

# Hard Drive and File System Configuration in Dual Boot Linux Environment

Presentation #81

David A. Border

Assistant Professor

Bowling Green State University

Bowling Green, Ohio

419-372-2461

Fax 419-372-6066

[dborder@bgnet.bgsu.edu](mailto:dborder@bgnet.bgsu.edu)

# Composition

- Topics assume Intel Architecture Machines
- Physical Media
  - IDE
  - SCSI
- BIOS constraints
- Partitions and Boot Dynamics
- Partitions Topology
- Resident File Systems / Structures
- Linux in Windows

# Physical Media

- ATA
  - Integrating controller onto physical device
- SCSI
  - Bus Interface Controller resides apart from physical device

# ATA

- AT (advanced Technology) Attachment
- Foundation
  - Based on AT bus
    - 40 pin ATA interface
      - 44 pin interface option
      - 68 pin interface option
      - eventual 80 conductor interface option
  - I/O Instruction set
    - standard AT fixed disk interface protocol
  - Standards published by ANSI

# ATA standards (past)

- ATA-1
  - Obsolete initial hard drive interface standard
  - 40 pin standard
  - Equivalent to vanilla IDE
- ATA-2
  - Obsolete hard drive interface standard
  - 40 pin standard plus others
  - Equivalent to EIDE (Enhanced IDE)

# ATA Standards (recent)

- ATA-4 w/ ATAPI definition included
  - DMA capability
  - Various names, incl. Ultra-DMA/33, UDMA
  - Up to 33 MB/sec
- ATA-5
  - DMA capability
  - Various names, incl. Ultra DMA/66
  - Increased noise immunity available
    - 80 wire / 40 pin
    - Up to 66 MB/sec

# ATA Standards (current items)

- ATA-6
  - Ultra ATA/100, UDMA/100
  - 100 MB/sec
- ATA-7
  - Serial ATA
  - Up to 300 MB/sec

# ATA Distinctives

- Controller on Drive
- Physical Hard Drive Features Relevancy
  - Cylinder | Head | Sector (CHS)
    - 16 bit field Cylinders (65536 max)
    - 4 bit field Heads (16 max)
    - 8 bit field Sectors (255 max)
  - Consequence of CHS limits on ATA-1 -> ATA-5
    - 128 GB limitation per drive
  - Circumvented
    - somewhat by LBA (Large Block Addressing)



# ATA Device Topology

- Internal to PC, typically Motherboard based
  - Via 33 MHz IDE Channels
  - Via 66 MHz IDE Channels
- Maximum of two ATA devices per cable
  - Drive 0 (master)
  - Drive 1 (slave)
- Drive number selection
  - Jumper assigned drive number
  - Cable default
    - Mid-connector = drive 0
    - End-connector = drive 1

# ATA Pricing - Performance

- Lowest Priced per Byte PC Hard Drive Device
- Typical RPMs in Mass Consumer Market
  - 5400 RPM, 7200 RPM
- Max. Transfer Rates
  - 33 MBs (UDMA 33 w/ 40 pin connector)
  - 66 MBs (UDMA 66 w/ 80 wire conductor)
  - 100 MBs (ATA100 w/ 80 wire conductor)

# IDE and Linux

- Drive 0 on the Primary Connector
  - /dev/hda
- Drive 0 on the Secondary Connector
  - /dev/hdc
- Drive 1 on the Primary Connector
  - /dev/hdb
- Drive 1 on the Secondary Connector
  - /dev/hdd

# SCSI

- Small Computer System Interface (SCSI)
- Versatile Platform
  - Hard drives
  - Tape drive
  - Scanners
- Versatile Topology
  - External/Internal to PC case
  - Multiple active devices on a single cable
- Cross Platform
  - Intel and Non-Intel based machines

# SCSI Standards (obsolete)

- SCSI 1 (SCSI I)
  - Narrow (8 bits)
    - $5 \text{ MB/s} = 5 \text{ MT/s} * 1 \text{ B/T}$
- SCSI 2 (SCSI II)
  - Wide (2 Bytes)
    - $10 \text{ MB/s} = 5 \text{ MT/s} * 2 \text{ B/T}$
  - Wide and Fast
    - $20 \text{ MB/s} = 10 \text{ MT/s} * 2 \text{ B/T}$

# SCSI Standards (recent)

- SCSI 3 (or just SCSI)
  - At Narrow (8 bits)
    - $20 \text{ MB/s} = 20 \text{ MT/s} * 1 \text{ B/T}$
  - At Wide (UltraSCSI) (16 bits)
    - $40 \text{ MB/s} = 20 \text{ MT/s} * 2\text{B/T}$

# SCSI Standards (current)

- Ultra2 SCSI
  - $80 \text{ MB/s} = 40 \text{ MT/s} * 2 \text{ B/T}$
  - Includes SCA-2 80-pin (Hot swap) and other connector styles.
- Ultra3 SCSI
  - $80 \text{ MB/s} = 80 \text{ MT/s} * 1 \text{ B/T}$  (narrow)
  - $160 \text{ MB/s} = 80 \text{ MT/s} * 2 \text{ B/T}$  (wide) (Ultra160)
- Ultra320
  - $160 \text{ MB/s} = 160 \text{ MT/s} * 1 \text{ B/T}$  (narrow)
  - $320 \text{ MB/s} = 320 \text{ MT/s} * 2 \text{ B/T}$  (wide)

# General SCSI Distinctives

- SCSI devices contain individual device controllers
  - Host Adapter interfaces w/ PC bus
- SCSI Sector Addressing
  - Physical addressing method irrelevant
  - Logical sector addressing employed
- Multiple SCSI Devices Supported Per Adaptor



# Device Ids

- Narrow SCSI
  - 3 bit ID Number
- Wide SCSI
  - 4 bit ID Number
- Boot device normally ID 0
- Host (bus) adapter normally ID 7

# SCSI Device Topology

- Devices Daisy Chained
  - Internal devices on a single ribbon cable
  - External devices on multiple cables
    - two connectors on device
  - Proper termination of transmission line an absolute necessity
    - active
    - passive
- Host adapter may be at end of chain or elsewhere

# SCSI Pricing - Performance

- More Costly than ATA
- Theoretically Faster Performance
- 2 TB Limit versus 128 GB ATA Limit
- Wider Device Product Line
- Hot Swappable Device

# SCSI and Linux

- Arranged in Order of Detection
- Alphabetic Letter Appended to "sd"
- Lowest ID SCSI Device First
  - /dev/sda
- Subsequent SCSI Devices, ordered by ID
  - /dev/sdb
  - /dev/sdc

# BIOS Constraints on Boot Drives

- Both IDE and SCSI affected
- Infamous INT 13H
  - Addresses ALL drive types by CHS via:
    - 10 bit Cylinder field
    - 8 bit Head field
    - 6 bit Sector field
  - Limits Cylinders to 1024 on boot
  - If all fields used to maximum
    - 16515072 sectors addressable
    - Maximum physical size 8.4 GB

# Older BIOS Limits

- A Nasty Intersection of
  - The INT 13H address fields
    - 10 bits C, 8 bits H, 6 bits S (63 max)
  - Original IDE standard address scheme
    - 16 bits C, 4 bits H, 8 bits S (255 max)
- Results
  - OS calls via INT 13H limited to
    - 10 bits C, 6 bits S (63 max)
  - BIOS calls to IDE interface limited to
    - 4 bits H
  - 20 bits, equivalent to a 528.4 MB limit

