



# VM Device Support Overview

01/20/2005

Steve Wilkins  
[wilkinss@us.ibm.com](mailto:wilkinss@us.ibm.com)

# Agenda

- Overview / Levels of Support
- Dedicated OSAs
- QDIO Protocol
- Dedicated FCP
- IOASSIST and QIOASSIST
- Tape Support
- DASD Support
- MDC
- VDISK Support
- Emulated Device Support
- Guest Coupling
- Dynamic I/O
- LCSS Support
- Device Support Removed
- Trademarks



# Overview

- Various areas of device support provided by VM will be presented
- Both device support provided for guest operating systems and for managing and running the VM hypervisor (commonly known as 'native' support) will be covered; many times there are differences.
- This support information is intended for audiences not deeply familiar with VM device support; geared for those new to VM.
- Presentation is new; suggestions for improvement are encouraged!
- Acknowledgment goes to all those on the VM development team that provided input to this presentation; Eric Farman, Jim McGinniss, Lisa Reese, Bill Stephens, Patty Stump, Rick Tarcza, Brian Thompson, Steve Wilkins, and John Yacynych.

# Levels of Device Support

VM provides device support with varying degrees of involvement from the Control Program (CP)

## ■ Dedicated or Attached

- The guest has exclusive use of the entire real device
- CP involvement is minimal

## ■ Virtualized

- CP presents an image of a real device to multiple virtual machines, or
- CP subdivides a real device so that a portion appears real to the guest (e.g., minidisk)
- CP is more involved with a virtualized device. This allows services such as minidisk caching to be offered.

## ■ Emulated / Simulated

- CP emulates a piece of real hardware to the system and its guests as some other type of hardware, or
- CP completely simulates a device for a guest without the existence of real hardware
- Examples are emulated 9336 DASD, virtual CTCAs, Virtual Disks (VDISKs), Guest LANs, virtual Unit Record devices (reader, punch, printer), and virtual consoles.
- CP involvement is the highest with emulated and simulated devices

# VM Control Blocks for Device Support

- RDEV
  - CP control block representing a real device
  - RDEV can be attached or dedicated for exclusive guest use
  - RDEV can belong to the VM system to be virtualized (or subdivided) for guest use
  - RDEV can belong to the VM system for its own native use (e.g., paging volume)
  
- VDEV
  - CP control block that represents the guest's view of the device
  - VDEV appears to be a real device to the guest virtual machine
  - Behind VDEV is dedicated or virtualized RDEV (or simulated device with no RDEV)
  
- MDISK
  - Represents subdivision of a real disk (RDEV) and is paired with 1 or more VDEVs
  - Used to manage multiple guest links to the same subdivision (called a Minidisk)
  
- EDEV
  - An emulated 9336 model 20 DASD that corresponds to a real SCSI LUN (disk)
  - For native CP use (e.g., paging, spooling) or for virtualizing a SCSI LUN as a 9336 to a guest

# Dedicated OSAs

## OSA Express

- Provides IP ethernet connectivity to and from a physical LAN
- Data transmission is done by an OSA Adapter Card
- Requires an OSA Express CHPID (channel path) to be installed
- Each CHPID contains 240 or more (depending on hardware level) subchannels

## HiperSockets

- Provide IP ethernet connectivity only between Logical Partitions (LPARs) without a physical LAN or hardware adapter
- HiperSockets devices are defined by IOCP (or Dynamic I/O)
- Data transmission is done by hardware millicode

 OSA Express or HiperSockets subchannels are dedicated to a guest virtual machine for its sole use by the CP ATTACH command or DEDICATE user directory statement; used by guest running VM TCP/IP stack.

## Queued Direct I/O Protocol (QDIO)

OSA Express and HiperSockets subchannels use QDIO protocol rather than traditional z/Architecture Start Subchannel I/O for IP Datagram transmission

QDIO protocol provides:

- Memory mapped type of data transmission
- Direct communications between a guest TCP/IP stack and OSA adapter
- Data transmitted via set of QDIO queues allocated in guest storage
- Guest TCP/IP stack can allocate up to 4 output queues and 1 input queue

# IP QDIO Connection

Minimum of 3 QDIO capable subchannels (devices) are required to establish an IP QDIO connection by guest IP stack

Dynamically (Class B Privilege)

