z/VM Version 5 Release 4

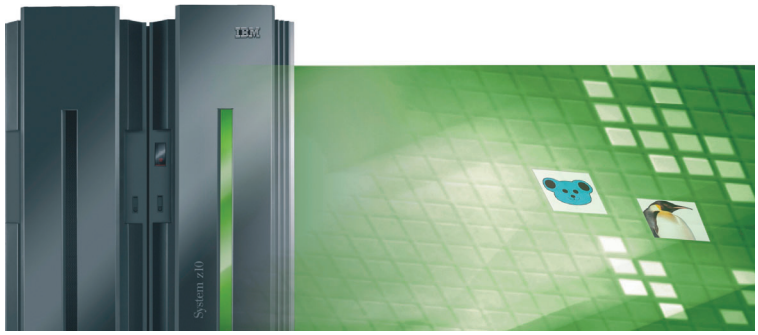

Flexible solutions for the competitive advantage

Highlights

- ***Increased flexibility with support for new z/VM®-mode logical partitions***
- ***Dynamic addition of memory to an active z/VM LPAR by exploiting System z® dynamic storage-reconfiguration capabilities***
- ***Enhanced physical connectivity by exploiting all OSA-Express3 ports including support for QDIO data connection isolation***
- ***Capability to install Linux® on System z from the HMC without requiring an external network connection***
- ***Enhancements for scalability and constraint relief***
- ***Operation of the SSL server in a CMS environment***
- ***Systems management enhancements for Linux and other virtual images***

Building successful virtual enterprises

The IBM z/VM hypervisor can help clients extend the business value of mainframe technology across the enterprise by integrating applications and data with exceptional levels of availability, security and operational ease. z/VM virtualization technology is designed to allow clients to run hundreds to thousands of Linux servers on a single mainframe with other IBM System z operating systems, such as z/OS®, or as a large-scale Linux-only enterprise server solution. z/VM V5.4 can also help to improve the productivity of hosting non-Linux workloads such as z/OS, z/VSE™ and z/TPF.

IBM System z

IBM System z (IBM System z10™, IBM System z9® and IBM eServer™ zSeries®) platforms offer a range of servers designed to be integrated into a robust, flexible infrastructure. IBM is expanding on the success and

widespread acceptance of the IBM System z9 platform to make it available to a broader set of clients with different computing needs with the System z10 Enterprise Class (z10 EC), System z10 Business Class (z10 BC), System z9 Enterprise Class (z9 EC), and the System z9 Business Class (z9 BC). This revolutionary brand responds to unprecedented demand by providing high levels of performance and scalability.

The z10 EC, with its advanced combination of reliability, availability, serviceability, security, scalability, and virtualization, delivers the technology that can help define this framework for the future. The z10 EC delivers improvements to performance, capacity and memory that can help enterprises grow their existing businesses while providing a cost-effective infrastructure for large-scale consolidation.

With increased capacity, the z10 EC virtualization capabilities can help to support more virtual servers than any other platform—hundreds or thousands of virtual servers in a single 2.83 square meter footprint. When consolidating on System z, you can create virtual servers on demand, achieve network savings through HiperSockets™ (internal LAN), provide security to enable and support new and existing applications, help improve systems management of virtual servers and, most importantly, consolidate software from distributed processors onto fewer processors.

The IBM System z10 Business Class (z10 BC) delivers innovative technologies for small and medium enterprises that give you a whole new world of capabilities to run modern applications. Ideally suited as the cornerstone of your new enterprise data center, this competitively priced server delivers unparalleled qualities of service to help manage growth and reduce cost and risk in your business. The z10 BC further extends the leadership of System z by delivering expanded granularity and optimized scalability for growth, enriched virtualization technology for consolidation of distributed workloads, improved availability and security to help increase business resiliency, and just-in-time management of resources. The z10 BC is at the core of the enhanced System z platform and is the new face of System z. The z10 BC Model E10 is designed to provide up to 1.5 times the total system capacity for general-purpose processing, nearly four times the available memory and over 40% more configurable processors than the z9 BC Model S07.

The System z environment, with self-configuring and self-healing attributes, provides new functions and features to meet the challenges of businesses. IBM mainframes provide reliability, security, scalability, virtualization and availability.

Put the power of System z environments combining partitioning and z/VM virtualization technology to work for you to help realize the benefits of workload isolation and resource sharing, including the:

- *Reliability, availability and serviceability of System z*
- *Flexibility to create as many as 60 LPARs on the z10 EC and z9 EC*
- *Increased memory of up to 248 GB on z10 BC*
- *Enhancements to Crypto Express2 and Crypto Express2-1P*
- *WWPN prediction tool for preplanning and setup of a SAN for a System z10 environment*
- *Ability to virtualize each LPAR into hundreds or more virtual machines*
- *Ability to virtualize processor, communication, memory, storage, I/O and networking resources*
- *Maximizing resources to achieve high system utilization*
- *Advanced dynamic resource allocation*
- *High-speed communications among LPARs and guests with IBM HiperSockets*
- *Advanced systems management, administration and accounting tools*

z/VM Version 5 (V5)

z/VM V5 offers attractive levels of price/performance, functional capabilities and hardware exploitation that help increase the advantages of deploying

Linux solutions on the mainframe. You can add capacity to System z servers for hosting Linux workloads by configuring them with Integrated Facility for Linux (IFL) processors.

When Linux is a guest of z/VM, it is designed to allow the capability to run hundreds to thousands of Linux images on a single System z server. These Linux images can be deployed on standard or IFL processors with z/VM V5. z/VM V5 requires IBM z/Architecture® (64-bit) and operates on the IBM z10 EC, z10 BC, z9 EC and the z9 BC, the zSeries 990 (z990), zSeries 890 (z890), zSeries 900 (z900) and zSeries 800 (z800).

z/VM V5 offers an ideal platform for consolidating select UNIX® and Linux workloads on a single System z server by providing a virtualization environment for hosting other IBM mainframe operating systems as guests, including Linux on System z, z/VSE and z/TPF.

z/VM V5.4 is designed to extend its IBM System z virtualization technology leadership by exploiting more capabilities of the IBM System z servers including:

- *Increased flexibility, with support for the new z/VM-mode logical partitions, allowing all System z processor-types (CPs, IFLs, zIIPs, zAAPs, and ICFs) to be defined in the same z/VM LPAR for use by various guest operating systems*
- *Capability to install Linux on System z from the HMC that eliminates the need for any external network setup or a physical connection between an LPAR and the HMC*
- *Enhanced physical connectivity by exploiting all OSA-Express3 ports, helping to service the network and reducing the number of required resources*
- *Nondisruptive dynamic addition of memory and processors to an active z/VM LPAR by exploiting System z dynamic reconfiguration capabilities*
- *Port isolation security that provides the ability to restrict guest-to-guest communications within a Virtual Switch (VSWITCH) by exploiting OSA-Express QDIO data connection isolation*

The operation and management of virtual machines has been enhanced with new systems-management APIs, improvements to the algorithm for distributing a guest's CPU share among virtual processors, and usability enhancements for managing a virtual network.