# Other BIOS Limits

- IDE Large Addressing Scheme
  - Increases Sector specification to 8 bits
  - Keeps the Head specification at 4 bits
  - Keeps the Cylinder specification at 10 bits
    - 22 bits, equivalent to 2.139 GB

# Yet More BIOS Limits

- Large Block Addressing
  - Purpose
    - Provides an IDE sector addressing scheme similar to SCSI while using INT 13H
- Translation Operation
  - Calls to BIOS based on CHS
  - BIOS translated to another address, say LBA
    - Prefers 16 Head, 63 Sector configuration
- BIOS/OS CHS Translation Difficulties
  - Recognized at 2.1 GB, 4.2 GB and elsewhere



# Existence of Partitions

- Division of Physical Hard Drive
  - Used in DOS and Wintel machines
  - Used in Unix boxes
  - Used in OS/2

# DOS and Windows need Partitions

- Existing Advantages
  - Helps DOS/Wintel to overcome limits on larger drives
    - Multiple 2 GB DOS/Wintel Partitions
  - Allows use of different Wintel filesystems
    - FAT (FAT16, FAT32 etc.)
    - NTFS
    - Etc
  - Allows use of special features
    - Hibernation mode

# Linux Makes Use of Partitions

- Existing Advantages
  - Separate filesystems
    - Access speed
    - Organization
      - e.g. Read-only partitions
    - Backup routine
    - Security

# Linux and IDE Partitions

- Numeric Follows Device Name
  - IDE Drive 0 on primary connector, partition 1
    - /dev/hda1
  - IDE Drive 0 on primary connector, partition 4
    - /dev/hda4
  - IDE Drive 0 on secondary connector, partition 2
    - /dev/hdc2

# Linux and SCSI Partitions

- Partitions are Numeric Suffices
  - 1<sup>st</sup> Partition, on 2<sup>nd</sup> detected SCSI Device
    - /dev/sdb1
  - 2<sup>nd</sup> Partition, on 3<sup>rd</sup> detected SCSI Device
    - /dev/sdc2

# Partitions are an OS Advantage

- Multiple Win OSes
- Multiple Non-Win OSes
- Win / Linux dual boot

# Boot Drive

- Contains Master Boot Record (MBR)
  - Contains the boot partition record
  - MBR outside of all other partitions
  - 1 sector in length
  - Begins at CHS: 0, 0, 1
  - Boot machine code
    - Held in first 446 bytes of sector
  - Descriptors for Primary Partitions
    - Held in last 66 bytes of sector
  - Limited to four primary partitions
    - Partition can be extended type
      - » Point to another partition record elsewhere

# MBR Boot Code

- MSDOS FDISK version
  - Not specific to booting Microsoft
  - Attempts to boot "active" partition
    - Loads OS boot from boot sector into memory
    - Code begins at boot sector relative address zero
    - Executes OS boot code
- Boot Code Initial 62 Bytes
  - Contains Reserved BIOS Parameter Block



# MBR Partition Descriptor

- 66 Byte entry
  - Begins at rel. address 01BE hex (446 dec.)
  - 64 Bytes devoted to descriptor
    - 16 Bytes - Partition #1
    - 16 Bytes - Partition #2
    - 16 Bytes - Partition #3
    - 16 Bytes - Partition #4
      - MSDOS Extended partition goes here
  - 2 Byte "signature" follows (0AA55 hex)

# Boot Schemes

- Boot to primary partition - typical
- Boot to extended partition
  - The extended partition must contain the machine boot code

# Over Writing MBR

- Over writing first 466 Bytes
  - Over written by other boot loaders
    - Boot Magic
      - A Wintel based application
    - Lilo
      - A Linux based application
    - Grub
      - A Linux based application
    - Windows installation always clobbers MBR
  - Writes over the Machine Boot Code
    - Typically it is the pre-existing MSDOS FDISK written code

# Saving the MBR with Linux


- Example: IDE boot device without partition table
  - `dd if=/dev/hda of=/bootmbr bs=446 count=1`
- Example: IDE boot device with partition table
  - `dd if=/dev/hda of=/bootmbr bs=512 count=1`

# Boot Loader Precedence

- Commercial Utilities (Boot Magic etc.)
  - If present
    - They should remain in MBR
    - Linux loaders
      - Reside on floppy
      - Reside within first sector of Linux boot partition
        - » MBR conditionally directed to Linux boot partition
  - If not present
    - For Linux boot from Hard Drive
      - Linux loader (lilo, grub) must be in MBR
    - Boot from floppy
      - Retain old MBR contents
      - Linux effectively hidden from view

# Example Boot Loader GUI

## Raster Software Vigo



# Partition Topology

- Primary Partition
  - 1<sup>st</sup> sector is a boot sector
  - Contains partition record for itself
- Extended Partition
  - 1<sup>st</sup> sector is a boot sector
  - Contains partition record current logical partition
  - Contains links to additional logical partitions

# Wintel Partitioning

- Three types acknowledged
  - DOS
    - Wintel limit: one DOS partition/drive
  - EXT DOS
    - Wintel limit: one EXT DOS partition/drive
  - NON DOS
    - Contents invisible to Wintel file management



# Linux Partitioning

- Various Partitions Types Recognized
- Minimum of Two Partitions for Linux OS
  - Linux OS partition
  - Linux Swap partition

# Extended Partition Topology

- Encapsulated Topology
  - Single contiguous disk region
    - Contains all logical disks
- Linear Topology
  - Non-contiguous disk regions
    - Logical disks linked together

# Mixed Filesystem Example

Disk /dev/hda: 255 heads, 63 sectors, 1650 cylinders  
Units = cylinders of 16065 \* 512 bytes

Device	Boot	Start	End	Blocks	Id	System
/dev/hda1		1	765	6144831	b	Win95 FAT32
/dev/hda2	*	766	1650	7108762+	85	Linux extend
/dev/hda5		766	790	200781	82	Linux swap
/dev/hda6		791	1325	4297356	83	Linux
/dev/hda7		1326	1650	2610531	83	Linux

# Partition Filesystems

- Bootable OS Filesystems
  - OS boot sector lies within 8 GB limit
- Data Filesystems
  - Linux
    - ext
    - ext2fs
    - ext3fs
  - Windows 9x, NT, beyond
- Non PC environments
  - MAC

