



WAVV 2005

# Using z/VM in a SCSI Environment

Monday, May 23, 2005

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

# Agenda

- Overview
- FBA Emulation
- C SCSI Driver Stack
- System Structure
- Configuration Elements
- Commands
- Configuration File Statement
- HCD
- Install
- IPL
- Dump and Service
- Performance
- FCP LUN Access Control
- Last Call ... for Q&A
- Reference
- Trademarks



## Overview

- VM provides native support for SCSI disks for paging, spooling, and other system data
- Support is provided by configuring SCSI disk LUNs to VM as 9336 FBA 512-byte block DASD
- VM guests (such as CMS, GCS, and VSE) may keep data on SCSI disk LUNs without requiring the guest software to have SCSI disk drivers
- IPL, Dump, and Service from/to SCSI disk LUNs is provided to achieve a SCSI-only VM environment
- SCSI-only as well as mixed SCSI and ECKD environments are supported
- Available with z/VM 5.1.0 (9/2004)

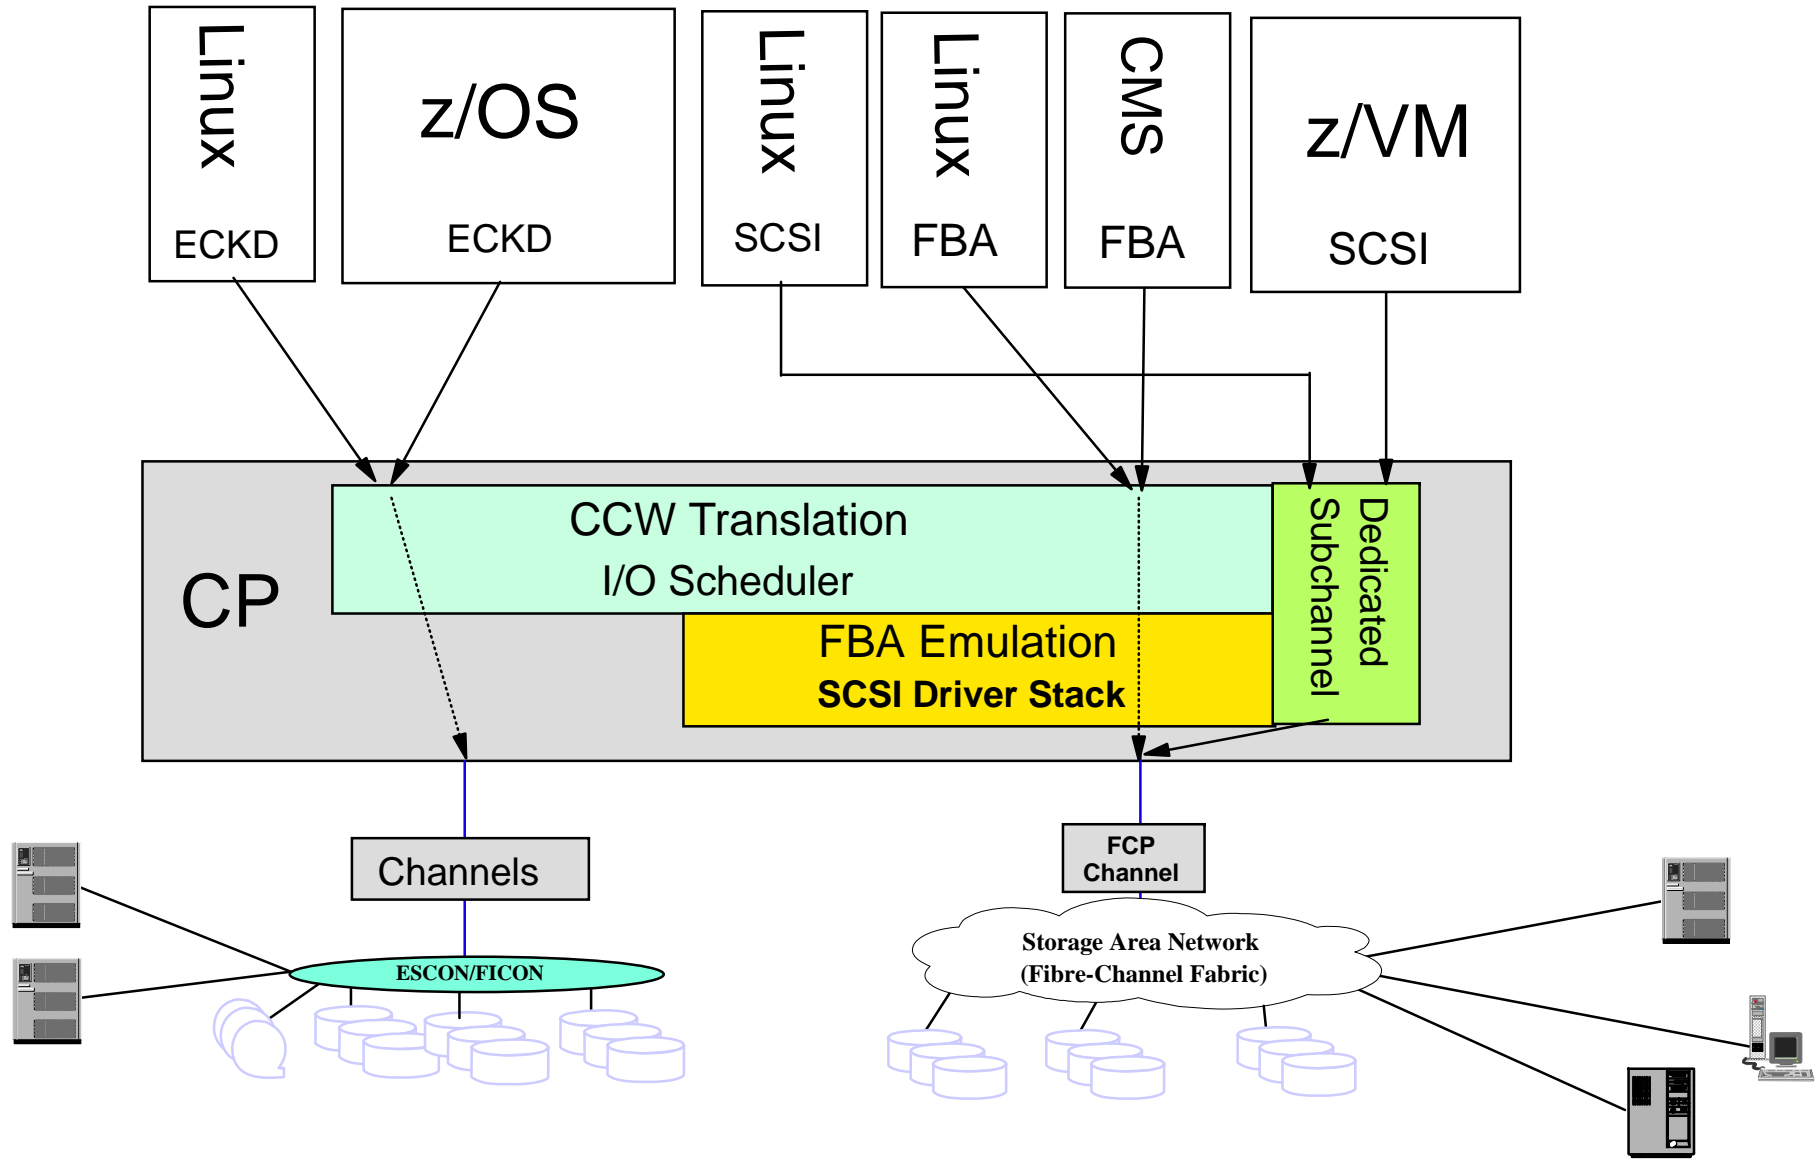
## FBA Emulation

- SCSI Disks are emulated as 9336 Model 20 FBA DASD
- FBA Emulation is used to reduce effort in supporting SCSI disks as well as allowing any operating system or application that supports a 9336 to utilize SCSI disks without change
- Emulated 9336 disk sizes supported are:
  - ▶ 1TB for CP with APAR VM63534
    - ✓ Exception: CP PAGE, SPOL, and DRCT allocations must remain below the 64GB mark on a CP formatted volume since internal addressing of these slots is limited to  $2^{24}$  4K pages.
  - ▶ 381GB for CMS/GCS
    - ✓ This includes software functions dependent on CMS functions such as DIRMAINT MOVE, COPY, ERASE and DFSMS MOVE, COPY, CHECK.
- VM officially supports IBM TotalStorage 1750 (DS6000), 2105, and 2107 (DS8000) SCSI disks as emulated 9336 DASD. However, other SCSI disks may also work since a generic SCSI driver is provided in addition to the TotalStorage drivers.
  - ▶ VM63534 is required for 2107 support; VM63535 is required for 1750 support.