Security capabilities of z/VM V5.4 provide an upgraded LDAP server at the function level of the z/OS V1.10 IBM Tivoli® Directory Server and enhancements to the RACF Security Server to create LDAP change log entries in response to updates to RACF group and user profiles, including user passwords and password phrases. The z/VM SSL server now operates in a CMS environment, rather than requiring a Linux distribution, thus allowing encryption services to be deployed more quickly and helping to simplify installation, service, and release-to-release migration.

Enhanced scalability and constraint relief

z/VM V5.4 provides additional enhancements to help enhance its scalability and provide storage (memory) constraint relief. These include:

- *Moving specific CP data structures out of the first 2 GB of main storage reduces contiguous-frame allocation searches for requests in the first 2 GB and may help improve performance and scalability for systems with more than 2 GB of real storage, particularly those supporting large virtual storage environments.*
- *Removing the constraint that restricted Discontiguous Saved Segments (DCSS) to residing below 2047 MB in virtual storage. Also, this allows stacking DCSSs to grow the amount of shared memory in aggregate beyond 2 GB.*
- *Using the dynamic reconfiguration function of z/VM, reserved storage can be assigned and accessed when the guest operating system decides to use it.*

z/VM V5 already provides support to help enhance scalability and ease constraints:

- *z/VM V5.3 provided:*
  - *Support for significantly more real memory, up to 256 GB, twice the size supported by z/VM V5.2, and more than 1 terabyte (TB)[1] of total virtual memory in use by guests, benefiting customers with large amounts of real memory, which may help reduce or eliminate the need to spread large workloads across multiple z/VM images.*
  - *Support for customer growth by allowing up to 32 real processors to be configured in a single z/VM system on an IBM System z server, a 33% increase from z/VM V5.2.*
  - *Support for the Collaborative Memory Management Assist (CMMA) on the z10 EC, z9 EC and z9 BC, by which z/VM and Linux guests exchange information to optimize their use and management of memory.*
  - *Improved memory management algorithms to help benefit paging workloads with large memory environments.*
  - *Processors dynamically added to or removed from a z/VM LPAR in reserve without preplanning with the PTFs for APARs VM64249, VM64323, and VM64389.*
  - *Enhanced memory utilization using Virtual Machine Resource Manager (VMRM) between z/VM and Linux guests assists in managing memory contention in the z/VM system.*
  - *Support for the Hyper Parallel Access Volume (HyperPAV) function of the IBM System Storage DS8000® series.*
  - *Enhanced IBM FlashCopy® support allows the specification of up to 12 target minidisks, determines the status of FlashCopy requests, and exploits hardware asynchronous cache destage and discard that makes the FlashCopy appear synchronous to the virtual machine.*
  - *Support for the IBM System Storage SAN Volume Controller (SVC) Storage Engine 2145 allows Linux on System z guests of z/VM V5 (all releases) to access the IBM System Storage disk subsystems, including the DS4000® series and OEM SCSI disk devices supported by the SVC.*
- *Prior releases provided:*
  - *Improved scalability with the control program (CP) now using memory locations above 2 GB for a much broader set of functions. These improvements can offer constraint relief for large real-memory, virtual-server environments that are memory-intensive.*
  - *Enhancements to CP to increase the number of Linux and other guest virtual machines that can be managed concurrently.*

[1]The actual amount of usable real and virtual memory is dependent on the amount of real memory in the z/VM logical partition, the hardware server model, firmware level and configuration, and the number of guests and their workload characteristics.

Virtualization technology enables Linux and other guests

With z/VM and IFL processors, a low-cost, flexible environment can be created to test and develop on Linux while simultaneously running Linux production applications. z/VM V5 support for IFL processors is designed to run Linux workloads without increasing the IBM software charges for z/OS, z/VM, z/VSE, TPF or z/TPF operating systems and applications running on System z standard processors. Only Linux or OpenSolaris workloads in an LPAR or Linux guests of z/VM V5 can operate on IFL processors.

z/VM V5.4 provides virtualization technology enhancements in support of Linux and other guests, including:

- *Removing one of the consequences of stopped virtual processors for virtual processor prioritization by performing share redistribution whenever a virtual processor is started or stopped and no longer includes stopped virtual processors in the calculation of how much share to distribute to each virtual processor.*
- *Providing additional support for Linux on System z guests using Dynamic Storage Reconfiguration with the PTF for APAR VM64524 for V5.4 that:*
  - *Allows a more compatible operation when running a second level on z/VM as when running directly under LPAR*
  - *Displays configured, standby, and reserved values for each storage element via the QUERY VIRTUAL STORAGE command*
- *Capability to dump Linux on System z guests to FCP-attached SCSI disks:*
  - *More guest virtual memory can be dumped because SCSI disks can be larger than ECKD disks.*
  - *Dumping on SCSI disks avoids the need to convert a* **VMDUMP** *into a Linux tool format.*
  - *Allows the same SCSI dump mechanism to be used when running Linux in a logical partition and in a z/VM virtual machine.*

Previous releases of z/VM provided virtualization technology enhancements that included:

- *New guest support for specialty processors, zAAPs and zIIPs, including:*
  - *Simulation support–z/VM guest virtual machines can create virtual specialty processors on processor models that support the same types of specialty processors but don't necessarily have them installed. Virtual specialty processors are dispatched on real central processors (CPs). Simulating specialty processors provides a test platform for z/VM guests to exploit mixed-processor configurations. This allows users to assess the operational and CPU utilization implications of configuring a z/OS system with zIIP or zAAP processors without requiring the real specialty processor hardware.*

— *Virtualization support—z/VM can create virtual specialty processors for virtual machines by dispatching the virtual processors on specialty processors of the corresponding type in the real configuration. Guest support for zAAPs and zIIPs may help improve your total cost of ownership by allowing available zAAP and zIIP capacity not being used by z/OS LPARs to be allocated to a z/VM LPAR hosting z/OS guests running Java™, IBM DB2®, and other zIIP- and zAAP-eligible software.*

