

z/VM Version 5 Release 2



Flexible solutions for the competitive advantage.

Highlights

- **Enhanced scalability for memory-constrained environments**
- **Improved management for Linux® and other virtual images**
- **Improved performance of native SCSI disk I/O**
- **Exploitation of the System z9™ Enterprise Class (z9 EC), formerly the IBM System z9 109 (z9-109), and the IBM System z9 Business Class (z9 BC)**

Building successful virtual enterprises

Using virtualization technology as a foundation, z/VM® offers function and technology that take advantage of IBM Virtual Machine (VM) capabilities. z/VM Version 5 (V5) benefits from more than three decades of innovation and refinement, providing users with the ability to respond to rapidly changing market requirements more quickly and easily than by using additional physical servers. z/VM offers the ability to host a large number of Linux images on a single mainframe, and provides an operational environment that is designed for on demand computing: highly flexible, adaptable, and efficient.

IBM System z

IBM System z™ (System z9 and @server™ 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 two new models: the System z9 Enterprise Class, formerly the IBM System z9 109, and the System z9 Business Class. This revolutionary brand responds to unprecedented demands by providing high levels of performance and scalability.

The System z environment, with self-configuring and self-healing attributes, provides new functions and features to meet the challenges of on demand business. 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 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 HiperSockets™*
- *Advanced systems management, administration and accounting tools*

z/VM Version 5 (V5)

z/VM V5 offers new levels of price/performance, functional capabilities, and hardware exploitation that increase the attractiveness 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) engines.

Using Linux as a guest of z/VM allows you to run hundreds of Linux images on a single System z server. These Linux images can be deployed on standard processor engines or on IFL engines with z/VM V5. z/VM V5 operates on the IBM z9 EC and the IBM z9 BC, the IBM @server zSeries 990 (z990), zSeries 890 (z890), zSeries 900 (z900), zSeries 800 (z800), and such other servers as IBM may specify.

z/VM V5 offers an ideal platform for consolidating select UNIX®, and Linux workloads on a single System z server for hosting other IBM mainframe operating systems as guests, including Linux on System z, z/OS®, z/OS.e, VSE/ESA™, z/VSE™, TPF, and z/TPF, as well as z/VM itself.

Virtualization technology enables Linux and other guests

With z/VM and IFL engines, 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 engines is designed to run Linux workloads without increasing the IBM software charges for z/OS, z/OS.e, z/VM, VSE/ESA, z/VSE, TPF, or z/TPF operating systems and applications running on System z standard engines. Only

Linux workloads in an LPAR or Linux guests of z/VM V4 or V5 can operate on the IFL engines.

Engine-based Value Unit pricing for z/VM V5 replaces the per-engine 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 engine 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.

z/VM V5.2 provides virtualization technology enhancements in support of Linux and other guests, including:

- *Improved scalability with the control program (CP) now using memory locations above 2GB for a much broader set of functions. Guest I/O buffers for both standard I/O and Queued Direct Input/Output (QDIO), used for networking and Fibre Channel Protocol (FCP) devices, were moved below 2GB when an I/O operation was initiated. I/O can now be done using buffers anywhere in real memory, and QDIO structures can reside above 2GB, as can most CP control blocks. These improvements offer constraint relief for large-real-memory virtual-server environments that are memory-intensive.*
- *z/VM V5.2 provides improved memory management for Linux on System z with the PTF for APAR VM63856*
- *Small Computer System Interface SCSI disks supported are those within an IBM TotalStorage® subsystem when it is connected to a fibre-channel (FC) fabric via FCP channels of System z. SCSI disks can be used directly by a guest operating system when an FCP sub-channel is dedicated to a guest that provides its own SCSI device support such as Linux on System z. z/VM SCSI support allows a Linux server farm to be deployed on z/VM in a configuration that includes only SCSI disks. ECKD™ disks are no longer required. 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 are also supported. z/VM SCSI support is based on emulation of the IBM Fixed Block Architecture (FBA). z/VM supports FCP-attached SCSI disks with capacities of nearly 1TB (2,147,483,640 512-byte blocks) for CP volumes and 381 GB for CMS and GCS volumes.*
- *Improved 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. The OSA-Express2 Gigabit Ethernet and 1000BASE-T Ethernet features now have the capability to provide channel connectivity from System z operating systems to the IBM Communication Controller for Linux (CCL) on System z with the introduction of Open Systems Adapter for the Network Control Program (OSA NCP, CHPID-type OSN) supporting the Channel Data Link Control (CDLC) protocol. OSA-Express2 OSN can help to 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. OSA-Express OSN support is exclusive to the z9 EC and z9 BC.*

