

z/VM Version 5 Release 3



Flexible solutions for the competitive advantage

Highlights

- **Enhanced memory utilization supporting up to 256 GB real storage and up to 32 Processor Units (PUs) to help relieve workload constraints**
- **Comprehensive security with a new Lightweight Directory Access Protocol (LDAP) server and IBM RACF® feature, including support for password phrases (passphrases)**
- **Guest support enhancements, including an IBM z/OS® testing environment for the simulation and virtualization of zAAP and zIIP specialty engines (processors)**
- **Management enhancements for Linux® and other virtual images**
- **Support for selected features of the IBM System z10™ Enterprise Server (z10 EC)**
- **Delivery of RSCS FL530 as a priced, optional IPLA feature**
- **Systems management of IBM z/VM® guests from the Hardware Management Console (HMC)**

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 the capability for 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 V5.3 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 EC, System z9 Enterprise Class (z9 EC) and the System z9 Business Class (z9 BC). This revolutionary brand responds to unprecedented demands by providing high levels of performance and scalability.

The z10 EC offers huge capacity to enable large scale consolidation* that can drive up to:

- *80% reduction in energy consumption and costs*
- *85% reduction in floor space*
- *80% reduction in labor costs*

With increased capacity, the z10 EC virtualization capabilities can help to support more virtual servers than any of our competitors—thousands of x86 virtual servers in a single 2.83 square meters 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 to fewer and consolidated processors.

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*
- *Ability to virtualize each LPAR into hundreds or more virtual machines*
- *Ability to virtualize processor, communication, memory, storage, I/O and networking resources*
- *Help with 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 helps 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, 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/OS.e, z/VSE and z/TPF.

* Comparison is versus x86 Blade servers without virtualization, reflecting a current-day consolidation. Reductions will vary by the number and age of the x86 servers being consolidated.

z/VM V5.3 provides support for larger logical partitions (LPARs) to improve scalability and to facilitate growth. A single z/VM partition can now be configured with up to 256 GB of real storage (memory), twice the size supported by z/VM V5.2 and up to 32 Processor Units (PUs), a 33 percent increase over the previous release. z/VM V5.3 and Linux on System z collaborate to make more informed choices about how memory is managed. This level of cooperation can allow z/VM to run more virtual servers in the same amount of memory.

z/VM V5.3 provides an increased focus on security capabilities with the introduction of an LDAP server and associated client services for a more comprehensive security solution on z/VM. Security is also enhanced to support the use of password phrases in z/VM through a new RACF feature, more security-rich TCP/IP sessions and enhanced data protection by exploiting drive-based data encryption of the IBM System Storage™ TS1120 Tape Drive.

z/VM V5.3 extends its world-class virtualization technology by providing guest support for IBM System z Application Assist Processors (zAAPs) and

System z9 Integrated Information Processors and System z10 Integrated Information Processors (zIIPs), the Modified Indirect Data Address Word (MIDAW) facility and ASCII consoles. Manageability, reliability and usability of virtual networks have also been improved.

Engine-based Value Unit pricing for z/VM V5 replaces the per-processor pricing model applicable to z/VM V4. Engine-based Value Unit pricing is designed to provide a lower entry price and a decreasing price curve as hardware capacities and workloads grow, which may help improve price/performance. Engine-based Value Unit pricing is designed to help you:

- *Add capacity and workload with an incremental, lower price*
- *Manage software costs better*
- *Aggregate licenses acquired across machines that are part of your enterprise*

The z/VM V5 pricing model makes it more feasible to add z/VM virtualization technology to a standard processor environment (compared to the pricing models of z/VM V3 and V4). z/VM V5 requires z/Architecture (64-bit) and provides additional support and exploitation opportunities for the thousands of users who have built enterprise-wide

automation and infrastructure enhancements on the VM platform in support of their applications, database systems and on demand business solutions.

Enhanced scalability and constraint relief

z/VM V5.3 provides considerable enhancements to help enhance its scalability and provide storage (memory) constraint relief. These include:

- *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. This can benefit customers with large amounts of real memory and 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.*

¹The 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.