- *Usability enhancements for the VSWITCH and guest LAN environments including enhanced ease-of-use for Virtual LAN (VLAN) and promiscuous mode configuration changes, as well as a new capability to configure a native VLAN ID identifier.*
- *Guest use of Modified Indirect Data Address Words (MIDAWs) to allow more flexibility and performance in certain channel programs, as an alternative to data-chained channel-command words (CCWs). This allows guest operating systems to exercise their code-paths just as they would on the real machine during, for example, pre-production testing of z/OS systems.*
- *Guest access to the system ASCII console to facilitate recovery of the guest during an emergency.*
- *Additional enhancements to Small Computer System Interface (SCSI) disk support for Linux users including Point-to-Point Fibre channel links, dynamically determined preferred paths for emulated FBA devices (EDEVICEs) on SCSI disks in an IBM System Storage DS6000™, faster formatting of EDEVICEs on SCSI disks, and display of additional SCSI device characteristics.*
- *Support for the OSA-Express2 and OSA-Express3 Open Systems Adapter for NCP to help eliminate the requirement to have any form of external medium (and all related hardware) for communications between the host operating system and the CCL image.*
- *Installation of z/VM from a DVD to an IBM System Storage SCSI disk or to 3390 DASD.*
- *An IBM HyperSwap™ function that allows the virtual devices associated with one real disk to be swapped transparently to another. HyperSwap can be used to switch to secondary disk storage subsystems mirrored by Peer-to-Peer Remote Copy (PPRC).*
- *A dynamic virtual machine time-out capability enables a guest operating system to specify an action to be taken by the z/VM Control Program (CP) if the guest becomes unresponsive.*
- *Enhancements to the VMRM provide the infrastructure to support more extensive workload and systems resource management features.*
- *Virtual IBM FICON® CTCA devices for guest operating systems enhancing previous virtual-CTCA capabilities by adding the FICON protocol as an option. Guests use virtual CTCAs to communicate among themselves within a single z/VM system image, without the need for real FICON CTCAs.*
- *Support for real and virtual integrated 3270 console devices. Real-device support enables this device, provided through the Hardware Management Console (HMC), to be used as the z/VM system operator console.*
- *Virtual Coupling Facility (CF) support was enhanced to allow z/VM systems to run as second-level (or higher) guests while simulating z/OS coupled sysplexes. This allows the testing of a z/OS Parallel Sysplex® environment at any guest level.*

Exploiting new technology

z/VM provides a highly-flexible test and production environment for enterprises deploying the latest business solutions. Enterprises that require multi-system server solutions will find that z/VM can help them address the demands of their businesses and IT infrastructures with a broad range of support for operating system environments such as z/OS, VSE/ESA™, z/VSE, TPF, z/TPF, CMS and Linux on System z. The ability to support multiple machine images and architectures enables z/VM to run multiple production and test versions of System z operating systems, all on the same System z server. z/VM can help simplify migration from one release to another, facilitate the transition to newer applications, provide a test system whenever one is needed and consolidate several systems onto one physical server. z/VM can also be used to enable access to the latest storage and processor architectures for systems that lack such support. Technological enhancements in z/VM are designed to exploit z10 EC, z10 BC, z9 EC, z9 BC, z990 and z890 servers including:

- *Support for specialty processors to allow all processor types (CPs, IFLs, zIIPs, zAAPs, and ICFs) on a System z10 to be defined as virtual CPUs, all in the same z/VM LPAR.*
- *Guest exploitation of the System z10 at the level of System z9.*
- *Exploitation of selected functions of the System z10 including:*
  - *Dynamic I/O configuration to define, modify and delete a Coupling channel using InfiniBand® link, CHPID type CIB, when z/VM V5.4 is the controlling LPAR for dynamic I/O.*
  - *Processors dynamically added to or removed from a z/VM LPAR in reserve without preplanning.*
  - *Recognizing all four ports in V5.4 on a System z10 OSA-Express3 Gigabit Ethernet (GbE) and 1000BASE-T Ethernet feature and two ports on the GbE and 1000BASE-T 2P features on the z10 BC providing more physical connectivity to service the network and reduce the number of required resources (I/O slots, I/O cages, and fewer CHPIDs to define and manage).*
  - *Port isolation security that provides the ability to restrict guest-to-guest communications within a Virtual Switch (VSWITCH) by exploiting OSA-Express QDIO data connection isolation with the PTFs for APARs PK67610 and VM64463 and required minimum MCLs.*
  - *Hardware Configuration Definition (HCD) provides I/O device information from the input/output definition file (IODF) for the World-Wide Port Name (WWPN) prediction tool (available from Resource Link™) with the PTF for APAR VM64579 for V5.3 and V5.4.*
  - *TCP/IP and VSWITCH gain the performance benefit of OSA-Express3 on the z10.*
- *z/VM V5.4 can be installed in an LPAR and both V5.4 and Linux on System z can be installed in a virtual machine from the Hardware Management Console (HMC) DVD drive on the IBM System z10.*
- *Additional systems management enhancements are provided by the HMC and Support Element (SE) 2.10.1 exploiting the z/VM System Management APIs to allow selected virtual resources to be defined and managed.*
- *Support for the IBM FlashCopy SE feature on the IBM DS8000, which provides a space-efficient snapshot capability that can greatly reduce the storage capacity needed for point-in-time copies with the PTF for VM64449 for V5.4.*

- *Prior releases included:*
  - *Facilities to dynamically add and delete logical partitions using CP's Dynamic I/O command interface and the z/VM HCD/HCM support when operating on the z10, z9, z990 and z890 servers.*
  - *I/O-configuration definition and dynamic-I/O configuration for up to 60 LPARs on the z10 EC and z9 EC.*
  - *Support the OSA-Express Integrated Console Controller (OSA-ICC) helping to eliminate the requirement for external console controllers.*
  - *Virtual switch exploitation of Layer 2 for OSA-Express, OSA-Express2, OSA-Express3 and link aggregation support for OSA-Express2 and OSA-Express3 devices.*
  - *Support for OSA-Express3 Gigabit Ethernet (GbE), 10 GbE, and 1000Base-T Ethernet on a z10.*
  - *Support the System z capability to cascade two FICON directors within a fibre-channel fabric. z/VM and its guests can take advantage of this enhanced and simplified connectivity, which is particularly useful in disaster recovery and business continuity procedures.*
  - *Support for FICON Express2 and FICON Express4 (4 Gigabit/second) on the z10 and z9 can help increase channel capacity and performance.*
  - *z/OS and Linux for System z guest support for Crypto Express2.*
  - *Support for the On/Off Capacity on Demand (On/Off CoD) and the Capacity Backup Upgrade (CBU) functions on System z servers, including functional enhancements that allow z/VM to recognize and report changed processor configuration and capacity settings on a z10, z9, z990 or z890.*
  - *Support for native control unit modes for the DS6000 (1750 control unit) and DS8000 (2107 control unit) series.*
  - *Features of the DS6000 and DS8000 supported by z/VM include:*
    - *Parallel Access Volumes (PAVs) as minidisks for guest operating systems that exploit the PAV architecture and provide the potential benefit of PAVs for I/O issued to minidisks owned or shared by guests that don't exploit PAVs.*
    - *Dynamic Volume Expansion simplifies disk management by allowing the dynamic increase of a DS8000 volume size to accommodate application data growth with the PTFs for APARs VM64305 and VM64354.*
    - *FlashCopy V2, designed to enable business continuance solutions with the delivery of new FlashCopy functions and enhancements is intended to help improve business efficiency, along with FlashCopy performance improvements that may help to help minimize operational disruption.*
    - *Guest use of Peer-to-Peer Remote Copy Extended Distance (PPRC-XD), designed to copy full volumes of data in non-synchronous mode.*
    - *Capability to define and operate FCP-attached SCSI disks with capacities of approximately 1 TB (2,147,483,640 512-byte blocks) for CP volumes and 381 GB for CMS and GCS volumes.*

