



z/VM Version 5 Release 2



Chuck Morse
morsec@us.ibm.com



Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

APPN*	GDPS*	POWERPC*	Virtual Image Facility
CICS*	Geographically Dispersed Parallel Sysplex	PR/SM	Virtualization Engine
DB2*	HiperSockets	Processor Resource/Systems Manager	VisualAge*
DFSMS/VS	HyperSwap	QM*	VM/ESA*
DFSMS/VM*	IBM*	RACF*	VSE/ESA
DirMaint	IBM eServer	Resource Link	VTAM*
Distributed Relational Database Architecture*	IBM logo*	RMF	WebSphere*
DRDA*	IBMlink	RS/6000*	z/Architecture
e-business logo*	Language Environment*	S/390*	z/OS*
ECKD	MQSeries*	S/390 Parallel Enterprise Server	z/VM*
Enterprise Storage Server*	Multiprise*	System 370	z/VSE
Enterprise Systems Architecture/390*	On demand business logo	System 390*	zSeries*
ESCON*	OS/390*	System z9	zSeries Entry License Charge
FICON*	Parallel Sysplex*	Tivoli*	
GDDM*	Performance Toolkit for VM	Tivoli Storage Manager	
	POWERS	TotalStorage*	

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.
Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries.
Linux is a trademark of Linux Torvalds in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation in the United States and other countries.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply. All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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



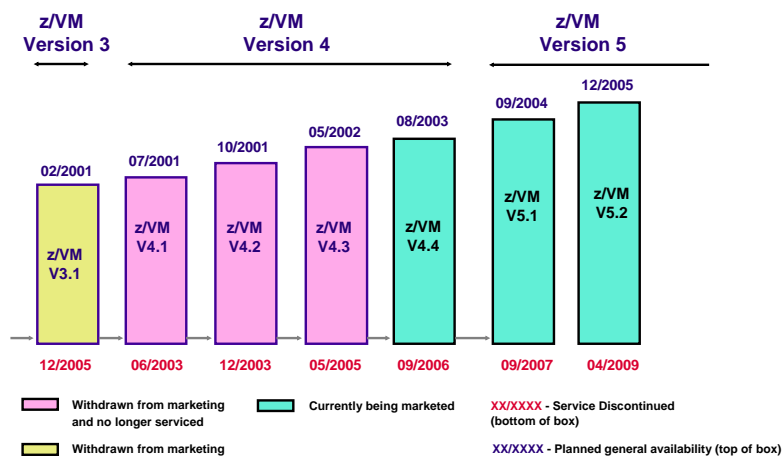
Agenda

- z/VM® Evolution
- Servers Supported by z/VM
- z/Architecture™ Support
- z/VM Version 5 Release 2 (V5.2) Enhancements
 - Virtualization Technology and Linux Enablement
 - Network Virtualization and Security
 - Technology Exploitation
 - Systems Management Improvements
- Optional Features of z/VM V5.2



z/VM Evolution

Expand your opportunities with z/VM



Servers Supported by z/VM Version 5 Release 2



System z9 109
(z9-109)

zSeries 990
(z990)

zSeries 890
(z890)

zSeries 900
(z900)

zSeries 800
(z800)

- Or Equivalent processors

z/VM Version 5 Release 2 Availability - 12/16/05

z/Architecture Supported by z/VM

- **Requires 64-bit capable processor**
- **Exploits central storage greater than 2 GB for guest and virtual disk (VDISK) pages and for minidisk cache**
- **Runs 64-bit capable guests with greater than 2 GB of virtual storage (z/OS, z/OS.e, z/VM, and Linux for System z9 and zSeries)**
- **Enhanced exploitation of large real memory providing 2 GB real-storage constraint relief**

z/VM Version 5 Terms and Conditions

■ International Program License Agreement (IPLA)

- Program Use License
 - One-time charge (OTC) for standard or IFL engines
 - Engine-based Value Units
 - Service by mail, fax, and e-mail only under basic warranty
- Subscription and Support (S&S)
 - Comparable service as traditional ICA products
 - Not required but highly recommended
 - Must decline when ordering if not desired
 - Annual renewable charge per engine-based value unit
 - Adds telephone support
 - No additional charges for updates, new versions and releases

■ SoftwareXcel available for an additional charge

■ IPLA applies to z/VM base code and the optional features

- DirMaint, RACF for z/VM, and the Performance Toolkit for VM


z/VM Version 5.2 Enhancements

■ Virtualization enhancements for Linux and other guests

- Enhanced exploitation of large real memory
- Improved performance of z/VM use of SCSI disk I/O

■ Networking virtualization and security enhancements

- Improved problem determination for guest LANs and virtual switches
- Enhanced dynamic routing capabilities with new MPRoute server
- SSL server support for additional Linux distributions



