Good news, everyone!

mTCP

User Documentation

Version: 2026-10-10

M. Brutman (mbbrutman@gmail.com)
http://www.brutman.com/mTCP/mTCP.html

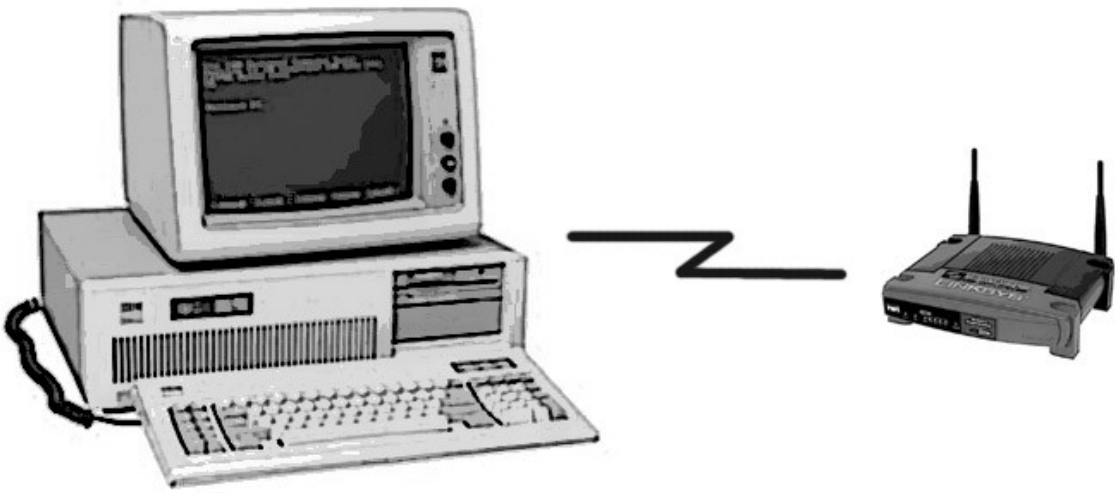


Table of Contents

Introduction and Set up

Introduction....9
- What is mTCP?....9
- Features....9
- Tested machines/environments....10
- Licensing....11
- Packaging....12
  - Binaries....12
  - Documentation....12
- Support and contact information....13
- Disclaimer....13
- And remember ….13

Set up....14
- Introduction....14
- Hardware set up....14
  - Packet drivers....14
  - Loading the packet driver....15
  - Sample packet driver command lines....16
  - Finding packet drivers....17
- Creating the mTCP configuration file....17
  - Creating the minimal configuration file for use with DHCP....18
  - Creating a minimal configuration file using static a networking configuration....19
- DNS resolving and hosts file support....19
  - DNS resolving....19
  - Hostsfile support....20
  - How DNS and the hostsfile work together....20
- Networking related configuration parameters....21
- Other configuration parameters....23
- Setting the TZ environment variable....24
- Advanced set up topics....26
  - The MTU setting....26
  - How the configuration file works....28

mTCP and Virtual Environments....29
- VirtualBox....29
- VMWare Player....29
- QEMU....30
- DOSBox-X....30
- 86Box....30

Checking for updates....31

mTCP Programs

DHCP client....33
- Introduction....33
- Set up instructions....33
- Using DHCP....33
- DHCP lease duration....34
- DHCP failures....35

DHCP without a real-time clock ... 35
Configuration parameters ... 35
DNSTest ... 37
Introduction ... 37
Set up instructions ... 37
Using DNSTest ... 37
Recursive vs. Non-recursive DNS lookups ... 37
FTP client ... 39
Introduction ... 39
Using FTP ... 39
ASCII vs. BINARY transfers ... 40
FTP transfer modes ... 41
DOS filename limitations ... 41
Filenames and Quoting ... 42
Command line editing ... 42
Reading from STDIN ... 42
FTP Performance ... 43
Configuration Parameters ... 43
FTP server ... 45
Introduction ... 45
Special hardware or software requirements ... 45
Before you fire it up ... 45
Quick Start Instructions ... 45
FTP server background information ... 45
Active connections vs. Passive connections ... 45
UserIDs, Anonymous FTP, and permissions ... 46
Configuration parameters ... 46
Required parameters ... 46
Optional parameters ... 47
Choosing the number of clients ... 49
Setting up the password file ... 49
Setting up Anonymous FTP ... 51
DOS filesystem limitations ... 52
Using mTCP FTP server with a firewall ... 53
Choosing a control port to listen on ... 53
Set up your firewall to forward the control port ... 53
Set up your firewall to forward data ports ... 54
Tell mTCP FTP server your public IP address ... 54
Starting, running and ending the FTP server ... 54
Starting ... 55
Running ... 55
Ending ... 56
Notes on specific FTP clients ... 56
Implemented RAW FTP commands ... 57
HTGet ... 58
Introduction ... 58
Using HTGet ... 58
When the server sends content ... 59
Return codes ... 60
Differences with the original WATTCP HTGet ... 61
HTTPServ ... 62
Introduction ... 62
Special hardware or software requirements ... 62

Using HTTPServ....63
Console messages and controls....64
Serving static content....64
Serving pre-compressed static content....65
Protecting files and directories....66
Serving to the public Internet....67
Security, Security, Security....67
Your ISP probably hates private servers....68
Configuring your home firewall....68
Telling people your server name....69
Time synchronization....70
Enabling HTTPS connections....70
Miscellaneous tips....70
Choosing files to cache in memory....70
Log file considerations....71
Using disk caching software....71
HTTPServ Configuration options....71
HTACCESS Files....77
Map Directive....77
FileAuth Directive....78
GeneratedPage Directive....79
DirectoryIndex Directive....80
IndexAllowed Directive....80
Special server functions....80
HTTPServ Status Page....80
HTTPServ Clients Page....81
File cache functions....81
Remote shutdown function....82
Sample files for exposing the special server functions....83
IRCjr....85
Introduction....85
Special hardware or software requirements....85
Set up....85
Starting IRCjr....86
Using IRCjr....86
Screen layout....87
Sending IRC commands....87
Sending CTCP commands....87
IRCjr Hotkeys....88
Session/Window handling....88
Logging sessions....89
Using attributes (bold, reverse, etc) and mIRC color codes....89
Selecting your display mode and size....91
Removing CGA Snow....91
Unicode support....92
Optional Configuration Parameters....92
Netcat....96
Introduction....96
Using Netcat....96
Netcat Options....96
Usage....96
Netcat default operation....97
Redirecting STDIN and STDOUT from/to files....97
Netcat options (details)....98

Keyboard Handling.... 100
Special keys:.... 100
Environment variables.... 101
Examples (and Quick start!).... 101
Network printing.... 102
Limitations.... 102
NetDrive.... 103
Introduction.... 103
DOS installation and usage.... 104
Primary commands.... 105
CONNECT.... 105
STATUS.... 106
DISCONNECT.... 107
Checkpoint commands.... 107
LIST_CP.... 108
MARK_CP.... 108
GOTO_CP.... 108
Passthrough: Running other networking code with NetDrive.... 109
Setting the DOS Verify mode.... 110
Error handling.... 111
Avoid using “Abort”.... 111
CHKDSK errors.... 111
If the server goes away.... 112
Changing the error handling behavior.... 112
Performance tuning: Read-ahead and Large Writes.... 112
Read-ahead.... 112
Finding a good read-ahead value with ND_CACHE.... 114
Large Writes.... 115
Known issues.... 116
DOS issues.... 116
PCjr BIOS keyboard handling.... 116
Ping.... 117
Introduction.... 117
Using Ping.... 117
Return codes.... 118
Notes.... 118
Broadcast ping.... 118
Ping timer resolution.... 118
Ping responses from other mTCP programs.... 118
PktTool.... 119
Introduction.... 119
Using PktTool.... 119
Scanning for packet drivers.... 119
Displaying packet driver statistics.... 120
Listening for packets on the network.... 120
Using the -v option:.... 121
Using the -raw option:.... 121
Return codes.... 124
Services.... 125
Introduction.... 125
Using Services.... 125
Implementation notes.... 126
SNTP (Simple Network Time Protocol).... 127

Introduction....127
Using SNTP....127
Examples....127
Return codes....128
Miscellaneous notes....128
Spdtest....129
Introduction....129
Using Spdtest....129
Telnet....131
Introduction....131
Special hardware or software requirements....131
Using Telnet....131
Telnet protocol features and limitations....132
Uploading and Downloading files....133
Protocols....133
Limitations on file transfers....134
Notes for Linux/Unix....134
Notes for Synchronet....135
Shelling to DOS....135
ANSI Emulation Notes....136
Unicode support....136
Skipping Unicode support....137
BBSes and Unicode....137
Code pages....137
Terminal Types....137
Monochrome Adapter Users....138
Removing CGA Snow....138
Color rendering....138
Strict DOS ANSI.SYS Mode....138
Special note: Telnet BBSes and MUDs....139
[Enter] key handling....139
RLE Graphics....139
Sixel Graphics....140
Advanced set up....140