- *Preferred paths for I/O operations to devices attached to a DS6000 1750 control unit that automatically switches the data path used to help improve overall performance.*
- *Support for the 65,520 cylinder (55.7 GB) 3390 Model 54 volume to help relieve addressing constraints, improve disk resource utilization and improve storage administrator productivity by providing the ability to consolidate multiple disk volumes into a single address.*

- *Support for the IBM System Storage Enterprise 3592 Tape Controller Model J70 and C06 and 3592 Tape Drive Models J1A, E05, and E06, which are designed to provide new levels of performance and attachment capabilities for System z. z/VM, including DFSMS/VM™, also support Write Once Read Many (WORM) data cartridges. Support for the Model E06 requires the PTFs for APARs VM64458 and VM64459 for z/VM V5.2 and later, including DFSMS/VM.*

Systems management

Enhancements in systems management, some of which help provide **self-configuring**, **self-managing** and **self-optimization** features in z/VM V5.4 including:

- *Providing enhancements to the z/VM Virtual Systems Management Application Programming Interface (API) for ease of use in creating and managing large numbers of Linux and other virtual images running on z/VM:*
  - *Sockets-based servers have been updated to communicate among themselves via shared queues, helping to improve the speed of request processing, enhance error recovery, and increase the scalability of the Systems Management API.*
  - *New APIs and enhancements to existing APIs to help manage virtual images.*

Previous releases of z/VM provided systems management enhancements that included:

- *A new z/VM Virtual Systems Management API sockets-based server in V5.3 to support the API. The sockets-based server is multitasking capable and supports both AF_INET and AF_IUCV socket requests.*
- *An interface in z/VM V5.3 to allow basic z/VM systems management functions to be performed from the Hardware Management Console (HMC) without having to establish network connections and reducing complex configuration of the system.*
- *Assisting network administrators to help manage virtual network performance, find and solve virtual network problems and plan virtual network growth by establishing a method of providing Simple Network Management Protocol (SNMP) data for virtual networking devices.*
- *Repackaging Remote Spooling Communications Subsystem (RSCS) V3.2.0 (5684-096) and preinstalling it on z/VM V5.3 system DDRs. RSCS is available as a priced, optional feature and is available for both IFL and standard processor configurations. RSCS V3.2.0 has been withdrawn from marketing effective May 26, 2008 and service support is planned to be discontinued effective October 31, 2009.*
- *Enhancing the guest LOGON process by providing a new COMMAND statement in a virtual machine definition or profile to configure the virtual machine.*

- *Hardware Configuration Manager (HCM) and Hardware Configuration Definition (HCD) components to create and manage your I/O configuration, providing a comprehensive, easy-to-use I/O configuration-management environment similar to that available with z/OS.*

Performance management

Performance Toolkit for VM provides enhanced capabilities for a z/VM systems programmer, operator or performance analyst to monitor and report performance information. The Performance Toolkit is functionally equivalent to RealTime Monitor (RTM) and Performance Reporting Facility (PRF). The Performance Toolkit for VM provides:

- *Full-screen mode system console operation and management of multiple z/VM systems*
- *Post-processing of Performance Toolkit to produce VM history files and of VM monitor data captured by the MONWRITE utility*
- *Viewing of performance monitor data using either Web browsers or PC-based 3270 emulator graphics*
- *Monitoring TCP/IP for VM*
- *Processing of Linux performance data obtained from the Resource Management Facility (RMF™) Linux performance gatherer (rmfpms). Linux performance data obtained from RMF can be viewed and printed similar to the way that VM data is viewed and presented*

Enhancements to the Performance Toolkit for VM feature in z/VM V5.4 include:

- *Display an optional "banner" page that can be customized and is presented prior to displaying the Performance Toolkit logon page that provides access to the Web interface*
- *Provide updated displays and reports to support new monitor data for the dynamic memory upgrade enhancement*
- *Enable the SYSTEM (FCX102) and PRIVOPS (FCX104) reports to be generated using monitor data, allowing them to be produced in BATCH and MONSCAN modes*

Previous releases of z/VM V5 provided Performance Toolkit enhancements included:

- *Supporting passphrases when accessing the Performance Toolkit using the Web interface*
- *Changing the service process for the Performance Toolkit from a full-part replacement MODULE to service by individual object parts, reducing the size of the service deliverable*
- *Correct display of the z10 EC and z9 processor models with the PTF for APAR VM64369 for z/VM V5.3 and V5.2*
- *High-level Linux reports based on application monitor records from Linux*
- *New reports for SCSI disks, updated control blocks and new monitor data*
- *Capability to not have to shut down the Performance Toolkit server and restart it when adding new VM systems within the enterprise from which to retrieve performance data*

Directory and security management

Directory Maintenance Facility (DirMaint)

IBM DirMaint™ FL540 is an optional, priced feature of z/VM V5.4 and can be licensed for IFL processors. DirMaint is designed to provide efficient and highly-secure interactive facilities for maintaining the VM system directory. Directory management is simplified by the DirMaint command interface and automated facilities. DirMaint provides a command corresponding to every VM directory statement, including those for Cross System Extensions (CSE) clusters. DirMaint error checking validates directory changes and permits only authorized personnel to make changes.

DirMaint FL540 includes support for:

- *End-user authentication using the end-user's password phrase, if one is defined in an external security manager (ESM)*
- *An end-user to use DirMaint to set a password phrase in the ESM*
- *Simplified DirMaint configuration, enabling user and resource management to be coordinated with the RACF Security Server*

DirMaint supports the Systems Management APIs. A directory manager exit routine can be supplied by another directory manager to invoke its underlying functions in place of those provided by DirMaint.