■ Technology exploitation

- Support for the System z9 109:
 - Up to 60 logical partitions
 - Enhanced performance assists for guests using OSA-Express, FCP, and HiperSockets
 - Support for OSA-Express2 OSA for NCP (OSN)
 - Crypto Express2 Accelerator for SSL acceleration
 - Improved FCP channel utilization and sharing
- Dynamic addition and deletion of logical partition names
- Support for the IBM TotalStorage DS4000 Midrange Disk systems

■ Systems management improvements

- Service and installation enhancements
- Improved management for Linux and other virtual images
- Simplified user administration with the coordination of DirMaint and RACF changes
- Improved DirMaint directory management performance
- Performance Toolkit for VM enhancements

Virtualization Technology and Linux Enablement

Enhanced Scalability Support for Memory-constrained Environments

- CP enhanced to improve performance and scalability for environments with high demand for storage below the 2 GB line
 - I/O can be done from buffers residing anywhere in memory
 - QDIO structures and most CP control blocks may now reside above the 2 GB line
- Storage above 2 GB included in dumps
 - CP hard/soft abend dumps and SNAPDUMPs
 - Stand-alone dumps of z/VM systems or VMDUMP of a virtual machine in z/Architecture mode
- TCP/ IP for z/VM exploitation of 64-bit Diagnose x'98'
 - Enhanced QDIO device driver to use I/O buffers above 2 GB when possible
 - Helps increase system reliability by reducing contention for storage below 2 GB
- Block I/O (Diagnose x'250')
 - Virtual machines can specify parameter addresses and I/O buffers with addresses above 2 GB
 - IBM is working with its Linux distribution partners to exploit this function in future Linux on System z9 and zSeries distributions or service updates
- **Potential benefit:**
 - Constraint relief for large-real-memory virtual-server environments that are memory-intensive

Improved Performance of SCSI Disk I/O



- QDIO efficiency improvements
 - Chaining I/O requests allows more data to be moved with a single I/O request
 - Decreases CP overhead
 - **Potential benefit** - significantly lowers the number of real SCSI I/O operations
- Paging and spooling optimization
 - No longer uses FBA emulation
 - SCSI commands are built and sent directly to the SCSI stack
 - **Potential benefit** - improves efficiency and performance
- Improved FBA emulation efficiency reducing the number of SCSI I/O requests
 - More efficient FBA emulation for reading CMS I/O buffers
 - **Potential benefit** - significantly improves SCSI I/O performance for CMS
 - More efficient FBA emulation of LOCATE CCWs
 - **Potential benefit** - significantly improve SCSI I/O performance for Linux guests, including the length of time to IPL a Linux guest
 - More efficient FBA emulation of FBA padding function
 - **Potential benefit** - significantly improve SCSI I/O performance for format functions

Network Virtualization and Security

Improved Problem Determination for Guest LANs and Virtual Switches

- LAN “sniffer” to capture network traffic on a z/VM guest LAN or virtual switch (VSWITCH) in both Linux and traditional environments:
 - Native Linux tracing capability on a guest LAN or VSWITCH
 - Traffic can be traced and analyzed by existing Linux tools, such as tcpdump
 - Linux guest must be authorized
 - The authorized guest can then use CP commands or the Linux device driver (when available) to put the guest NIC in "Promiscuous Mode"
 - IBM is working with its distribution partners to exploit this function in future Linux for System z9 and zSeries distributions or service updates
 - Native z/VM tracing capability on a guest LAN or VSWITCH
 - CP TRSOURCE command has been enhanced to trace and record the data transmissions, which can be analyzed with a new tool IPFORMAT
 - Only available to users with Class C privileges
- RACF for z/VM has been updated to control promiscuous mode authorizations
- **Potential benefit:**
 - Allows capture of network data to resolve virtual networking problems

Enhanced Dynamic Routing Capabilities With a New MPRoute Server

- OMPROUTE V1.7 module has been ported from z/OS to z/VM V5.2
- New MPRoute server includes support for:
 - IPv6 dynamic routing including RIPng and OSPF
 - IPv4 OSPF authentication using MD5 (cryptographic authentication)
 - Improved IPv4 VIPA support
 - Receiving RIPv1 and RIPv2 on same link
 - Up to 16 equal cost paths to a single destination
- RouteD server will be removed from a future release of z/VM
 - Utility (RTD2MPR EXEC) is supplied to assist in migration from RouteD to MPRoute
- **Potential benefits:**
 - Additional protocols are supported
 - Greater efficiency may be achieved within an IP network, and manual network routing table updates may be reduced or eliminated