Advanced Topics

Unicode support....143
Introduction....143
Set up....143
Limitations....144
Debugging problems....146
Introduction....146
Common symptoms and causes....146
Hardware and packet driver problems....147
I/O ports....147
Hardware Interrupts (IRQ)....147
Shared Memory....148
Cabling selection....148
Bad cabling or network connection....148
Packet driver....149
Virtual machine problems....149
Collecting mTCP traces....149

Checksum errors....151
PPP and SLIP Set up....152
Introduction....152
Prerequisites....152
Picking addresses for the connection....152
Option 1: Establish the PPP connection to the Linux machine....153
Option 2: Establish the SLIP connection to the Linux machine....154
Configuring mTCP....156
Routing beyond the Linux machine....157
SLIP connections: Enable Proxy ARP....157
PPP and SLIP: Enable IP forwarding....157
Raspberry Pi WiFi users: Enable promiscuous mode on wlan0....157
Advanced options....158
DOS Idle/Power control....159
Introduction....159
How mTCP minimizes CPU usage....159
Notes on specific power saving programs....159
POWER....159
FDAPM....160
DOSIDLE....160
Disabling power control....160
Connectivity problems with specific machines....161
Firewalls....161
Ethernet card acceleration options....161
Application specific problems....161
Network Address Translation (NAT)....161

NetDrive Server

NetDrive Server....164
Introduction....164
Running the server....164
Flags for serving images....164
Serving standard disk images....166
Serving Journaled disk images....167
Serving Session scoped disk images....167
Disk images....167
Disk image formats (FAT variations)....168
Creating disk images....169
Manipulating disk images with Linux....170
Manipulating disk images with Windows....170
Advanced disk images: Session scoped images and Journaled images....171
Session scoped disk images....171
Journaled disk images....172
Other operations on disk images....174

Bonus Materials!

Make your own floppy disk jacket!....176
Old style customer comment card!....177

Introduction and Set up

Introduction

What is mTCP?

mTCP is a set of TCP/IP applications for personal computers running DOS. The applications are designed to run well on older 16 bit x86 compatible computers. Applications include:

- a DHCP client for auto-configuring your machine for use on your network
- an FTP client for transferring files across your network (or further!)
- an FTP server that supports multiple connected clients, anonymous FTP and more
- HTGet for fetching files or web pages from web servers using HTTP
- HTTPServ for serving files on the web using HTTP 0.9, 1.0 or 1.1
- IRCjr, an Internet Relay Client for chatting on IRC networks
- the Netcat utility for sending and receiving data across the network
- the Ping command which is used for checking basic network connectivity
- the PktTool utility for working with packet drivers and sniffing packets on the network
- SNTP, a Simple Network Time Protocol client for syncing your clock with Internet servers
- SpdTest for figuring out how fast your system/network card is
- a Telnet client for connecting to Unix systems, BBS systems, and anything else that supports Telnet
- NetDrive, a device driver that lets you use network attached storage as a drive letter under DOS

The TCP/IP code is compiled into each application allowing for each application to choose the features that they need. This approach allows for better performance and per-application customization that is not possible with other approaches, such as a DOS TSR.

The TCP/IP code takes a "framework" approach to applications, allowing you to quickly get to the important parts of your application without getting bogged down in the details of initialization, error handling, etc.

This document (“mTCP User Documentation) describes the mTCP programs and is intended for end users. A separate document (“mTCP Developer Documentation”) is intended for people wanting to use the mTCP TCP/IP library for their own projects.

Features

Some of the design goals of mTCP are:

- Configuration flexibility: A lot of features are enabled or disabled at compile time using #defines so that the TCP/IP code can be tailored to each application.
- High performance: Even on the slowest machine dating back to 1981 one can get raw TCP socket performance well in excess of 70KB/sec. (This is dependent upon the Ethernet card being used.)
- Small space requirements: The library is compact without missing key features. Buffer sizes are configurable. Most of the applications run comfortably on a 256KB machine and many will run with less memory.
- Robustness: The library and applications are pretty rigorously tested. Some of the applications have been left running for months at a time with no memory leaks or crashes. The correctness of the TCP/IP and other protocols are checked against a variety of other machines and operating systems.
- Usability: DHCP configuration makes it easy to use on a modern network. Command line options and the configuration file are not complicated.

Some of the more advanced features of the mTCP TCP/IP library are:

- Automatic detection and re-transmit of lost packets using measured round-trip times.
- Support for listening sockets to write server applications.
- DNS resolving (using UDP) and a local hosts file.
- TCP zero window support.
- IP fragment reassembly and automatic sending of UDP fragments.
- Extensive tracing to help find bugs or performance problems.

And some of the current limitations:

- Minimal support for IP and TCP header options
- Only one gateway may be configured
- Only one nameserver may be configured

In general, I implemented a set of features that makes sense for a small machine. There is enough implemented to allow the machine to inter-operate well with a variety of other machines and not violate specifications.

Tested machines/environments

mTCP is not particularly fussy about what it runs on. As long as you have some variant of DOS and a machine (physical or emulated) to run it on, mTCP should work.

General requirements:

- An IBM PC compatible with an 8088 or better CPU. Virtual machines and emulators are included.
- 128KB to 384KB of RAM available depending on the program.
- DOS 2.1 or newer. (DOS 2.0 probably works but it has not been tested.)
- A network card or device that has a “class 1” packet driver. “Class 1” generally means Ethernet but there are packet drivers for Token Ring, ARCNet, SLIP and PPP that emulate Ethernet and thus will work with mTCP.

Here are some of the machines, network cards, and environments that I have tested mTCP with:

Machines:

- IBM PCjr (4.77Mhz 8088 class machine)
- IBM PC, IBM PC XT and IBM Portable PC (4.77Mhz 8088 class machine)
- Compaq Portable (4.77Mhz 8088 class machine)
- IBM PS/2 Model 25 (10Mhz 8086 class machine)
- IBM AT (8Mhz 80286 class machine)
- IBM PS/2 L40SX (25Mhz 80386 class machine)
- Generic 80386-40 clone
- Compaq Deskpro 433i (66Mhz 80486DX-2 class machine)
- Generic Pentium 133 clone
- Aptiva E3N (300Mhz AMD K6-2 class machine)
- Celeron 1100 clone machine

CPUs: 8088, 8086, V20, 80286, 80386, 80486, Pentium, Celeron, etc ...

Video cards: Monochrome, CGA, EGA and VGA

Ethernet cards: The only requirement is a class 1 (Ethernet) packet driver that. Cards that I have personally tested are:

- 3Com 3C503 (8 bit ISA)
- 3Com 3C509 (16 bit ISA)
- Novell or Eagle NE1000 (8 bit ISA)
- Novell or Eagle NE2000 (16 bit ISA)
- Danpex EN-2200T (16 bit ISA, NE2000 clone)
- Davicom DM9008F chipset based cards (16 bit ISA)
- D-Link DE220PCT (16 bit ISA)
- SMC 8216T (16 bit ISA)
- Western Digital/SMC 80x3 series (8 and 16 bit ISA)
- LinkSys LNE100 (PCI)
- 3Com 3C905 (PCI)
- Intel EtherExpress 8/16 (8 and 16 bit ISA)

Ethernet devices not on a card or serial port devices:

- Xircom parallel port attached Pocket Ethernet adapters (PE, PE2 and PE3)
- Compex Paraport parallel port attached Ethernet adapter
- D-Link DE-620 PocketLAN Adapter (parallel port attached)
- Generic 8250/16550 PC UART using SLIP or PPP to a Linux machine.
- “WiFi Modems” using SLIP firmware. These devices use a SLIP connection to the DOS machine and eliminate the need for a Linux machine to do packet routing.
- SCSI Ethernet devices (DaynaPORT SCSI/Link)

Operating Systems:

- IBM PC-DOS and MS-DOS version 2.1 or higher.
- FreeDOS 1.0, 1.1, 1.2 and 1.3. (Future versions are assumed to work.)

Virtual environments:

- VirtualBox
- VMWare
- DOSBox-X
- 86Box
- Qemu

Licensing

Early versions of mTCP were distributed in binary only form. This made sense as I was doing a lot of development and things were churning quickly.

On May 27th 2011 I open-sourced the code and made it available under the terms of the GNU General Public

License version 3. Most versions since then have been open source. Please see the "copying.txt" file in those distributions for the full text of the license that covers those versions.

This release (2026-10-10) is being released with full source code. The source code is available from:

- The mTCP home page: https://www.brutman.com/mTCP/mTCP.html
- GitHub: https://github.com/mbbrutman/mTCP

There are many copies and forks of the code, some of which add useful programs or features. Those changes are the work of their respective authors. As per the GPL Version 3, any code directly using mTCP should also be licensed using GPL 3.

If you are going to use mTCP in a project please considering adding it as a dependency instead of forking it. Adding it as a dependency will help you by keeping your project up to date as mTCP updates.

As the copyright holder of mTCP I am able to offer alternative licenses, including for commercial use. Please contact me if you need that option.

Packaging

mTCP is distributed as a set of binaries, this documentation, and source code. The binaries and documentation are in separate files because it is assumed that you will want to read the documentation on a more modern machine with a PDF viewer.