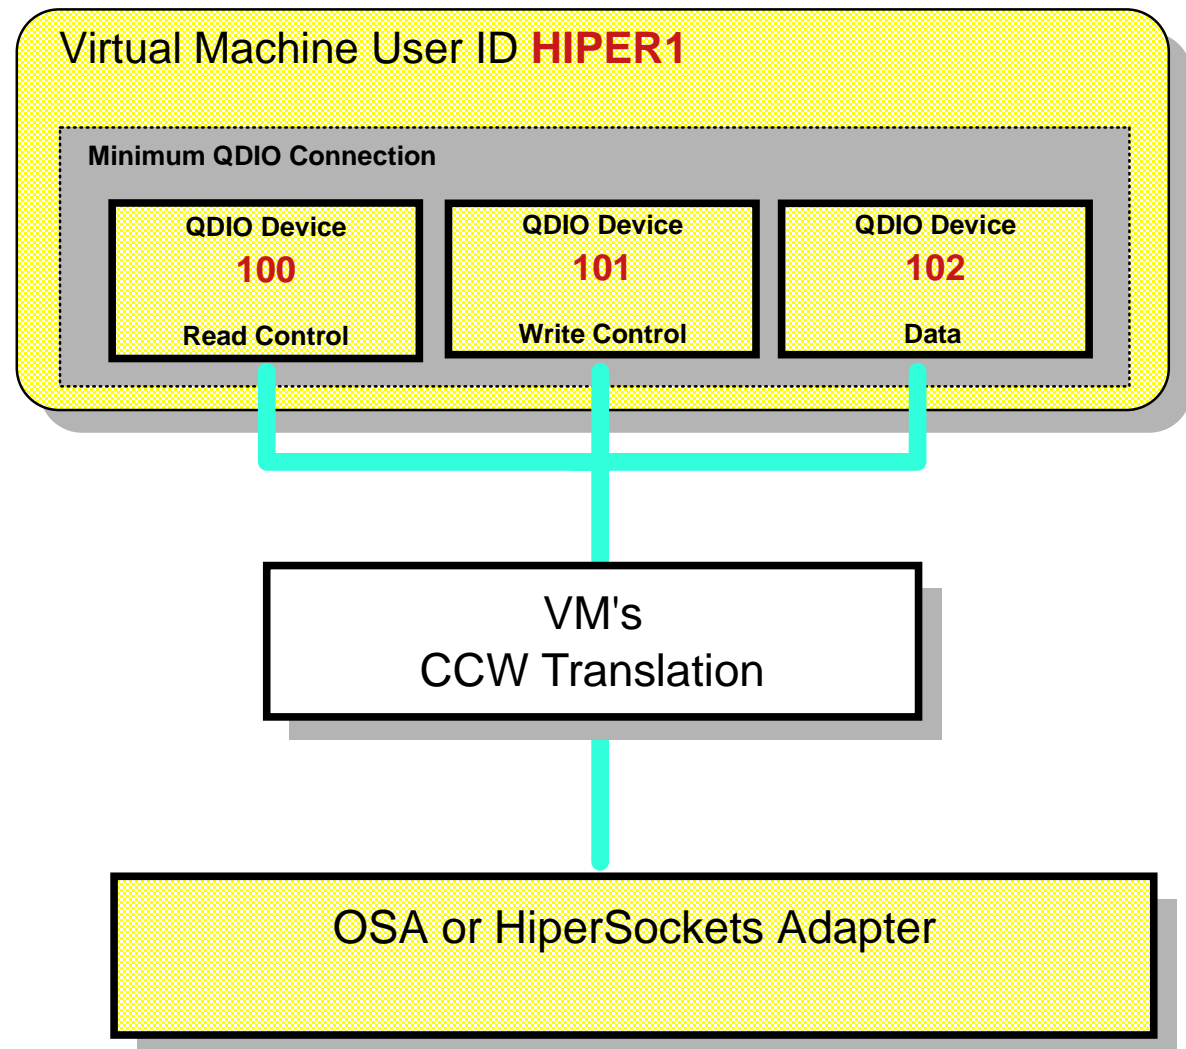
**ATTACH 100-102 TO HIPER1**

At LOGON via User's Directory Entry

**DEDICATE 100 100**

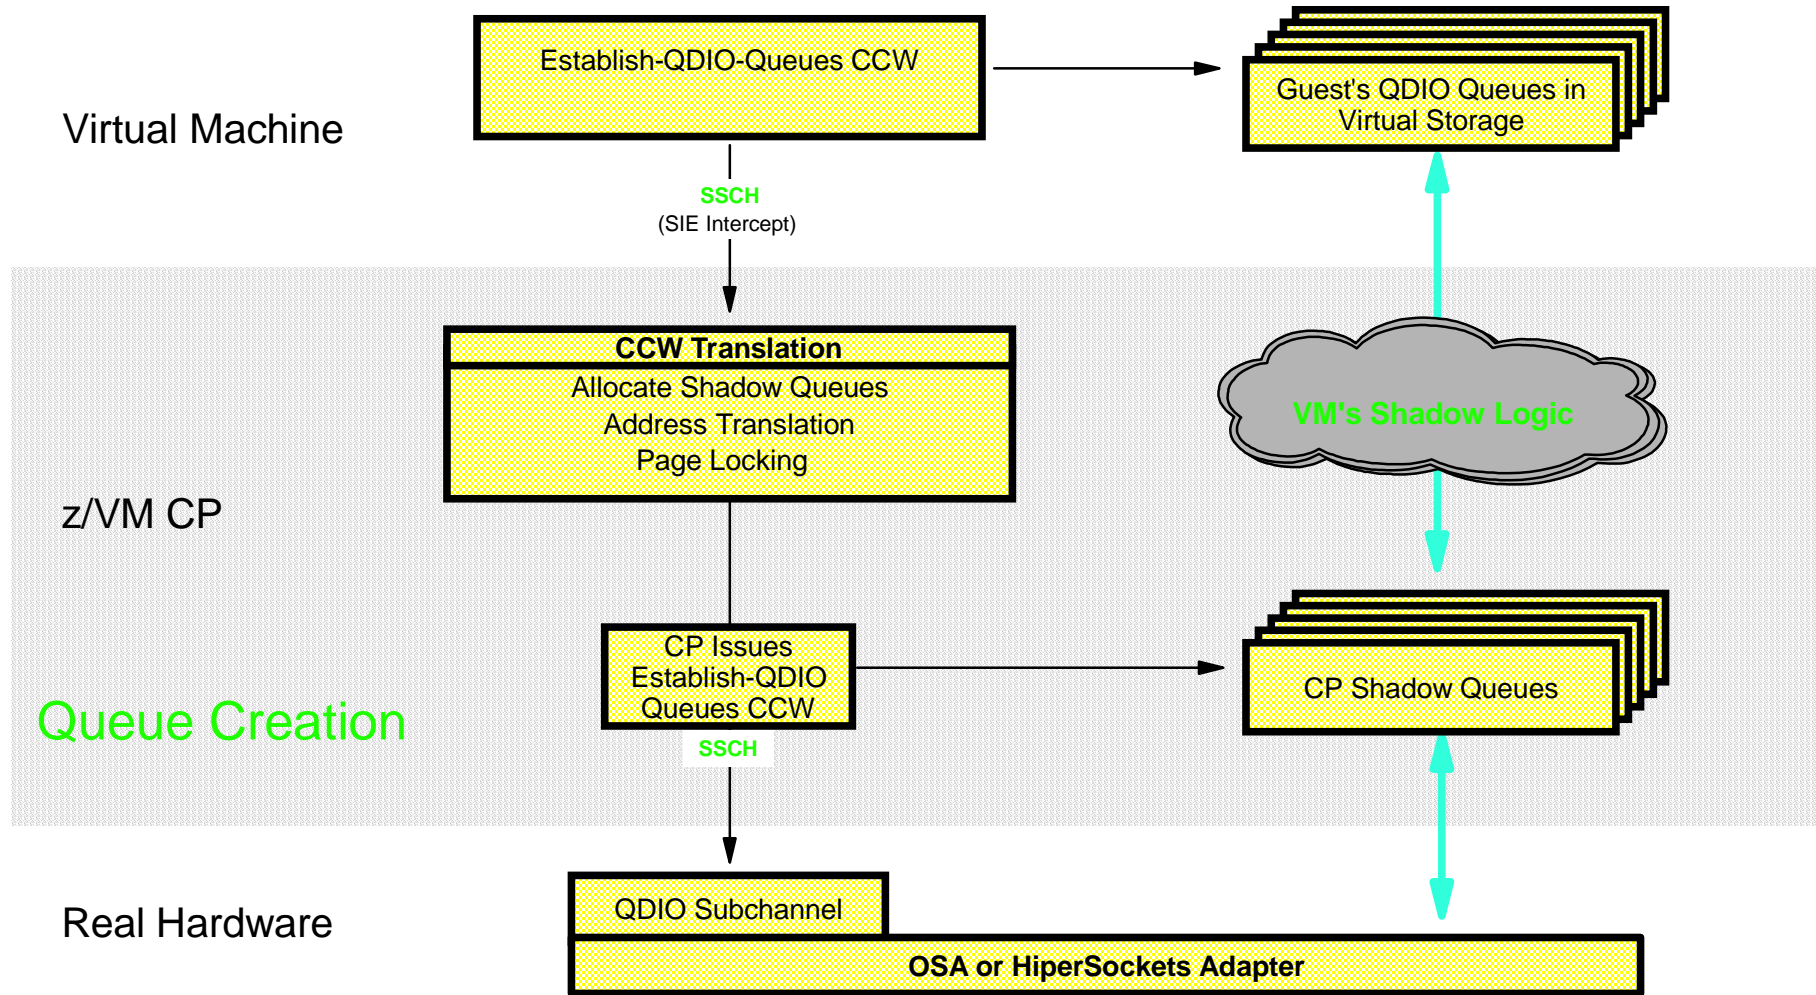
**DEDICATE 101 101**

**DEDICATE 102 102**





# VM Shadow Queues Processing



# Displaying QDIO Connection Status

## QUERY VIRTUAL OSA or QUERY {device number}

```

OSA 0100 ON OSA 0100 SUBCHANNEL = 000B
    0100 QDIO-ELIGIBLE
OSA 0101 ON OSA 0101 SUBCHANNEL = 000C
    0101 QDIO-ELIGIBLE
OSA 0102 ON OSA 0102 SUBCHANNEL = 000D
    0102 QDIO ACTIVE
    0102 INP + 01 IOCNT = 000002B1    ADP = 008 PROG = 000 UNAVAIL = 120
    0102 OUT + 01 IOCNT = 00000000    ADP = 000 PROG = 000 UNAVAIL = 128
    0102 OUT + 02 IOCNT = 000002B0    ADP = 000 PROG = 000 UNAVAIL = 128
    0102 OUT + 03 IOCNT = 00000000    ADP = 000 PROG = 000 UNAVAIL = 128
    0102 OUT + 04 IOCNT = 00000000    ADP = 000 PROG = 000 UNAVAIL = 128
  
```

+ indicates the guest and VM's shadow queues are synchronized. A - indicates the queues are not synchronized.

Number of QDIO buffers transferred on the queue since it was activated.

Number of QDIO buffers waiting for the adapter to process.

Number of QDIO buffers being used by the TCP/IP stack.

## Dedicated FCP

- Introduced with z/VM 4.3.0
- Provides connectivity to Fibre channel Storage Area Network (SAN) via FCP Adapter Card (i.e., provides access to SCSI devices)
- Requires an FCP CHPID (channel path) to be installed
  - ▶ 240 FCP subchannels per CHPID (480 on z990)
  - ▶ Each subchannel provides same hardware interface as real machine
- Dedicate one or more FCP subchannels to a virtual machine with the CP ATTACH command or DEDICATE user directory statement
- Guest uses QDIO protocol to transfer 'SCSI talk' over a FCP subchannel
  - ▶ Each subchannel allows 1 Input Queue and 1 Output
  - ▶ Guest is responsible for QDIO and SCSI protocol drivers
  - ▶ VM QDIO shadow queues processing for FCP subchannels similar to OSA subchannels
- CP QUERY VIRTUAL **FCP** / QUERY {virtual device number} commands provide output similar to QUERY VIRTUAL **OSA** / QUERY {virtual device number}

## Guest IPL from SCSI

Allows guest virtual machine to IPL from SCSI device using FCP subchannel

- Exploits dedicated FCP subchannel support delivered in z/VM 4.3.0; available in z/VM 4.4.0 (and up)
- Special Machine Loader brought into guest storage from Support Element (SE) to perform operating system load from SCSI device on Storage Area Network (SAN)
