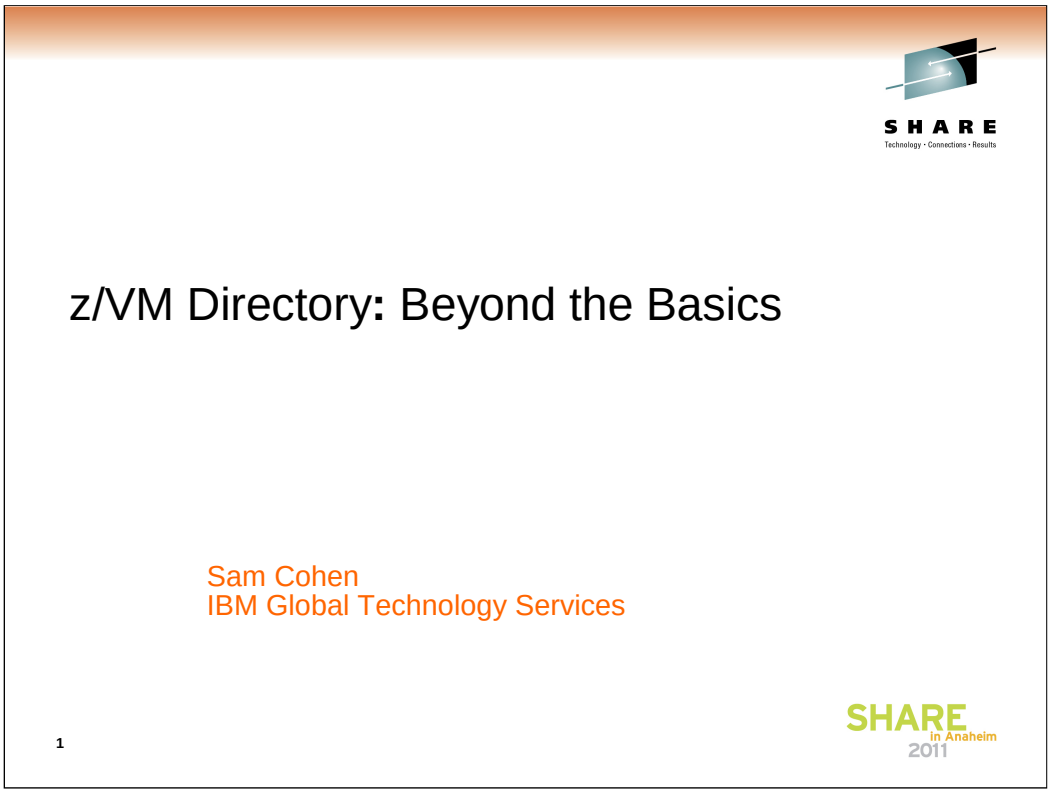

Trademarks

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

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*, AS/400®, e business(logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

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

2

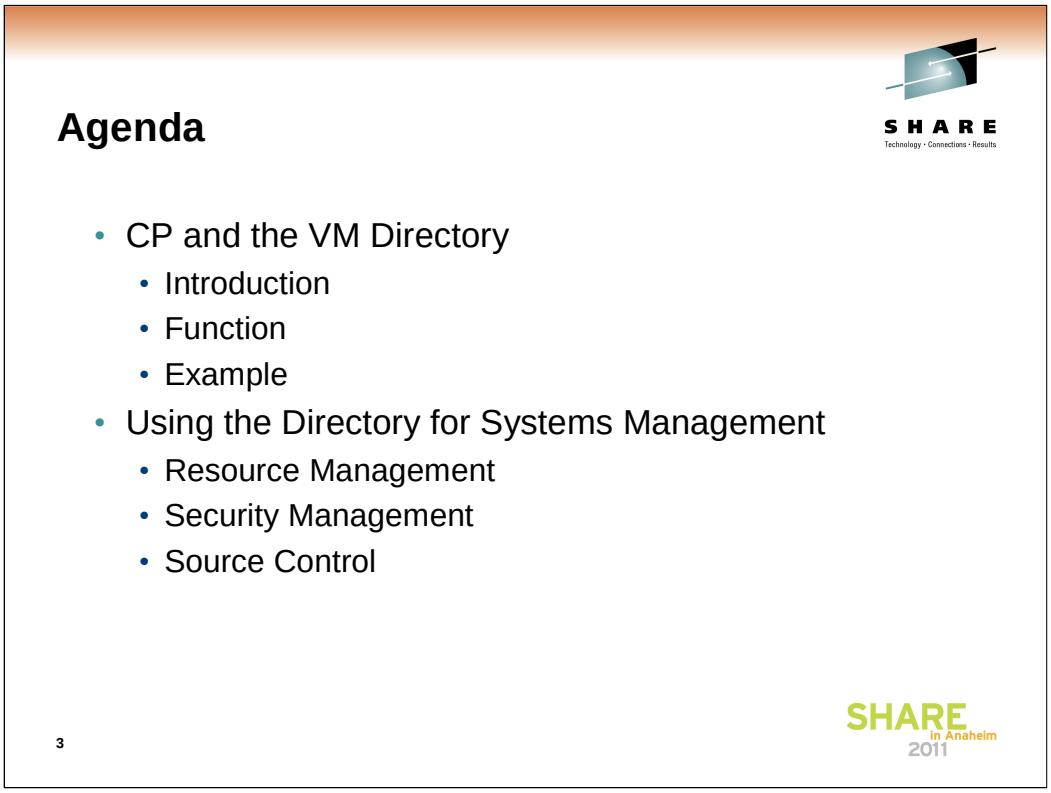

This talk is focused on the role of the VM directory, especially how it can help with Systems Management