- *Improved memory management algorithms to help benefit paging workloads with large memory environments with the PTF for APAR VM64349. This enhancement may be more beneficial with the faster processor speeds of the System z10 EC.*
- *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 is designed to reduce the number of alias-device addresses needed for parallel I/O and transparently provide the potential benefits of HyperPAV volumes for minidisks owned or shared by guests that do not specifically exploit HyperPAV volumes, such as Linux and CMS.*
- *Enhanced IBM FlashCopy® support allows the specification of up to 12 target minidisks, determine the status of FlashCopy requests and exploit hardware asynchronous cache destage and discard that makes the FlashCopy appear synchronous to the virtual machine. In addition, z/VM has reduced the number of FlashCopy hardware related error conditions that can be reflected to the guest for the z/VM FlashCopy command response back to the guest.*

- *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. The SVC can be used to provide SCSI devices as emulated FBA devices for use by CP and guest operating systems by z/VM V5.3 and z/VM V5.2 with the PTF for APAR VM64128.*

Previous releases of z/VM V5 also provided support to help ease constraints that included:

- *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.*

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/OS.e, z/VM, z/VSE, TPF or z/TPF operating systems and applications running on System z standard processors. Only Linux workloads in an LPAR or Linux guests of z/VM V5 can operate on the IFL processors.

z/VM V5.3 provides virtualization technology enhancements in support of Linux and other guests, including:

- *New guest support for specialty processors, zAAPs and zIIPs, includes:*
 - *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 corresponding specialty processors of the same 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™ and IBM DB2®.*
- *Usability enhancements for the virtual switch (VSWITCH) and guest LAN environments including enhanced ease-of-use for Virtual LAN (VLAN) and promiscuous mode configuration changes and 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.*

Previous releases of z/VM V4 and V5 also provided virtualization technology enhancements that included:

- *SCSI disk support that allows a Linux server farm to be deployed on z/VM in a configuration that includes only SCSI disks, installation of z/VM from DVD to a SCSI disk, IPL from a SCSI disk using the VM Stand-Alone Program Loader (SAPL) and VM system dumps to a SCSI disk.*
- *Enhanced performance of FCP-attached SCSI disks for both system and guest use. This includes performance enhancements for:*
 - *QDIO efficiency*
 - *Paging/spooling optimization*
 - *FBA emulation efficiency*
- *Support for the OSA-Express2 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 a 3390 DASD.*

- An IBM HyperSwap™ function so that the virtual devices associated with one real disk can be swapped transparently to another. HyperSwap can be used to switch to secondary disk storage subsystems mirrored by Peer-to-Peer Remote Copy (PPRC).
- Dynamic virtual machine timeout 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 for guest operating systems. 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 or z/OS.e 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 such operating system environments as z/OS, z/OS.e, 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, z9 EC, z9 BC, z990 and z890 servers including:

- Guest exploitation of the System z10 EC at the level of System z9 functionality with the PTFs for APARs VM64180 and VM64242
- Exploitation of selected functions of the System z10 EC including:
 - Dynamic I/O configuration to define, modify and delete a Coupling using InfiniBand® link, CHPID type CIB, when z/VM V5.3 is the controlling LPAR for dynamic I/O
 - Processors dynamically added to or removed from a z/VM LPAR in reserve without preplanning with the PTFs for APARs VM64249, VM64323 and VM64389