- *Support for the PCI Express Cryptographic Coprocessor (PCIXCC) and the Crypto Express2 features can help protect your assets. z/VM V5 provides z/OS and Linux on System z guest support for these z890 and z990 features and the Crypto Express2 feature of the z9 EC and z9 BC. z/VM support includes:*
 - *Dedicated-queue support for clear-key and secure-key cryptographic functions for z/OS guests*
 - *Shared-queue and dedicated-queue support for clear-key cryptographic functions for Linux guests.*
 - *On the z9 EC and z9 BC, the Crypto Express2 feature can be configured as a coprocessor or as an accelerator for Secure Sockets Layer (SSL) acceleration.*
- *Additional guest support enhancements provided in V4 or V5 are:*
 - *Installation of z/VM from a DVD to an IBM TotalStorage SCSI disk emulated as a FBA device or to a 3390 DASD.*
 - *A 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.*
 - *Improvements to the CP scheduler increase the number of Linux and other guest virtual machines that can be managed concurrently.*
 - *Enhancements to the Virtual Machine Resource Manager (VMRM) provide the infrastructure to support more extensive workload and systems resource management features. IPLing from SCSI disks attached to FCP channels for Linux and other guest operating systems that contain the required SCSI device support when z/VM is running in System z environments.*
 - *Virtual 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 helps 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 z9 EC, z9 BC, z990, and z890 servers including:

- *N_Port identifier virtualization (NPIV) support for FCP channels is designed to allow the sharing of a single physical FCP channel among operating-system images, whether in logical partitions or virtual machines. This new function offers improved FCP channel utilization and sharing among operating- system images, joining ESCON® and FICON in offering channel-sharing through virtualization. This may help to reduce your hardware requirements and facilitate infrastructure simplification. NPIV is exclusive to the z9 EC and z9 BC and is applicable to all FICON features supported on the z9 EC and z9 BC when configured as CHPID type FCP. z/VM V5.2, V5.1 and V4.4 enable guest use of NPIV when FCP subchannels are dedicated to a guest. In addition, z/VM V5.2 and V5.1, which support the use of SCSI FCP devices, provide for CP use of NPIV. z/VM V5.1 requires the PTF for APAR VM63744 for CP system use of NPIV. z/VM V5.1 cannot be installed from DVD to SCSI disks when NPIV is enabled. z/VM V5.2 can be installed from DVD to SCSI disks when NPIV is enabled. With the PTF for APAR VM63952 and when using NPIV, guest operating systems and z/VM users can query hardware-defined virtual port names.*
- *z/VM V5.2 is designed to provide 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 z9 EC, z9 BC, z990, and z890 servers.*
- *Extend dynamic-I/O configuration support to 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 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 and OSA-Express2 with z/VM V5.2 and V5.1 with the PTFs for APARs VM63538 and PQ97436*
- *Support for OSA-Express2 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 (4Gigabit/second) on the z9 EC and z9 BC can help increase channel capacity and performance
- 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 z9 EC, z9 BC, z990 or z890
- z/VM 5.2 supports selected functions of the z9 EC and z9 BC with the PTF for APAR VM63952, including:
 - 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 aid in debugging of wild branches during virtual machine execution, and the Store-Clock-Fast Facility to reduce overhead of Store Clock instructions and CP program tracing
 - CP TRACE support for new instructions available with the z9 EC and z9 BC.
- z/VM V5.2 supports the IBM TotalStorage DS6000 and DS8000 series emulated as 3990 control units. With the PTF for APAR VM63952, z/VM V5.2 supports 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 DS8000 and DS6000 supported by z/VM V5.2 include:
 - Parallel Access Volumes (PAVs) as minidisks for guest operating systems that exploit the PAV architecture with the PTF for APAR VM63952. In addition, the APAR will provide the potential benefit of PAVs for I/O issued to minidisks owned or shared by guests that don't exploit PAV's
 - 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.
 - Peer-to-Peer Remote Copy Extended Distance (PPRC-XD), designed to copy full volumes of data in non-synchronous mode. PPRC-XD. It is supported for guest operating systems on z/VM V4.4 and later. z/VM V4.4 and later versions are also designed to support PPRC Version 2 (V2) for guest operating systems.
- z/VM V5.2 also provides the following DS6000 and DS8000 support:
 - Capability to define and operate FCP-attached SCSI disks with capacities of nearly 1TB (2,147,483,640 512-byte blocks) for CP volumes and 381 GB for CMS and GCS volumes. PTFs for DirMaint APAR VM63700 and DFSMS/VM® APAR VM63664 are required to support the nearly 1TB CP SCSI disks on z/VM 5.1.