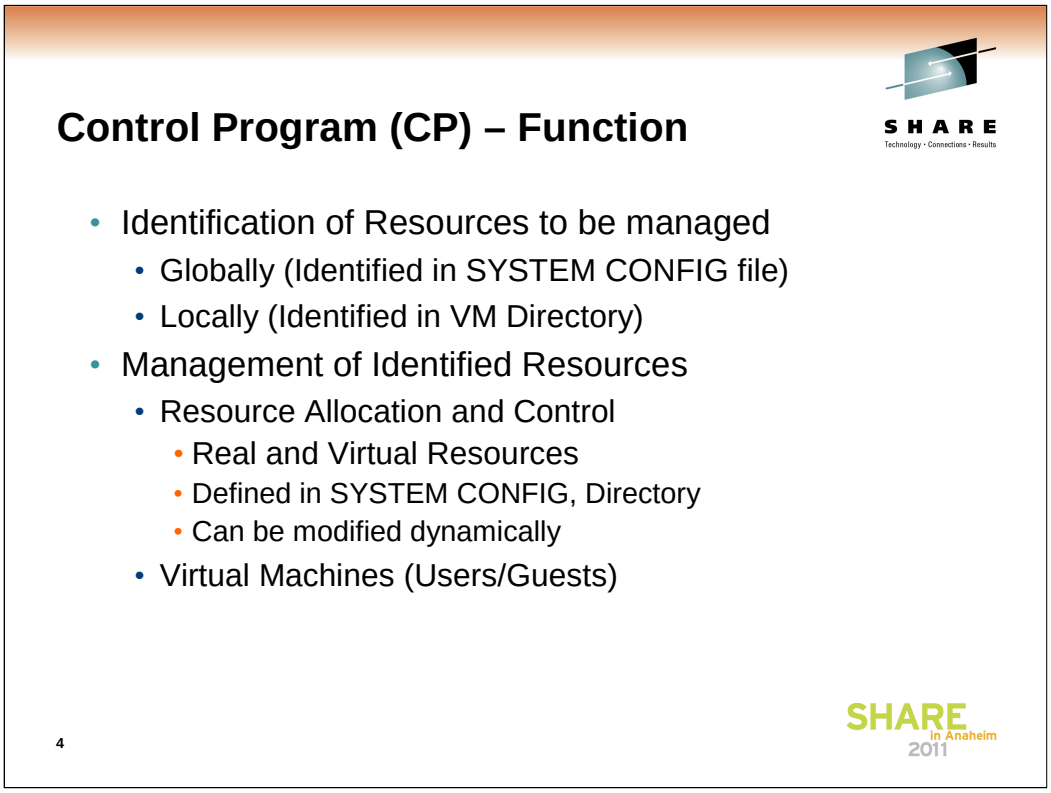

Although a Virtual Machine uses real computer resources, it has no idea what the "real" world looks like, only what it can see. Everything in a virtual machine looks real to the guest operating system.

A Virtual Machine uses real hardware resources, but even with dedicated devices (like a tape drive), the virtual address of the tape drive may or may not be the same as the real address of the tape drive. Hence, a virtual machine only knows virtual hardware that may or may not exist in the real world.