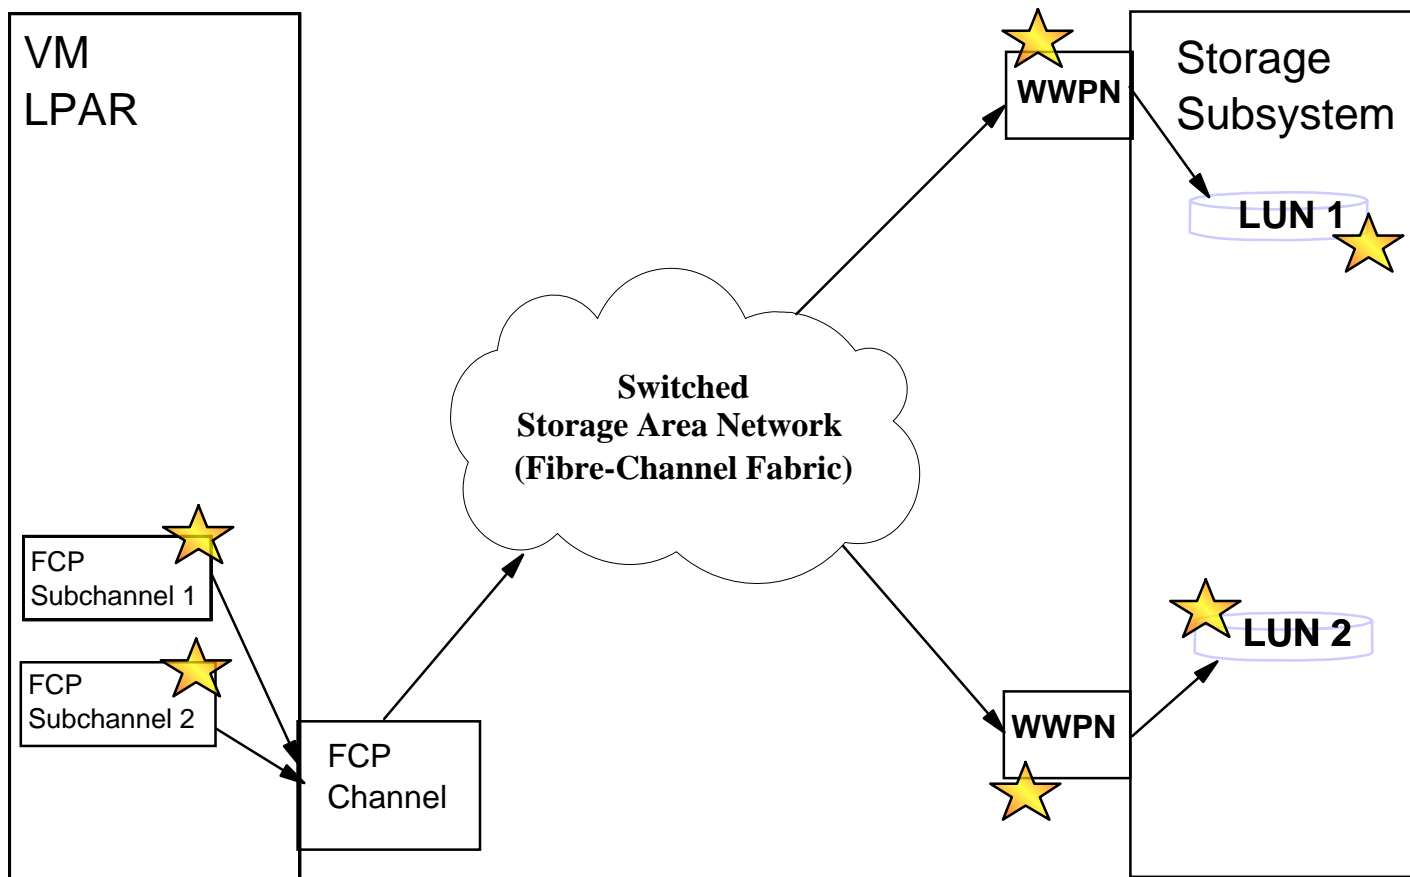
## C SCSI Driver Stack

- A SCSI driver stack written in the C programming language has been coded to drive I/O to SCSI LUNs over FCP subchannels
- This is the first time for C in the CP Nucleus of VM. Therefore, C infrastructure (i.e., linkage and memory stack) has been added with this project. This will allow future CP projects to be written in C.
- This C SCSI driver stack exists in the real I/O layer of the CP component of VM (see next slide).
- An FBA channel program emulator intercepts real Start I/O (SSCH) requests from the real I/O dispatcher of VM and *transforms* these channel programs into API calls to the C SCSI driver stack
- Ending status from the C SCSI driver stack is in turn *transformed* into appropriate z/Architecture and FBA status to be returned to the originator of the FBA I/O request

# System Structure



# Configuration Elements



★ VM configuration element

- SCSI disk is referred to as a LUN
- VM addresses a LUN via a path made up of 3 components
- A FCP subchannel (specified in IOCP), target WWPN on the disk controller, and a LUN name make up the 3 components
- VM allows 8 unique paths per LUN, each with the 3 components
- Figure shows 2 LUNs, each with a single path

# Commands

- New SET EDEVICE command provided to configure a SCSI disk LUN to the VM system as an emulated 9336 FBA DASD
- New QUERY EDEVICE command provided to obtain information related to SCSI disk LUNs defined to VM as emulated FBA DASD
- New DELETE EDEVICE command provided to remove a SCSI disk LUN from the VM system that has been previously defined as an emulated FBA DASD



## Commands ...

Privilege Class: B

```

>>-SET--EDEvice--rdev----->
>-.Type--FBA--ATTRIBUTES--1750-.--| Paths |-.-><
|          | -2105- | | -ADD PATH---- |          |
|          | -2107- | | -DELETE PATH-' |          |
|          | -SCSI-' | |          |          |
'-CLEAR-----'

```

*Paths:*

```

<-----<
|--FCP_DEV--nnnn--WWPN--xxxxxxxxxxxxxxxx--LUN--1111111111111111-->
>-----|
|          | -PREferred---- |
|          | -NOTPREferred-' |

```

**Notes:** You can specify a maximum of 8 "paths" to the device.  
 PREF/NOTPREF only valid for 1750.

## Commands ...

### SET EDEVICE Tips:

- ▶ ATTRIBUTES keyword tells VM which SCSI driver to use
- ▶ FCP\_DEV is the real device number associated with an FCP subchannel that has connectivity to the LUN