Binaries

There are two flavors of the binaries available to choose from:

- Standard binaries: These are normal DOS EXE files that consume about 1.2MB on disk.
- UPX compressed binaries: These are the standard binaries compressed using “Ultimate Packer for eXecutables” (http://upx.sourceforge.net/). UPX compressed binaries take less space on disk (744KB) and take just a little longer to get running. These are ideal for systems where space is tight, like on floppy disks.

Once running both sets of binaries are equivalent; there are no performance or feature differences between them.

Documentation

This PDF file is the documentation for mTCP. There may be additional information available at the mTCP home page, but everything that you need to get started should be here.

The documentation used to be provided in TXT files, one for each program and several others for set up and miscellaneous topics. At the time I decided to consolidate the documentation in this PDF file there were 20 TXT files consuming 328KB of space. That became unwieldy. Moving to PDF has enabled me to improve the formatting and usability of the documentation.

Documentation for developers who want to use mTCP for their projects can be found in a separate PDF file (“mTCP Developer Documentation”).

I consider the documentation as important as the code. If you find a problem with the documentation or have a suggestion to make it better please let me know.

Support and contact information

Links to the latest mTCP code and documentation can be found on the home page at https://www.brutman.com/mTCP/mTCP.html. There is a low-traffic mailing list that can be found at https://groups.google.com/g/mtcp. If you have questions, comments, feature requests or bug reports I can be reached directly at mbbrutman@gmail.com. I also check github for issues reported there.

If you are having a problem getting the mTCP applications running please try to give me a good idea of your machine setup, your network configuration, and the symptoms you are experiencing. If I don't know the answer off the top of my head we can use the trace facility and try to figure it out. (See the chapter on Debugging for how to do some basic debugging and generating traces.)

If you have a feature request for mTCP please send me an email with what you are looking for.

Disclaimer

While I make nearly every attempt to ensure that my code is free from bugs and safe to run on your system, I am not perfect and mistakes will happen. You are encouraged to test this code in isolation on a machine that you can afford to reinstall. While that is extreme, it is probably the only prudent thing to do with any new code.

Use this, and anything else that I write at your own risk. I take no liability for anything that might happen to you, your computer, your network, house, descendants, houseplants, etc.

And remember …

Share and enjoy. mTCP has been a fun project for me. I hope it is fun for you too.

If you are enjoying mTCP and are in a position to do so, then consider doing something charitable; there are a lot of great charities out there. Personally, I have never seen an animal shelter that was over-funded.

Set up

Introduction

mTCP is both a TCP/IP stack and a set of applications that use that TCP/IP stack. Before you can use the mTCP applications some set up is needed. You need to:

- Make sure your Ethernet card or networking device is ready to use.
- Tell mTCP where to find a configuration file.
- Place some networking values in the configuration file.
- Add some more values to the configuration file depending on what programs you want to use.

Most of this work is required to be done just once when you first set up mTCP.

Hardware set up

All networking depends on having a connection from one computer to another. There are many ways to make that connection – serial ports, parallel ports, Ethernet, Token Ring, etc. It would be too complex for any program to try to know how to use every possible connection type so some simplifying assumptions are usually made.

Packet drivers

Note: For the purposes of this discussion serial ports using SLIP or PPP packet drivers that emulate Ethernet are equivalent to Ethernet; mTCP is not aware that a serial port being used. The same goes for Token Ring, ARCnet, or other packet drivers that emulate Ethernet.

mTCP is designed to work with Ethernet cards. But instead of having code to talk directly to every possible Ethernet card mTCP has code to work with "packet drivers."

A packet driver is a device driver that provides a simplified programming interface for Ethernet cards. Ethernet cards are often very different from each other; a packet driver is used to hide those differences. With packet drivers networking programs become simpler to write because instead of trying to know the details of every card, programs just need to know how to talk to packet drivers. The full packet driver specification can be found at http://crynwr.com/packet_driver.html.

In particular, mTCP is designed to work with "class 1" packet drivers. Class 1 packet drivers provide the packet driver API for Ethernet. Packet drivers for non-Ethernet hardware such as Token Ring, ARCNet or serial ports can also be considered class 1 if they make the device that they support look like Ethernet. Class 1 packet drivers also exist to convert ODI or NDIS drivers to the packet driver API – these are commonly known as "shims."

Packet drivers that do not provide a class 1 interface such as such as class 3 (native Token Ring) or class 8 (native ArcNet) are not compatible with this version of mTCP. It is possible to modify mTCP to use them but there have been very few requests to do that.

For a general discussion on packet drivers see http://www.brutman.com/Dos_Networking/ .

Loading the packet driver

Packet drivers are usually loaded from the DOS command line as terminate and stay resident utilities. Each packet driver is specific to a model of Ethernet card so every one will be slightly different. However most have some common options that are required.

Software interrupt

A packet driver implements a high level programming interface for working with a network device. The packet driver must be told where to put the entry point to that programming interface so that higher level programs like mTCP applications can find it. The technical term for that entry point is "software interrupt". When you load a packet driver you tell it which software interrupt to install its services at.

Software interrupts are generally written as hexadecimal numbers and are limited in range between 0x60 and 0x7F. You need to pick one that is not in use already but since most DOS software does not use software interrupts 0x60 is usually available.

The next set of parameters are often required to tell the packet driver where the Ethernet card is located in your system

I/O port

Each Ethernet card (or emulated Ethernet card) needs an I/O port address that can be used to send the card data and commands. The I/O port is usually set by a jumper on the card or a software configuration utility. It is necessary to specify the correct I/O port address; if the address is incorrect the packet driver will not be able to control the Ethernet card.

Hardware interrupt

A hardware interrupt is a signal from your Ethernet card to the computer that lets the computer know when the Ethernet card has new data that needs to be processed. When a new packet arrives the Ethernet card will send the signal using the hardware interrupt, waking up the loaded packet driver. The packet driver will then do the work to move the data from the Ethernet card to the main memory of the computer, and possibly signal any higher level software (mTCP) that new data is available.

The hardware interrupt number that the Ethernet card uses is usually set by a jumper on the card or a software configuration utility. The hardware interrupt number is often required on the packet driver command line so that the packet driver knows which hardware interrupt it should be listening to. If you specify the wrong hardware interrupt number on the packet driver command line it will look like you can not receive any data, as the Ethernet card will be signaling that new work is available on one hardware interrupt while the packet driver will be listening on a different interrupt number.

Memory window address

Some Ethernet cards use shared memory to make the transfer of data between the Ethernet card and the computer faster. This shared memory is usually 8KB to 16KB in size and cards that use it are usually much faster that cards that use only an I/O port address. (Note that this is for a RAM chip, not an option ROM.)

If your Ethernet card uses shared memory there should be hardware jumpers or a software configuration utility to set the shared memory address. It is also possible that the shared memory address used by the

hardware is set by the packet driver when the packet driver loads. Please check the packet driver for your card to see which method it uses for setting the shared memory location, if applicable.

Note for ISA bus systems: It is extremely important that any hardware resources that you use are not shared with another active piece of hardware and that they are correct. This is a general ISA bus rule. Using the wrong hardware interrupt number, I/O port address or shared memory region will result in flaky system behavior or crashes.

Sample packet driver command lines

Here are some sample packet driver command lines with an explanation of their parameters.

NE2000 and compatible:

ne2000 is the packet driver for "NE2000" and compatible network cards. (NE2000 is a 16 bit ISA card that is the successor to the NE1000 card, which was a popular 8 bit ISA card.)

```
ne2000 0x60 3 0x300
```

On this command line we are telling the ne2000 packet driver to install its services at software interrupt 0x60, to listen to hardware interrupt 3, and to communicate with the NE2000 card at I/O port 0x300.

SMC / Western Digital 8000 series:

smc_wd is a packet driver for the SMC/Western Digital 8003 series Ethernet cards.

```
smc_wd 0x61 0x03 0x280 0xD000
```

Here the smc_wd packet driver is being told to install its services at software interrupt 0x61, to listen to hardware interrupt 3, and to communicate with the card at I/O port 0x280. The last parameter tells the packet driver to set the card up to use the memory at D000:0000 as a shared memory area for data buffers between the card and the computer.

Xircom Pocket Ethernet 3:

The Xircom PE3 is interesting in that it is an Ethernet device but it attaches to your computer using the standard parallel printer port. While this limits the speed of the Ethernet connection, it does make it easy to add Ethernet to computers that can't use an ISA Ethernet card. More information about the Xircom PE3 Ethernet device can be found online at http://www.brutman.com/Dos_Networking/xircom_pe3.html.

```
pe3pd SINT=62 LPT=2 INT=5 NON
```

On this command line we are telling the pe3pd packet driver to install its services at software interrupt 0x62, to find the Xircom device on the second parallel port, to listen to hardware interrupt 5 (which should be associated with that parallel port), and to not attempt to use the parallel port as a bi-directional parallel port. Note the different way in which the software interrupt is specified and the extra parameter that is specific to this device.