- CP IPL command initiates Machine Loader process. Recognizes FCP subchannel (not SCSI load device).
- CP SET LOADDEV command (or LOADDEV user directory statement) defines SCSI load device and extended parameters
- Extended parameters:
  - ▶ Port Name - target WWPN associated with load device
  - ▶ LUN - logical unit name of load device
  - ▶ Boot Program Number and Boot Record Block Address - defines offset of IPL dataset on load device
  - ▶ SCP Data Parameters - Operating System (OS) parameters
- SET LOADDEV is issued before IPL command. Use CP QUERY LOADDEV command to display settings.

# IOASSIST and QIOASSIST

- **Hardware I/O Assist defined:**

- ▶ Allows guest I/O *instructions* to go directly to the *device* without the need for VM to intercept the instructions, translate addresses, perform validity checks, or perform other simulation.
- ▶ Allows I/O *interruptions* to go directly to the *guest* without the need for VM to take the interrupts, translate addresses, perform validity checks, or perform other simulation.
- ▶ Improves guest I/O performance to near native performance by avoiding I/O instruction interceptions and I/O interrupts for assisted subchannels.
- ▶ There are two flavors. I/O Assist (IOASSIST) on releases before 5.1.0, and Queued Direct I/O Assist (QIOASSIST) on releases 4.4.0 and later.

# IOASSIST and QIOASSIST

## ■ Hardware I/O Assist controls:

- ▶ IOASSIST - Enabled by default for dedicated devices attached to fixed-page, preferred guests (V=R or V=F). Not applicable for pageable guests (V=V).
- ▶ QIOASSIST - Enabled by default for dedicated devices attached to pageable guests. Not applicable for preferred guests.
- ▶ Disabled on a device basis with operand NOIOAssist/NOQIOAssist on the DEDICATE directory statement or on the CP ATTACH command
- ▶ Enabled/Disabled on a guest basis with CP SET IOASSIST command or CP SET QIOASSIST command

# IOASSIST compared to QIOASSIST

## I/O Assist

- Available in releases prior to 5.1.0
- Not available in 5.1.0
- Benefits at most 6 preferred guests (V=R and V=F)
- Requires dedicated devices
- Does not support QDIO subchannels (OSA and FCP)
- Must be running in Basic mode

## Queued Direct I/O Assist

- Not available in releases prior to 4.4.0
- Available in 4.4.0 and up
- Benefits all V=V guests with QDIO subchannels (OSA and FCP)
- Requires dedicated devices
- Supports QDIO subchannels
- Supports running in LPAR mode