On z/VM systems with large user directories, changes made to the user directory should be faster in z/VM V5.2 than in previous releases because a change to be made without requiring reprocessing of the entire directory results in less processing time to make the change. Performance improvements depend on the type of directory changes being made and the size of the VM directory being changed.

z/VM V5.2 integrated the directory management functions of DirMaint with the security management functions of RACF. DirMaint can be configured to automatically notify RACF whenever important changes are made to user definitions and the resources they own. This configuration reduces the administrative effort and skills needed to deploy and manage users and their resources when DirMaint and RACF are used together. By eliminating the need to manually define and manage z/VM resources in RACF, the possibility of incomplete or incorrect RACF configuration is reduced. Functions that are coordinated by this DirMaint support include:

- *User creation, deletion and changes*
- *Password management*
- *POSIX segment management*
- *Access Control Interface (ACI) group management*
- *Profile creation and deletion for selected VM functions*

RACF on z/VM

RACF is an optional, priced feature of z/VM V5 and can operate on standard and IFL processor configurations. RACF works with system features of z/VM to help provide improved data security for an installation. RACF is designed to help meet the need for security by providing:

- *Flexible control of access to protected resources*
- *Protection of installation-defined resources*
- *Ability to store information for other products*
- *Choice of centralized or decentralized control of profiles*
- *Transparency to end users*

With z/VM V5.3, the stand-alone RACF for VM V1.10.0 (5740-XXH) product was repackaged with all service applied and was renamed the RACF Security Server for z/VM. It is delivered as a release-specific priced, optional feature and is the base for all future RACF enhancements on z/VM. It works with the existing functions and features of z/VM to provide improved discretionary and mandatory access controls, separation of duties, and auditability capabilities of z/VM. The RACF Security Server feature includes support for mixed-case passwords and password phrases.

Working together, z/VM V5.3, and later, and the RACF Security Server for z/VM FL530 and later, feature support passwords that are longer than eight characters, called password phrases. A password phrase may contain mixed-case letters, numbers, blanks and special characters, allowing for an exponentially greater number of possible combinations of characters than traditional passwords. To ease migration from passwords to password phrases, the RACF Security Server for z/VM continues to support traditional eight-character passwords. An installation exit is also provided to help enable customers define rules governing the length and content of password phrases.

z/VM V5.3 introduced new user authentication, authorization and auditing capabilities with the inclusion of an LDAP server and associated client utilities. The z/VM LDAP server has been adapted from the IBM Tivoli® Directory Server for z/OS, to be delivered in z/OS V1.8. The z/VM LDAP server is designed to provide:

- *Multiple concurrent database instances (referred to as backends)*
- *Interoperability with LDAP V2 or V3 protocol-capable clients*
- *LDAP Version 2 and Version 3 protocol support*
- *Native authentication using Challenge-Response Authentication Method (CRAM-MD5), DIGEST-MD5 authentication and Simple (unencrypted) authentication*
- *Root DSE information master/slave and peer-to-peer replication*
- *Ability to refer clients to additional directory servers*
- *Capability to create an alias entry in the directory to point to another entry in the directory*
- *Access controls on directory information*
- *Change logging*
- *Schema publication and update*
- *SSL communication (SSL V3 and TLS V1)*
- *Client and server authentication using SSL/TLS*

Executing in a CMS virtual machine, LDAP is integrated into the base of z/VM V5.3 as a subcomponent of TCP/IP.

LDAP upgrade and RACF password change logging

In order to help maintain cross-platform consistency, the z/VM V5.4 LDAP server introduced in z/VM V5.3 has been upgraded in V5.4 to the function level of the z/OS V1.10 IBM Tivoli Directory Server. The z/VM RACF Security Server FL540 feature has been enhanced to create LDAP change log entries in response to updates to RACF group and user profiles, including changes to user passwords and password phrases. This update enables password changes made on z/VM to be more securely propagated to other systems, including z/OS, using applications such as the IBM Tivoli Directory Integrator.

Common criteria certification

IBM has received certification of z/VM V5.3 with the RACF Security Server optional feature from the

German Federal Office of Information Security (Bundesamt für Sicherheit in der Informationstechnik [BSI]) for conformance to the Controlled Access Protection Profile (CAPP) and Labeled Security Protection Profile (LSPP) of the Common Criteria standard for IT security, ISO/IEC 15408, at Evaluation Assurance Level 4, augmented by flaw remediation procedures (EAL4+). z/VM V5.4 has not been evaluated for conformance, but is designed to meet the same standards.

Tape encryption

z/VM is designed to support drive-based data encryption with the IBM System Storage TS1120 Tape Drive (machine type 3592, model E05) and the IBM System Storage TS1130 Tape Drive (machine type 3592, model E06 with the PTF for APAR VM64459 for V5.3 and V5.4) to help protect data on tape in a cost-effective way. Encryption of tapes by z/VM itself requires that the IBM Encryption Key Manager be running on another operating system, using an out-of-band (such as TCP/IP) connection to the tape control unit. z/VM native support includes encryption for DDR and SPXTAPE, as well as for guests that do not provide their own encryption enablement (for example, CMS and Linux for System z). z/VM enables encryption of tapes by guests (such as z/OS) that have the ability to control the tape-encryption facilities themselves and to optionally run the Encryption Key Manager. Encryption Re-Key support provides the capability to update a previously encrypted tape cartridge with a new set of Key Encryption information with the PTF for APAR VM64260. This allows for continuous protection of tape cartridge data, even as the encryption certificates that were used to create them are changed or replaced.

DFSMS/VM FL221 with the PTFs for APARs VM64062 and VM64458 for V5.3 and V5.4 supports locating encryption-capable 3592 tape drives in an Enterprise Automated Tape Library for a z/VSE guest running on z/VM.

Conformance with industry standards

z/VM V5.3 and later provide Secure Sockets Layer/Transport Layer Security (SSL/TLS) support for industry-standard secure FTP (RFC 4217), TELNET (draft specification #6) and SMTP (RFC 3207) sessions. This support includes new socket APIs to permit a Pascal or Assembler client or server application to control the acceptance and establishment of TCP sessions that were encrypted with SSL/TLS. Data transmission on a connection can now begin in clear text and at some later point be made available in secure text, thus reducing the need to dedicate a separate port for secure connections.

Networking with z/VM

z/VM network LAN simulation supports networking communication among virtual machines without the need for real IQD or OSD channels, much as z/VM simulates channel-to-channel adapters for communication among virtual machines without the need for ESCON, FICON or other real channel-to-channel connections.