Emulated Ethernet using SLIP on a serial port:

The EtherSLIP packet driver emulates an Ethernet (class 1) packet driver. Under the covers it establishes a Serial Line Internet Protocol (SLIP) connection to another machine, giving you a point-to-point connection to that machine. SLIP connections are much slower than Ethernet connections but you can use them on almost every PC made, as most PCs have at least one serial port.

```
ethersl 0x60 3 0x2f8 9600
```

This tells EtherSLIP (ethersl.com) to install at software interrupt 0x60 on COM2 (IRQ3, IO port 0x2f8) at 9600 bps.

See the PPP and SLIP Set up for more information on SLIP connections.

Finding packet drivers

Packet drivers are specific to models of Ethernet cards. Older ISA bus cards almost always had packet drivers available for them. Many older PCI cards also had packet drivers available too.

If you have the original driver diskette or CD-ROM that came with your Ethernet card that might have the packet driver included with it. Here are some other places to look:

- http://crynwr.com/drivers/ has a collection of open source packet drivers for older cards
- http://www.georgpotthast.de/sioux/packet.htm has some packet drivers for newer Ethernet cards
- http://www.brutman.com/Device_Drivers/Device_Driver_Collection.html for a small, curated collection of packet drivers.

If your card does not have a packet driver available it might have an ODI driver instead. ODI drivers can be made to look like packet drivers using a program called ODIPKT.COM. See http://files.mpoli.fi/unpacked/software/dos/network/odi-shim.zip/ to find a copy of ODIPKT.COM.

If your card only has an NDIS style driver look for DIS_PKT.DOS. See http://ftp.uv.es/pub/msdos/network/dis_pkt.dos/ to find a copy of it. A discussion of how to use it can be found at http://wiki.freedos.org/wiki/index.php/Networking_FreeDOS_-_NDIS_driver_installation (Archive.org)

Creating the mTCP configuration file

The mTCP configuration file is a text file that tells mTCP applications the networking parameters and other configuration values that they need to run. You only have to spend time once to create the file; after that you may change it or add to it but the basics will remain the same.

mTCP finds the configuration file by reading the value of the MTCPCFG environment variable. The environment variable should be set to the full path and filename of the configuration file so that mTCP can find it even when you are not in the directory where you installed mTCP.

Here is an example:

(Assuming the configuration file is located in C:\MTCP\MYCONFIG.TXT)

```
set MTCPCFG=C:\MTCP\MYCONFIG.TXT
```

All of the lines in the mTCP configuration file are either blank, a comment, or a key/value pair used to tell mTCP a configuration parameter that it needs. A comment line starts with the '#' character. Blank lines and comment lines are ignored, while the other lines are used by the mTCP programs.

The next two sections describe the minimum configuration parameters needed for the configuration file. A sample configuration file is also provided with the programs.

Creating the minimal configuration file for use with DHCP

When using DHCP your DHCP server will provide the networking parameters mTCP needs. While your mTCP configuration file may have many configuration parameters, only one is required.

PACKETINT tells mTCP which software interrupt to use to communicate with the packet driver. Example:

```
PACKETINT 0x62
```

That tells mTCP that the packet driver is located at software interrupt 0x62. (The '0x' indicates that it is a hexadecimal number.)

No other configuration parameters are required. When you run the DHCP client it will talk to the DHCP server and add networking related parameters to the configuration file. The exact parameters depend on the DHCP server, but they generally include the assigned IP address, network mask, and gateway IP address.

Assuming you have an mTCP configuration file with PACKETINT specified in it and the MTCPCFG environment variable is pointing at that file, this is what will happen when you run the DHCP client:

- The DHCP program will load and look for an environment variable called MTCPCFG.
- The file pointed at by the MTCPCFG environment variable will be read.
- DHCP will look for the "PACKETINT" line in the file and see that it is supposed to use software interrupt 0x62 to communicate with a loaded packet driver.
- DHCP will make a connection to the packet driver so that it can send and receive packets.
- DHCP will use the packet driver to send and receive packets with your DHCP server. This is part of the process of negotiating a DHCP lease for your machine. (The lease is the right to use an IP address for a limited period of time.)
- DHCP will update the configuration file so that it looks like this when it is done:

```
DHCPVER DHCP Client version May  6 2022
TIMESTAMP ( 1652571761 ) Sat May 14 16:42:41 2022
PACKETINT 0x62
HOSTNAME_ASSIGNED DOSRULES
IPADDR 192.168.2.80
NETMASK 255.255.255.0
GATEWAY 192.168.2.1
NAMESERVER 192.168.2.1
LEASE_TIME 86400
```

Notice that two lines were added at the top and six more lines were added at the bottom. The six lines at the bottom are the networking parameters given to your machine by the DHCP server. The two lines at the top are added so that it is easy to determine that the DHCP program set the networking parameters. If your DHCP

server is configured for it you may also see a DOMAIN parameter added to the file.

While this configuration file will work for many programs, your configuration file will probably have more in it. If you have other configuration parameters, comments or blank lines in the file they will be preserved as is.

Creating a minimal configuration file using static a networking configuration

If you are in an environment without a DHCP server or if you do not want to use DHCP then you can set the networking parameters in the mTCP configuration file by hand. The basic steps are the same except instead of using DHCP to add the networking parameters you will put them in the file when you create it.

Here is a sample file:

```
PACKETINT 0x62
HOSTNAME zoidberg
IPADDR 192.168.1.10
NETMASK 255.255.255.0
GATEWAY 192.168.1.1
NAMESERVER 192.168.1.1
```

You will notice that this sample is similar to what DHCP would have created except that it is missing DHCPVER, TIMESTAMP, and LEASE_TIME - those are specific to the DHCP client. PACKETINT tells mTCP which software interrupt to use to communicate with the packet driver, and it is required for all mTCP configuration files. The remaining lines are networking parameters.

Notes:

- If you choose to use a static configuration please ensure that other machines on the network do not try to use the same address. Most DHCP servers can be told reserve specific address ranges so that they do not hand out an address that will conflict with your statically configured machines.
- As an alternative to using a static configuration, you can also tell most DHCP servers to reserve an address for a specific Ethernet card MAC address. This allows you to use the DHCP client while ensuring you get a stable, reserved IP address.
- To convert a DHCP generated configuration file to a static configuration file just delete the DHCPVER and TIMESTAMP lines at the top of the file. (You should also delete other DHCP related configuration parameters to be complete, but that is not necessary.) Similarly, to convert a static configuration to a DHCP configuration just run the DHCP CLIENT – it will add those lines back and update the networking parameters.

DNS resolving and hosts file support

mTCP supports hostname resolution using the Domain Naming System (DNS) and a local hosts file. The relevant configuration options are HOSTNAME, HOSTNAME_ASSIGNED, DOMAIN, NAMESERVER, NAMESERVER_PREFERRED, and HOSTSFILE and their descriptions can be found in the next section (“Networking related configuration parameters”).

DNS resolving

To enable DNS use the NAMESERVER or NAMESERVER_PREFERRED configuration options.

Most mTCP programs require you to use a recursive nameserver. Using a recursive nameserver keeps the mTCP code simple. The nameserver provided by your router or Internet Service Provider will generally work.

Public DNS services such as Google DNS (8.8.8.8 or 8.8.4.4) and Cloudflare DNS (1.1.1.1) also work well.

Usually all you need is the NAMESERVER configuration option. If for some reason you want to use a different nameserver than what your DHCP server provides you can use the NAMESERVER_PREFERRED configuration option.

The ability to resolve machine names that are on your local network using DNS depends on your DNS and DHCP setup. Most home routers include a DHCP server and a DNS server, so that local machines can be resolved using DNS queries. (The router will update its internal DNS table each time it hands out a DHCP lease.) This can work, even if you do not have a domain name configured on the router.

If resolving local machine names does not work then consider using the hostsfile support described next.

Hostsfile support

A local hosts file allows you to specify IP addresses for specific host names. This takes precedence over DNS and it can be used when DNS is not available or does not know about specific machines. The format is as follows:

```
# Comment line: Any line starting with # is a comment.
#
# Format: <IP_Address> <name> [<alias 1> <alias 2> … <alias n>]
#
# Samples:
192.168.2.30 zoidberg.planetexpress.org zoidberg
192.168.2.31 calculon.badacting.org calculon
10.0.0.2 fry
```

The IP_Address and name are required while the aliases are optional. mTCP searches the file from top to bottom so put your most commonly used names at the start of the file.

Use the HOSTSFILE configuration parameter to tell mTCP where to find the file.

How DNS and the hostsfile work together

Here are the general rules for how mTCP resolves host names to IP addresses:

- If you provide a numerical IP address that will be used as no lookup is required.
- If the name is in the hosts file the address in the hosts file will be used.
- If a nameserver is configured, a DNS query will be sent to the nameserver.

Example: Trying to ping “leela”, which is a machine on your local network:

- mTCP will search the hosts file (if available) for something called “leela.”
- If it is not found or there is no hosts file then a DNS search for “leela” is done.
- If your specified DNS server is your local router, it might return an address for “leela” if it knows about a machine with that name on your network. (Many home routers that issue DHCP leases also include a nameserver that can do this.)
- If you specified a public DNS service then the lookup is going to fail, because “leela” is not a fully qualified name.