- TCP/IP and VSWITCH gaining the performance benefit of OSA-Express3 10 GbE on the z10 EC
- Additional PTFs must be applied to support the z10 EC:
 - EREP support requires the PTF for APAR VM64367
 - CMS IOCP support requires the PTF for APAR VM64302
 - HCD support requires the PTF for APAR VM64020
- N_Port identifier virtualization (NPIV) support for FCP channels that is designed to allow the sharing of a single physical FCP channel among operating-system images, whether in logical partitions or virtual machines. This function offers improved FCP channel utilization and sharing among operating-system images, joining IBM ESCON® and FICON in offering channel-sharing through virtualization. This may help to reduce your hardware requirements and facilitate infrastructure simplification
- Provided 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 EC, z9 EC, z9 BC, z990 and z890 servers
- Extend dynamic-I/O configuration can allow channel paths, control units and devices to be dynamically added, changed and deleted in multiple Logical Channel SubSystem (LCSS) configurations and transparently share internal and external channels across LCSSs
- Handle 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-Express2 and OSA-Express3 Gigabit Ethernet (GbE)
- 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 EC, z9 EC and z9 BC can help increase channel capacity and performance
- z/OS and Linux for System z guest support for Crypto Express2
- Support 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 EC, z9 EC, z9 BC, z990 or z890
- Guest support for SIGP Conditional-Emergency-Signal and Sense-Running-Status orders and for Program-Event-Recording 3

- *CP exploitation of Program-Event-Recording 3, providing access to the guest breaking-event-address register to help aid in debugging of wild branches during virtual machine execution and the Store-Clock-Fast Facility to help reduce overhead of Store Clock instructions and CP program tracing*
- *Support for the DS6000 and DS8000 series in their native control unit modes. That is, the DS6000 will be supported as a 1750 control unit and the DS8000 as a 2107 control unit*
- *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 providing the potential benefit of PAVs for I/O issued to minidisks owned or shared by guests that don't exploit PAV's.*
- *Support for Dynamic Volume Expansion simplifying disk management by allowing for the dynamic increase of a DS8000 volume size in order 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 and is intended to help improve business efficiency, along with FlashCopy performance improvements that may help to help minimize operational disruption.*
- *Guest support for 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 1750 control unit to automatically switch the data path used to help improve overall performance on the DS6000*
- *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 3592 Tape Drive Models J1A and E05, 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*

Systems management

Enhancements in systems management, some of which help provide **self-configuring, self-managing** and **self-optimization** features in z/VM V5.3 including:

- *Providing enhancements to the z/VM Virtual Systems Management Application Programming Interface (API) for System z platform provisioning applications (such as IBM Director and programs developed by non-IBM solution providers) for ease of use in creating and managing large numbers of Linux and other virtual images running on z/VM. In z/VM V5.3, a new sockets-based server supports the API. The sockets-based server is multi-tasking capable and supports both AF_INET and AF_IUCV socket requests.*
- *Providing 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 any 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 for providing Simple Network Management Protocol (SNMP) data for virtual networking devices. A pre-configured subagent and exit routine are provided in z/VM V5.3 to supply bridge Management Information Base (BRIDGE-MIB) data, as documented in RFC 1493, for the z/VM Virtual Switch (VSWITCH). This subagent, through the use of a Network Management System client can acquire BRIDGE-MIB data for the z/VM virtual switch. In addition, this support provides a programming interface to obtain information about virtual networks.*
- *Repackaging Remote Spooling Communications Subsystem (RSCS) V3.2.0 (5684-096) and preinstalling on z/VM V5.3 system DDRs. RSCS Function Level 530 (FL530) is now available as a priced, optional feature and is available for both IFL and standard processor configurations. RSCS V3.2.0 is planned to be withdrawn from marketing on May 26, 2008.*
- *Enhancing the guest LOGON process by providing a new COMMAND statement in a virtual machine definition or profile to configure the virtual machine.*

A provision of the U.S. Government's Energy Policy Act of 2005 and similar legislation enacted by the governments of Canada and Bermuda extends Daylight Saving Time (DST) by four weeks, beginning in 2007. Starting in March 2007, Daylight Savings Time in the United States, Canada and Bermuda will begin on the second Sunday in March and end on the first Sunday in November. New sample system configuration file statements will be shipped with z/VM to support this change. The Language Environment PTF for APAR VM64117 must be applied to z/VM V5.1 and V5.2.