- ▶ All 16 digits of the WWPN and LUN must be specified. Otherwise, the VM command parser will add leading zeros causing the path to be invalid. For example, LUN 5150 must be specified as 5150000000000000.
- ▶ An EDEVICE must be varied off-line to ADD, DEL, or CLEAR path information.
- ▶ An EDEVICE takes up a slot in VM's real device space. That is, there can't exist another real device with the same real device number.
- ▶ Once an EDEVICE is defined, it is managed on VM like a real 9336 FBA. CP commands such as VARY, ATTACH, and QUERY execute as if the emulated disk were a real FBA. This applies also to user directory and system configuration file statements.

# Commands ...

Privilege Class: B

```

                                <-----<
>>-Query-EDEVice-.--rdev-----.-.-><
      | '-rdev1-rdev2-' '-DETAILS-' |
      '-ALL-----'
```

# Commands ... Query edevice responses

**q edev 607**

**EDEV 0607 TYPE FBA ATTRIBUTES 2105**

**Ready;**

**q edev 608 details**

**EDEV 0608 TYPE FBA ATTRIBUTES 2107**

**PATHS:**

**FCP\_DEV: 8100 WWPN: 5005076300C604DA LUN: 5137000000000000**

**FCP\_DEV: 8200 WWPN: 5005076300C604DB LUN: 5137000000000000**

**Ready;**

**q edev 609 details**

**EDEV 0609 TYPE FBA ATTRIBUTES 1750**

**PATHS:**

**FCP\_DEV: B100 WWPN: 4004066300C2023B LUN: 0011000000000000 PERF**

**FCP\_DEV: B200 WWPN: 4004066300C2023C LUN: 0011000000000000 NOTPERF**

**Ready;**

## Commands ...

Privilege Class: B

```
>>--DELEte--EDEVice--.-rdev-----.---><  
      '-rdev1-rdev2-'
```

Note: An EDEVICE must be varied off-line to use the DELETE command.

## Configuration File Statement

New SYSTEM CONFIG file statement performing same function as SET EDEVICE command

```
>>-EDEVice--rdev---TYpe--FBA--ATTRibutes--.-1750-.---| Paths |--><
                                     |-2105-|
                                     |-2107-|
                                     '-SCSI-'
```

### *Paths:*

```
<-----<
|--FCP_DEV--nnnn--WWPN--wwwwwwwwwwwwwwww--LUN--1111111111111111-->
>-----|
                                     |-PREFered----|
                                     '-NOTPREFered-'
```

**Notes:** You can specify a maximum of 8 "paths" to the device.  
 PREF/NOTPREF only valid for 1750.

## HCD

HCM GUI interface to VM's HCD support will allow an EDEV to be created with up to 8 unique paths.