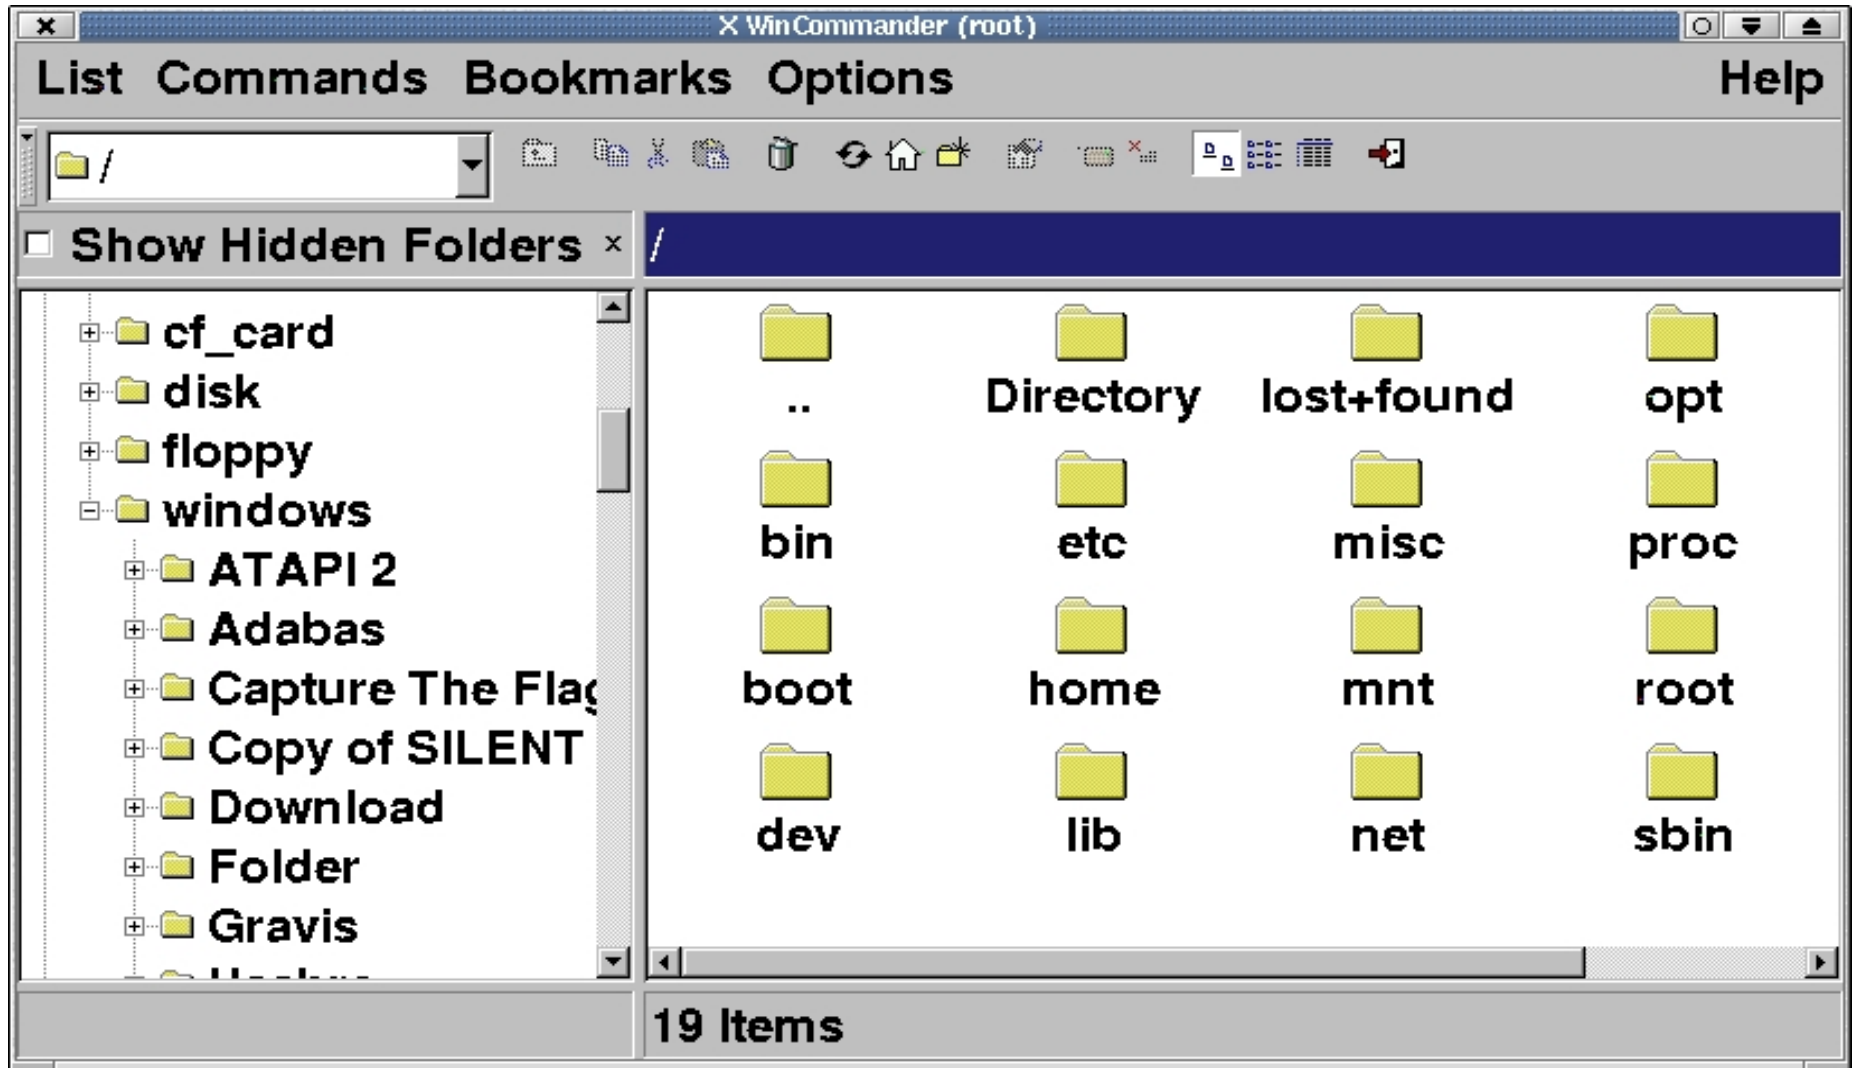
# Partitioning Transparency

- Foreign Partitions Transparent to Linux
  - Windows 9x
  - Windows NT
  - OS/2
  - MAC
- Foreign Partitions Transparent to Wintel
  - None