Guest LANs

Guest LANs can be defined to function as OSA-Express QDIO transport media, in addition to HiperSockets transport media. Guest LANs can be defined to function as OSA-Express QDIO transport media supporting Layer 2 mode on z/VM. In this mode, each guest on the guest LAN is referenced by its Media Access Control (MAC) address and data is transported and delivered in Ethernet frames.

Virtual machines (z/VM and other guest operating systems) in the z/VM guest LAN environment can define and use simulated OSA-Express devices that support both the IPv4 and IPv6 protocols. IPv6 support allows the z/VM TCP/IP stack to be configured for IPv6 networks connected through OSA-Express adapters operating in QDIO mode. The stack can be configured to provide static routing of IPv6 packets and to send IPv6 router advertisements. In addition, support is provided to help application developers create socket applications for IPv6-based communications.

Virtual Switch (VSWITCH)

z/VM provides the capability to deploy virtual IP switches in the guest LAN environment. The z/VM virtual switch eliminates the need for virtual machines acting as routers to provide IPv4 connectivity to a physical LAN through an OSA-Express adapter. Virtual routers consume valuable processor cycles and require additional copying of the data being transported. The virtual switch can help alleviate these problems and also provides centralized network configuration and control. These controls allow the z/VM guest LAN administrator to more easily grant and revoke access to the network and to manage the configuration of VLAN segments. VSWITCH provides enhanced failover support for less disruptive recovery after some common network failures, helping to improve business continuity and infrastructure reliability and availability.

The virtual switch support also provides a transport option to define it as operating in ETHERNET Layer 2 mode. In this mode, each port on the virtual switch is referenced by its Media Access Control (MAC) address instead of by an Internet Protocol (IP) address. Datagrams are transported and delivered in Ethernet frames, providing the ability to send and receive protocol-independent traffic for both IP and non-IP applications for both IPv4 and IPv6 networks.
VSWITCH support for IEEE 802.3ad link aggregation and failover support is designed to allow OSA-Express2 or OSA-Express3 ports that are associated with a virtual switch to be grouped and used as a single “fat pipe.” This helps increase bandwidth and provides more seamless failover in the event of a link failure and requires associated OSA-Express2 support on the IBM z9 EC and z9 BC or OSA-Express3 on a z10 EC and z10 BC. In z/VM V5.4, a z/VM TCP/IP stack connected to a VSWITCH operating in Layer 2 mode is now supported in Link Aggregation configurations.

z/VM Network Virtualization

z/VM provides authorization capabilities for z/VM guest LANs and virtual switches by using RACF or any equivalent External Security Manager (ESM) that supports this new authorization function. It is designed to provide ESM-centralized control of authorizations and Virtual Local Area Network (VLAN) assignment.

z/VM provides improved problem determination for a z/VM guest LAN or a virtual switch (VSWITCH) by virtualizing a LAN sniffer to capture network traffic. This capability can help an administrator (or owner of the guest virtual machine) capture network data to help resolve virtual networking

problems. Procedures are provided to capture and process the data for both Linux and traditional VM environments:

- *Native Linux tracing capability on a guest LAN or VSWITCH*
  - *When a Linux guest is deployed, traffic can be traced, recorded and analyzed by existing tools directly from the guest virtual machine. This Linux guest must be authorized to use this capability through CP commands. The authorized guest can then use CP commands or the Linux device driver (when available) to put the guest NIC in "Promiscuous Mode."*
- *Native z/VM tracing capability on a guest LAN or VSWITCH*
  - *LAN traffic can be traced, recorded and analyzed using native z/VM facilities. This function is only authorized to users with Class C privileges.*

z/VM exploits IEEE VLAN technology to help ease the administration of logical groups of users so that they can communicate as if they were on the same physical LAN. VLANs help increase traffic flow and may help reduce overhead, allowing the organization of networks by traffic patterns rather than by physical location. To support VLANs, z/VM provides:

- *Functions to enable membership in a VLAN for OSA-Express adapters (in QDIO mode) and HiperSockets adapters that support IEEE 802.1q*
- *Virtual QDIO and HiperSockets network interfaces support VLAN frame tagging as described in IEEE 802.1q*
- *Management and control of the VLAN identifiers that can be used by guest virtual machines*
- *Simplified networking administration and management of VLANs with support for Generic Attribute Registration Protocol (GARP) VLAN Registration Protocol (GVRP) using OSA-Express2 or OSA-Express3 on z/VM*

z/VM V5.3 and later support port isolation security that provides the ability to restrict guest-to-guest communications within a VSWITCH and OSA-Express QDIO data connection isolation with the PTFs for APARs VM64463 and PK67610 for V5.3 and V5.4. The System z QDIO data connection isolation function is exclusive to z/VM and is designed to adhere to security guidelines, including the Health Insurance Portability and Accountability Act (HIPAA) by providing the ability to isolate QDIO data connection communications on a shared OSA port in support of multi-tier security zones. QDIO data connection isolation is available for OSA-Express2 and OSA-Express3 features on a System z10 EC and z10 BC and OSA-Express2 features on a System z9 EC and z9 BC with required minimum MCLs.

z/VM recognizes all four ports on an OSA-Express3 Gigabit Ethernet (GbE) and 1000BASE-T Ethernet feature and two ports on the GbE and 1000BASE-T 2P features on the z10 BC. There are

two PCI Express (PCI-E) adapters per feature, with two ports per adapter, for a total of four ports per feature or two per 2P feature. Activating all four ports (or two on the 2P) on an OSA-Express3 feature provides more physical connectivity to service the network and reduces the number of required resources (I/O slots, I/O cages, fewer CHPIDs to define and manage). With double the port density, reduced latency, and improved throughput, OSA-Express3 is designed to satisfy the demands of bandwidth-hungry applications. Four-port and two-port exploitation is supported by z/VM V5.4 and with the PTFs for APARs VM64277 and PK50120 for z/VM V5.2 and V5.3. If the PTFs are not applied, only two of the four ports (or one of the two) will be "visible" to z/VM V5.2 and V5.3.

The usability of managing virtual networks, including the z/VM virtual switch, has been enhanced in V5.4. Specifically, the QUERY LAN and QUERY VSWITCH commands have new USERID and RDEV operands that can provide a subset of the output. In addition, detailed transmission counters can be turned on and queried for a VLAN-aware VSWITCH.