Define Device <==> OS Configuration Parameters

Device Number: 8891      OS Config: CT1

Number of devices: 1      Type: VM

Type: FBASCSI      Description: TCFT

OK

Cancel

Help

Parameter	Value	Req	Description
ATTR	2105	<input checked="" type="checkbox"/>	Name of SCSI Device attribute set
FCPDEV	2A00	<input checked="" type="checkbox"/>	Associated FCP device
FCPDEV -2	2A01		Associated FCP device
FCPDEV -3			Associated FCP device
FCPDEV -4			Associated FCP device
FCPDEV -5			Associated FCP device
FCPDEV -6			Associated FCP device
FCPDEV -7			Associated FCP device
FCPDEV -8			Associated FCP device
WWPN	5005076300C204DA	<input checked="" type="checkbox"/>	World Wide Port Number
WWPN -2	5005076300CE04DA		World Wide Port Number
WWPN -3			World Wide Port Number
WWPN -4			World Wide Port Number

# HCD

## GUI continued

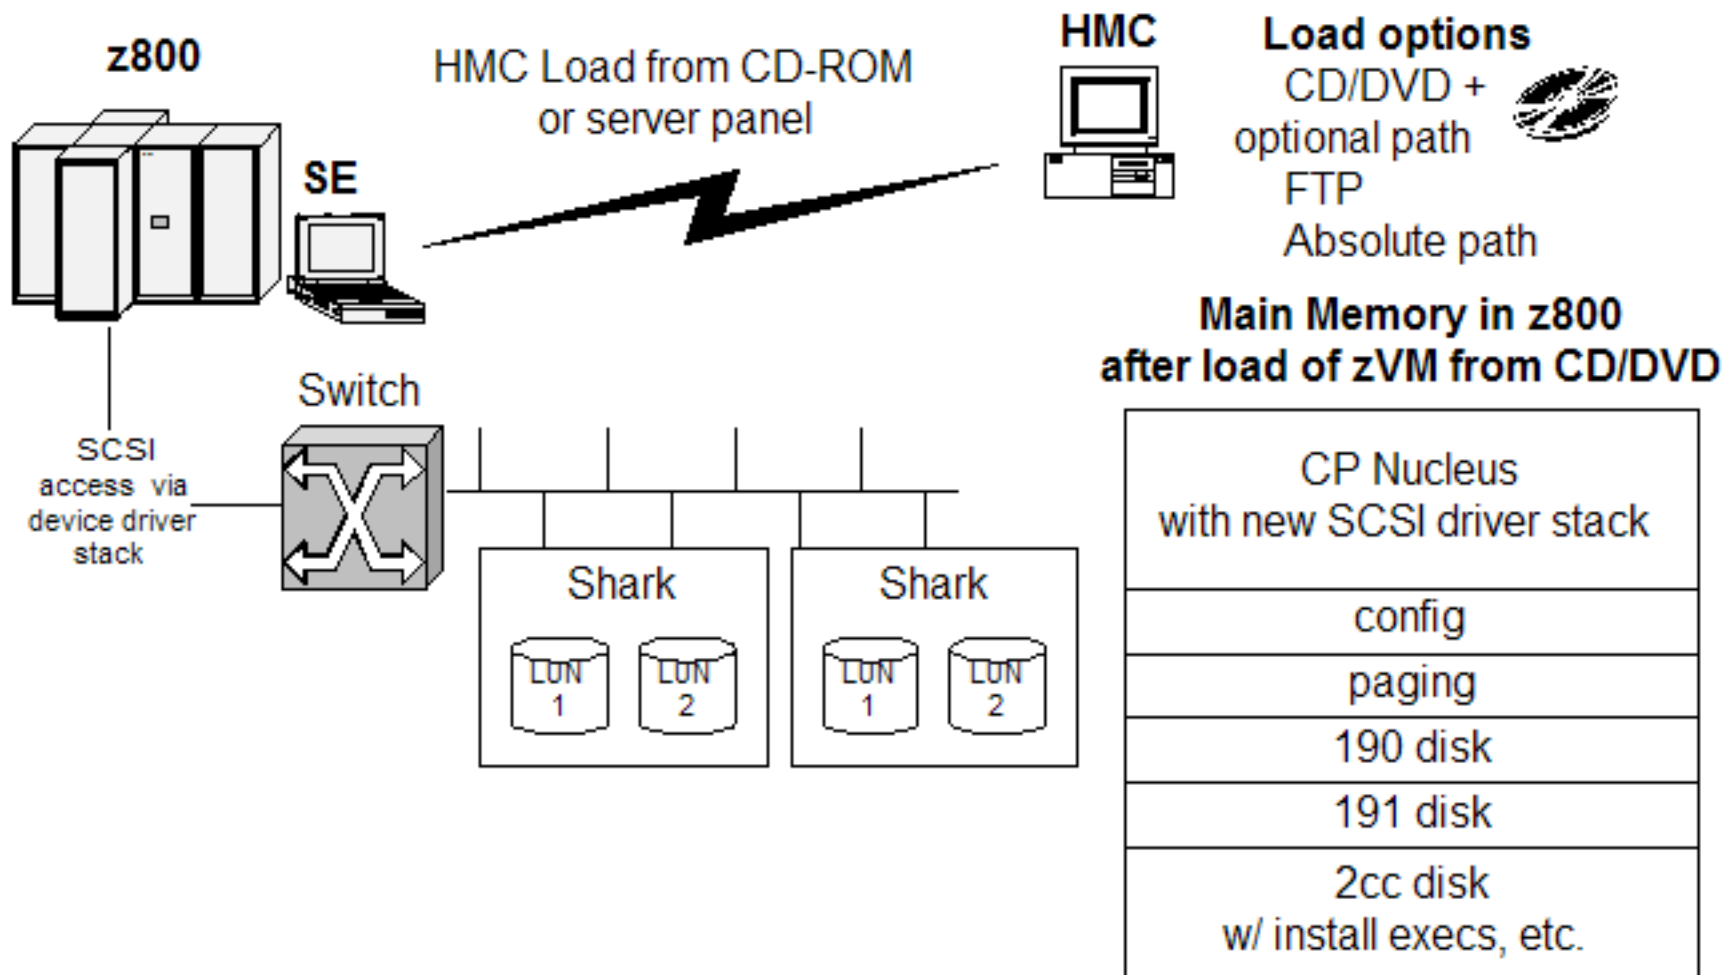
WWPN -5	<input type="text"/>	World Wide Port Number	?	▲
WWPN -6	<input type="text"/>	World Wide Port Number	?	
WWPN -7	<input type="text"/>	World Wide Port Number	?	
WWPN -8	<input type="text"/>	World Wide Port Number	?	
LUN	E15E 000000000000	Logical Unit Number	?	
LUN -2	E15C 000000000000	Logical Unit Number	?	
LUN -3	<input type="text"/>	Logical Unit Number	?	
LUN -4	<input type="text"/>	Logical Unit Number	?	
LUN -5	<input type="text"/>	Logical Unit Number	?	
LUN -6	<input type="text"/>	Logical Unit Number	?	
LUN -7	<input type="text"/>	Logical Unit Number	?	
LUN -8	<input type="text"/>	Logical Unit Number	?	
OFFLINE	<input type="checkbox"/>	Device considered online or offline at IP_	?	▼