# Mounting Windows

- Sample Mount Command
  - `mount -t vfat /dev/hda1 /mnt/windows`

# Filesystem View from Linux

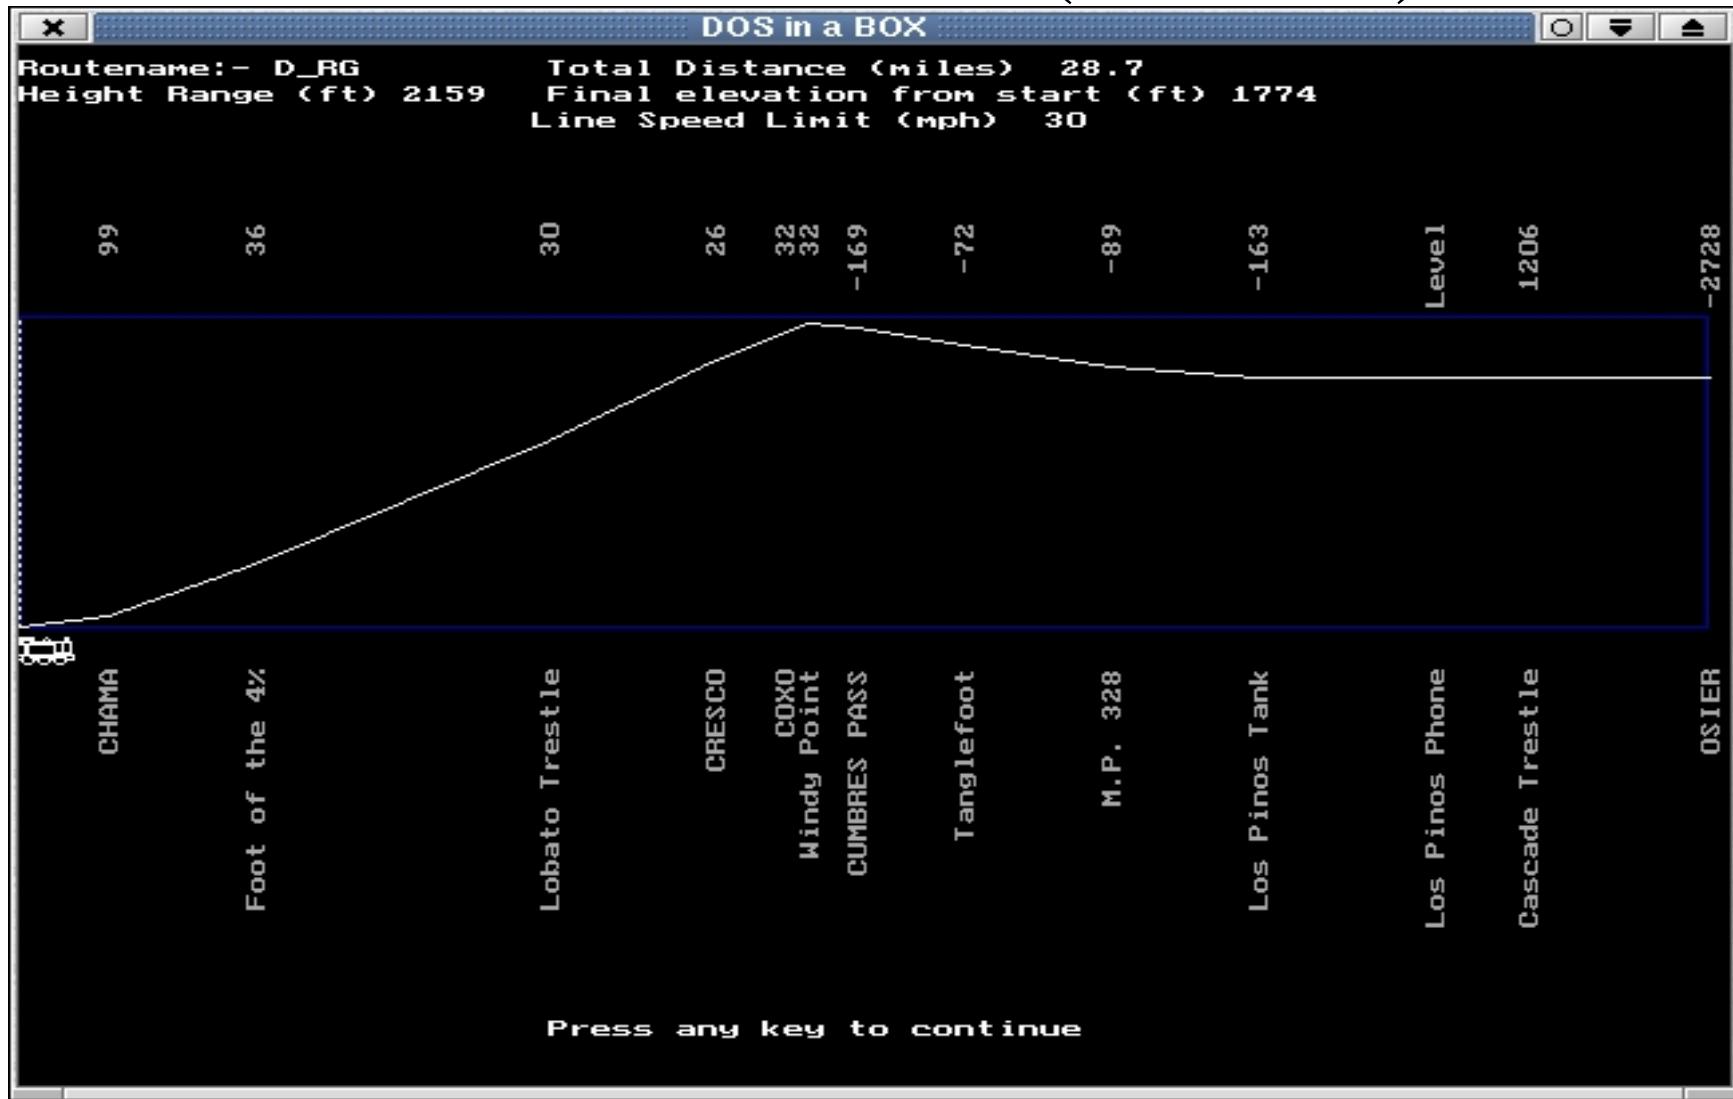


# Sample Linux Features

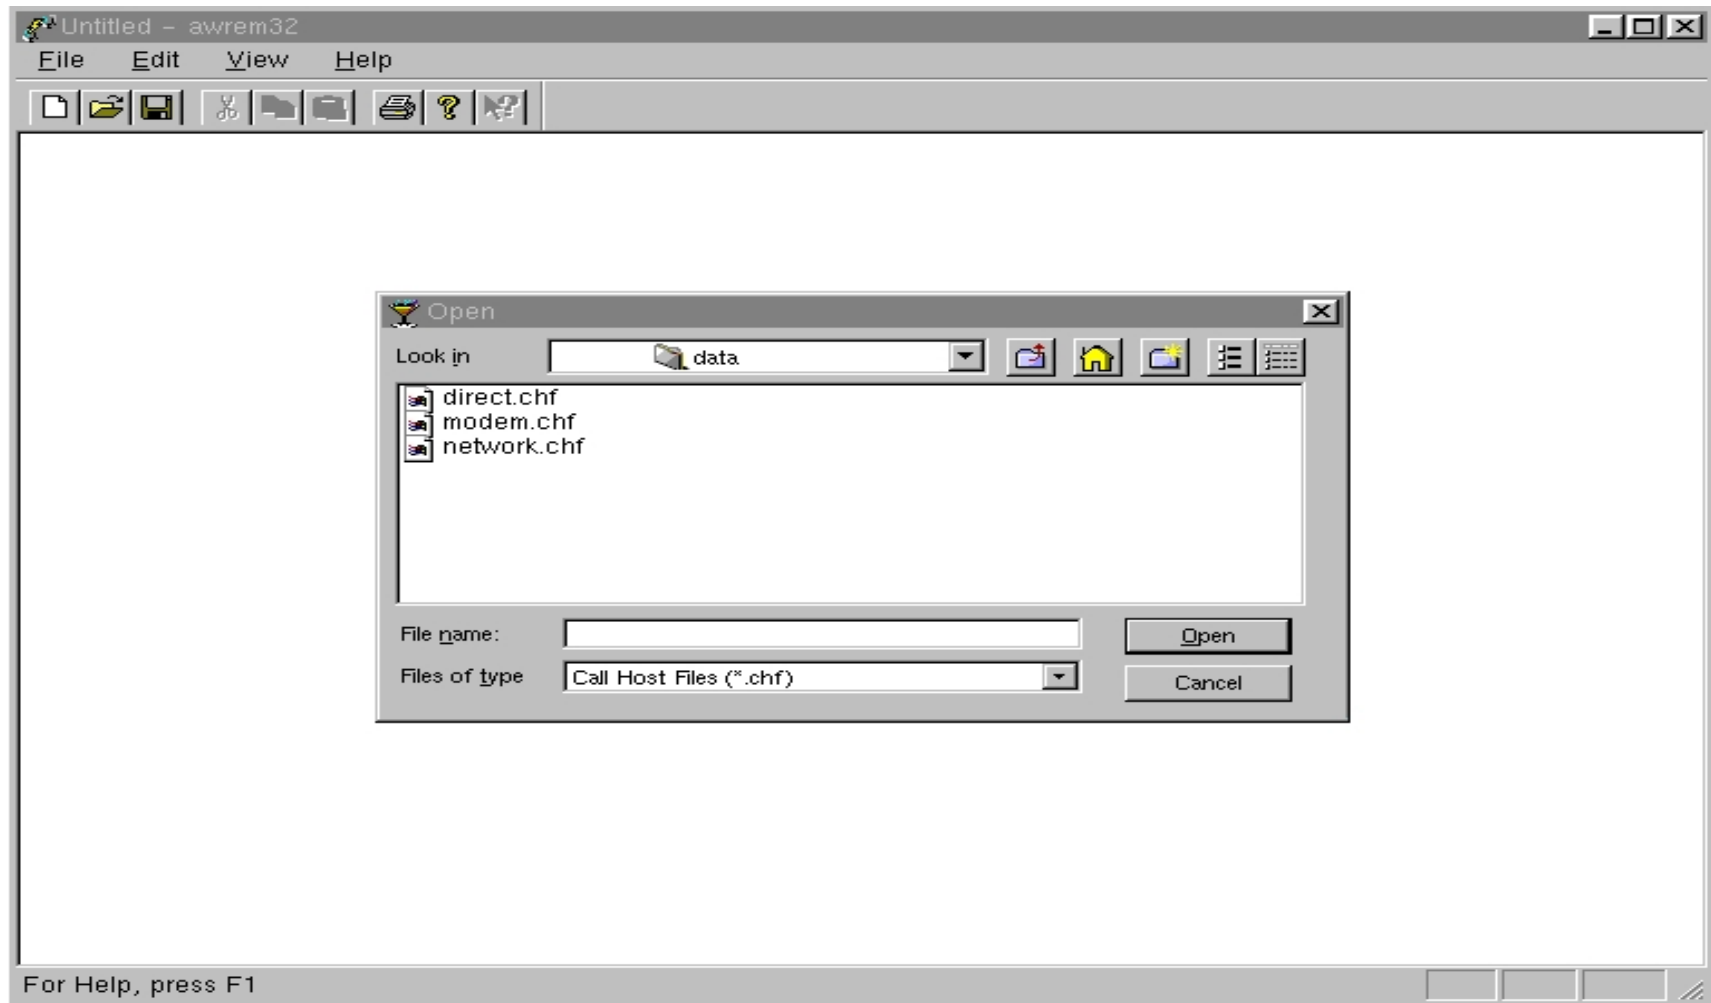
- DOS emulation
  - XDOS
- DOS/Windows emulation
  - Wine
- Office Suite Compatibility
  - StarOffice
- Resource Sharing
  - Samba