## Tape Support

- VM uses tape drives for many native system maintenance functions:
  - ▶ Backing up disk volumes with the DASD Dump Restore utility (DDR)
  - ▶ Backing up system spool data with the CP SPXTAPE command
  - ▶ Backing up CMS Minidisk files with the CMS TAPE command
  - ▶ Storing system trace data with the CP TRSAVE command
  - ▶ Managing volumes from a robotics tape library with the DFSMS/VM RMS Service Virtual Machine (SVM)
- A tape drive can be dedicated for use by a guest operating system with the SINGLEUSER option of the CP ATTACH command or DEDICATE user directory statement
- A tape drive can be *serially* shared by multiple guest operating systems with the MULTIUSER option of the CP ATTACH command (or DEDICATE statement)
  - ▶ Specifically for guests that manage their own assignment of tape drives, such as z/OS JES 2 Static Vary and Tape Auto-Switching
  - ▶ Works with 3480, 3490, 3590, 3592 tape drives
  - ▶ Not intended for CMS



## Supported Tape Units / Libraries

- 3480/3490 Tape Subsystems
- 3590/3592 TotalStorage Enterprise Tape Systems
- 3494 TotalStorage Enterprise Automated Tape Library
- 3494 TotalStorage Virtual Tape Server (VTS)
- 3495 Tape Library Data Server
- ✚ Native Support for 3494 and 3495 provided in conjunction with DFSMS/VM
- ✚ Currently only ESCON/FICON connection supported (for all tapes)



### Key Differences between drive and cartridge types:

	<b>3480</b>	<b>3490</b>	<b>3590E</b>	<b>3590H</b>	<b>3592</b>
<b>Raw Capacity</b>	<b>200MB</b>	<b>800MB</b>	<b>20.5GB</b>	<b>30.8GB</b>	<b>300GB</b>
<b>Bandwidth</b>	<b>3MB/s</b>	<b>4MB/s</b>	<b>14MB/s</b>	<b>14MB/s</b>	<b>40MB/s</b>

# 3592 Tape Support

VM service required for latest enterprise tape subsystem:

- ✓ Base CP support - VM63325 (z/VM 3.1.0, 4.3.0)
- ✓ Base RMS support - VM63353 (DFSMS/VM FL221)
- ✓ WORM CP support - VM63461 (z/VM 3.1.0, 4.3.0, 4.4.0)
- ✓ WORM RMS support - VM63460 (DFSMS/VM FL221)



3592 R/W Cartridge



3592 WORM Cartridge

4-drive 3592 subsystem with cartridge stackers



## DASD Support

- Chief device for maintaining VM system and guest data
- Can be dedicated to guest operating system with CP ATTACH command or DEDICATE user directory statement
- Can be attached to system with CP ATTACH rdev TO SYSTEM command or automatically at IPL with system configuration file statements SYSTEM\_RESIDENCE, CP\_OWNED, or USER\_VOLUME\_LIST for:
  - ▶ Checkpoint data
  - ▶ Warmstart data
  - ▶ Paging, spooling data
  - ▶ Guest Minidisks
  - ▶ Guest Temp disks

# VM TotalStorage Disk Support

- IBM TotalStorage Models supported
  - ▶ IBM TotalStorage Enterprise Storage Server (ESS) 750
  - ▶ IBM TotalStorage Enterprise Storage Server (ESS) 800
  - ▶ IBM TotalStorage DS6000
  - ▶ IBM TotalStorage DS8000
  
- These models support connections:
  - ▶ ESCON (except DS6000)
  - ▶ FICON
  - ▶ FCP
  
- ESCON/FICON supported as 3990 Model 3 or 6 Controller with:
  - ▶ 3390 Model 2, 3, and 9 DASD
  - ▶ 3390 Model 2 and 3 DASD in 3380 track compatibility mode
  
- ▶ FCP supported as emulated 9336 model 20 DASD

## ESS Model 750

- Smaller and less expensive than traditional Model 800
- Peer-to-Peer Remote Copy (PPRC) and FlashCopy v2 supported
- Parallel Access Volumes (PAV) support available for enhanced zSeries performance
- upgradeable to the ESS Model 800
- Supported by z/VM 5.1.0 as well as z/VM 3.1.0, 4.3.0, and 4.4.0




# DS6000 and DS8000

- New storage systems announced on October 12, 2004
- DS6000:
  - ▶ Exceptional price and performance breakthrough in a modular package that redefines enterprise-class storage
  - ▶ Shares advanced software features with ESS 750, ESS 800, and DS8000
- DS8000:
  - ▶ New standard delivering massive scalability and performance -- helping to lower costs
    - Up to 6 times the performance of the ESS 800 for increased response time and usable capacity (up to 192 TB of storage)
    - 20% smaller footprint than ESS 800
  - ▶ Storage system LPARs allow two independent storage systems within a single array
- z/VM support details:
  - ▶ z/VM 4.4.0 and 5.1.0
  - ▶ DS8000 with CP APAR VM63534 (available now)
  - ▶ DS6000 with CP APAR VM63635 (available TBD)
  - ▶ DIRMAINT APAR VM63656 required for both (available now)




# Flashcopy

- Instant copy of source disk to target
- Guest support via CCW channel program support
  - ✦ Restriction: dedicated or fullpack minidisks
- Native support via CP FLASHCOPY command
- Flashcopy Version 1 (v1) hardware
  - ▶ No concurrent copies (one at a time per source/target pair)
  - ▶ No extent relocation (full volume only)
- Flashcopy Version 2 (v2) hardware
  - ▶ Dataset Relocation (data on different CCHH track offsets when moved to target)
  - ▶ Multiple Relationships (source CCHH track can have many target tracks, but target track can have only one source track)
  - ▶ Cross LSS (source and target can be in different logical control units)



# Peer-to-Peer Remote Copy (PPRC)


- Continuous copy of source disk changes to target
- Guest support via CCW channel program support
  - ✦ Restriction: dedicated or fullpack minidisks
- Native support via ICKDSF Release 17 running in CMS virtual machine
- PPRC hardware
  - ▶ Base PPRC (changes from source to target synchronous to initiating I/O)
  - ▶ PPRC-XD (changes from source to target asynchronous to initiating I/O; allows extended distances between source and target)
  - ▶ PPRC Cascading (3 volume PPRC link; synchronous link between source and intermediate volume; asynchronous link between intermediate and target volume; allows extended distances with added data integrity)
  - ▶ PPRC Async (asynchronous link between source and target; software managed consistency groups allow failover to point-in-time consistent copy in case of disaster)
- CP QUERY DASD DETAILS command displays PPRC volume status (Primary, Secondary, or Cascading Intermediate) when a PPRC link is active for a disk





