IBM z/VM

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

IBM® z/VM® delivers extremely high levels of security, scalability, and efficiency, providing a robust foundation for on-premises private cloud computing. z/VM virtualization technology is designed to run hundreds to thousands of Linux® servers and IBM Cloud Paks® based on Red Hat® OpenShift® Container Platform 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 architectures provides a highly flexible production, development and test environment for IBM Z and LinuxONE operating systems. This simplifies enterprise solutions and infrastructure upgrades in a timely manner, provides a test environment whenever needed, and deploys and consolidates several systems onto one physical server.

IBM 7.2, IBM z/VM 7.1 and IBM z/VM 6.4 provide support for all IBM z15™ (z15) and IBM LinuxONE III servers.

IBM z/VM 7.2

IBM z/VM 7.2 provides the hypervisor for hosting enterprise-class virtual servers to exploit the IBM Z and IBM LinuxONE advantages in scalability, performance, high availability, and security.

While cloud computing has become the standard use model for IT services, the IT infrastructure continues to be the foundation for every IT service, including strong virtualization technology, such as z/VM.

Highlights

- Foundation for private cloud on IBM Z and IBM LinuxONE
- Host IBM Cloud Paks based on Red Hat OpenShift in z/VM virtual machines
- Host many virtual servers running different operating systems
- Virtualize and share resources with very high levels of utilization
- Continuous Delivery model enabling fast benefit from new functions

With z/VM 7.2, IBM continues to offer the choice to deploy product enhancements via the z/VM Continuous Delivery model for faster adoption and client benefit. New z/VM capabilities will be delivered in the service stream of the latest release as Small Programming Enhancements (new Function APARs), thus providing the flexibility to select and deploy new capabilities immediately, along with moving from one release to another on a regular two-year cadence. Details about new z/VM functions and subscribing to availability alerts are found at: www.vm.ibm.com/newfunction.

IBM z/VM 7.2 enhancements include:

Centralized Service Management, a simplified process for applying and deploying maintenance across multiple non-Single System Image (SSI) z/VM systems in a timely manner. Centralized Service Management provides support to deploy service to multiple systems, regardless of geographic location, from a centralized primary location that manages distinct levels of service for a select group of z/VM systems. This centralized service process keeps track of available service levels and manages whatever files are needed to supply a client-defined service level to a managed system. Centralized Service Management supports servicing both base z/VM and preinstalled features.

Even greater availability and data redundancy for Geographically Dispersed Parallel Sysplex (GDPS) environments and improved direct access storage device (DASD) scalability for clients exploiting Multiple Subchannel Set Multi-Target Peer-to-Peer Remote Copy (MSS MT-PPRC). This function enables a disk device to be the primary disk for up to three secondary devices in each of up to three alternate subchannel sets.

An enhanced test and debug environment provided by z/VM ADJUNCT support. z/VM ADJUNCT support enables the manipulation of a guest operating system running in a primary configuration from an adjunct configuration of the same virtual machine. The companion environment can display and modify the primary's memory, making it useful for observing and manipulating a guest system that is running there.

Roll up of timely, client-driven functions[1] previously delivered in the service stream of z/VM 7.1:

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 customer-defined guest NIC importance.

EAV Paging: Paging capacity for z/VM partitions can be defined on Extended Address Volume (EAV) devices. This can reduce 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, IBM 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. Defining more logical processors for workloads on each LPAR may mean that fewer LPARs are required to support the same workloads.

Dynamic Crypto: 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: 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.

RACF® Multifactor Authentication for z/VM: This functionality enables a z/VM system with an External Security Manager to authenticate a z/VM user ID via a non-password token. This provides greater security by requiring an additional form of proof to help avoid an exposure if one token becomes compromised.

TLS/SSL Certificate Verification: This functionality allows authentication of client certificates, host name validation and extraction of fields from a certificate. In addition, new APIs are provided to allow fields to be extracted from a client or server certificate programmatically.

CMS Pipelines SSL/TLS Enhancements. This functionality enhances the CMS Pipelines TCP/IP stages and allows applications to establish secure connections using SSL/TLS. Both implicit SSL/TLS connections (e.g., HTTPS) and explicit SSL/TLS connections (e.g., FTP) are supported. A sample stage is provided that implements a client for FTP with SSL/TLS.

z/VM 7.2 includes a new Architecture Level Set (ALS) that requires an IBM z13, IBM z13s, IBM LinuxONE Emperor, IBM LinuxONE Rockhopper, or later server.

Planned Continuous Delivery enhancements for z/VM 7.2 in 2020[1,2] are:

CP New Feature Interrogation API: This functionality will provide an interface for interrogating whether a named capability is present in the running CP system. This enhancement will simplify scripting and automation written by customers and vendors that is sensitive to the function level of CP.

FlashCopy® Preserve Mirror Support: This functionality will provide new operands for z/VM FlashCopy commands to allow managing FlashCopy relationships between PPRC primary devices without affecting the PPRC status of the target devices. This will avoid PPRC suspend conditions on FlashCopy target devices.