# Install

- For customers without a **non-SCSI** Enterprise tape drive (e.g., 3490 or 3590), install is done via DVD from the Hardware Management Console (HMC).
- HMC software with DVD Load and Integrated 3270 console support is required. A 512MB LPAR is also required.
- Install uses a special HMC hardware interface to bring in the VM starter system. The starter system uses a RAM disk to get itself and the install procedure started. Install writes the VM system and other files out to SCSI disk using FBA emulation (see next slide).
- SCSI install is done to 4 disks; 2 1GB LUNs and 2 3.5GB LUNs
- Install is business as usual for customers with an Enterprise tape drive
- Second-level install is also supported for customers without an Enterprise tape drive but requires electronic transfer (via FTP) of the install files to the first-level system. This transfer is done automatically from the install DVD by the installation utility.
- Install from DVD is also available for 3390 models 3 and 9. 2 install DVDs are available; 1 for SCSI, 1 for 3390.

# Install ...



# IPL

- Once installed, VM can be IPLed from a SCSI LUN using the Load panel on the SE or HMC
- Stand-Alone Program Loader (SAPL) continues to be the VM IPL interface
- SAPL contains its own generic, stripped down SCSI stack so that it can read the VM load module from disk into memory. Once SAPL passes control to VM, FBA emulation is used to do all SCSI I/O.
- VM IPL parameter PDVOL must be specified for a SCSI IPL. It indicates the EDEVICE number of the SYSRES.
  - ▶ System stops with a wait state 6505 if not specified
  - ▶ Can be entered on the SAPL screen (PDVOL=xxxx) or defined when SAPL is set up ('burned in')

## IPL

**Load**

CPC: P000F12B  
Image: ZFCP4

Load type:  Normal  Clear  SCSI  SCSI dump

Store status

Load address: 5C00

Load parameter: SYSG

Time-out value: 060 60 to 600 seconds

World wide port name: 5005076300CE93A7

Logical unit number: 5732000000000000

Boot program selector: 0

Boot record logical block address: 0000000000000000C8

OS specific load parameters

OK Reset Cancel Help

- Select SCSI radio button
- Load address is FCP subchannel
- Load parameter is console address for SAPL
- World wide port name (WWPN) and Logical unit number (LUN) address the target load volume
- Boot program selector and Boot record logical block address specify offset of IPL deck on load volume (use 0 and 0000000000000000C8)

## Second-Level IPL

- Second-level IPL of z/VM 5.1.0 from a SCSI LUN works via existing support that went into z/VM 4.4.0 for Linux Guest IPL from SCSI
- Setup WWPN and LUN name of target load volume with CP SET LOADDEV command
- Setup Boot Program Selector and Boot Record Logical Block Address with SET LOADDEV command
- Use same Boot Program Selector and Boot Record Logical Block Address as first-level IPL (0 and 0000000000000000C8)
- IPL virtual address of FCP subchannel to kickoff the load

## Second-Level Example

```
att 50aa *
FCP 50AA ATTACHED TO STEVO 50AA
Ready; T=0.01/0.01 13:16:20

q v fcp
FCP 50AA ON FCP 50AA CHPID 40 SUBCHANNEL = 000E
50AA QDIO-ELIGIBLE QIOASSIST-ELIGIBLE
Ready; T=0.01/0.01 13:16:24

set loaddev portname 50050763 00c20b8e lun 52410000 00000000 bo
otprog 0 br_lba 00000000 000000C8
Ready; T=0.01/0.01 13:16:33

q loaddev
PORTNAME 50050763 00C20B8E LUN 52410000 00000000
BOOTPROG 0 BR_LBA 00000000 000000C8
Ready; T=0.01/0.01 13:16:38

IPL 50aa <Hit ENTER and the IPL starts...>
```

# SALIPL

- The SALIPL utility continues to be used to setup SAPL for IPL
- SALIPL now writes to blocks 5-207 if the device is a 9336 (emulated or not). SALIPL previously wrote to blocks 5-31.
  - ▶ Affects size of RECOMP area for a CMS minidisk containing SAPL, or
  - ▶ Affects placement of allocations such as PAGE and SPOL when SAPL is put on a CP formatted volume
- SALIPL must run in a virtual machine to setup a SCSI disk for IPL. SALIPL can be run against either:
  - ▶ A virtual device, such as a fullpack minidisk, on an emulated 9336
  - ▶ A virtual FCP subchannel where new SALIPL parameters WWPN and LUN designate the target SCSI disk
- The SYSRES device and device containing the PARM area must be the same for a SCSI IPL. This isn't the case for an ECKD IPL.