## Parallel Access Volumes (PAV)

- Base volume shadowed by 1 or more Alias volumes
- Looks like multiple separate, real DASD to host operating system
- Allows concurrent I/O to Base data
- Greater I/O performance (throughput) by reducing I/O queuing contention
- Base and Alias volumes supported for VM guests as dedicated DASD
- Guest access via CP ATTACH command or DEDICATE user directory statement
- Supported by z/OS (WLM) and z/Linux (LVM) guests



## MiniDisk Cache (MDC)

- MiniDisk Cache is a host based, full track, write through cache for partial pack and full pack minidisks (linked by guest virtual machines)
- When one record of a track is read, the entire track is brought into the cache
- Records are kept in the cache based on *last reference*. Records not referenced are aged out of the cache.
- Works well with XSTORE, though XSTORE is not required
- Since MDC is centralized at the host, it is always benefiting the most active disks at any given moment in time
- A *Fair Share* insertion algorithm prevents one user from monopolizing the cache
- By default, MDC is active for all minidisks

## Managing MDC

- CP SET MDC command
  - Turn MDC OFF/ON for the system, a RDEV, or Minidisk
  - FLUSH cache for a given RDEV or Minidisk
  - Turn cache insertion OFF/ON for a specified userid
  - Alter the amount of main storage or XSTORE available to MDC
- CP QUERY MDC command to display MDC settings
- User Directory controls
  - Turn MDC OFF for a user's minidisk with NOMDC option on MINIOPT statement for the Minidisk
  - Disable *Fair Share* limit for a guest by coding NOMDCFS on the user's OPTION statement. Good option for critical server or production virtual machines.
- RDEVICE system configuration file statement
  - Force MDC OFF for all minidisks on real device
  - Set default state of all minidisks on real device to MDC OFF
  - Set default state of all minidisks on real device to MDC ON (default except when real device is shared across LPARs)

## Virtual Disk in Storage (VDISK) Support

- A VDISK is a fully simulated FBA 9336 model 20 DASD. All data resides in main memory.
- VM simulates all I/O commands, interrupts, and CCW's associated with the VDISK.
- Since the VDISK data is in volatile storage, it should be used for applications that make use of temporary files (e.g., assembly, VSE lock file, or sort applications)
- Any guest OS that supports FBA DASD can use VDISKs
- A VDISK can be created via the CP DEFINE VFB-512 command or by the MDISK user directory statement. VDISKs created by the MDISK statement are shareable by multiple guests.
- CP SET VDISK command or FEATURES VDISK system configuration file statement can be used to control amount of main memory used by the system and each guest for VDISKs

## Emulated Device Support

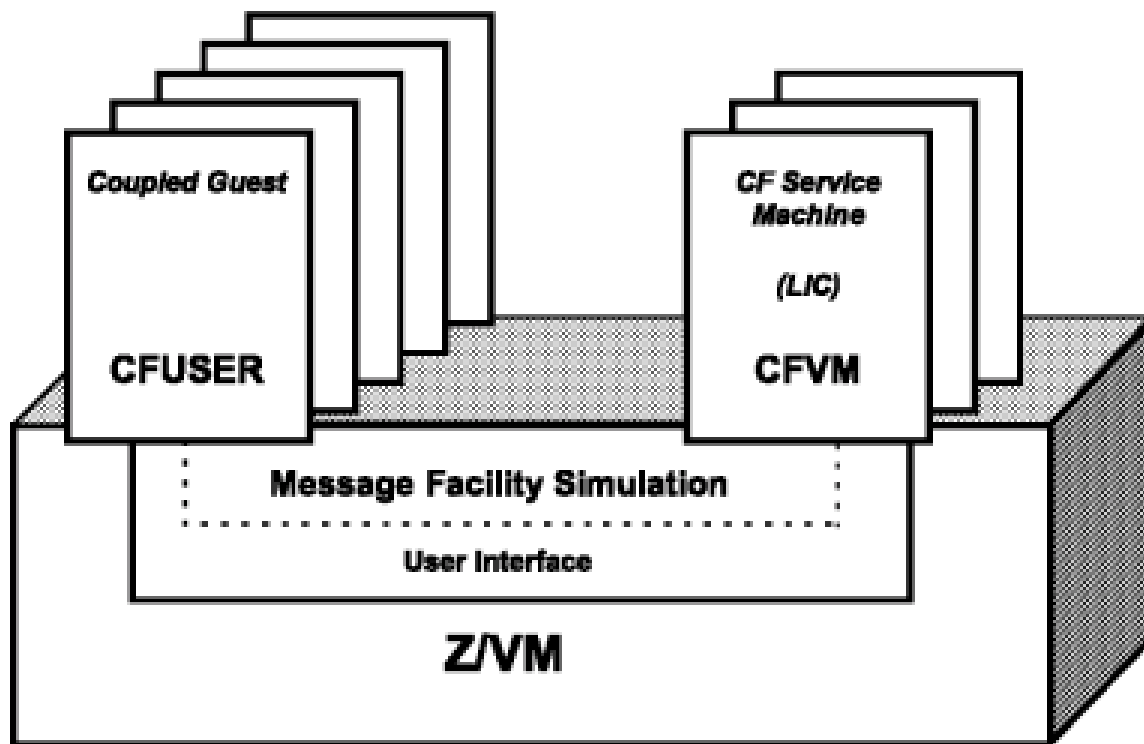
- z/VM 5.1.0 emulates FCP-channel-attached SCSI disks as 9336 Model 20 FBA DASD
- VM and its guests see an emulated device as a real 9336-20
- Emulated SCSI disks can be used for system data such as Paging, Spooling, Directory, and Minidisks
- ECKD DASD is no longer needed to run VM
- Managing emulated devices
  - ▶ Use CP SET EDEVICE command or EDEVICE system configuration file statement to configure a SCSI disk LUN to the VM system as an emulated 9336
  - ▶ Use CP QUERY EDEVICE command to obtain information related to an emulated 9336
  - ▶ Use CP DELETE EDEVICE command to remove an emulated 9336 from the VM system

## Guest Coupling

z/VM Guest Coupling Simulation Support is the software that simulates the Coupling Facility hardware and the software required to run a z/OS sysplex environment in guest virtual machines