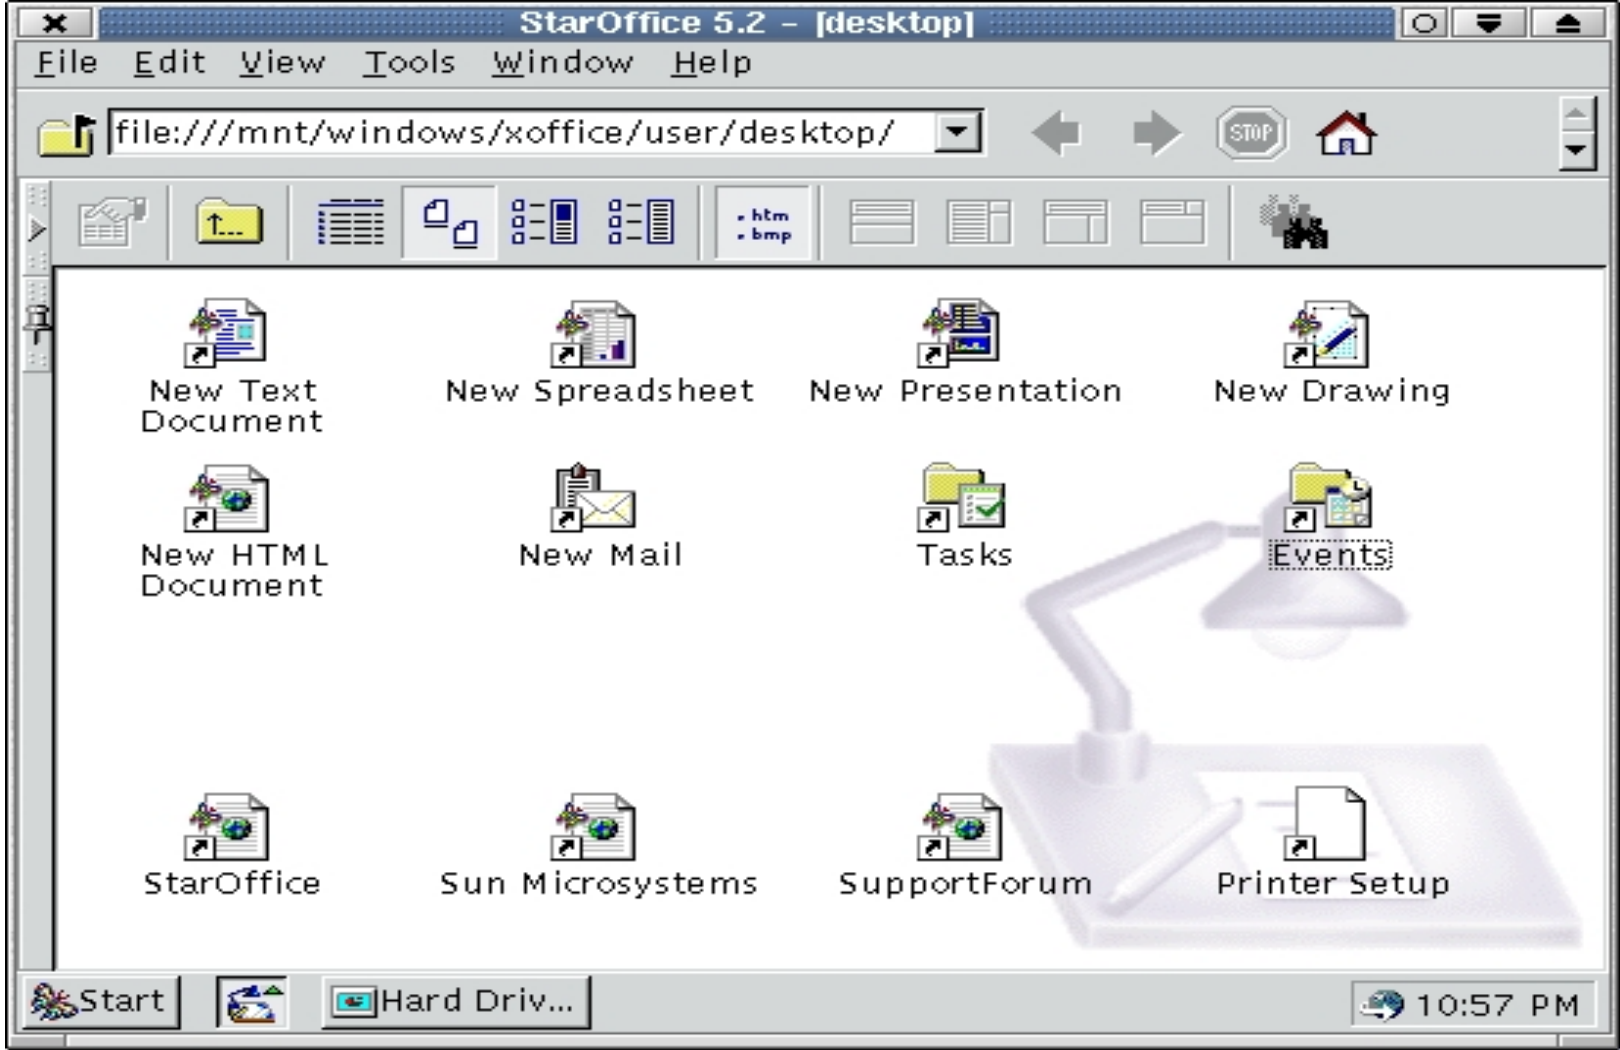
# DOS in a Box (XDOS)