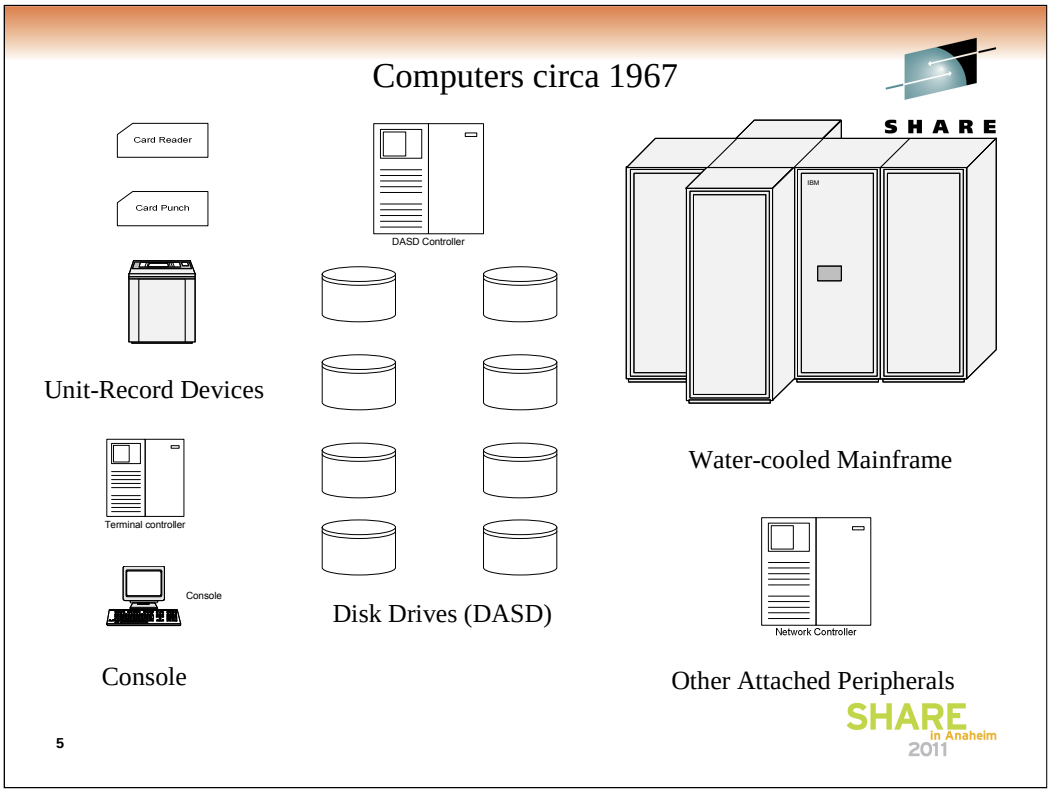

So, what made up a computer in 1967? A computer had:

1) A Water-Cooled mainframe computer, with a certain amount of memory (also known as Central Storage)
2) Unit-Record Devices…card reader, card punch and line printer
3) A console for interfacing with the computer
4) Drum or disk drives (known as Direct Access Storage Devices or DASD)
5) Other devices attached via channel cables

Role of the VM Directory

- Define the characteristics of a virtual machine
  - Based on the definition of a 1967 computer system
- Allocate real and virtual resources
- Isolate each virtual machine from every other virtual machine
  - Unless allowing specific connectivity or visibility to another virtual machine's resources

6

Resource Management in Directory

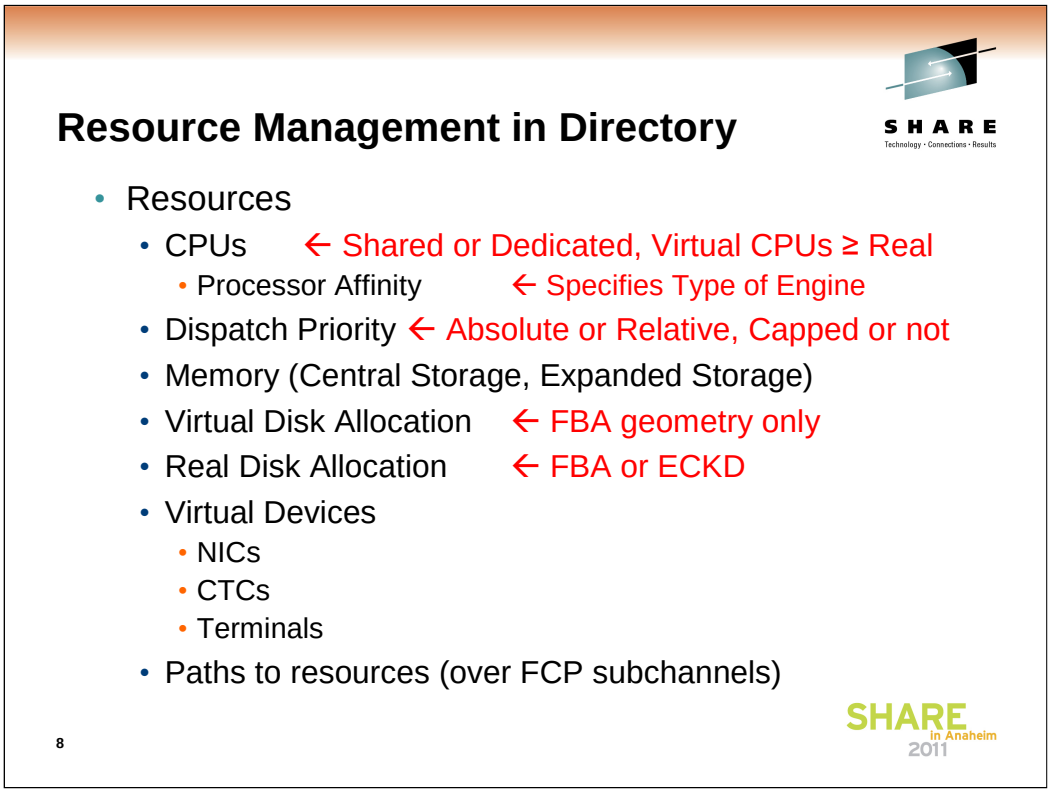

CPUs can be shared or dedicated. You can allocate more virtual processors than actually exist (This would be a case where you are not dedicating CPUs, since there is not a 1:1 relationship.)