Setting a domain name in the configuration changes the search behavior slightly. When the domain name is set it can be added to a hostname to form a fully qualified name to search for.

Example: Your domain name is set to “planetexpress.org” and you ping “leela”:

- mTCP will search the hosts file for “leela” or “leela.planetexpress.org.” (The hosts file is always searched first.)
- If that doesn’t work a DNS search for “leela.planetexpress.org” is done. (A search for “leela” is not done, as it is assumed that if you used just a hostname that it was shorthand for the hostname in your domain.)

If you always use fully qualified names (hostname and domain) there is no ambiguity about what you are looking for, whether you set a domain name or not. Adding the domain name onto a hostname during a search only happens if you do not use a fully qualified name.

Networking related configuration parameters

What follows is a detailed description of the networking related parameters that are valid to use in the mTCP configuration file.

`PACKETINT <x> [, <y>]`

Required. <x> is the software interrupt number that mTCP will use to talk to the packet driver. This has to match the software interrupt number that you told the packet driver to use. It is specified in hexadecimal notation. (e.g.: 0x60)

<y> is an optional secondary software interrupt that the NetDrive device driver can use to provide “pass-through” support, allowing other mTCP programs to run at the same time as NetDrive. See Passthrough: Running other networking code with NetDrive for more information

`IPADDR <x.x.x.x>`

Required; Automatically added or updated when using DHCP.

This is the assigned IP address of the machine. IP addresses need to be unique on your network. If two machines try to share the same IP address neither machine will be able to communicate reliably.

`NETMASK <x.x.x.x>`

Required; Automatically added or updated when using DHCP.

This is the assigned network mask for your local area network. The network mask is used to determine if a target machine for a packet is on the same local area network as your machine or on a remote network. Packets that are destined for a remote network will be sent to the machine specified by the GATEWAY configuration parameter.

`GATEWAY <x.x.x.x>`

Optional, but recommended; Automatically added or updated when using DHCP.

This is the IP address of the machine that will be used to forward packets destined for remote networks. This is usually the IP address of a router. If this is not provided or it is set incorrectly you will only be able to send and receive packets that are on your same local network.

`NAMESERVER <x.x.x.x>`

Optional, but recommended; Automatically added or updated when using DHCP.

This is the IP address of the name server which is responsible for translating human readable machine names (e.g.: www.brutman.com) to IP addresses. If this value is incorrect then converting machine names to IP addresses will not work.

`NAMESERVER_PREFERRED <x.x.x.x>`

Optional. Instructs mTCP to use this nameserver instead of the one provided by your DHCP server. (This is useful if you do not have control over your DHCP server.)

DHCP will not overwrite this preference.

`HOSTNAME <hostname>`

Optional. This is the host name that you are requesting for your machine. The host name can be up to twenty characters and should conform to internet standards. (Avoid punctuation and you should be fine. Hostnames are not case sensitive.) A DHCP server does not have to honor this request. The DHCP assigned name will be provided in the HOSTNAME_ASSIGNED parameter.

If you do not provide a hostname mTCP will unhelpfully provide “DOSRULES” as a default.

`HOSTNAME_ASSIGNED <hostname>`

Automatically added or updated when using DHCP. Unless there is a problem, this will generally match your requested HOSTNAME.

`DOMAIN <domainname>`

Optional; Automatically added or updated when using DHCP.

This is the domain name for your machine. If you are using DHCP and the DHCP server is configured correctly, this will be provided for you. If not, you may provide the domain name.

`HOSTSFILE <x:\path\filename>`

Optional; This is the filename of your hosts file. By convention it is usually “HOSTS.TXT” but you can name it anything that you would like to. You should specify a full path and filename so that the mTCP programs can find it even when you are not in the same sub-directory.

`MTU <x>`

Optional, but recommended. MTU stands for “Maximum Transmission Unit”, which is the largest number

of bytes that can be sent by your networking hardware in a single packet. If not specified the default MTU is set to 576.

For Ethernet this value is typically 1500 bytes, but that value might not work well on mixed networks where other networks have smaller MTU sizes. The default setting of 576 should be safe for almost all environments. SLIP and PPP users may have a smaller value depending on their SLIP or PPP connection.

Please see the section on MTU in "Advanced Topics" for a more detailed discussion.

Other configuration parameters

The mTCP configuration file is used for more than network configuration. Other mTCP applications may read it to look for their configuration settings.

Here are the rules for the content of an mTCP configuration file:

- Blank lines are OK.
- Comments start with a # character in the first position of the line.
- Lines should never exceed 160 characters.
- The DHCP client may add or alter DHCP related networking parameters. All other lines in the file will be preserved as-is.
- Other mTCP programs generally read but do not alter the file. They will skip lines that they do not understand. (In the future it is possible that other mTCP programs besides the DHCP client might alter the file.)

Here is a sample of a more complex mTCP configuration file:

```
DHCPVER DHCP Client version Sep 26 2026
TIMESTAMP ( 1790457167 ) Sat Sep 26 14:12:47 2026

# Parms for my machine
#
packetint 0x60, 0x61
mtu 1500
hostname DOS
dhcp_lease_threshold 1800
dhcp_lease_request_secs 100000

hostsfile c:\mtcp\hosts.txt

# IRCjr parms
#
ircjr_nick Zoidberg
ircjr_user ZoidBerg
ircjr_name John F Zoidberg

# FTP parms
#
ftp_connect_timeout 15
ftp_tcp_buffer 8192
ftp_file_buffer 16384

# DHCP generated settings will appear here
#
```

```
HOSTNAME_ASSIGNED DOS
IPADDR 192.168.2.102
NETMASK 255.255.255.0
GATEWAY 192.168.2.1
NAMESERVER 24.159.193.40
LEASE_TIME 86400
```

Note that the DHCP client adds/replaces the first two lines of the file and the last six lines of the file. The rest of the lines in the file are kept as is for the other applications.

Configuration parameters for each application are documented by that application. The sample file above is just a sample; see the individual programs to see which configuration parameters they support.

Setting the TZ environment variable

Some of the mTCP programs such as HTGet, the HTTP server and the SNTP client need to know the time zone you are located in. Before running these programs you should set the 'TZ' environment variable in DOS. (It's a good practice to have this variable set, so consider adding it to your AUTOEXEC.BAT file.)

The following is reproduced from "The TZ Environment Variable" in the "Watcom C Library Reference" manual. (Some light edits were made to fix some details.)

The TZ Environment Variable

The TZ environment variable is used to establish the local time zone. The value of the variable is used by various time functions to compute times relative to Coordinated Universal Time (UTC) (formerly known as Greenwich Mean Time (GMT)).

The time on the computer should be set to the local time. Use the DOS time command and the DOS date command if the time is not automatically maintained by the computer hardware.

The TZ environment variable can be set (before the program is executed) by using the DOS set command as follows:

```
SET TZ=PST8PDT7
```

The value of the TZ environment variable should be set as follows (spaces are for clarity only):
std offset dst offset , rule

The expanded format is as follows:

stdoffset[dst[offset][,start[/time],end[/time]]]

std, dst three or more letters that are the designation for the standard (*std*) or summer (*dst*) time zone. Only *std* is required. If *dst* is omitted, then summer time does not apply in this locale. Upper- and lowercase letters are allowed. Any characters except for a leading colon (:), digits, comma (,), minus (-), plus (+), and ASCII NUL (\0) are allowed.

offset indicates the value one must add to the local time to arrive at Coordinated Universal Time (UTC). The *offset* has the form:

hh[:mm[:ss]]

The minutes (*mm*) and seconds (*ss*) are optional. The hour (*hh*) is required and may be a single digit. The *offset* following *std* is required. If no *offset* follows *dst*, summer time is assumed to be one hour ahead of standard time. One or more digits may be used; the value is always interpreted as a decimal number. The hour may be between 0 and 24, and the minutes (and seconds) - if present - between 0 and 59. If preceded by a "-", the time zone will be east of the Prime Meridian ; otherwise it will be west (which may be indicated by an optional preceding "+").

rule indicates when to change to and back from summer time. The *rule* has the form:

date/time,date/time

where the first *date* describes when the change from standard to summer time occurs and the second *date* describes when the change back happens. Each *time* field describes when, in current local time, the change to the other time is made.

The format of *date* may be one of the following:

Jn The Julian day n (1 <= n <= 365). Leap days are not counted. That is, in all years - including leap years - February 28 is day 59 and March 1 is day 60. It is impossible to explicitly refer to the occasional February 29.

n The zero-based Julian day (0 <= n <= 365). Leap years are counted, and it is possible to refer to February 29.

Mm.n.d The d'th day (0 <= d <= 6) of week n of month m of the year (1 <= n <= 5, 1 <= m <= 12, where week 5 means "the last d day in month m" which may occur in the fourth or fifth week). Week 1 is the first week in which the d'th day occurs. Day zero is Sunday.

The *time* has the same format as offset except that no leading sign ("+" or "-") is allowed. The default, if time is omitted, is 02:00:00.