# Wine - Establishing PCanywhere Remote Host Connection



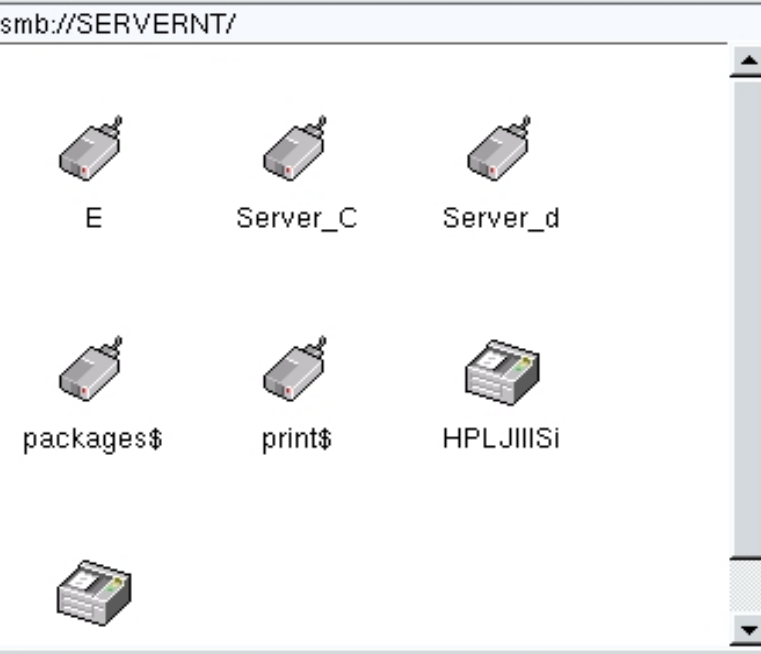
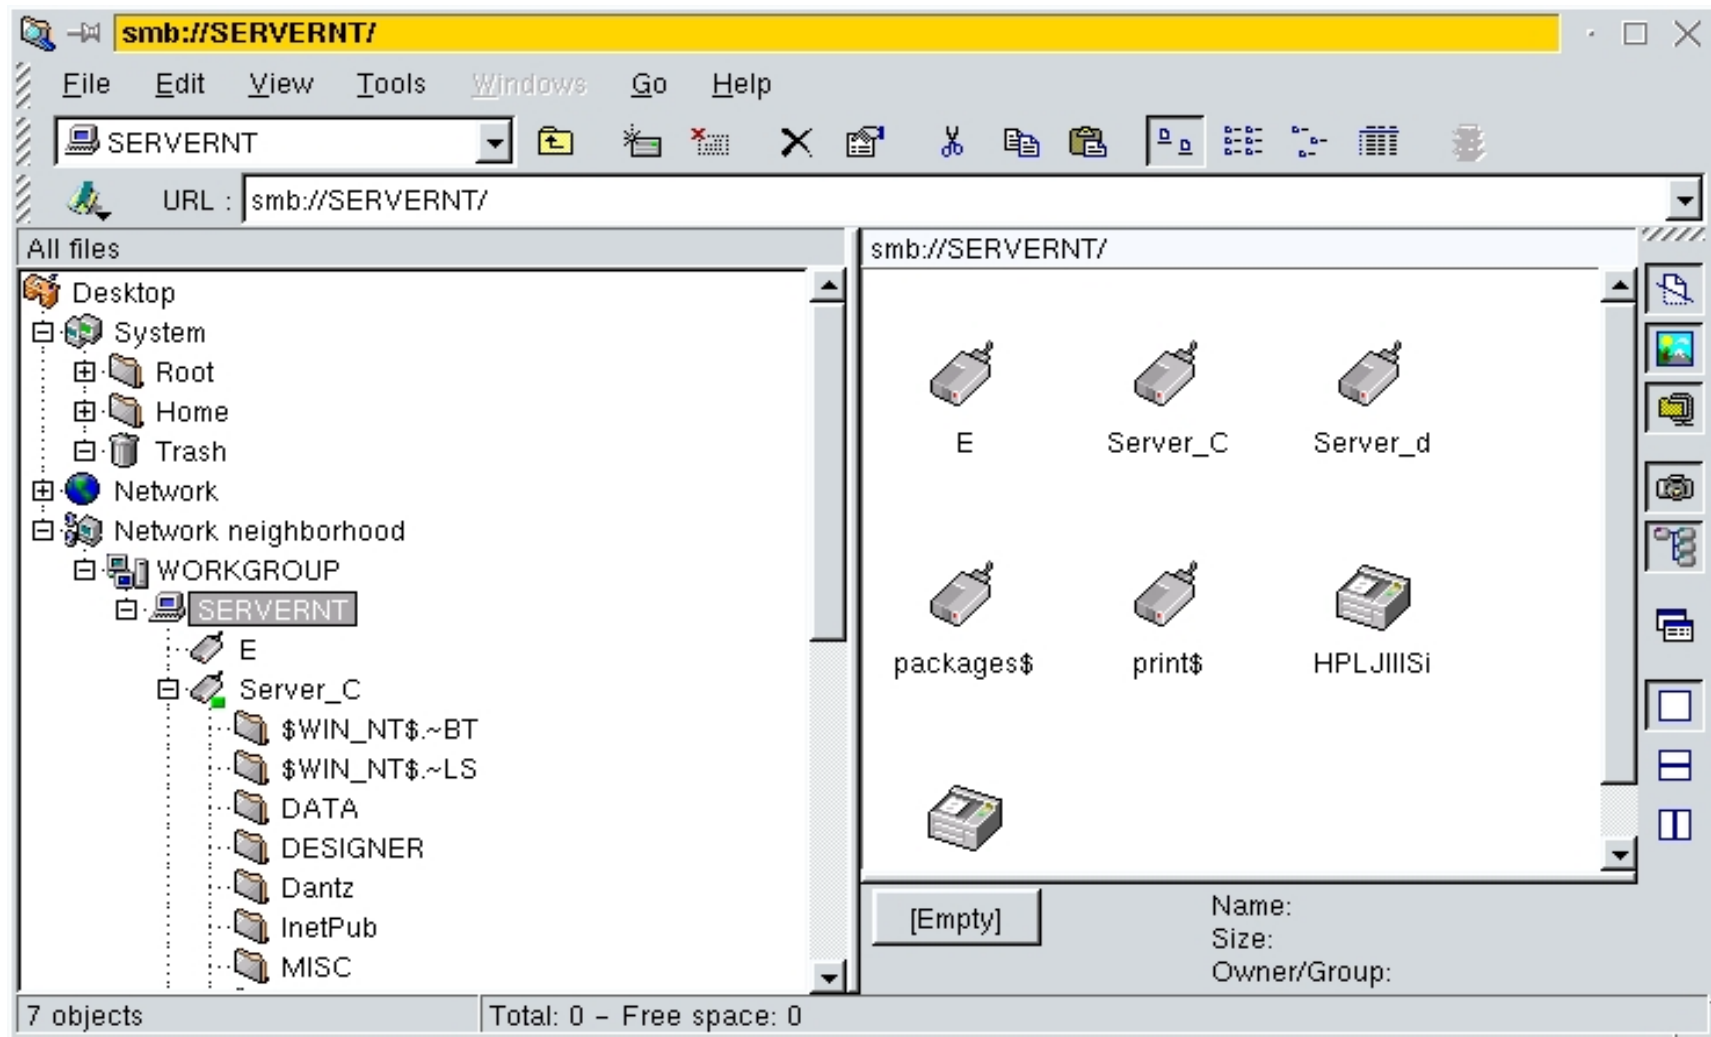
# Star Office