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

- *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 toolkit is an optional, priced feature derived from the FCON/ESA program (5788-LGA). 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 for 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*
- *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.*

The toolkit can monitor TCP/IP for z/VM and can also process Linux performance data. Previous releases of z/VM V5 included enhancements for new high-level Linux reports based on application monitor records from Linux, new reports for SCSI disks, updated control blocks and new monitor data and provided the capability to not have to shut down the Performance Toolkit server and restart it when adding new VM systems within the enterprise to retrieve performance data.

Enhancements to the Performance Toolkit for VM feature in z/VM V5.3 include:

- *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*
- *New or updated displays and reports to support new V5.3 functions*
- *Correct display of the z10 EC and z9 processor models with the PTF for APAR VM64369 for z/VM V5.3 and V5.2*

Directory and security management

The IBM Directory Maintenance Facility (DirMaint™) FL530 is an optional, priced feature of z/VM V5.3 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 FL530 supports the new Systems Management Application Programming Interfaces made available in V5.3. 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 new 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 as 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 has been repackaged with all service applied and is now called the RACF Security Server for z/VM FL530 feature. It is delivered as a release-specific priced, optional feature and operates only with z/VM V5.3. The new RACF Security Server feature includes support for mixed-case passwords and password phrases.

This new feature will be the base for all future RACF enhancements on z/VM and is expected to work 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 new RACF Security Server feature has also been updated to interoperate with the new z/VM LDAP server.

IBM is currently in evaluation for Common Criteria Certification of z/VM V5.3 with the RACF Security Server optional feature 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 (EAL4). IBM no longer intends to evaluate z/VM V5.2.

Working together, z/VM V5.3 and the RACF Security Server for z/VM FL530 feature support the use of 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 8-character passwords. An installation exit is also provided to help enable customers to define rules governing the length and content of password phrases.

z/VM V5.3 introduces 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 in the base of z/VM V5.3 as a subcomponent of TCP/IP.

z/VM is designed to support drive-based data encryption with the IBM System Storage TS1120 Tape Drive (machine type 3592, model E05) 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 for their own encryption (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 a 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 PTF for APAR VM64062 supports locating encryption-capable 3592 tape drives in an Enterprise Automated Tape Library. This DFSMS/VM support provides tape-encryption capabilities for a z/VSE guest running on z/VM.

z/VM V5.3 adds 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. In order to enable enforcement of enterprise requirements for strong encryption on network connections (128 bits or higher), the z/VM SSL server has been enhanced to more easily allow weak cipher suites to be excluded.

Networking with 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, z/OS.e, 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.

TCP/IP for z/VM can support tens of thousands of users and communicate with multi-vendor systems within your enterprise via your intranet or with external systems via the Internet.

TCP/IP for z/VM allows users to send messages, transfer files, share printers and access remote resources across a broad range of systems from multiple vendors.

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 broadcast support is provided for the HiperSockets environment when using Internet Protocol version 4 (IPv4). Applications that use the broadcast function can propagate frames to all TCP/IP applications.

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.

The z10 EC, 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, 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 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.*

z/VM exploits IEEE Virtual Local Area Network (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 VLAN, 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 for 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.*

The guest LAN support simulates IP Networking communication among virtual machines without the need for real IQD or OSD channels, much as 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 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.

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.

The virtual switch (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 provides a transport option to define it as operating in 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. A new port isolation security mechanism provides the ability to restrict guest-to-guest communications within a VSWITCH with the PTF for APAR VM64281.

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.

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-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.

TCP/IP and VSWITCH can also gain the performance benefit of OSA-Express3 10 GbE on the z10 EC.

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 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 V5.2 and later has been adapted from z/OS V1.7 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. The ROUTED and BOOTP servers have been removed from z/VM V5.3.

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-2 or OSA-3 devices (in Ethernet mode), Virtual IP Addresses (VIPAs) and addresses for which PROXYARP services are being provided through a takeover-eligible device.