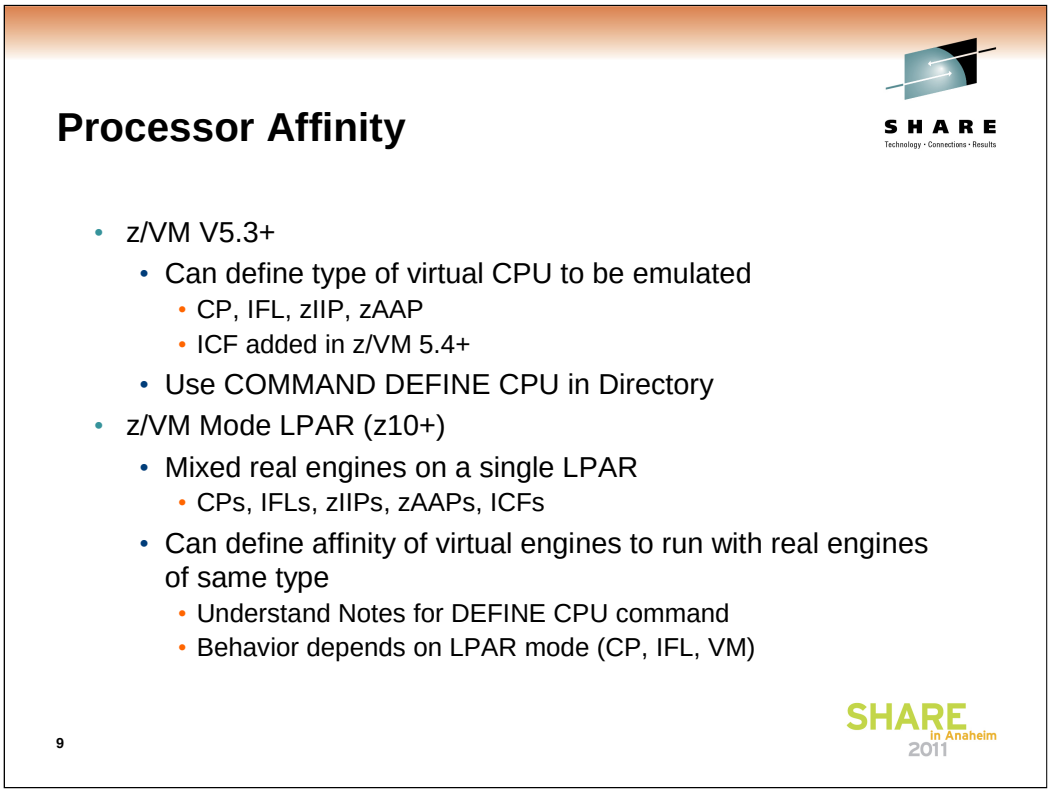

Hardware systems prior to the z10 did not allow you to mix CPs and IFLs in the same LPAR. z/VM mode LPARs allow mixed real engine types to be in a single LPAR. This may allow you to reduce the number of LPARs.

You can issue the DEFINE CPU command within a virtual machine, but placing it in the directory allows the virtual machine environment to be defined prior to the IPL of an operating system.

Real Disk Allocation

- Disk space can be contiguous subsets of physical disk (termed “minidisk”; similar to disk partitions in the x86 world)
- Disk space can exist on “traditional” ECKD devices or Open Systems (SCSI) devices
- When used by z/VM components (CP, CMS, etc), disk space on SCSI is presented as emulated FBA disks (EFBA)
- Minidisks can be shared between multiple virtual machines
- A minidisk cannot span physical disks
- There is NO Volume Table of Contents (VTOC) on a CP-managed disk that is suballocated
  - If cylinder 0 is part of a minidisk definition, the user of that space may put a VTOC on cylinder 0, but it will not be used by CP
- A CP-owned disk never has a VTOC

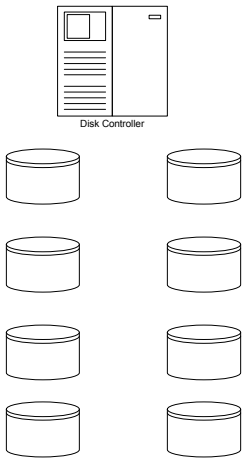

10

DEDICATE vs. MDISK

- A DEDICATEd device is available only for the use of a single Virtual Machine
- A minidisk can be shared between multiple virtual machines
- A minidisk can be an entire physical disk (called a “full-pack” minidisk)
- Allocating less than a full-pack minidisk requires that the volser be unique and the volume be attached to SYSTEM
- Same volser on multiple disks?
  - Use DEDICATE or MDISK with DEVNO specification

11

Security Management in Directory

12

SHARE in Anaheim 2011

Security Management in Directory

- Passwords
  - Logon
  - Minidisk
- CP Commands
- Shared Code (DCSS) Access
- Inter-User Communication (IUCV)
- Restricted Device Access
  - Dedicated Device
  - VCTC/NICDEF