SSL Support for Additional Linux Distributions

- Support will be provided 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)
- Upgraded SSL server also includes:
 - A variety of industry-standard encryption algorithms, including DES, triple-DES, RC2, and RC4, with keys up to 128 bits in length. Hashes are provided by SHA-1 and MD5
 - Certificate activation and removal without server restart.
 - Federal Information Processing Standard (FIPS 140-2) operational mode support
- **Potential benefits:**
 - Additional Red Hat Package Manager (RPM) packages supporting more distributions
 - Allows a new certificate to be added or an existing one to be deactivated while existing SSL-secured sessions are active, removing the need to shut down and restart the SSL server when certificates are added or deleted
 - SSL server may be operated in FIPS mode, restricting connections to those that employ FIPS-approved cipher suites

Technology Exploitation

Support for System z9 109

- z/VM V5.2 Only
 - Enhanced Performance Assists for z/VM Guests
- z/VM V5.2 and V5.1 only
 - Up to 24 Processors per VM System Image
 - OSA-Express2 Open Systems Adapter for NCP (OSN)
 - Crypto Express2
- z/VM V5.2, V5.1 and V4.4
 - Basic Processor Support
 - Up to 60 LPARs
 - N_Port ID Virtualization (NPIV)
- Compatibility support for currently supported z/VM releases requires PTFs for APARs
 - VM63646
 - VM63721
 - VM63743
 - VM63744
 - VM63722
 - OA11650
 - VM63740




Enhanced Performance Assists for z/VM Guests

- QDIO Enhanced Buffer-State Management (QEBSM)
 - Two new machine instructions reduce overhead of hypervisor interception
- Host Page-Management Assist (HPMA)
 - Interface to z/VM paging-storage management designed to allow machine to assign, lock, unlock page frames without hypervisor assistance
- Applicable to the following:
 - First-level guests of z/VM V5.2
 - HiperSockets (CHPID type IQD)
 - All OSA features (CHPID type OSD)
 - All FICON features (CHPID type FCP)
- Complements performance assists introduced in z/VM V4.4
- Supported by V5.2 on z9-109, z990, and z890
- **Potential benefit:**
 - Guest operating systems can initiate QDIO operations directly to channel without interception by z/VM

OSA-Express2 Open Systems Adapter for NCP (OSN)

- Designed to help eliminate dependencies on hardware
 - 3745/3746, ESCON, Token-Ring
- OSA-Express2 (Gigabit and 1000BASE-T Ethernet)
 - New CHPID type – OSN
- Provides support for IBM Communication Controller for Linux
- Appears to the operating systems as an ESCON-attached channel connected to a 3745 device type
- Traffic can flow LPAR-to-LPAR
- z/VM V5.2 enables VTAM® on z/VM and other guests to access OSA-Express CDLC channels
 - V5.1 requires PTFs for APARs VM63722 (CP) and VM63721 (HCD/HCM)
- **Potential benefits:**
 - Allows system administrators to configure, manage, and operate their CCL NCPs just like their real 374x NCPs
 - 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



Crypto Express2

- Can be configured as
 - Coprocessor:
 - Shared-queue and dedicated-queue support for clear-key cryptographic functions for Linux guests
 - Dedicated-queue support for clear-key and secure-key cryptographic functions for z/OS guests
 - Accelerator for SSL acceleration:
 - Designed to support clear-key RSA operations for Linux and z/OS guests
 - Offloads compute-intensive RSA public-key and private-key cryptographic operations employed in the SSL protocol
 - Supported by V5.2 and also V5.1 with the PTF for APAR VM63646

N_Port ID Virtualization (NPIV)

- Multiple operating system images can now concurrently access the same or different SAN attached devices (LUNs) via a single, shared FCP Channel
 - Can improve channel utilization
 - Less hardware required
 - Helps reduce complexity of physical I/O connectivity
- Supported by V5.2 and also V4.4 and V5.1 with the PTF for APAR VM63744 for CP use of NPIV
 - V5 cannot be installed from DVD to SCSI disks when NPIV is enabled and a future enhancement is planned to provide this capability in z/VM V5.2.
- IBM also intends to provide a future enhancement to z/VM V5.2 for NPIV such that guest operating systems and VM users can obtain virtual port names.
- IBM is working with its distribution partners to exploit this function in future Linux on System z9 and zSeries distributions or service updates.