Some examples are:

TZ=PST8PDT7,M3.2.0/02:00:00,M11.1.0/02:00:00

Pacific Standard Time is 8 hours earlier than Universal Coordinated Time (UTC). Standard time and daylight saving time both apply to this locale. By default, Pacific Daylight Time (PDT) is one hour ahead of standard time (i.e., PDT7). The rule says that daylight savings time starts on the second Sunday of March at 2:00 AM and ends on the first Sunday of November at 2:00 AM.

TZ=NST3:30NDT1:30

Newfoundland Standard Time is 3 and 1/2 hours earlier than Coordinated Universal Time (UTC). Standard time and daylight saving time both apply to this locale. Newfoundland Daylight Time is 1 and 1/2 hours earlier than Coordinated Universal Time (UTC).

TZ=CET-1:00

Central European Time is 1 hours later than Coordinated Universal Time (UTC). If Daylight saving time is being observed (CEST) the hour offset changes to -2:00.

Note: The Watcom C language run-time library had the correct defaults for the *rule* part of the TZ environment variable when it was written years ago. US legislation has since changed, making it necessary to provide the full *rule* and not assume any defaults. It will undoubtedly change again.

If a timezone is not set the default timezone is Eastern Standard or Eastern Daylight time, using obsolete rules. You probably want to set the TZ variable explicitly.

Some of you might note that the time offset is opposite of what Windows uses. Windows defines the offset has how far behind UTC you are, while the POSIX standard defines the offset as how many hours need to be added to get to UTC. The Watcom C language run-time is using the POSIX standard, hence the difference from Windows. If you are west of GMT the offset is positive and if you are east of GMT the offset is negative.

Advanced set up topics

The MTU setting

In general, for best performance set MTU to 1500. Keep reading if you want to know why …

MTU stands for 'Maximum Transmission Unit' and it tells the TCP/IP stack how big of a packet it can send on the local network. Bigger packets are generally more efficient to work with.

In a world where all networks are built using the same technology the MTU setting would not be needed. But that is not our world. Ethernet predominates, but there are other physical transports out there that you might not be aware of that you are using. Your Ethernet traffic can be transmitted over many different things that don't look like Ethernet before it reaches its final destination.

TCP/IP uses 'fragments' to deal with this problem. If a packet is too big to pass from one network to another the gateway machine will break the packet into fragments suitable for the next network. It is possible that a single packet will be broken into fragments multiple times over the course of its journey. The receiving side is responsible for gathering all of the fragments for an IP packet together and reassembling them.

While this process works it is not ideal. It causes a performance loss as the packets are fragmented by the gateways and a bigger performance loss when the fragments reach their final destination and have to be reassembled. In addition to the performance loss, it requires quite a bit of memory too. And losing a fragment hurts performance even more because all of the other fragments for that packet have to be thrown away.

In the real world fragments are rare. Most modern servers using TCP sockets do some form of path MTU

discovery to determine the largest packet that can be sent without causing fragments to be created. For a TCP socket it works very well and you will have to work hard to try to get a server to send you a packet that needs fragmenting.

On the other hand, that mechanism is not applicable to UDP packets. A server sending a UDP packet has no choice about how to 'chunk' the data to avoid fragmentation; it has to send what the application says to send. Most UDP based programs avoid sending large packets that will be fragmented, but there are exceptions. (NFS and programs sending streaming video data are two notable exceptions.)

mTCP supports the reassembly of fragmented packets, with some restrictions. There is a limit on the total size of the fragmented packet that is set at compile time, and there is a limit to the number of fragments that can be reassembled at the same time. Currently those limits are set at 2048 bytes per packet, and 4 packets at a time. While low, those limits will protect you against a server on a standard Ethernet that goes through a misconfigured gateway that causes fragments to be created. (UDP applications that require larger packets can specify that at compile time.)

Even though mTCP can reassemble fragments there are still all sorts of fragment related dangers floating around out there: In particular, some TCP/IP servers set the 'Don’t Fragment' bit on their packets. If the packet needs fragmentation at a gateway and this bit is set the packet effectively gets dropped at the gateway. mTCP will have no knowledge and no way to get around the dropped packets, other than mysterious timeouts from the application. While the server might be trying to avoid additional overhead caused by fragmentation, it is really doing a disservice by forcing these packets to be dropped. (I'd rather get the data slowly than not get it at all.)

The good news is that you have some control over the problem.

Applications that use TCP sockets will tell the remote side what their MSS (Maximum Segment Size) is. (MSS is based on MTU.) This happens during the socket connect and once the remote side knows your MSS it will not knowingly send you a packet that needs to be fragmented. (Something in the middle might cause fragmentation, so it is still possible. MTU path discovery is supposed to avoid that danger by actively probing all gateways between the server and your machine.)

If you are having problem with dropped packets on a TCP socket, set your MTU to 576. That is guaranteed to work on all gateways and all routers. This is also the default for mTCP if you do not specify an MTU size.

UDP should not have this problem; if an application knowingly sends a large packet it will not turn on the Don’t Fragment bit in the IP header. Doing so would effectively break the application as the data would never be able to get out. On a large UDP packet fragmentation is the only solution; there is no concept of MSS.

For best performance set MTU has large as you can. For Ethernet that value is usually 1500 bytes. This minimizes the overhead of the protocol headers and minimizes interrupt handling.

If you are having strange connection issues you might be bumping into a fragmentation problem. Many times there is no warning message because the packets just simply don't arrive, and the ICMP messages get sent directly to the server, not the mTCP code. If this happens set your MTU to 576 and try again.

SLIP and PPP connections generally have smaller MTU settings which may cause fragmentation. This is unavoidable; consider it part of the joy of using low speed point-to-point links. Your local gateway that is providing the SLIP or PPP connectivity will generate fragments that mTCP will have to reassemble.

How the configuration file works

At the beginning of time, mTCP used environment variables for all of the configuration. This worked well until the IRC client came along and I had to start increasing the size of the DOS environment space. Then came the DHCP client - I couldn't find a way to set environment variables from within the DHCP client and make them stick for the next program without writing a file. At that point the configuration file was born.

When the DHCP client starts up it reads the PACKETINT and HOSTNAME parameters from the configuration file and ignores the rest. (HOSTNAME is optional.) It then tries to communicate with the DHCP server on your network. If it gets an address from the DHCP server it will write a new configuration file. All of the contents of the old configuration file are copied, except for the HOSTNAME_ASSIGNED, IPADDR, NETMASK, GATEWAY, NAMESERVER, and LEASE_TIME parameters which get updated with new values at the end of the file.

Having DHCP write a configuration file has an interesting side effect. It means that at its core, every mTCP application assumes static network addressing. They all read the configuration file to get the network parameters, and they can't tell if the network parameters were set by a human or a machine.

At start-up each application looks for the MTCPCFG environment variable and tries to read the configuration file. It makes one pass to get the networking parameters that it needs to initialize the TCP/IP stack.

If an application has other configuration parameters that it needs to read it will reopen the configuration file and scan for them. The format for a parameter is a single token that represents the parameter name, a space for a delimiter, and then the value of the parameter. The entire rest of the line is assumed to be the parameter. Lines must be under 160 characters in total length.

mTCP and Virtual Environments

If you want to try DOS and mTCP but don't have an old computer to dedicate to it, don't worry! All of this runs fine in virtual and emulated environments. Here are some notes to get you started.

If you find that your virtual machine is chewing up a lot of CPU please read the section on DOS Power control. Using the wrong power/idle management program will reduce your CPU usage but also hurt the usability of mTCP.

VirtualBox

VirtualBox is an open source program that creates a virtual x86 machine within your host operating system. You define the amount of memory, the size of the virtual hard drive, the amount of video memory, networking support, and a few other attributes to create a virtual machine. Then you boot the virtual machine and install DOS on it, just like you would do on a bare machine that you are installing for the first time. The virtual machine emulation takes advantage of modern CPU features making it extremely fast - almost as fast as running directly on the hardware without the host operating system in the way. Virtual machines slow down when doing I/O operations but it is not enough to worry about.

VirtualBox has a variety of network cards that it can emulate and it can make your virtual machine look like it is directly on the network ("bridged"), behind a firewall ("NAT", or Network Address Translation), and a few other options. I use "bridged" mode and the AMD PCNet device, which is the same device used by VMWare Player. If you use a different emulated device be sure to find the correct packet driver for it.

VirtualBox versions 6 and earlier had some limitations which made it occasionally slightly annoying to use. VirtualBox 7 seems to have fixed many of those problems. The only minor problem I know about is that VirtualBox does not emulate the PC speaker, so beeps generated by the programs are silent. I use VirtualBox for much of my testing and it generally works well.

VirtualBox main site	http://virtualbox.org/
AMD PCNet packet driver	(Search for the AMD PCNet PCNTPK.COM packet driver)

VMWare Player

VMWare Player is a closed source (but free) virtual x86 machine that is very similar in operation to VirtualBox. (VMWare predates VirtualBox by quite a bit.) You define your virtual machine, install DOS, and then run it side by side with your host operating system.