The z10 EC, z10 BC, z9 EC, z9 BC, z990 and z890 servers are designed to include:

- *Virtualized adapter interruptions: This function can be used with V=V (pageable) guests. With TCP/IP stack enhancements, adapter interruptions can be used for OSA-Express channels and TCP/IP for VM can benefit from this performance assist for both HiperSockets and OSA-Express adapters. z/VM provides support for enhanced performance assists to allow adapter interruptions to be passed directly to z/VM guests for OSA-Express, FCP and HiperSockets operating on a z10 EC, z10 BC, z9 EC, z9 BC, z990 or z890. These assists include:*
  - *QDIO Enhanced Buffer-State Management (QEBSM)–two new hardware instructions designed to help eliminate the overhead of VM-Hypervisor interception for cooperating guest operating systems that initiate QDIO operations.*
  - *Host Page-Management Assist (HPMA)–an interface to the z/VM paging-storage management function designed to allow page frames to be assigned, locked and unlocked without z/VM Hypervisor assistance, primarily benefiting the QEBSM environment.*
- *TCP/IP stack improvements for OSA-Express increases the number of TCP/IP stacks that can share an OSA-Express (from 84 to 160) and is supported by z/VM to connect more virtual machines to an external network.*
- *Support for more TCP/IP stacks with OSA-Express2 and OSA-Express3 is supported by z/VM to help enable the number of connections (TCP/IP stacks) to be increased up to 640. This new capability allows additional connections to virtual machines, particularly Linux images.*

TCP/IP for z/VM

TCP/IP for z/VM with your System z server can communicate and share data with multi-vendor systems via your intranet and the Internet. Applications can be shared transparently across z/VM, z/OS, UNIX and other environments. TCP/IP can be characterized as providing functions and services that can be categorized as follows:

- *Connectivity and gateway functions that handle the physical interfaces and routing of data*
- *Server functions that provide a service to a client (for example, sending or transferring a file)*
- *Client functions that request a certain service from a server anywhere in the network*
- *Network status and management functions that detect and solve network problems*
- *Application Programming Interfaces (APIs) that allow you to write your own client/server applications*

TCP/IP is used to build interconnections between networks (including the Internet) through universal communication services. To allow communication among networks, addresses are assigned to each host with a network connection.

z/VM-based TCP/IP servers and clients can exploit Gigabit Ethernet, 1000BASE-T Ethernet, Fast Ethernet, Token-Ring and ATM networks through the OSA-Express Adapter using QDIO. QDIO can help improve performance through a highly efficient data transfer architecture that can reduce TCP/IP path lengths. Data can be directly exchanged with an I/O device without using traditional I/O instructions. Using QDIO can help an application achieve the full performance potential of a high-speed network.

Also available in z/VM V5.4, the Queued Direct Input/Output (QDIO) device driver of the TCP/IP stack can now operate in Layer 2 (of the Open Systems Interface (OSI) reference model) mode. The stack now deploys its own unique locally defined Media Access Control (MAC) addresses instead of sharing the OSA MAC address. It sends and receives Ethernet frames when transferring data through the OSA feature. Layer 2 support can be employed with an OSA-Express, OSA-Express2, or OSA-Express3 feature when configured in Queued Direct I/O (QDIO) mode (CHPID type OSD). This new support allows connecting a z/VM TCP/IP stack to an IPv4 or IPv6 Layer 2 Virtual Switch (VSWITCH). You can also connect a z/VM TCP/IP stack to a Layer 2 guest LAN or connect a z/VM TCP/IP stack (via a real OSA-Express) to a physical LAN segment in Layer 2 mode. By default, TCP/IP now automatically restarts OSD and HiperSockets devices without requiring operator, network, or system-programmer intervention. This can help provide faster recovery of connectivity to the system.

TCP/IP is designed to support the z/Architecture HiperSockets function for high-speed communication among virtual machines and logical partitions within the same IBM mainframe. The HiperSockets function allows virtual machines and logical partitions to communicate internally over the memory bus using the internal-queued-direct (IQD) channel type in System z servers.

z/VM provides TCP/IP and guest LAN support for HiperSockets using IPv6 protocol.

TCP/IP for z/VM includes support for File Transfer Protocol (FTP) and Trivial File Transfer Protocol (TFTP). FTP and TFTP clients running on z/VM or other systems can access files residing anywhere on the Internet. z/VM provides FTP support for access to the VM Shared File System (SFS), Byte File System (BFS) and CMS minidisk file system, as well as TFTP support for the BFS.

The multi-protocol dynamic routing server (MPRoute) implements Open Shortest Path First (OSPF) and Routing Information Protocol (RIP), providing a powerful alternative to TCP/IP static routing. When properly configured, a z/VM host running the MPRoute server can become an active OSPF or RIP network router, providing network access to z/VM virtual networks. The MPRoute server in z/VM V5.3 and later has been adapted from z/OS V1.8 and supports the following protocols:

- *For IPv4, MPRoute implements the OSPF protocol described in RFC 1583 (OSPF Version 2) and the RIP protocols described in RFC 1058 (RIP Version 1) and in RFC 1723 (RIP Version 2)*
- *For IPv6, MPRoute implements the IPv6 OSPF protocol described in RFC 2740 (OSPF for IPv6) and the IPv6 RIP protocol described in RFC 2080 (RIPng for IPv6)*

The MPRoute server removes the limit of four equal-cost paths, allowing the generation of up to 16 equal-cost routes to a destination, thus providing improved load-balancing support. MPRoute is the only dynamic routing server supported by TCP/IP for z/VM FL530 and later. The ROUTED and BOOTP servers have been removed from z/VM V5.3.

TCP/IP for z/VM V5.4 dynamically discovers the Maximum Transmission Unit (MTU) size of a given IPv4 or IPv6 Internet/intranet path. This helps to reduce fragmentation along the path, thus minimizing wasted network resources and helping to improve throughput by sending datagrams at the path MTU size. This function is automatically enabled for all IPv6 links and optionally can be enabled for IPv4 links.

Virtual IP Addressing (VIPA) can increase the reliability and availability of TCP/IP in the event of a network or interface failure. With VIPA, hardware link fault tolerance is supplied for both inbound and outbound TCP/IP communications on z/VM, which can provide automatic recovery of hard link failures and network traffic splitting. Virtual IP addressing support in the TCP/IP stack has been extended in z/VM V5.3 to support IPv6 addresses. It is now

possible to enable and configure a virtual device for IPv6 and to associate real IPv6-capable network adapters with a specific IPv6 virtual link for determining the source address used in outgoing packets. VIPA support is designed to improve the capability of the TCP/IP stack to maintain connections in the event that a real network device fails.

Failover support for IPv4 and IPv6 devices has also been enhanced in z/VM V5.3. When the z/VM TCP/IP stack has two (or more) QDIO or LCS Ethernet devices on the same network and one device is stopped or fails, another device is designed to take over responsibility for traffic destined for the failing device (or any devices the failing device had previously taken over). This failover support includes OSA-Express devices (in QDIO Ethernet or LCS Ethernet mode), OSA-Express2 or OSA-Express3 devices (in Ethernet mode), Virtual IP Addresses (VIPAs) and addresses for which PROXYARP services are being provided through a takeover-eligible device.