Dynamic Addition and Deletion of LPAR Names

- V5.2 provides facilities to dynamically define and delete logical partitions using
 - CP's Dynamic I/O command interface
 - z/VM HCD/HCM support
- Supported on z9-109, z990, and z890
- **Potential benefit:**
 - Ability to add meaningful Logical Partition (LPAR) names to your configuration without a Power-On Reset

Support for DS4000 Midrange Disk Systems

- V5.2 SCSI support extended to support the lower-priced DS4000 (formerly FASiT) Midrange Disk Systems
- Added the Redundant Disk Array Controller (RDAC) driver providing multipath attachment with failover support
 - Designed to continue accessing storage on attached devices even if an element in the path (for example, an adapter, cable, or switch) fails



Systems Management Improvements

Enhanced Systems Management APIs

- z/VM V4.4 introduced a set of APIs to allocate/manage resources for guests
 - Functions are invoked through Remote Procedure Calls (RPC)
- V5.2 provides:
 - New APIs
 - Creating and updating of the LOADDEV directory statement in support of SCSI IPL
 - Obtaining the time when a virtual image was activated
 - Defining Local Tags by the Directory Manager
 - Scanning the directory for a specified pattern
 - Enhancements made to existing APIs:
 - Larger disk sizes (up to 2,147,483,640 blocks)
 - Capability to create persistent virtual switch definitions
 - New definition parameters for virtual switch APIs
- Requires a directory manager
 - DirMaint has been enhanced to support the new APIs
- **Potential benefit:**
 - Applications can be more easily written by solution providers to help administrators, especially those who lack in-depth VM knowledge, manage large numbers of virtual images running in a single z/VM system

Service and Installation Enhancements

- Further automation of the local modification process
 - Adds the capability to rework local modifications and provide support for local service
- Simplified migration of the preinstalled z/VM products
 - Allows the disks associated with the preinstalled products on your first-level z/VM V5.1 system to be made available to your second-level z/VM V5.2 system and then transfers the following types of files:
 - Customized files
 - Local modifications
 - Service
 - User-created files residing on selected disks

z/VM Version 5 Product Changes

- HCD/HCM upgraded to new level
- Withdrawal of CD-ROM installation media
- Publications
 - z/VM Collection Kit available on DVD (supplied with order)
- Functions Removed
 - System Administration Facility
 - Support for Server-Requester Programming Interface (SRPI)
 - Tivoli® Storage Manager
 - Tivoli Storage Manager™ V5.3 not supported on z/VM
 - Tivoli Storage Manager V5.3 or Tivoli Storage Manager Extended Edition V5.3 is recommended when running in a Linux guest environment
- Device support withdrawn
 - IBM 2741 and TWX Terminal Model 33/35 (TTY), or their equivalents, as virtual consoles
 - IBM 3705 Communication Controller
 - IBM 3720 Communication Controller
 - IBM 3725 Communication Controller
 - IBM 8232 LAN Channel Station

Performance Toolkit for VM FL520

- Licensed as an IPLA optional feature of z/VM V5
- Based on FCON/ESA
 - Replaces both PRF and RTM in V5.1
- New function level for V5.2
 - Supports new/changed CP Control blocks
 - New System Execution Space report
 - New System Storage reporting
 - Eliminate Detailed User Storage report
 - Expand QDIO report based on new Monitor Data
- Helps simplify performance analysis and resource management on your z/VM system by analyzing system monitor data and produces performance reports and history files, including:
 - System resource utilization, transaction response time, and throughput
 - Resource utilization by userID, DASD activity, and channel utilization

Directory Maintenance Facility (DirMaint) FL510

- Licensed as an IPLA optional feature of z/VM V5
 - Will only run on z/VM V5
- Provides interactive facilities for maintaining the z/VM user directory
- Support for z/VM V5.2 applied to the DirMaint feature on the V5.2 system DDRs and system image (DVD) and includes support for
 - All new/updated directory statements
 - Large Volumes (up to 56520 Cylinders)
 - New/updated Systems Management APIs
 - Faster Directory Updates for Large Directories
 - Updates can be made without re-writing the entire directory
 - Coordination with RACF for changes related to
 - User creation, deletion and changes
 - password management
 - POSIX segment management
 - ACI Group management

RACF for z/VM Feature

- Licensed as an IPLA optional feature of z/VM V5
 - Will only run on z/VM V5.2 or later
- RACF helps 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
 - Exits for installation-written routines
- Required support for guest LAN sniffer support is applied to the RACF feature supplied with the V5.2 system DDRs and system image (DVD)

Statements of Direction

- IBM intends to provide future enhancements to z/VM supporting the following System z9 functions:
 - System and guest exploitation of HiperSockets supporting the IPv6 protocol
 - Improved memory management between z/VM and Linux for System z9 and zSeries
 - Simplified networking administration and management of VLANs with support for GARP VLAN Registration Protocol (GVRP) using OSA-Express2
 - Capability to allow guest operating systems and z/VM users to query virtual port names when using N_Port ID Virtualization.
- 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).
- IBM intends to provide IBM Director support for Linux on System z9 and zSeries, extending IBM virtualization technology leadership with the exploitation of system management and virtual server deployment functions based on the Common Information Model (CIM) standard. IBM Director is a key component of the IBM Virtualization Engine.
- IBM intends to provide exploitation of the IBM TotalStorage Parallel Access Volume (PAV) feature for z/VM system data and guest data residing on VM minidisks in a future release of z/VM.
- IBM 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.