [z/VM Cloud Connector](#).



## Hardware Currency

*z/VM supported using IBM Dynamic Partition Manager:* IBM z/VM is a supported environment using IBM Dynamic Partition Manager (DPM) for Linux-only servers with SCSI storage. DPM 3.1 provides support for ECKD DASD. This simplifies system administration tasks for a more positive experience with IBM Z or LinuxONE.

*z-Thin Provisioning<sup>4</sup>:* DS8880 z-Thin Provisioning and Extent Space-Efficient (ESE) volumes are supported. This enables guest support for thin-provisioned volumes and allows CPOWNER volumes be defined on thin-provisioned volumes.

*Dynamic Simultaneous Multithreading (SMT) Level:* Once z/VM has been IPLed with multithreading enabled in the system configuration file, the SET MULTITHREAD command can be used to switch non-disruptively between one and two activated threads per IFL core. Performance of a system and workload with one active thread per core is comparable to that of the same system and workload with multithreading disabled. Thus, the Dynamic SMT Level capability allows the benefit of evaluating multithreading for a workload without requiring an outage to enable or disable SMT. The SET MULTITHREAD command is allowed only when the system has been enabled for multithreading in the system configuration file, which can specify activating either one or two threads per core. It is not possible to revert to a non-SMT configuration without an IPL. SMT-enabled configurations are restricted to 32 or 40 cores, depending on the z/VM software level, even when operating in single-threaded mode due to the logical processor addressing limit of 64 or 80.

*I/O architecture enhancements on the z13 (Driver D27) and z13s:* A user defined identifier (UID) can be assigned to a real PCI function to more accurately indicate equivalent functions between different LPARs and for exploitation by guest operating systems, and in particular by Linux. z/VM supports dynamic I/O and guest use for PCIe UID support.

*Shared Memory Communications Direct (SMC-D):* The SMC-D protocol provides fast, low-latency LPAR-to-LPAR traffic using Direct Memory Access over firmware-provided Internal Shared Memory (ISM) devices. Supported for z/VM guest exploitation, SMC-D and ISM are designed to use shared memory areas to provide low-latency, high-bandwidth, cross-LPAR connections for applications. SMC-D is expected to provide substantial performance, throughput, response time, and CPU consumption benefits compared with standard TCP/IP communications over HiperSockets™. z/VM supports dynamic I/O and guest use of the Internal Shared Memory (ISM) PCI function type.

## Installation and Serviceability

*Enhanced upgrade in place:* The upgrade in place process allows upgrading an existing system to a new release of z/VM with minimal impact to the running system. The upgrade in place process allows the upgrade from z/VM V6.2 or V6.3 to z/VM V6.4 and positions a system for releases beyond z/VM V6.4. Upgrade in place is supported for a member of a z/VM SSI cluster as



well as for a non-clustered z/VM system. The only release supported for upgrading to z/VM V7.1 is V6.4.

*Determine installed service:* Enhancements to CP and VMSES/E enable you to determine if specific CP service is built into the CP nucleus (load module). This CPSERVICE option on the CP QUERY command allows queries based on APAR, PTF, or local modification identifiers of the nucleus that is currently running.

*3590 and 3592 tape formats not supported:* z/VM is available on DVD and electronic delivery. In addition, z/VM service for all releases will no longer be orderable in 3590 or 3592 tape format.

*Installing z/VM V6.4 on a z14* requires updated installation media available since August 25, 2017.

*Dump to tape* is not supported with z/VM V7.1.

*ESA/390 removal:* z/VM V6.4 enhancements enable hypervisor initialization and termination, the Stand-Alone Program Loader (SAPL), DDR, Stand-Alone Dump, and stand-alone utilities to run entirely in z/Architecture mode.

*Architecture Level Set (ALS):* z/VM V7.1 requires IBM zEnterprise EC12 or BC12 or later systems, and z/VM V6.4 requires IBM zEnterprise 196 and IBM zEnterprise 114 and later systems. See the appropriate preventive service planning (PSP) bucket for the minimum microcode level and any required updates.

## Foundation for on-premise Cloud

The z/VM hypervisor extends the capabilities of the IBM Z and IBM LinuxONE enterprise platforms from the standpoint of sharing hardware assets, virtualization facilities, and communication resources. Together with [IBM Wave for z/VM](#), the comprehensive administration solution for z/VM-based virtual Linux server environments, z/VM makes it easier to receive the maximum value from large-scale virtual server hosting on IBM Z and LinuxONE. This includes software and personnel savings, operational efficiency, power savings, and optimal qualities of service. This virtualization technology is designed to enable organizations to run thousands of Linux servers on a single server with the highest degrees of efficiency and elasticity. New services can be launched quickly and easily. In addition, development teams can rapidly deploy and remove test environments.

With z/VM and Linux, IBM Z and LinuxONE provides a highly scalable, secure and efficient on-premise cloud infrastructure for the hybrid cloud architecture.

<sup>1</sup> Often referred to as “New Function APARs”

<sup>2</sup> Please see detailed information at: [www.vm.ibm.com/newfunction](http://www.vm.ibm.com/newfunction), including the required PDFs and APARs

<sup>3</sup> Planned support availability with PTFs in 2019. Statements by IBM regarding its plans, directions, and intent are subject to change or withdrawal without notice at the sole discretion of IBM. Information regarding potential future products is intended to outline general product direction and should not be relied on in making a purchasing decision.

<sup>4</sup> Support available with PTFs

<sup>5</sup> A PTF is required to get this capability with z/VM V6.4; it is in the base of z/VM V7.1.

## Why IBM?

As you transform your business and differentiate yourself in a trust economy, IBM remains your partner.

We have the total expertise in systems, software, delivery and financing to help you create a secure, open and intelligent foundation for the future.

Our experts can help you configure, design and implement z/VM on IBM z15, not only as your on-premise cloud infrastructure, but always optimized for your needs.

## For more information

To learn more about the IBM z/VM offering, please contact your IBM representative or IBM Business Partner, or visit the following website: [ibm.com/it-infrastructure/z/zvm](https://ibm.com/it-infrastructure/z/zvm)

To keep informed and about the latest z/VM capabilities and news visit: [www.vm.ibm.com/newfunction](https://www.vm.ibm.com/newfunction)

© Copyright IBM Corporation 2019.

IBM, the IBM logo, and ibm.com are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at <https://www.ibm.com/legal/us/en/copytrade.shtml>, and select third party trademarks that might be referenced in this document is available at [https://www.ibm.com/legal/us/en/copytrade.shtml#section\\_4](https://www.ibm.com/legal/us/en/copytrade.shtml#section_4).

This document contains information pertaining to the following IBM products which are trademarks and/or registered trademarks of IBM Corporation:  
IBM, IBM logo, DS8000, FICON, IBM FlashSystem, HiperSockets, LinuxONE, LinuxONE Emperor, LinuxONE Rockhopper, RACF, Storwize, IBM Z, z/Architecture, IBM z14, IBM z13, IBM z13s, zEnterprise, z/OS, z/VM, z/VSE

---



Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

---

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