# Samba

- Unix-Windows Resource Connections
  - Allows entry to "Windows Network Neighborhood"
  - Windows penetrates into Linux Resources
  - Linux penetrates into Network Resources
  - Common Internet File System Protocol
    - File and Print Services
    - Authentication and Authorization
    - Name Resolution
    - Service Announcement

# Samba Connections

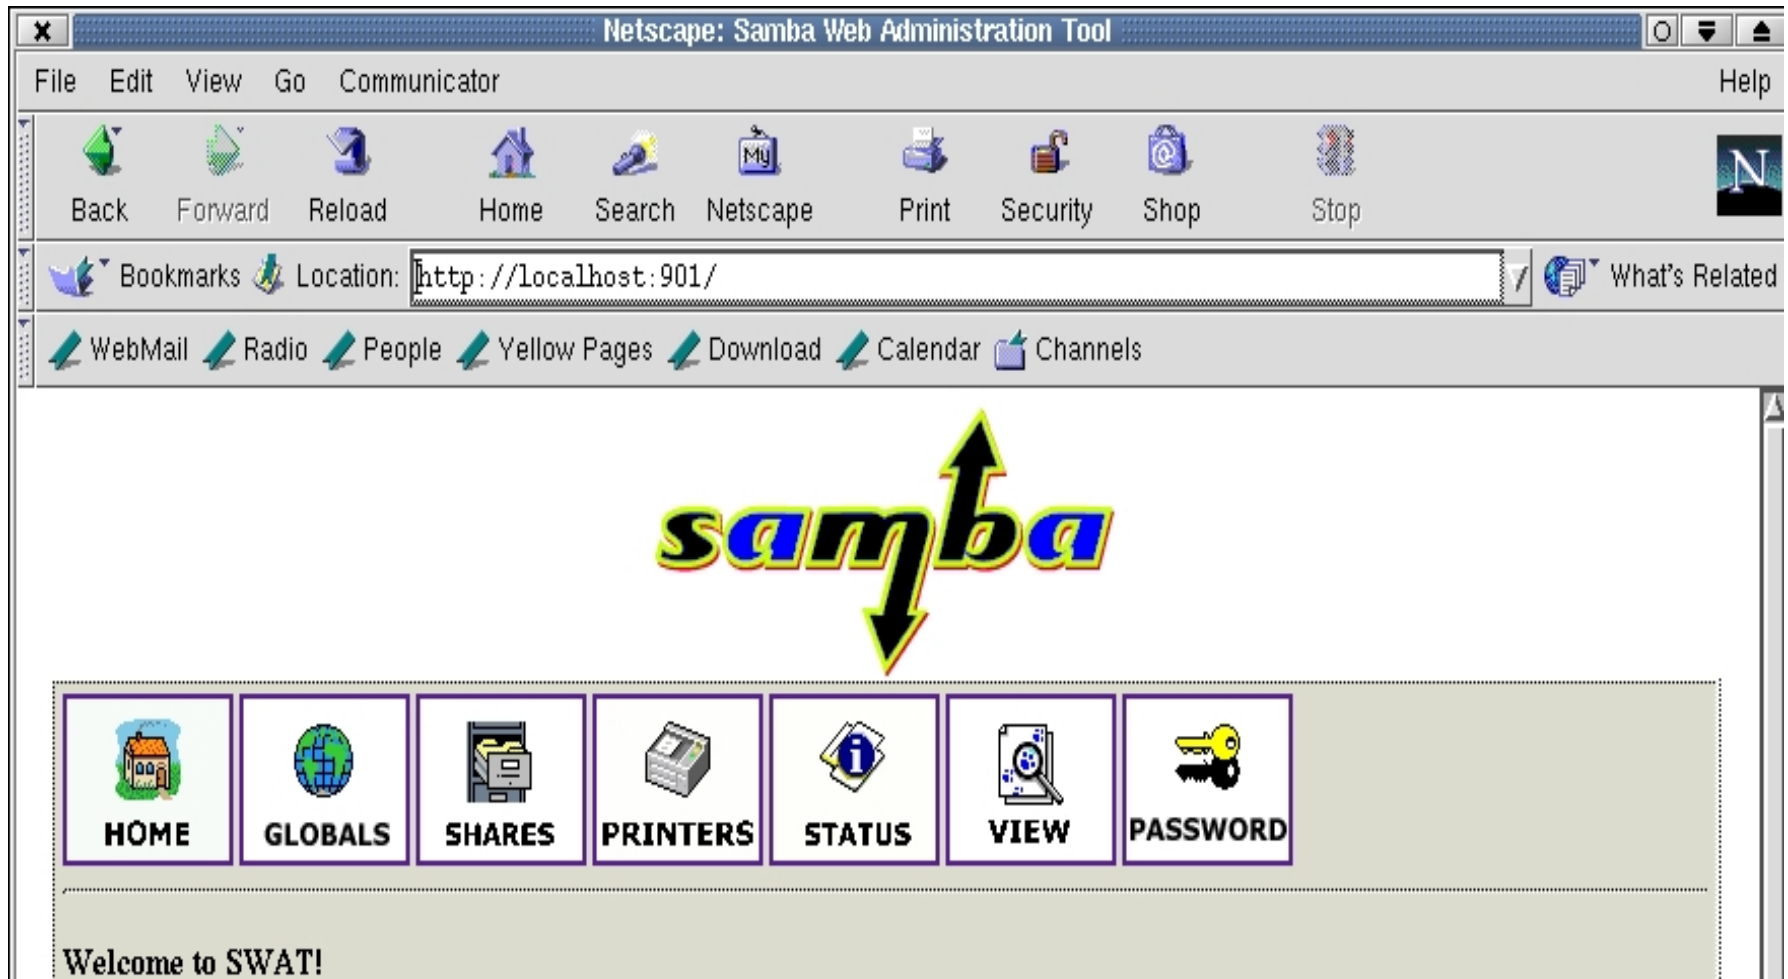


7 objects Total: 0 - Free space: 0

# Linux Samba Commands

- Assumptions
  - A Windows machine resource exists
  - Example
    - Windows machined named TECH
    - Resource has been SHARED
    - Resource has been named "DOCUMENTS"
    - Appropriate permissions are known to you
  - Use smbmount command to connect
    - Smbmount //TECH/DOCUMENTS /mnt/win\_doc