- Simulates advanced CF (Coupling Facility) and MF (Message Facility) hardware functions
- Provides pre-production testing platform for a coupled-system installation
- Not intended for production use since its single points of failure negate the intent of the z/OS parallel sysplex environment
- No special hardware needed (no coupling links and no external coupling facilities)
- Up to 32 User Coupling Facility virtual machines can be coupled within a simulated sysplex
- Each of these User virtual machines can be coupled with up to 8 Service Coupling Facility virtual machines
- Each simulated Service Coupling Facility virtual machine can connect to up to 7 peer simulated Service Coupling Facility virtual machines

# Guest Coupling



Three components:

- Service Machines
- Coupled Guests (User Machines)
- Simulated Message Facility

CFUSER and CFVM are specified on the OPTION user directory statement

## Guest Coupling Simulation Commands

### From class G CFUSER:

- DEFINE MSGPROC - Define virtual message processor and associated message devices for communication with a CFVM
- DETACH MSGPROC - Detach the specified virtual message processor and associated message devices
- SET MSGFACIL OFF - Remove message facility capability from the issuing virtual machine. Detaches **all** message processors and message devices.
- QUERY VIRTUAL MSGDEVICES - Display all virtual message devices in virtual machine configuration
- QUERY VIRTUAL MSGPROC - Display all virtual message processors associated with the virtual machine

### From any class A or B user (or class G CFUSER with restrictions):

- DEFINE CFLINK - Define links between CFVM coupling facility virtual machines
- SET CFLINK - *Logically* connects or disconnects a coupling facility link between CFVM virtual machines
- QUERY CFLINKS - Display information about coupling facility links between CFVM virtual machines
- RESTART MSGPROC - Reload and start a specific CFVM service virtual machine



## Dynamic I/O Configuration

- Function allows modification of real hardware I/O definitions (IOCP) without an outage
  - ▶ Define, Modify, and Delete - CHPIDs (channel paths), Control Units, and Devices
  - ▶ Manage I/O definitions for all LPARs on a CEC from one LPAR
  - ▶ Controlled by VM's dynamic I/O command interface or HCD/HCM support (two are mutually exclusive)
  
- Function allows modification of VM software view of I/O definitions
  - ▶ Define, Modify, Delete, and Clear - RDEV control blocks
  - ▶ Controlled by VM's dynamic I/O command interface or HCD/HCM support (two are mutually exclusive)