TCP/IP for z/VM provides numerous self-protection functions. An SSL server is available to facilitate security-rich and private conversations between z/VM servers and external clients. With z/VM support for SSL, a VM server can communicate with a secure client without a change to the server itself. The SSL server supplied with z/VM supports 40-bit, 56-bit and 128-bit encryption/decryption services.

z/VM V5.4 now has a CMS-based SSL server with the PTFs for APARs PK65850, PK73085, PK75268, VM64540, VM64519, and VM64570 replacing the z/VM V5.3 Linux-based SSL server. Instead of an SSL server that requires a Linux distribution, a CMS-based SSL server may enable encryption services to be deployed more quickly and helps make installation, service, and release-to-release migration simpler. Other enhancements to the z/VM V5.4 SSL server include:

- *Network-free SSL server administration*
  The SSL server can be managed without requiring a network connection between the SSL server administrator and the SSL server.
- *New encryption and decryption engine*
  The SSL server uses z/OS V1.10 System SSL technology for encryption, decryption, and certificate management.
- *New certificate-management services*
  The System SSL gskkyman utility is now used to manage the SSL server certificate database. New services available for the SSL server include certificate renewal, certificate signing, and certificate exportation with or without the private key. The gskkyman application also manages certificates for the z/VM LDAP server.

Previous releases of z/VM enhanced the z/VM SSL server allow the host Linux guest system to remain active after a critical error is encountered during server operations. Also, the SSLADMIN command has been enhanced to:

- *Allow the specification of the number of days that a self-signed certificate is to be valid*
- *Improve the management of SSL server LOG files*

In order to enable enforcement of enterprise requirements for strong encryption on network connections (128 bits or higher), the z/VM SSL server was enhanced to more easily allow weak cipher suites to be excluded.

A configuration wizard, IPWIZARD, automates the connection of a newly installed z/VM system to a TCP/IP-based network. This easy-to-use tool can help the z/VM installer provide IP configuration information such as host and domain names, IP addresses and subnet masks. This tool also generates an initial z/VM TCP/IP configuration and verifies that connectivity to the network has been established.

Once the initial IP network configuration has been created, a dynamic TCP/IP configuration tool, IFCONFIG, is available that can eliminate the need to learn the statement syntax of the z/VM TCP/IP server configuration file. This additional tool can optionally generate configuration statements for incorporation into the configuration file so that the changes may be made permanent.

The Network File System (NFS) V3 server allows applications and users from heterogeneous systems to access files stored in the VM Byte File System (BFS), Shared File System (SFS) and CMS minidisk file system. NFS support on z/VM is a natural extension of the VM file systems and enables Internet-based heterogeneous systems to use the vast DASD resources available on z/VM. Additionally, NFS permits z/VM to be a centralized, transparent file server for PC servers and workstations.

The z/VM NFS client gives CMS users and applications transparent access to data on remote systems that run NFS servers, including z/OS, Microsoft® Windows®, IBM AIX®, UNIX, Linux and VM. Mounting remote data on the BFS structure in a single virtual machine allows access by an NFS client.

The Simple Mail Transfer Protocol (SMTP) server, which includes TCP/IP mail services, is integrated with CMS mail functions. This can deliver a consistent method of mail and file transfer for TCP/IP and CMS users. The SMTP server provides service extension support, including acceptance and forwarding of MIME-formatted messages.

The Internet Message Access Protocol (IMAP) Server provides support for an IMAP Version 4rev1 (IMAP4rev1) mail server that runs on z/VM. This support allows you to utilize the strengths of z/VM (reliability, availability and security) for storing and serving electronic mail while allowing any IMAP4rev1 client to access and manipulate mail messages using the IMAP protocol as defined by RFC 2060. In z/VM V4.4, an IMAP user authentication exit removed prior restrictions on user ID and password-lengths and helped eliminate the need for every IMAP client to have a VM user ID and password. Authentication is handled by a user-written exit routine, providing greater flexibility for choosing authentication methods.

Access to 3270-based applications from UNIX and other systems is available with the TELNET TN3270 support provided by TCP/IP for z/VM. The VM

SSL server, along with an SSL-enabled TELNET client such as IBM Personal Communications, can be used to enable the appropriate level of security and privacy of TELNET session data as it travels over the Internet or your intranet. The z/VM V5.4 TCP/IP stack provides an IPv6-capable TELNET server and client. Because TELNET is written in Pascal, IPv6-capable Pascal APIs are also provided.

Users or applications can execute a command on a remote host and receive results based upon TCP/IP remote execution protocol (REXEC) and support from z/VM.

TCP/IP for z/VM allows you to print data from your z/VM system on remote printers in your TCP/IP network. It also delivers enterprise-wide network printer support with line printer router (LPR), line printer daemon (LPD) and TN3270E printer attachment. VM LPR, LPD and TN3270E print support has been incorporated into the RSCS print server. You can specify whether you want remote print data to be processed for delivery by TCP/IP or RSCS.

z/VM provides network management support with Simple Network Management Protocol (SNMP). Message Queuing (MQ) is a popular method for applications to interface with one another across heterogeneous systems. MQ communication requires client API support on the communicating platforms and a message queue manager (MQ server) somewhere in the network. The MQ server facilitates communication between applications without requiring them to actually connect to one another. The IBM MQSeries® Client code is supplied with z/VM. Therefore VM-based applications can interact over the Internet with other IBM WebSphere® MQ and MQSeries enabled applications and servers.

In addition, a new SNMPTRAP utility can be used to generate SNMP TRAP messages, providing unsolicited data messages that are sent by an SNMP agent to its SNMP managing system. These messages are usually used to inform the managing station about special conditions that have occurred either in an agent system or in the network.

Statements of Direction for z/VM V5.4

- **TCP/IP functions:** *IBM intends to withdraw support in a future z/VM release for the Kerberos Authentication System.*
- **MMC (Mainframe to Micro Channel) card:** *IBM intends to withdraw support in a future z/VM release for the MMC card, which enabled communication between the PS/2 (PWSCS) and VM (PWSCF or ISFC).*

For more information

To learn more about z/VM V5.4, visit: **ibm.com**/vm/

To learn more about the IBM System z environment, contact your IBM marketing representative, IBM Business Partner or visit: **ibm.com**/systems/z/

Integrated Marketing Communications,
Server Group
Route 100
Somers, NY 10589
U.S.A.

Produced in the United States of America
April 2009

ZSD03003-USEN-10