The networking works basically the same way as it does in VirtualBox - "bridged" makes your virtual machine look like it is directly connected to the LAN. "NAT" (Network Address Translation) makes it look like your virtual machine is hiding behind a firewall on your host operating system. "Host-only" is the most limited, and it allows you to communicate only with the host operating system. VMWare emulates some form of AMD PCNet Ethernet device.

I am using VMWare Workstation 16 Player Windows 10. Everything works perfectly, including the PC speaker emulation. Current versions require you to create an account at the Broadcom Support Portal, so I'm sticking with my slightly older version for now.

VMWare Player main site	https://www.vmware.com/products/desktop-hypervisor
AMD PCNet packet driver	(Search for the AMD PCNet PCNTPK.COM packet driver)

QEMU

I don't use QEMU regularly but I have done some basic testing with it and it works well.
QEMU is a bit more complicated to get started because it is command line driven. Here is a minimal command line that starts a floppy only system with AMD PCNet emulation:

```
qemu-system-i386 -fda Celeron_boot -net user -net nic,model=ne2000 -display curses
```

The floppy drive image ("Celeron_boot") had DOS 6.3, the PCNet packet driver, and just enough files and programs to test mTCP. (I used the curses display option to avoid having to use X-windows.)

QEMU home page	https://www.qemu.org/
AMD PCNet packet driver	(Search for the NE2000.COM packet driver.)

DOSBox-X

DOSBox-X is based on DOSBox, which is a DOS emulator that also emulates some specific graphics and sound cards. Unlike VirtualBox, VMWare Player or QEMU it is not a full virtual machine. This provides some nice opportunities for integration with the host operating system. DOSBox-X takes the DOSBox project and adds more features including networking support, making it usable for mTCP. It is a good way to get DOS running on a modern machine without having to fire up a full virtual machine.

While I don't have a lot of experience with DOSBox-X, I trusted it enough to use it to demonstrate mTCP for two days at a recent event.

DOSBox-X main site	https://dosbox-x.com/
NE2000 packet driver for use with DOSBox	(Search for the NE2000.COM packet driver.)

86Box

86Box is a low level x86 system emulator that emulates machines ranging from the first IBM PC to later PCI based systems. It focuses on accuracy and customization, making it a great test environment for performance and compatibility testing. I've used 86Box to test the performance of the CRC16 and Fletcher-32 Checksum routines in NetDrive on an emulated IBM PC AT, and to test EGA specific portions of Telnet.

86Box home page	https://86box.net/
Packet drivers: 3Com EthernLink II, NE1000, NE2000, AMD PCNet, Western Digital 8003, etc.	

Checking for updates

mTCP is still being actively developed so it is a good idea to check for updates once in a while. mTCP now has a fairly simple utility for letting you know if a newer version is available.

To check to see if an update is available use the “services” command with the version option:

```
C:\MTCP>services version

C:\MTCP>echo off

mTCP Services by M Brutman (mbbrutman@gmail.com) (C)opyright 2022
  Version: Jul  1 2022

Checking to see if your version is up to date.
Connecting to brutman.com on port 8088 ...

Server response:

Congrats! Your version of mTCP (2022-07-01) is still up-to-date.
```

If your version is older a different message advising you of the current version and updates in that version will be displayed.

The services command uses Netcat under the covers, so mTCP must already be set up and working before you can use it. If DNS is working and you can connect to the outside world then this utility will work for you.

See “Services” for more information on the services command.

mTCP Applications

DHCP client

Introduction

A DHCP client is a small utility that helps you automatically configure the networking parameters for a machine, allowing it to use your network. This is an alternative to "static" addressing where you set the network parameters for your machine manually.

Most networks have a router that has a DHCP server built in, making DHCP the preferred way to configure a machine for networking. This DHCP client can be used to automatically configure the networking parameters for all of the mTCP applications and with some cleverness it can be used for other networking stacks as well.

Set up instructions

DHCP requires:

- A valid mTCP configuration file. This can be a minimal file that just specifies the software interrupt to use to communicate with the packet driver.
- The MTCPCFG environment file is set to the full path and filename of the configuration file.
- The correct packet driver for your Ethernet card is loaded.

The details can be found in the section titled "Set up."

Additional configuration parameters may be added to the mTCP configuration file to control DHCP behavior; those are described later.

Using DHCP

The DHCP command line format is:

```
dhcp [options]
```

where options are:

`-help`	Show a help message
`-retries <n>`	Retry n times before giving up (default=3)
`-timeout <n>`	Set timeout for each attempt to n seconds (default=10)

When DHCP runs it will send a series of UDP packets and try to negotiate a DHCP lease with a DHCP server on your network. If DHCP is successful it will set a return code of 0, show the network settings it received on the screen, and write the new network settings to the mTCP configuration file.

If DHCP fails it will set a return code of 1.

You can check the return code in a batch file using ERRORLEVEL. Here is a simple example:

```
ne2000 0x60 3
set mtcpcfg=e:\mtcp\tcp.cfg
```

```
set tz=PST8PDT7,M3.2.0/02:00:00,M11.1.0/02:00:00
dhcp -retries 3 -timeout 20
echo off
if errorlevel 1 goto Errmsg
echo DHCP worked
goto End
:Errmsg
echo DHCP failed!
:End
```

You should run this program at least once a day, and before running any other mTCP programs. This program will write the settings it gets into a configuration file that the other mTCP programs will look for and use. (You do not need to run DHCP before each mTCP application.)

Note: If you are going to use SNTP to sync your system time then you should set the TZ environment variable before running DHCP. This ensures that the timestamp written by DHCP is accurate. If you set the TZ environment variable after running DHCP the lease time calculations may be off by hours because of the difference from the default timezone and the timezone specified in the TZ environment variable.

DHCP lease duration

Remember that a DHCP lease has an expiration time; your router will not let you use an address forever. The normal lease duration is a few hours but it can be much shorter than that.

mTCP does not automatically renew DHCP leases. While that is possible, it would require every program to have the DHCP client built into them, making every program larger. It is assumed that your DHCP server will assign a lease that is long enough to let you use mTCP programs without having to renew the lease. If lease renewal is going to be an issue then consider running the DHCP client again before the lease expires, or move to static addressing. You can also request longer leases from the DHCP server.

DHCP stores the time the lease was granted and the expiration time in the mTCP configuration file. When an mTCP program starts it will check the current time and the lease expiration time. It will issue a warning and refuse to start if the lease is going to expire within an hour. (This is your cue to run the DHCP client again.)

If your DHCP server issues short leases (one hour or less) this will be a problem, as the warning will keep popping up and you won't be able to start any mTCP programs. The work-around is to set an mTCP configuration parameter that reduces the threshold for the lease expiration warning. The name of the configuration parameter is DHCP_LEASE_THRESHOLD and the value is the number of seconds to set the threshold to. For example

```
DHCP_LEASE_THRESHOLD 1800
```

will change the threshold to 1800 seconds so that if your DHCP lease expiration is further than 30 minutes out into the future you will not get a warning.

Also, if you are going to use SNTP then set the TZ environment variable before running DHCP. (See DHCP without a real-time clock for an explanation.)

DHCP failures

DHCP is often the first thing that runs after a user does the basic mTCP configuration. As a result it is often the first thing to fail. The most common form of failure is a timeout. Below are some factors that can cause timeouts:

- a misconfigured card
- using the wrong packet driver
- bad parameters to the packet driver
- cabling/network issues

A detailed discussion of these factors can be found in the Debugging section.

In the unlikely event that you can not use DHCP you can still use mTCP applications; you will have to do the network parameter set up manually though. See the Set up section for details. (I have not yet met a DHCP server that mTCP does not work with so being forced to do a manual set up is unlikely.)

DHCP without a real-time clock

DHCP writes a timestamp with the information it gets after receiving a response from a DHCP server. The timestamp is used to determine when the lease is going to expire.

If the system time is updated after DHCP gets a lease it may cause mTCP to report that the DHCP lease has expired. The extreme example of this happens when an XT class system without a real-time clock first runs DHCP before the system date and time are set; the DHCP timestamp will be from 1980. If you then update the time by entering it manually or using SNTP the system time will jump ahead 45 years, making the DHCP lease look to be long expired.

If this happens just run DHCP again, which will rewrite the timestamp for the DHCP lease.

Configuration parameters

Normally configuration parameters are read-only and are used to set the behavior of a program. DHCP is unique in that it writes some configuration parameters to the mTCP configuration file. These parameters are networking related parameters:

- HOSTNAME
- HOSTNAME_ASSIGNED
- IPADDR
- NETMASK
- GATEWAY
- NAMESERVER
- LEASE_TIME
- DOMAIN

A detailed description of these parameters can be found in the "Table of networking related configuration parameters."

The HOSTNAME parameter is the requested host name for your machine, and it is an input (read-only) parameter. HOSTNAME_ASSIGNED is the name the DHCP server actually assigned to your machine. In general the DHCP server will honor your request and both of these will have the same value. If your requested hostname is invalid or there is a naming collision the DHCP server might assign something else or nothing at all; the behavior depends on the DHCP server.

All of the other parameters in the list are output parameters provided by the DHCP server.