- *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 (55.7GB) 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*
- *Supports the IBM TotalStorage 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. The 3592 Model J70 can also be used to attach IBM 3590 tape drives. z/VM V5.2, including DFSMS/VM, also support Write Once Read Many (WORM) data cartridges. The PTFs for APARs VM63811 and VM63812 are required to support the 3592 E05 (TS1120) Tape Drive.*

Systems management

Improvements in systems management, some of which help provide **self-configuring, self-managing** and **self-optimization** features in z/VM V4.4, V5, and later include:

- *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.*
- *Systems management APIs provide functions that may be called by applications to allocate and manage resources for guests running in z/VM virtual machines (virtual images). All enhancements to the systems management APIs in z/VM V5.1 were implemented using Version 2 (V2) of the systems management server. A new server security identification procedure is provided for special authorized clients, which can remove the need for a password, simplifying the logon process. New and enhanced APIs have been implemented in z/VM V5.2 using a new version of the z/VM systems management server, Version 3 (V3). APIs provided in earlier releases of z/VM, implemented using V1 or V2 of the server, can also function with the V3 server.*

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 the RealTime Monitor (RTM) and Performance Reporting Facility (PRF) optional priced features thereby eliminating the need to use separate products (RTM and PRF) to help manage performance. 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. z/VM V5.1 enhancements included new high-level Linux reports based on application monitor records from Linux and new reports for SCSI disks. The Performance Toolkit has been updated in V5.2 to support updated control blocks and new monitor data. The Performance Toolkit server does not have to be shut down and restarted when adding new VM systems within the enterprise to retrieve performance-data with the PTF for APAR VM63952.

Directory and security management

The Directory Maintenance Facility (DirMaint™) is an optional, priced feature of z/VM V4 and V5 and can be licensed for IFL engines. 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.

On z/VM systems with large user directories, changes made to the user directory using DirMaint FL510 with the PTF for APAR VM63733 should be faster than on previous releases. This enhancement allows a change to be made without requiring reprocessing of the entire directory, resulting 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. The PTF for APAR VM63952 is required for the DirMaint feature to support improved memory management for Linux guests and PAVs.

RACF® for z/VM is available as an optional, priced feature of z/VM V4 and V5 and can be licensed for IFL engines. RACF for z/VM works with system features of z/VM to help provide improved data security for an installation. RACF for z/VM 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*

In z/VM V5.2, the directory management functions of DirMaint can be integrated 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*

IBM intends to evaluate z/VM V5.2 with the RACF for z/VM 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).

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. TCP/IP broadcast support is provided for the HiperSockets environment when using Internet Protocol version 4 (IPv4) with z/VM V4.4 and V5.1, and later. Applications that use the broadcast function can now propagate frames to all TCP/IP applications.

z/VM V4.4 introduced IPv6 support for guest LANs. The z/VM guest LAN support for OSA-Express simulation in QDIO mode was updated for IPv6. Virtual machines (z/VM and other guest operating systems) in the z/VM guest LAN environment are able to define and use simulated OSA-Express devices that support both the IPv4 and IPv6 protocols. z/VM V5.1 enhanced its IPv6 support by allowing 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.

z/VM V5.2 provides TCP/IP and guest LAN support for HiperSockets using IPv6 protocol with the PTF for APAR VM63952.