13

SHARE in Anaheim 2011

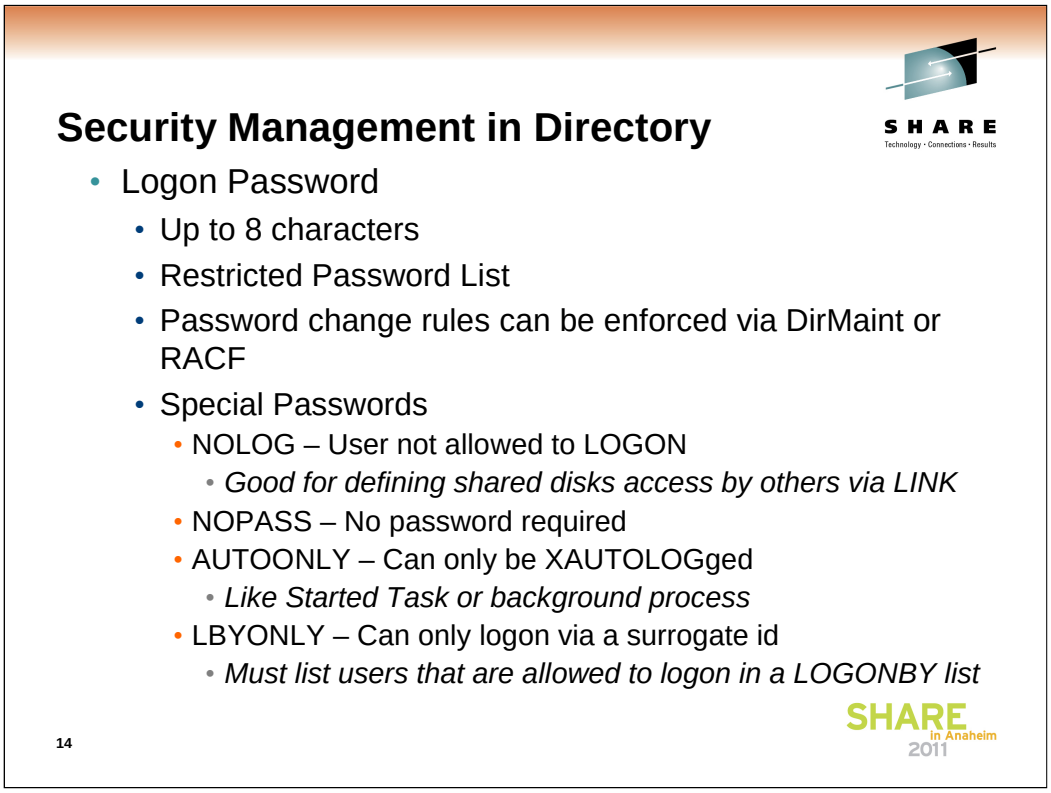

Logon Passwords are in clear text in the Directory Source. If this presents a security concern, you may need to install and use RACF and DirMaint to hide and encrypt logon passwords (they work together).

Security Management in Directory

- Minidisk Passwords
  - Used with CP LINK commands issued by other users
  - Positional (Read, Write, Multi-User)
  - “ALL” = UACC
  - No positional password means that no other user can access unless a LINK statement is in the Directory Entry for the other user

15

SHARE in Anaheim 2011

Security Management in Directory

- CP Commands
  - Commands are divided into classes (A-G)
  - Commands can be given new or additional classes
  - Secondary Console Interface Facility (SCIF)
    - Allows another user to issue commands on behalf of a guest
    - Entered on CONSOLE statement in Directory Entry
  - XAUTOLOG Entry
    - Allows identified user to issue XAUTOLOG command
      - *Like starting a task in z/OS or running a disconnected process in Linux*
  - Dynamic Allocation of Virtual Machine Resources
    - Can usually define/detach virtual resources
    - Can usually detach real resource
      - Detaching of real resources may cause your virtual machine to be reset by CP

16

Security Management in Directory

- Shared Code (DCSS) Access
  - A Discontiguous Shared Segment (DCSS) is similar in concept to LPA – code shared among many virtual machines located at only one place in real memory
  - A DCSS can be identified as a “restricted” segment
  - A virtual machine must be authorized in the VM Directory to access a restricted segment

17

Security Management in Directory

- Inter-User Communication (IUCV)
  - Applications that communicate between virtual machines via IUCV may restrict which machines can communicate with each other
    - Restrictions on both client and server virtual machines
  - Examples of IUCV usage:
    - DITTO/VM client communicating with DITTO/VSE server
    - Sending VSE console commands from a CMS User
    - Linux Terminal Server setup
    - Linux as a recording server (*LOGREC, *ACCOUNT, *SYMPTOM)