# Samba Administration

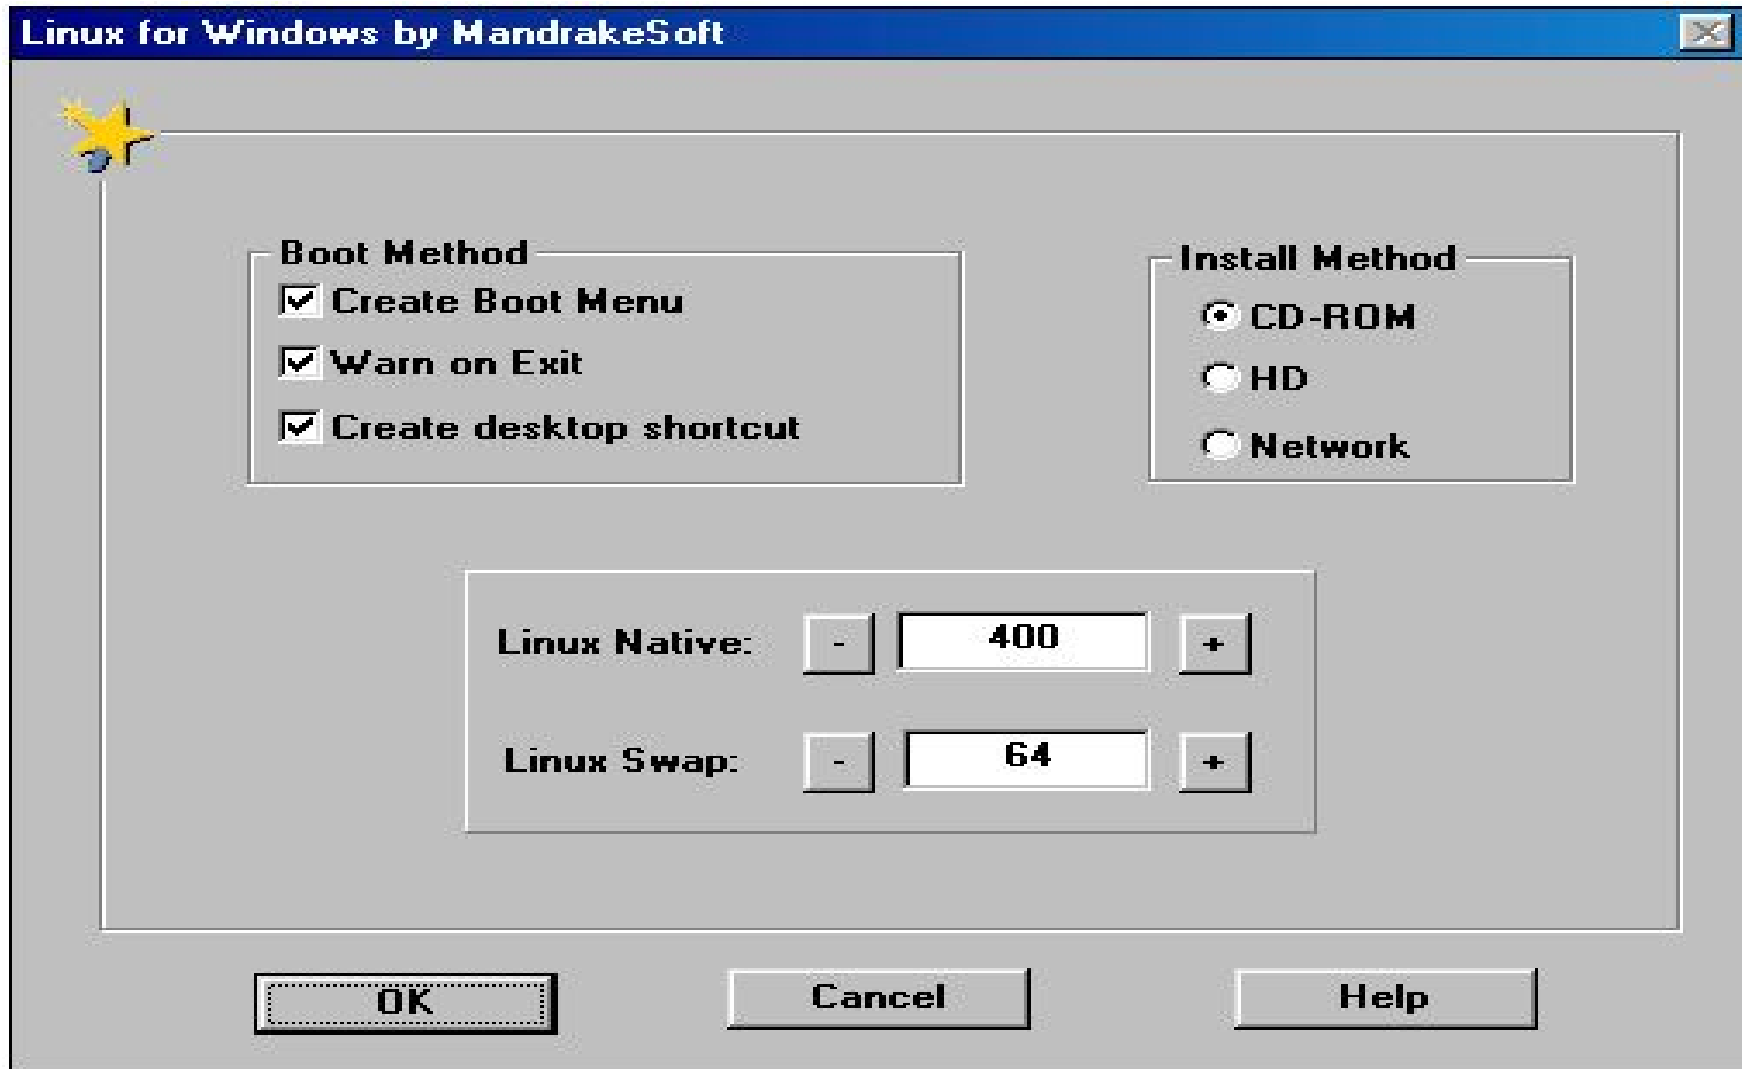


# Linux on Windows

- Avoids Linux on Assigned Partitions
  - Linux OS and Filesystem structures resides as a file on Windows
    - Denoted by a Windows Desktop Icon
- Pseudo-Linux Environment
  - Based on
    - Linux Image with Linux Filesystem and Swap
      - Windows Filesystem remains accessible
- Boot Method Selectable
  - Boot Menu Displayed at PC Startup/Restart
  - Boot into Linux from within Windows OS



# Setup of Linux on Windows



# Boot at PowerOn

Microsoft Windows 98 Startup Menu

---

---

1. Windows
2. linux

Enter a choice: 2

Time remaining: 03

# Boot within Windows



# Linux Browsers

- Netscape
  - Recent Releases include 4.75
- Konqueror
  - As seen in KDE 2.0
- Star Office
- File Managers

# Other Linux / Windows Mixes

- VNC
  - Virtual Network Computing (AT&T)
    - client
    - server
- VMWare
  - Multiple OSes running concurrently on single PC

# Conclusion

- Physical Environment Important
  - IDE and SCSI size, performance and topology
  - BIOS hard drive addressing limits
- Boot Mechanism
  - MBR
  - Boot Load Code
  - Partitions
- Filesystems
  - DOS/Windows/Linux
  - Interoperability