The 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 on z/VM V4.4 and V5.1. With the enhancement of the TCP/IP stack in z/VM V4.4 to use adapter interruptions for OSA-Express channels, TCP/IP for VM can benefit from this performance assist for both HiperSockets and OSA-Express adapters. z/VM V5.2 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 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 V5.2 and on V4.4 and V5.1 with PTFs for APARs PQ91421 and VM63524. This increase helps connect more virtual machines to an external network.
- *Support for more TCP/IP stacks with OSA-Express2* is supported by z/VM V5.2 and on V5.1 with PTFs for APARs PQ91421 and VM63524 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 V4.4, and later versions, exploit 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 on z/VM V5.2 with the PTF for APAR VM63952. This function is also supported on z/VM V5.1 with the PTFs for APARs VM63784 and PK08444.*

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 V5.2 and on V5.1 with the PTFs for APARs VM63538 and PQ97436. 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 V4.4 further enhanced its virtualization technology by providing 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 helps 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.

Improvements to the virtual switch in z/VM V5.1 provided enhanced failover support for less disruptive recovery after some common network failures, helping to improve business continuity and infrastructure reliability and availability.

With z/VM V5.2 and on V5.1 with the PTFs for APARs VM63538 and PQ97436, virtual switch support has been enhanced with a new 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.

z/VM V5.2 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 helps 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 or z/VM tracing capability on a guest LAN or VSWITCH 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 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. A new MPRoute server has been adapted from z/OS V1.7 to operate on z/VM V5.2 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 new 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. IBM recommends that MPRoute is the routing server for use with TCP/IP for z/VM and plans to remove the ROUTED and BOOTP servers from a future release of z/VM. z/VM V5.2 is planned to be the last release in which these servers will be available.

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.

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. A Secure Sockets Layer (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 upgraded SSL server, introduced in z/VM V4.4, is further updated in V5.2 to provide support for:

- *SUSE SLES8 Service Pack 3 (31-bit)*
- *SUSE SLES9 Service Pack 2 (31-bit)*
- *SUSE SLES9 Service Pack 2 (64-bit)*
- *Red Hat Enterprise Linux AS V3 (31-bit)*
- *Red Hat Enterprise Linux AS V3 (64-bit)*

This upgraded SSL server includes:

- *A variety of industry-standard encryption algorithms, including DES, triple-DES, RC2, and RC4, with keys up to 128 bits in length.*
- *Certificate activation/removal without server restart*
- *Federal Information Processing Standard (FIPS 140-2) operational mode support*

Security of the TCP/IP stack was improved to help withstand Denial of Service (DoS) attacks. The overall security and auditability of the TCP/IP for z/VM stack and the integrity of the z/VM system have been improved by providing better controls, monitoring, and defaults.

z/VM 5.1 enhanced the authorization capabilities for z/VM guest LANs and virtual switches by using Resource Access Control Facility (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 helps 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, OS/390, Microsoft Windows, 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, added in z/VM V4.3, 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 MQSeries® Client code is supplied with z/VM. Therefore VM-based applications can interact over the Internet with other WebSphere® MQ and MQSeries enabled applications and servers.

For more information

To learn more about z/VM V5.2, 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 2006

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

Produced in the United States of America
April 2006
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, IBM @server, the IBM logo, the e-business logo, AIX, DFSMS/VM, DirMaint, ECKD, Enterprise Storage Server, ESCON, FICON, FlashCopy, HiperSockets, HyperSwap, Hypervisor, MQSeries, OS/390, Parallel Sysplex, Performance Toolkit for VM, RACF, RMF, System z, System z9, S/390, TotalStorage, VM/ESA, VSE/ESA, VTAM, 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 trademark of Linus Torvalds in the United States, other countries, or both.

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

Red Hat and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

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

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.

ZSD01145-USEN-03