IP Multicasting provides a more efficient means of transmitting the same data or messages to multiple users. A set of recipients can be selected and only one copy of the data is sent to the group. TCP/IP for z/VM supports multicasting in this manner, helping you save valuable network resources and users' time.

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 and requires a copy of Linux on System z to run. The z/VM V5.3 SSL server provides support for:

- *Novell® SUSE Linux Enterprise Server (SLES) 9 Service Pack 3 (64-bit)*
- *Novell SUSE Linux Enterprise Server (SLES) 9 Service Pack 3 (31-bit)*
- *Red Hat Enterprise Linux (RHEL) AS 4 Update 4 (64-bit)*
- *Red Hat Enterprise Linux (RHEL) AS 4 Update 4 (31-bit)*

The z/VM SSL server has been enhanced to 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*

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 VLAN assignment.

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, z/OS.e, 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.

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.

Statements of Direction for z/VM V5.3

- **Common Criteria Certification:**

IBM is currently in evaluation for Common Criteria Certification of z/VM V5.3 with the RACF Security Server optional feature 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 (EAL4). IBM no longer intends to evaluate z/VM V5.2.

- **RPC server support for the Systems**

Management API: IBM intends to withdraw support for the RPC/CSL interface from the Systems Management API server in a future z/VM release.

- **TCP/IP functions:** *IBM intends to withdraw support for the Network Database (NDB) system, Trivial File Transfer Protocol (TFTP) server, X25 interface (including the X25IBI server) and SNALINK server in a future z/VM release.*

- **3480 Distribution Medium:**

IBM intends to withdraw 3480 tape as a distribution medium in a future z/VM release. IBM plans to continue distributing z/VM on 3590 and 3592 tape and on DVD and to be available for electronic delivery from ShopzSeries.

- **z/VM LPAR enhancements:**

IBM intends to further enhance z/VM in a future release to exploit the new System z10 EC support for a new logical partition (LPAR) mode “VM,” exclusively for running z/VM LPARs. This new LPAR mode allows z/VM to utilize a wider variety of specialty processors in a single LPAR. For instance, in a VM mode LPAR, z/VM can manage Linux on System z guests running on IFL processors while also managing z/OS guests running on central processors (CPs) and zIIPs and zAAPs.

- **Additional support for managing z/VM systems:** *IBM intends to further enhance z/VM in a future release to exploit the new Hardware Management Console (HMC) interface that allows the installation of Linux on System z into a z/VM virtual machine. Additionally, future support is planned for z/VM and the HMC to provide z/VM hypervisor-configuration tasks.*

For more information

To learn more about z/VM V5.3, visit:

ibm.com/eserver/zseries/zvm/

To learn more about the IBM System z environment, contact your IBM marketing representative, IBM Business

Partner or visit: ibm.com/systems/z/



© Copyright IBM Corporation 2008

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

Produced in the United States of America
February 2008
All Rights Reserved

References in this publication to IBM products or services do not imply that IBM intends to make them available in every country in which IBM operates. Consult your local IBM business contact for information on the products, features and services available in your area.

IBM, z9, the IBM logo, the e-business logo, AIX, DB2, DFSMS/VM, DirMaint, DS4000, DS6000, DS8000, ESCON, eServer, FICON, FlashCopy, HiperSockets, HyperSwap, MQSeries, Parallel Sysplex, RACF, RMF, System z, System z9, System z10, System Storage, Tivoli, VSE/ESA, WebSphere, z/Architecture, z/OS, z/VM, z/VSE and zSeries are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both.

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

Microsoft and Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries or both.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries or both.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through The Open Group.

InfiniBand is a registered trademark of the InfiniBand Trade Association.

Other trademarks and registered trademarks are the properties of their respective companies.

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

IBM hardware products are manufactured from new parts or new and used parts. Regardless, our warranty terms apply. This equipment is subject to all applicable FCC rules and will comply with them upon delivery.

Information concerning non-IBM products was obtained from the suppliers of those products. Questions concerning those products should be directed to those suppliers.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of a specific Statement of General Direction.