## Dump and Service

- Support has been added to take an ABEND dump (i.e., hard, soft, SNAPDUMP, VMDUMP) to a SCSI LUN. Stand-alone dump to SCSI is not supported.
- Dumps can be submitted to VM Level 2 electronically via FTP for customers without a **non-SCSI** Enterprise tape drive (e.g., 3490 or 3590)
- Dumps can be copied to tape via FBA emulation for customers with an Enterprise tape drive using existing tools such as the CMS TAPE command
- Service is obtained electronically for customers without an Enterprise tape drive
- Service is applied by existing tools without change using FBA emulation (i.e., put service files on emulated 9336 and run existing tools)
- Service is business as usual for customers with an Enterprise tape drive or 2074 CDROM



# Performance

## Negatives

- Significant path length increase, so one needs to insure proper planning for enough processor cycles to handle the I/O rate
- Monitoring capabilities are not as strong as traditional ECKD in isolating problems at the channel, control unit, and device levels.
- FCP paths are not expected to be significantly faster than FICON (from a hardware perspective)
- Further increases the VM "overhead" numbers for Linux guests

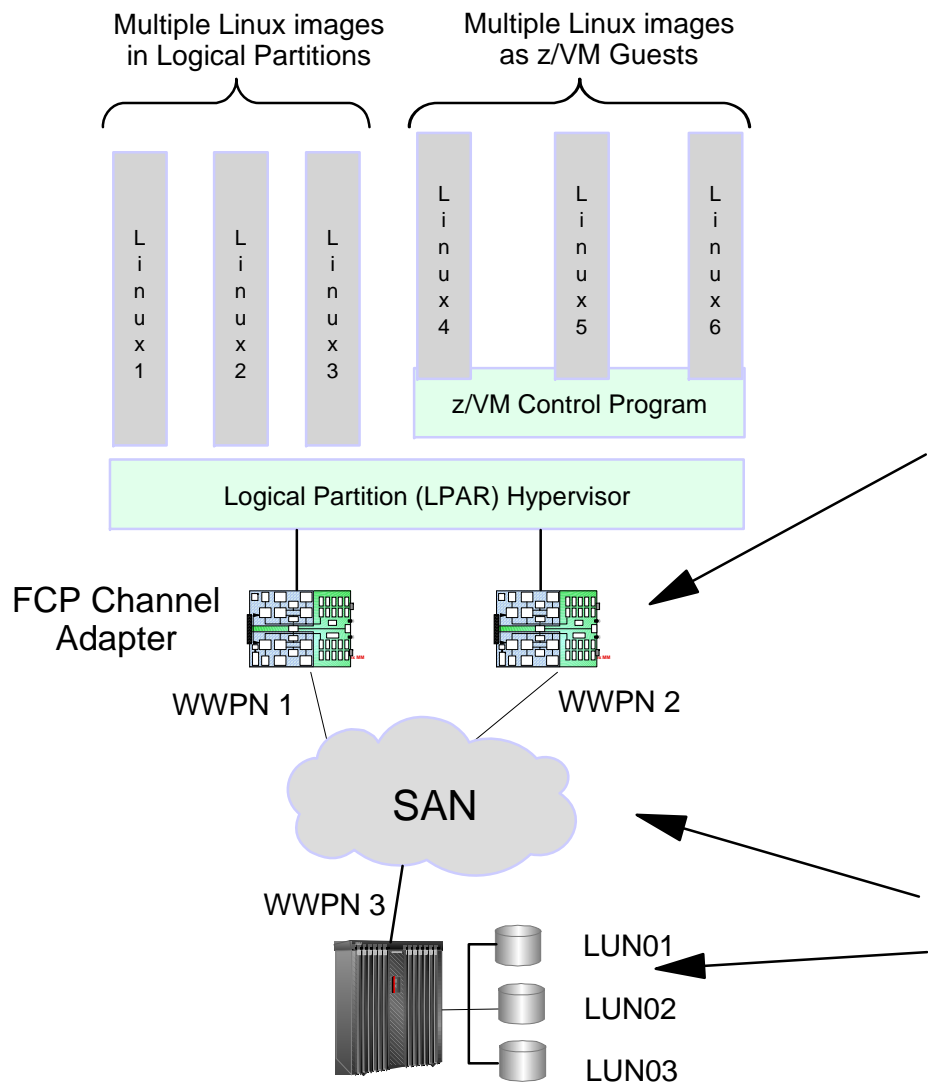
## Positives

- Allows exploitation of minidisk cache
- Allows minidisks on SCSI to be shared
- Improves performance measurement capabilities compared to FCP subchannels dedicated to a guest
- FCP paths are faster than ESCON
- Greater flexibility in size of minidisks based on SCSI devices tending to be larger in general
- Allows potential use of other tuning options not available for dedicated FCP subchannels:
  - ❖ SET THROTTLE
  - ❖ SET IOPRIORITY

# Performance Recommendations

- Continue to use dedicated FCP subchannels for Linux guests except when:
  - ❖ sharing minidisks can add disk and administrative savings
  - ❖ significant common data exists in read-mostly usage and minidisk cache can benefit
  - ❖ other VM minidisk management capabilities add significant value
- Continue to use traditional ECKD DASD for paging and spooling except when:
  - ❖ little paging and spooling activity exists
  - ❖ sufficient processor resources are available to handle increased path length
- Continue to use traditional ECKD DASD for CMS minidisks except when:
  - ❖ I/O to CPU ratios are low (minidisk cache helps lower this ratio)
  - ❖ sufficient processor resources are available to handle increased path length
- Apply APARs VM63534 (available) and VM63725 (not yet available) for significant performance improvements for CMS and Linux guests
- Consider the new z/VM native SCSI support when:
  - ❖ no ECKD DASD is available
  - ❖ it would facilitate moving off of ESCON channels to FCP
  - ❖ large minidisks are desired

# FCP LUN Access Control



- Problem: Without FCP LUN Access Control
  - ◆ LUN Masking and Zoning only allow access control per WWPN of FCP Channel Adapter
  - ◆ No individual access control/protection for Operating Systems sharing a FCP Channel Adapter
- Solution: With FCP LUN Access Control
  - ◆ Can specify separate permissions for each Linux LPAR, Linux VM guest, and VM hypervisor sharing a FCP Channel Adapter
  - ◆ Provides improved security and access protection

With LUN Access Control

Channel "policies" accesses for OSs sharing a FCP Channel Adapter based upon an Access Control Table (ACT)

LUN Masking and Zoning

Access to control unit ports and devices (LUNs) is controlled by zoning (in switches) and LUN masking (in controllers), based on identity of FCP Channel (WWPN).