SET EDEVICE Optional LUN Specification: This functionality will allow specification of the LUN operand on the SET EDEVICE command and the EDEVICE system configuration statement to be optional after the first path has been established.

z/VM Support for IBM z15 and IBM LinuxONE III

z/VM 7.2, z/VM 7.1, and z/VM 6.4 will support[3] z15 and LinuxONE III.

All versions will enable guest exploitation for:

- Synchronous execution support for on-chip data compression and deflate-conversion

- Enhanced sort acceleration and optimization

- 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 OSAExpress7S

- OSD CHPIDs, Crypto Express7S, FICON Express16SA FC and FCP CHPIDs, RoCE Express2, and Coupling Express3 LR adapters

- IBM Fibre Channel Endpoint Security between an z15 Model T01 or LinuxONE III Model LT1 and the IBM DS8900F

In addition, z/VM 7.2 and z/VM 7.1 support IBM System Recovery Boost in support of instant recovery, by providing temporary utilization of 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 7.2, z/VM 7.1, and z/VM 6.4 capabilities

z/VM 7.2, z/VM 7.1, and z/VM 6.4 provide support for all z15, z14, IBM z13® (z13), IBM z13s® (z13s) 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 7.2, z/VM 7.1, and z/VM 6.4. 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

Elliptic Curve Support[3,4]: 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 standard RSA or DSS algorithms.

ESM Authorization and Auditing of SMAPI Requests[3,4]: 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 existing SMAPI authorization mechanism.

Dump Scalability[4]: 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.

Encrypted paging support[3]: Ciphering is performed as data moves between active memory and a paging volume owned by CP. z/VM encrypted paging supports 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)[3]: 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[3]: 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 CPOWNED 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 also exploited for the SYSRES volume, 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[3]: 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[3]: z/VM manages internal spinlocks more efficiently and thereby reduces 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[3]: 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[3]: 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 LPAR 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

RSCS Query System Service[4]: This functionality provides a means to determine 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 about what service is installed and active.

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 deployment. 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 of 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 based on that response 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 EDEVICEs.[3]

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 Pipelines 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[3]: 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: NICDEF Security Controls introduce Directory Network Authorization capabilities 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 a 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.

Hardware Currency

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

z-Thin Provisioning[3]: DS8880 z-Thin Provisioning and Extent Space-Efficient (ESE) volumes are supported. This enables guest support for thin-provisioned volumes and allows CPOWNED 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 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 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 6.2 or 6.3 to z/VM 6.4 and positions a system for releases beyond z/VM 6.4. Upgrade in place is supported for a member of a z/VM SSI cluster as well as for a non-clustered z/VM system. Only 6.4 is supported for upgrading to z/VM 7.1.

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 for electronic delivery. z/VM service for all releases has not been orderable in 3590 or 3592 tape format since September 21, 2018. *Installing z/VM 6.4* on a z14 requires updated installation media.

Dump to tape is not supported with z/VM 7.1.

ESA/390 removal: z/VM 6.4 enhancements enable hypervisor initialization and termination, the Stand-Alone Program Loader, DDR, Stand-Alone Dump, and stand-alone utilities to run entirely in z/Architecture mode.

Foundation for private cloud

z/VM, together with Linux and IBM Cloud Paks based on Red Hat OpenShift Container Platform, IBM Z and LinuxONE provide a highly scalable, secure, and efficient on-premises private cloud infrastructure for your hybrid multicloud.

Realizing the benefits of cloud computing requires an infrastructure that delivers availability, reliability, security, and performance, while also providing strong virtualization technology. Virtualization is foundational to delivering Infrastructure as a Service, a basic building block for cloud.

To accelerate cloud adoption, IBM Cloud Infrastructure Center can help cloud administrators and cloud development teams to provide a consistent, industry-standard user experience for defining, instantiating, and managing the lifecycle of virtual infrastructure, and deployment of images. It also integrates with higher-level cloud automation tools, such as IBM Cloud Automation Manager or VMware vRealize Automation or vRealize Orchestrator.

z/VM, together with the comprehensive administration solution IBM Wave for z/VM, makes it easier to extract 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.

z/VM extends the capabilities of the IBM Z and IBM LinuxONE enterprise platforms from the standpoint of sharing hardware assets, virtualization facilities, and communication resources, and IBM continues to expand z/VM, evolving to meet the needs of IT organizations.

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 Z and IBM LinuxONE, not only as your on-premises private 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

To keep informed and about the latest z/VM capabilities and news visit:
www.vm.**ibm.com**/newfunction

[1] *Please see detailed information at: www.vm.ibm.com/newfunction, including the required PTFs and APARs.*

[2] *Planned support availability with PTFs in 2020.*

[3] *Support available with PTFs. Please see detailed information at: www.vm.ibm.com/newfunction.*

[4] *Not available with z/VM 6.4.*

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