18

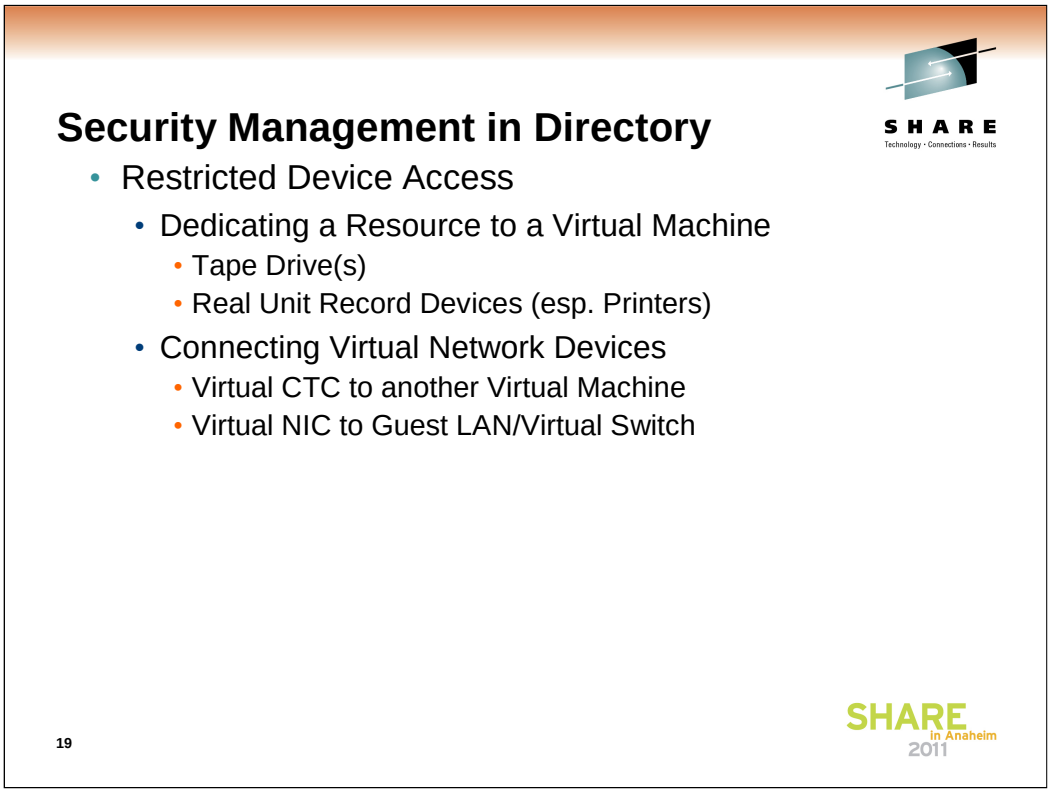

Command Classes restrict the ATTACH command to Class B users. The COUPLE command (for connecting Virtual NICs to a Guest Lan or VSwitch) is a general user command, since the administrator can limit who connects to a Guest Lan or VSwitch

Source Control in Directory

Source Control in Directory

- Can share source directory among multiple VM Systems
- Must compile source on each VM system
- SYSAFFIN statements provide granularity to:
  - Allow a user to exist only on certain systems
  - Allow a user to logon or not logon on certain systems
  - Allow a user to have different resources on different systems
- DirMaint has provision for managing single source

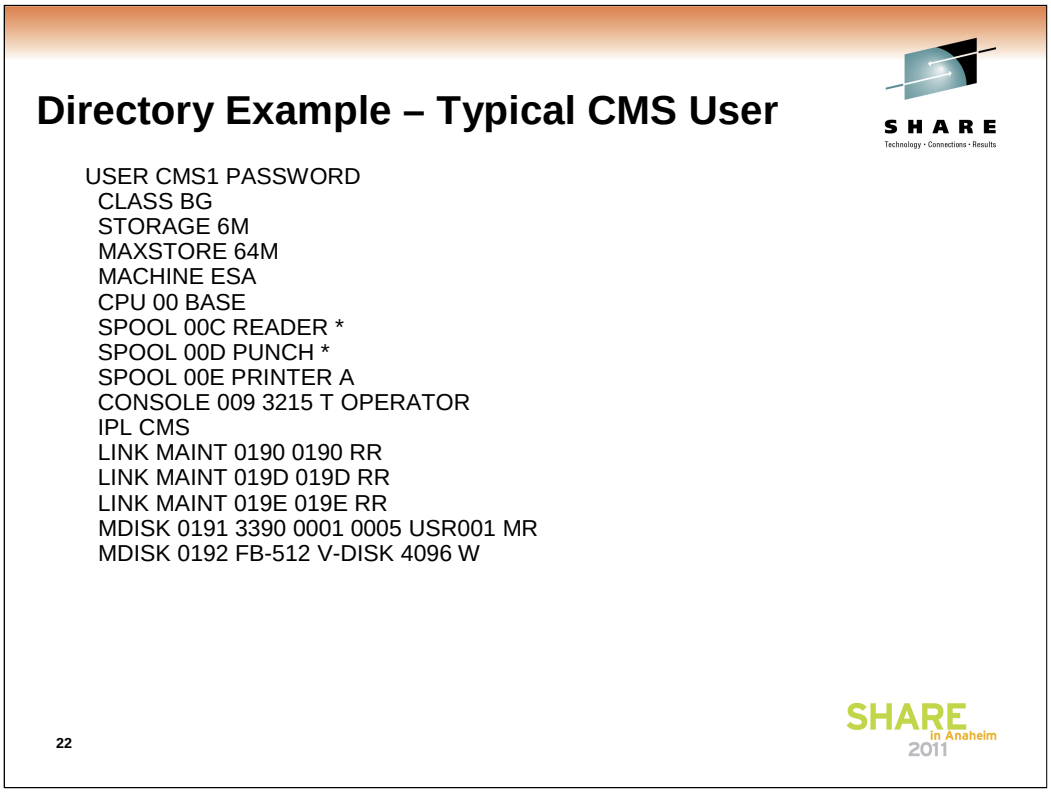

This is an example of a typical virtual machine setup for a user running the CMS operating system