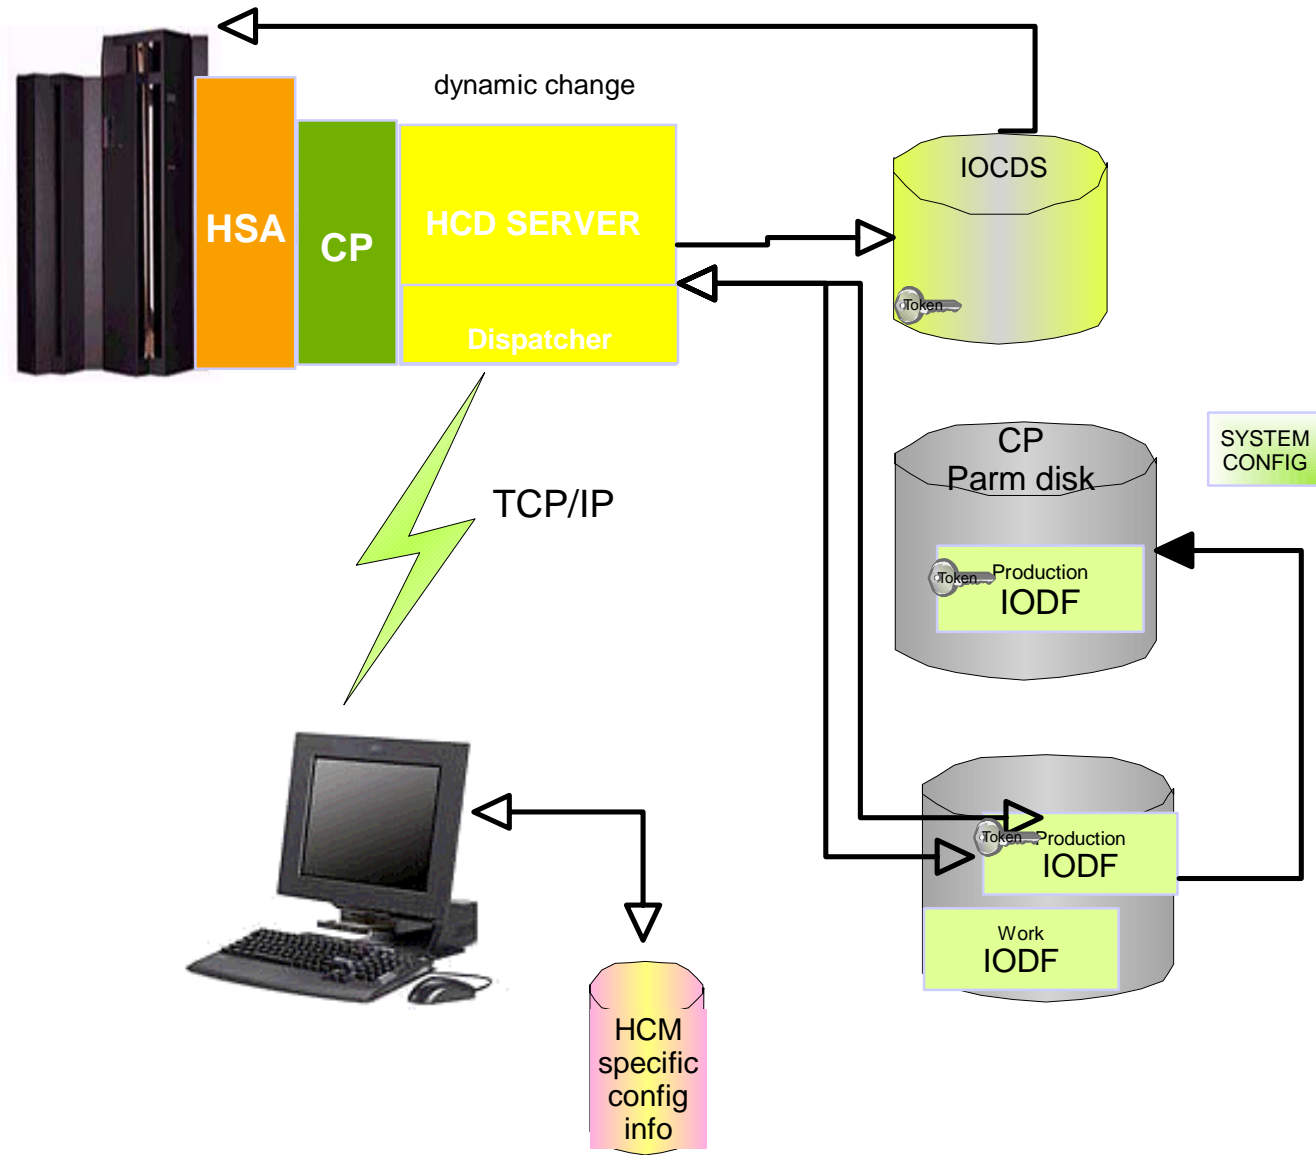
## Dynamic I/O Command Interface

- Enable hardware dynamic I/O command capabilities with the CP SET DYNAMIC\_I/O ON command or FEATURES ENABLE DYNAMIC\_I/O system configuration file statement
- Use CP DEFINE/MODIFY/DELETE CHPID/CU/DEVICE commands to change the hardware configuration one entity at a time
- VM software view of I/O definitions (RDEVs) created when devices are varied online via dynamic sensing or RDEVICE system configuration file statement
- Use CP SET RDEVICE command to create new software device definitions (RDEVs) or modify/clear existing device definitions after IPL

## Dynamic I/O HCD/HCM support

- HCD (Hardware Configuration Definition) runs in a VM service virtual machine and acts as an I/O configuration server that HCM (Hardware Configuration Manager running on a PC) can connect to via TCP/IP
- HCD uses an I/O Definition File (IODF) that contains the current hardware and (optionally) software configuration information. This information is graphically displayed via the HCM terminal.
- VM's IPL code recognizes when HCD is being used to control the I/O configuration (an IODF system configuration file statement is specified) and interprets data from the HCD-supplied I/O Definition File (IODF)
- Dynamic changes to the configuration are made by updating information on the HCM screens to create a target configuration IODF file and then activating that target IODF file via the ACTIVATE command on the HCD service virtual machine console
- HCD processes the ACTIVATE request by comparing the active production IODF file with the target production IODF file and creating a Configuration Change Block (CCB) which represents the hardware and/or software changes that need to be executed. The CCB is sent to CP to execute the change.
- Note: VM's dynamic I/O command interface is turned off and cannot be used when HCD is in control.

# HCD/HCM



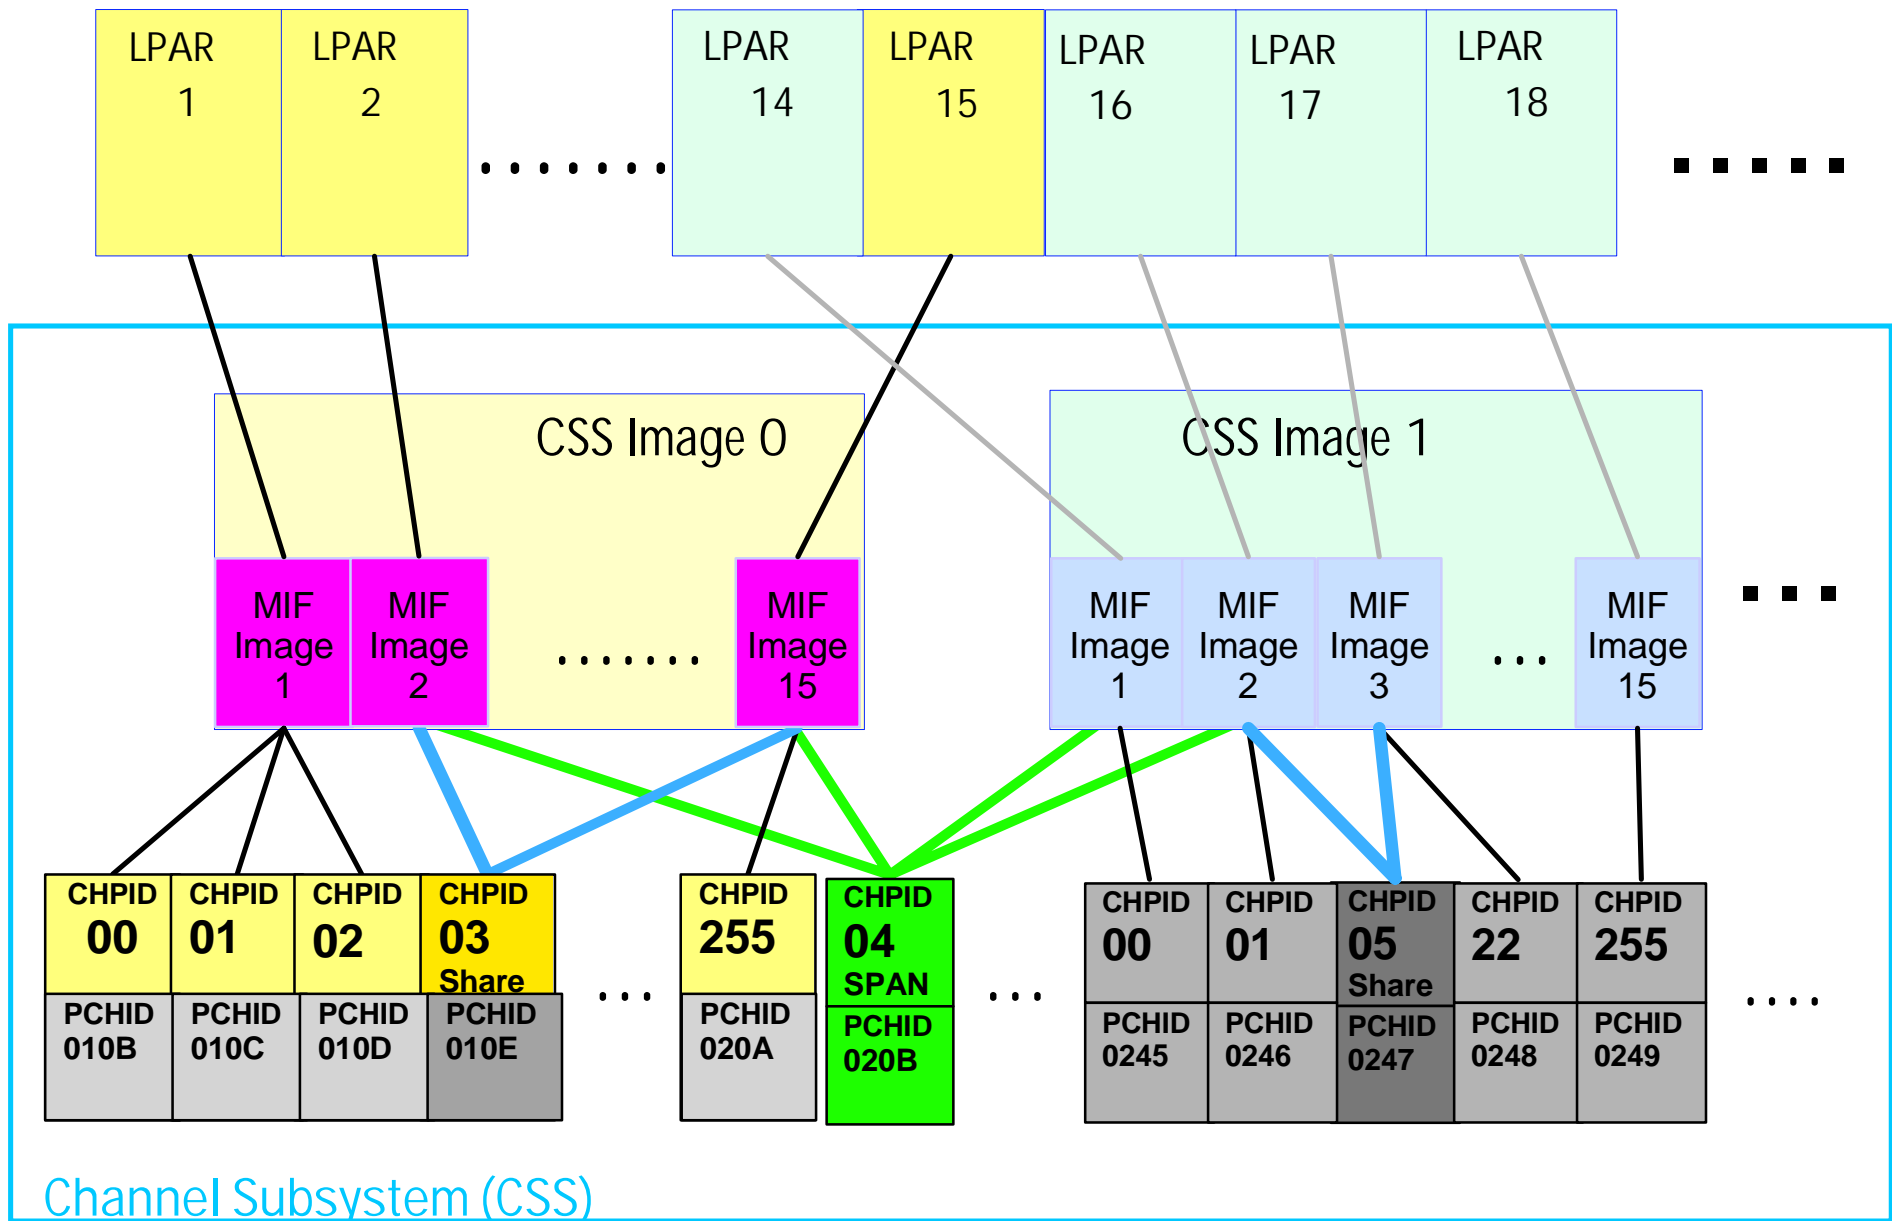
Dynamic changes are only allowed if the token of the current production IODF matches a system token.