LUN masking and zoning are industry standard methods of controlling access to target control units and LUNs

## Invocation and Use

- An Access Control Table (ACT) is built defining what LPARs and/or VM guest userids have access to specific target control unit WWPNs and LUNs
  - ACT is in XML format
  - Built on Linux using any text or XML editor
  
- ACT is downloaded to FCP Channel and activated via new Linux utility
  - *Configuration Utility for FCP LUN Access Control* is run from a Linux LPAR or Linux VM guest
  - Utility downloads / activates ACT thru special FCP subchannel associated with unit address xFC or xFD defined in IOCP
  
- Once ACT is downloaded / activated, FCP Channel monitors and controls the flow of I/O over specific FCP subchannels according to the policies in the ACT
  - When running under VM, VM tags a FCP subchannel as owned by a specific guest when the subchannel is attached to the guest
  - The reserved userid *SYSTEM* is used to tag VM hypervisor owned FCP subchannels (z/VM 5.1.0 only) and is used accordingly in the ACT
  - The FCP Channel generates an Access Control error when I/O targets a control unit port or LUN against the policies in the ACT

# Simple VM Guest XML Example

```

<ref_table>
  <wwpn_group id="shark_port">
    <wwpn wwpn="4711000000000000"/>
  </wwpn_group>

  <lun_group id="shark_lun" wwpn_ref="shark_port">
    <lun id="lun1" lun="5401000000000000"/>
  </lun_group>

  <os_id_group id="all_Linuxes">
    <os_id id="lin1" guest_id="Linux1"/>
    <os_id id="lin2" guest_id="Linux2"/>
  </os_id_group>
</ref_table>

```

Control Unit WWPN

LUN

Guests

```

<os_wwpn_rules>
  <os_wwpn_entry>
    <os_id_ref refid="all_Linuxes"/>
    <wwpn_ref refid="shark_port"/>
    <full_access/>
  </os_wwpn_entry>
</os_wwpn_rules>

<os_wwpn_lun_rules>
  <os_wwpn_lun_entry>
    <os_id_ref refid="all_Linuxes"/>
    <lun_ref refid="shark_lun"/>
    <rw_serial_mode/>
  </os_wwpn_lun_entry>
</os_wwpn_lun_rules>

```

Guests

## QUERY FCP AGENT

- Special FCP subchannel unit addresses xFC and xFD are known as "Agent" devices on VM
- One of these subchannels must be attached to the Linux guest for downloading the ACT to the FCP Channel via the Linux Configuration Utility
- Use QUERY FCP AGENT to find an Agent device on the associated FCP Channel (CHPID). Use the ATTACH command to dedicate the subchannel to the Linux guest where the Utility will run.

### Example:

**q fcp agent**

**FCP 8100 is an FCP Agent**

**Ready; T=0.01/0.01 09:56:46**

**q path 8100**

**Device 8100, Status ONLINE**

**CHPIDs to Device 8100 (PIM) : 66**

**Physically Available (PAM) : +**

**Online (LPM) : +**

**Legend + Yes - No**

**Ready; T=0.01/0.01 09:57:19**

(Optional) To verify that agent device is on proper Channel (or CHPID)

## Installation

- Configuration Utility for FCP LUN Access Control available on the web at IBM Resource Link: <https://www-1.ibm.com/servers/resourceLink/>
- Contains 2 RedHat Package Managers (RPMs); 1 for 31-bit Linux, the other for 64-bit Linux
- RPMs install the Utility and Access Control Table (ACT) input template
- User's Guide available from Resource Link with detailed install instructions, XML tag definitions, ACT examples, etc.
- VM support available with APAR VM63328 on 4.4.0 and 5.1.0

# Last Call ... for Q&A

# The End



Reference  
z/VM 5.1.0 Native  
SCSI Support

# HMC SE Requirements

- The Hardware Management Console (HMC) must be communicating with the Support Element (SE). The HMC can only communicate with versions of the SE that are equal to or lower than the HMC. For example, HMC version 1.8.0 can communicate with a SE at version 1.7.3, or 1.8.0, but it can not communicate with a SE at version 1.8.2.
  
- The following minimum SE levels are required to install from DVD:
  - ❖ zSeries 800 - Support Element (SE) version 1.7.3, Engineering Change (EC) J11213, change level 146 or higher must be active.
  - ❖ zSeries 890 - Support Element (SE) version 1.8.2. No Licensed Internal Code changes are required.
  - ❖ zSeries 900 - Support Element (SE) version 1.7.3, Engineering Change (EC) J11213, change level 146 or higher must be active.
  - ❖ zSeries 990 - Support Element (SE) version 1.8.0, Engineering Change (EC) J12560, change level 054 or higher must be active. Or, Support Element (SE) version 1.8.2. No Licensed Internal Code changes are required.

## FCP Firmware Requirements

- The following zSeries driver levels and FCP channel code levels are the minimum required to run z/VM 5.1.0 with SCSI disks:
  - ▶ zSeries 800 - Driver D3G with FCP Code 0.28, MCL J11233 #014
  - ▶ zSeries 890 - Driver D55 with FCP Code 3.04, MCL J13471 #003
  - ▶ zSeries 900 - Driver D3G with FCP Code 0.28, MCL J11233 #014
  - ▶ zSeries 990 - Driver D52 with FCP Code 2.07, MCL J12951 #004 for *GA2 level z990*
  - ▶ zSeries 990 - Driver D55 with FCP Code 3.04, MCL J13471 #003 for *GA3 level z990*
- Driver changes require a POR. FCP code changes on z890 and z990 at Driver D55 are concurrent. Otherwise, CHPID must be recycled off to on.