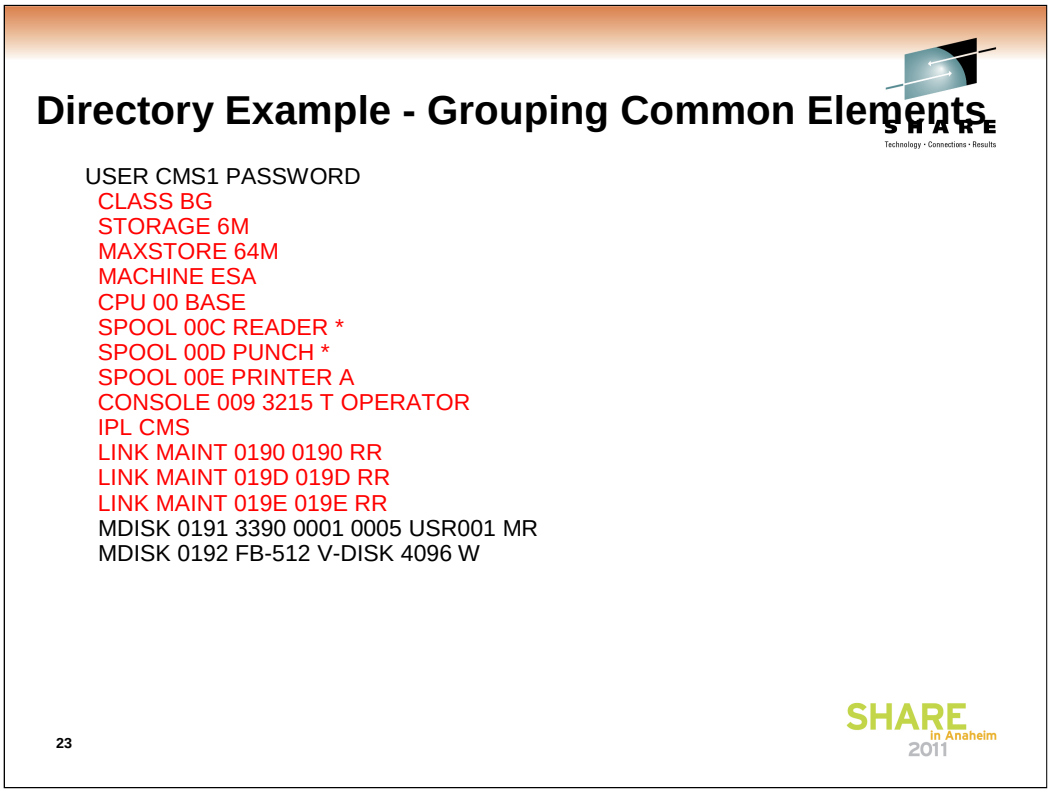

When you identify common elements for a group of virtual machines, they can be grouped into a common “profile”

Directory Example - Easier Directory Management

```
PROFILE CMSUSER
  CLASS BG
  STORAGE 6M
  MAXSTORE 64M
  MACHINE ESA
  CPU 00 BASE
  CPU 01
  SPOOL 00C READER *
  SPOOL 00D PUNCH *
  SPOOL 00E PRINTER A
  CONSOLE 009 3215 T OPERATOR
  IPL CMS
  LINK MAINT 0190 0190 RR
  LINK MAINT 019D 019D RR
  LINK MAINT 019E 019E RR

USER CMS1 PASSWORD
  INCLUDE CMSUSER
  MDISK 0191 3390 0001 0005 USR001 MR
  MDISK 0192 FB-512 V-DISK 4096 W

USER CMS2 PASSWORD
  INCLUDE CMSUSER
  MDISK 0191 3390 0006 0005 USR001 MR
  MDISK 0192 FB-512 V-DISK 4096 W

USER CMS3 PASSWORD
  INCLUDE CMSUSER
  IPL 190
  MDISK 0191 3390 0011 0005 USR001 MR
  MDISK 0192 FB-512 V-DISK 4096 W
```

24

SHARE in Anaheim 2011

By referring to a PROFILE in the virtual machine entry, you can update the profile and that change will be reflected each virtual machine using the profile. Entries in a profile can be overridden by explicit definitions in the virtual machine definition (note USER CMS3)

More Complex Directory Example – Linux Guest

```
PROFILE LNXGUEST
  CLASS BG
  STORAGE 256M
  MAXSTORE 1G
  MACHINE ESA 4
  COMMAND DEF CPU 00 TYPE IFL
  COMMAND DEF CPU 01 TYPE IFL
  SHARE REL 200
  OPTION APVIRT TODENABLE APPLMON
  XAUTOLOG LNXOPER
  LOGONBY LNXMAINT FRANK JIM
  IPL CMS PARM AUTOCR
  SPOOL 00C READER *
  SPOOL 00D PUNCH *
  SPOOL 00E PRINTER A
  CONSOLE 009 3215 T LNXOPER
  NICDEF E000 TYPE QDIO LAN SYSTEM VSWITCH1
  LINK LNXCMN 191 191 RR
  LINK LNXCMN 300 300 RR
  LINK LNXCMN 301 301 RR
```

```
USER LINUX1 LBYONLY
  INCLUDE LNXGUEST
  DEDICATE F000 F020
  DEDICATE F001 F021
  DEDICATE F002 F022
  MDISK 0100 FB-512 V-DISK 32768 W
  MDISK 0101 FB-512 V-DISK 65536 W
  MDISK 0200 3390 DEVNO 1445 MR
  MDISK 0202 3390 0 END SLES9A MR
  MDISK 0203 3390 1339 2000 USER01 MR

USER LNXCMN NOLOG
  MDISK 0191 3390 0001 0100 USER01 MR
  MDISK 0300 3390 0101 1000 USER01 MR
  MDISK 0301 3390 1101 0238 USER01 MR

USER LNXOPER LBYONLY
  INCLUDE CMSUSER
  LOGONBY FRANK JIM
  MDISK 0191 3390 0001 0150 USER02 MR
```

25

SHARE in Anaheim 2011

This example shows a virtual machine setup to run Linux using shared read-only minidisks as well as dedicated disks. These minidisks (contained in the profile) could contain executables (such as /boot and /usr) that can be shared between Linux instances and managed centrally.