The SYSTEM CONFIG file identifies the IODF to be used and optionally selects the OS configuration to be used.

## Logical Channel Subsystem (LCSS) Support

- The Logical Channel Subsystem (LCSS) hardware facility allows one or more sets of 256 physical channels paths (PCHIDS) to be represented as 1-to-4 logical channel subsystem images; each containing 256 logical channel paths (CHPIDs).
- Each logical partition (LPAR) on the CEC can still only be associated with one logical channel subsystem
- FICON, FCP, and *internal* (e.g., Coupling, HiperSockets) channels can be spanned across logical channel subsystems; allows sharing by logical partitions (LPARs) associated with different logical channel subsystems.
- Since a LPAR can still only be associated with a single logical channel subsystem, VM's support of the LCSS facility only involved updating dynamic I/O support:
  - ▶ PCHID parameter was put on the CP DEFINE CHPID and QUERY CHPID commands to define/display the 4-digit PCHID value

# LCSS+Spanning



## Device Support Removed in z/VM 5.1.0

- DASD (drives and controllers): 3370, 3375, 3380, Multiprise Internal Disk, 9332, 9335, 9336, 9340, 3830, 3880
  - ▶ Exception: RAMAC-emulated 3380 models J and K and 3390 DASD configured for 3380-track-compatibility are still supported
  - ▶ Exception: 9336 is still supported for VDISKS and emulated SCSI LUNs
- Optical: 3995 Optical Dataserver
- Tape (drives and controllers): 2440, 3420, 3422, 3424, 3430, 9348, 3803
  - ▶ Exception: OMA/2 CD-ROM emulating a 3422 is still supported
- Communications: All SDLC, BSC, and CETI ICAs

Refer to z/VM 5.1.0 General Information publication for complete listing of devices supported

# Trademarks

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

CICS*	Language Environment*	S/370
DB2*	MQSeries*	S/390*
DB2 Connect	Multiprise*	S/390 Parallel Enterprise Server
DB2 Universal Database	MVS	VisualAge*
DFSMS/MVS*	NetRexx	VisualGen*
DFSMS/VM*	OpenEdition*	VM/ESA*
e business( logo)*	OpenExtensions	VTAM*
Enterprise Storage Server*	OS/390*	VSE/ESA
ESCON*	Parallel Sysplex*	WebSphere*
FICON	PR/SM	z/Architecture
GDDM*	QMF	z/OS*
HiperSockets	RACF*	zSeries*
IBM*	RAMAC*	z/VM*
IBM(logo)*	RISC	

\* Registered trademarks of the IBM Corporation

The following are trademarks or registered trademarks of other companies.

Lotus, Notes, and Domino are trademarks or registered trademarks of Lotus Development Corporation.

Tivoli is a trademark of Tivoli Systems Inc.

LINUX is a registered trademark of Linus Torvalds.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., 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.

Penguin (Tux) compliments of Larry Ewing

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.

IBM considers a product "Year 2000 ready" if the product, when used in accordance with its associated documentation, is capable of correctly processing, providing and/or receiving date data within and between the 20th and 21st centuries, provided that all products (for example, hardware, software and firmware) used with the product properly exchange accurate date data with it. Any statements concerning the Year 2000 readiness of any IBM products contained in this presentation are Year 2000 Readiness Disclosures, subject to the Year 2000 Information and Readiness Disclosure Act of 1998.

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.