## TotalStorage Server Firmware Requirements

- 2105 Model F - LIC EC Identification Number: 2.3.2.108,  
EC Level: SC40723
- 2105 Model 800 - LIC EC Identification Number: 2.4.1.58,  
EC Level: SC40908
- 2107 (DS8000) - Latest microcode bundle
- 1750 (DS6000) - Latest microcode bundle

# Monitor Records

- **MRMTRDEV** (Device Configuration Data record in Monitor Domain)
  - ▶ Changed the reserved byte at the end of the record to a new flag byte called MTRDEV\_EDEVTYPE. This byte is a code associated with the type of SCSI device. For example, SCSI disk over FCP protocol. Also added the following information to the bottom of the record:
    - WWPN(s)
    - FCP subchannel(s)
    - LUN Address
    - Attribute Name
  
- **MRIODDEV** (Device Activity record in I/O Domain)
  - ▶ Added new flag byte called IODDEV\_EDEVTYPE at the end of the record. This byte is a code associated with the type of SCSI device. Added IODDEV\_RDEVDEV (real device number) to the end of the record. Previously, the record only contained RDEVSID (subchannel id) which isn't unique for SCSI devices. Adding RDEVDEV enables Monitor to uniquely identify the device.
  
- **MRIODVON** (Vary On Device record in I/O Domain)
  - ▶ Changed the reserved byte at the end of the record to a new flag byte called IODVON\_EDEVTYPE. This byte is a code associated with the type of SCSI device. Also added WWPN, FCP subchannel, and LUN path information.

# Monitor Records

- **MRIODDTD** (Detach Device record in I/O Domain)
  - ▶ Added IODDTD\_RDEVDEV (real device number) to the end of the record so that Monitor can uniquely identify the device.
  
- **MRIODVOF** (Vary Off Device record in I/O Domain)
  - ▶ Added IODVOF\_RDEVDEV (real device number) to the end of the record so that Monitor can uniquely identify the device.
  
- **MRSTOASS** - (Auxiliary Shared Storage Management record in Storage Domain)
  - ▶ Added STOASS\_RDEVDEV (real device number) to the end of the record so that Monitor can uniquely identify the device.
  
- **MRSTOATC** - (Page/Spool Area of a CP Volume record in Storage Domain)
  - ▶ Added STOATC\_RDEVDEV (real device number) to the end of the record so that Monitor can uniquely identify the device.
  
- **MRMTRPAG** - (Paging Configuration Data record in Monitor Domain)
  - ▶ Added MTRPAG\_RDEVDEV (real device number) to the end of the record so that Monitor can uniquely identify the device.

# Monitor Records

- **MRSEKSEK** - (Seek Data record in Seek Domain)
  - ▶ Added SEKSEK\_RDEVDEV (real device number) to the end of the record so that Monitor can uniquely identify the device.
  
- **MRIODSZI** (New I/O Domain Monitor Record 24)
  - ▶ This is the SCSI Device Activity Record and will be cut in addition to the existing MRIODDEV record. The first 20 bytes of the record will contain the standard header for monitor records. The rest of the record will contain the following:
    - RDEVDEV
    - Bytes in a block (512)
    - -> amount of time device is active
    - -> kbytes/sec transfer rate
    - -> #transfers to/from disk
    - -> #blocks read from disk
    - -> #blocks written to disk
    - -> #seek operations
    - -> I/O queue depth
    - Statistics on per path basis (same as -> arrows)

# Performance Toolkit

Updated to display performance stats from new and FBA 'emulated' Monitor data

IBM Performance Toolkit for VM

SCSI device performance data (GDLVMK4)  
Select a device for I/O device details

Command Refresh Systems Menu Help

Auto-Refresh

Interval 16:37:47-16:38:48, on 2004/11/18 (CURRENT interval

Device Number	Bytes per block	Activity/sec KByte	Trans fers	Seeks	Read	Written	Percent Connect
<u>150B</u>	512	0	.000	.000	.000	.000	0
<u>517C</u>	512	0	.000	.000	.000	.000	0
<u>5319</u>	512	0	.000	.000	.000	.000	0
<u>550D</u>	512	0	.000	.000	.000	.000	0
<u>6660</u>	512	0	292.4	.000	.000	22716	0
<u>6663</u>	512	0	168.4	.000	.000	13080	0

Example of format of EDEVs 6660 and 6663



## Performance Toolkit

Format  
example  
continued  
showing  
more  
output

Performance Details for a Single I/O Device (GDLVMK4)

IBM  
Performance  
Toolkit for VM

Command Refresh Systems Menu Return Help  Auto-Refresh

Interval 16:39:00-16:39:48, on 2004/11/18

Detailed Analysis for Device 6663 ( dedicated to PERFSVMK as 6663 )

Device type :	9336	Function pend. :	.0ms	Device busy :	0%
VOLSER :		Disconnected :	.0ms	I/O contention:	0%
Nr. of LINKs:	0	Connected :	14.4ms	Reserved :	0%
Last SEEK :	19871	Service time :	14.4ms	SENSE SSCH :	0
SSCH rate/s :	.0	Response time :	14.4ms	Recovery SSCH :	0
Avoided/s :	.0	CU queue time :	.0ms	Throttle del/s:	...

Status: ONLINE

Path(s) to device 6663: 7C  
Channel path status : ON

No active MDISKs on this volume - no SEEKs analysis possible

Reference  
FCP LUN Access  
Control

# FCP LUN Access Requirements

## ■ Hardware:

- ▶ z890 or z990 with driver D55 and patches:
  - EC J13471 LVL003
  - EC J13484 LVL039
  - EC J13486 LVL041
- ▶ FICON Express or FICON Express2 card with an associated port configured as a FCP Channel in IOCP
- ▶ FCP subchannel unit address xFC or xFD on the FCP Channel must be configured in IOCP to the LPAR where the Configuration Utility for FCP LUN Access Control will run

## ■ Software:

- ▶ Linux for zSeries instance either in a LPAR or running as a guest of VM to run the Configuration Utility for FCP LUN Access Control
- ▶ Linux Configuration Utility for FCP LUN Access Control from Resource Link
- ▶ Libxml 2.5.7, or any higher version from [www.xmlsoft.org](http://www.xmlsoft.org)
- ▶ When running under VM, PTFs for APAR VM63328:
  - UM31347 for z/VM 4.4.0
  - UM31348 for z/VM 5.1.0

# 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.