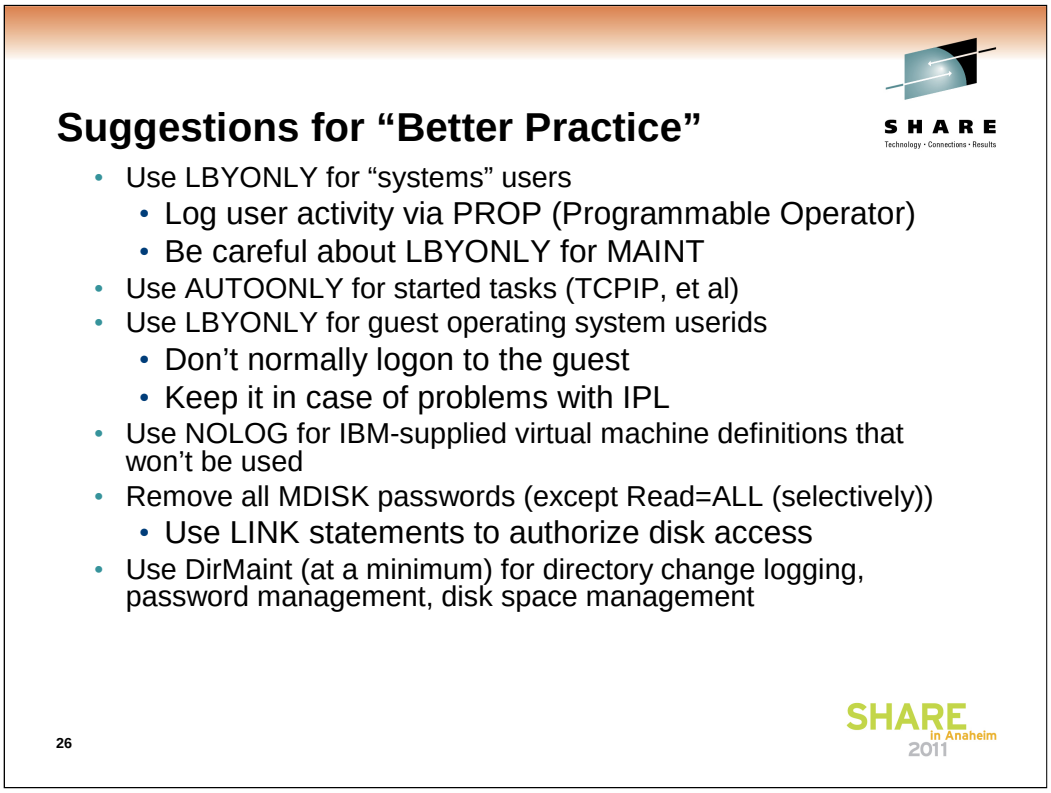

You may want to keep a password for TCPMAINT’s 191 for use with OBEYFILE commands. Be careful if you use LBYONLY with userid MAINT….if all the LOGONBY users have expired passwords, you may not be able to logon to MAINT.

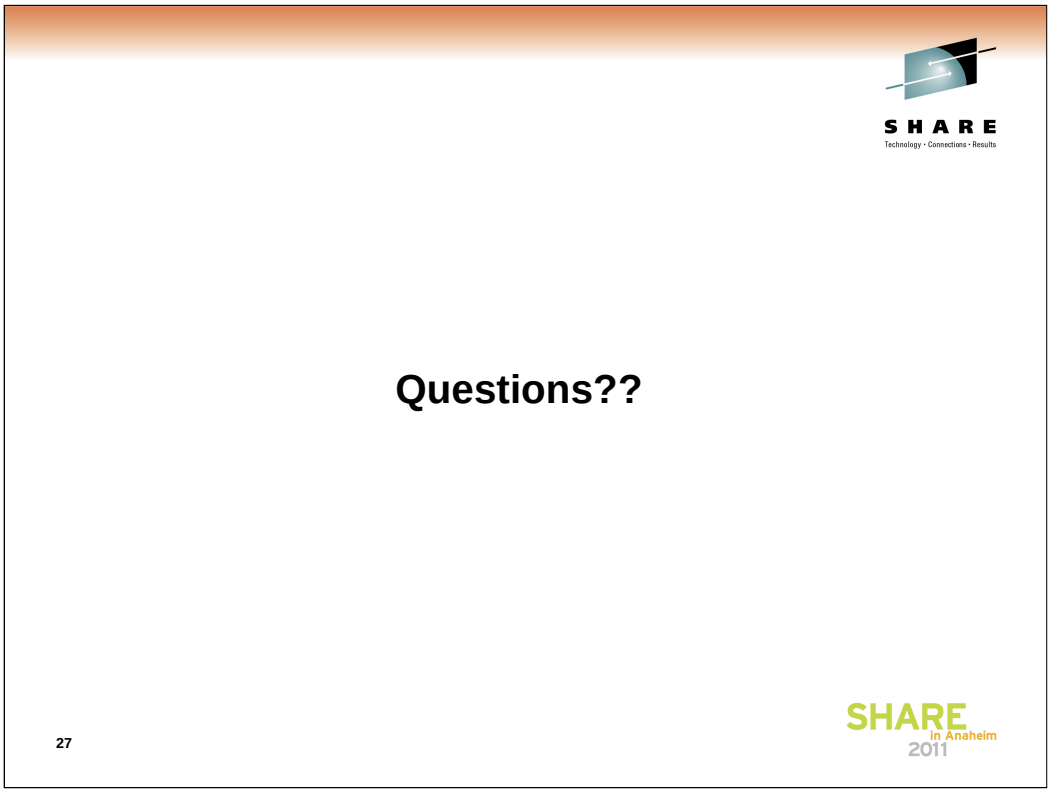
SHARE
Technology · Connections · Results
Questions??
27
SHARE
in Anaheim
2011