Additional parameters:

`DHCP_LEASE_REQUEST_SECS <x>`

This option lets you request a lease of a specific length from the DHCP server. This is useful for cases where your DHCP server might provide a short lease time (1 hour) but you know that you will be using the machine for several hours.

`DHCP_LEASE_THRESHOLD <x>`

This option sets when mTCP applications will warn you that your lease is about to expire. The unit is in seconds.

Technically this is not a DHCP parameter but it is related to DHCP operation so it is mentioned here. Please see DHCP lease duration to see how it works.

`NAMESERVER_PREFERRED <x.x.x.x>`

This option can be used to tell mTCP applications to use a different nameserver than the one provided by your DHCP server. For example, if you want to use Google DNS servers instead of your ISP's nameservers you could use the following:

```
NAMESERVER_PREFERRED 8.8.8.8
```

The NAMESERVER option set by your DHCP server will continue to be updated, but mTCP programs will use NAMESERVER_PREFERRED if it is set.

Technically this is not a DHCP parameter but it is related to DHCP operation so it is mentioned here.

DNSTest

Introduction

This is a utility for resolving machine names to IP addresses. It is useful for debugging DNS problems, as it lets you choose a DNS server to query right on the command line.

Set up instructions

DNS name resolving depends on the NAMESERVER configuration setting in the mTCP configuration file being set correctly. If you are unsure of what to set NAMESERVER to there are a few options:

- Use DHCP and let your DHCP server tell you what to use.
- Use a nameserver provided by your ISP.
- Use a public DNS server such as Google's (8.8.8.8 or 8.8.4.4) or Cloudflare's (1.1.1.1)

This program also lets you use any nameserver you choose via a command line option.

Using DNSTest

Usage is like this:

```
dnstest -name <your.machinename.here> [options]
```

Options are:

-help	Show a help message
-nameserver <x.x.x.x>	Use an alternative nameserver
-norecurse	Do not request a recursive DNS lookup
-timeout <n>	Timeout after <n> seconds
-verbose	Turn on verbose debug messages

When you run the program it sends a DNS query to your configured nameserver or the nameserver supplied on the command line. (Using the command line option allows you to test an alternative nameserver quickly.) Besides the nameservers listed above, you can try other nameservers such as the root servers. See https://www.iana.org/domains/root/servers for a listing of the current root servers.

Recursive vs. Non-recursive DNS lookups

A recursive DNS lookup is a search where the DNS server is supposed to provide the final answer to the client requesting the search. (It is recursive from the point of view of the DNS server.) All mTCP applications request a recursive search by default as it keeps the mTCP code simpler.

DNSTest is the only mTCP program that is capable of asking for a non-recursive DNS lookup. When this happens the DNSTest program is responsible for following the chain of nameservers through the Internet to find the final nameserver that can provide an answer. This is more complicated and it takes much longer, and is

generally not recommended but it can be interesting to look at the output. Use -verbose to see how it works. (The output from -verbose is generated by the mTCP trace facility, so set the LOGFILE environment variable to capture it in a file. See Collecting mTCP traces for the details on how to do that.)

Some nameservers will not do a recursive search for you. Others require it. In general, public DNS services will only allow recursive searches, while root nameservers and authoritative nameservers will not allow recursive searches. You will get an error message if the nameserver refuses your request.

FTP client

Introduction

This is an FTP client like many other FTP clients you have probably used before. Features include fairly high performance, low memory requirements, and compatibility with a wide selection of FTP servers.

Using FTP

FTP uses the following syntax:

```
ftp [-port <n>] <ftp_server_addr>
```

Options are:

`-help`	Show a help message
`-port <n>`	Connect using port <n> instead of port 21

You have to specify the FTP server address on the command line. After the control socket is opened you will be prompted for a user ID and password. After that you get a command line just like on other FTP command line clients.

Here are the commands:

dir [<filespec>]	directory list (full details)
ls [<filespec>]	directory list (names only)
pager <n>	pause ls or dir output after approximately n lines
cd [<directory>]	change to a directory on the server
cdup	move up to the server parent directory
pwd	print the current server directory
lcd [<directory>]	change the local directory
lmd [<directory>]	make a local directory
mkdir [<directory>]	make a directory on the server
rmdir [<directory>]	remove a directory on the server
ascii	change the file transfer mode to ASCII
image, binary or bin	change the file transfer mode to binary

get <server_file> [<new_local_file>]	get a file from the server
put <local_file> [<new_server_file>]	put a file on the server
delete <server_file>	delete a file on the server
prompt	toggle prompting for mget, mput and mdelete on or off
mget <server_filespec>	get multiple files matching server_filespec
mput <local_filespec>	put multiple files matching local_filespec
mdelete <server_filespec>	delete multiple files matching server_filespec
rename <from_file> <to_file>	rename a file on the server
xfermode [<new_mode>]	set PORT or PASSIVE transfer mode (see section below)
quote <remote command>	send a raw command to the FTP server to execute
quit	quit this program
shell	shell out to DOS
interactive	switch input from a script file to the keyboard

The program captures Ctrl-Break and Ctrl-C. If you use Ctrl-Break during a file transfer or directory listing the current operation will stop (prematurely) and you will be returned to the command line after a brief delay. If you use Ctrl-Break or Ctrl-C at the command line the program will end. (This might take up to 20 seconds depending on the state of the sockets, so please be patient.)

The shell command should be used with care. With the shell command you can do local directory listings, rename local files, look inside of files, and do most normal things. Keep in mind that while you are at the DOS command prompt the mTCP code is not servicing the socket connection - if you stay in DOS too long you might lose your connection to the FTP server.

ASCII vs. BINARY transfers

(This used to be common knowledge for FTP but it is slowly fading from the collective consciousness ...)

FTP clients and servers can transfer data in one of several modes. You need to ensure the correct mode is in use when you transfer a file to avoid data corruption.

- ASCII: The FTP server you are connecting to will treat the files being transferred as text files. You are on a DOS system so the end-of-line is marked by a CR/LF pair, while the server side probably uses something different. These conversions will be made for you by the server so you need to use BINARY mode if you do not want these conversions to be made.
- BINARY: The FTP server does not alter the content of the data in any way. This is usually what you

want to use, except for text files.

ASCII or BINARY mode is a server side function; every server has a different default so you should not assume what is set when you first connection. Use the ASCII or BIN commands to set the desired mode before you start transferring data to avoid data corruption.

In this version of the FTP client a “BIN” command will automatically be sent to the server immediately after you log in. BINARY mode is usually want people need and sending the BIN command automatically right after login will generally protect people from servers that set a default of ASCII. This feature can be disabled by setting FTP_BIN_CMD_STUFF to false in the mTCP configuration file.

FTP transfer modes

The xfermode command requires a little explanation.

In the beginning of time FTP servers always sent data by doing a connect to the FTP client, which was listening for the incoming connection on a default port. The FTP server in NCSA telnet does this. I will call this 'Classic' mode and you can tell this FTP client to handle data transfers this way.

'Classic' mode is only useful for ancient FTP servers that can't do anything else. Most modern FTP servers will not connect to an FTP client on the default ports and most firewalls/routers wouldn't allow it anyway.

'Port' mode is the current default for many clients. Instead of having the FTP server connect to the client on default ports, the client will tell the FTP server exactly what port it should use for the connection. This is slightly better than Classic mode, and should be usable on more FTP servers, firewalls and routers.

'Passive' mode is probably the best and most compatible mode to use with modern FTP servers. It turns things around by making the client connect to the FTP server to make data connections instead of having the FTP server connect to the client. (From the standpoint of the FTP server it is 'passive' and it waits for a connection from the client.) This will work with almost any firewall and router. This is the mode I'd recommend using too. Passive is the default mode for this client.

DOS filename limitations

As this FTP client runs under DOS, it is subject to DOS path and filename limitations.

When specifying a filename on the server you can use whatever filenames are legal for the server. This may require the use of quoting filenames (see below).

When dealing with local files you must obey the rules for DOS. For example, DOS filenames must fit the 8.3 naming convention and some characters are forbidden.

The FTP client will not try to check and reject invalid DOS filenames; it will just let DOS deal with them. Sometimes this results in filenames being truncated to fit the 8.3 format. Other times it might result in an error and an uncompleted operation because DOS refused to work with the filename. Please be aware of your filenames and try not to break DOS too badly.

One important exception: the FTP client will allow you to use either a forward slash or a backward slash as a directory delimiter for local DOS filenames.

Filenames and Quoting

DOS does not allow a filename to have a space in it. However, you may find that FTP servers running real operating systems do support filenames and directory names with spaces in them.

FTP supports adding quotes around filenames and directory names so that you can work with these servers. (You still have to provide a legal DOS filename for receiving files.)

If you need to specify a filename that has the double quote character in it use two double quote characters in a row. This will result in one double quote character being sent. If you have to deal with embedded spaces at the same time, add an extra set of quotes to enclose the entire parameter.

Examples:

You enter:	The server sees:
"Spacey name"	**Spacey name** *(All one parm)*
""QuoteChar""	**"QuoteChar"** *(All one parm)*
""Quote Char""	Uh oh .. you get **"Quote** as one parm and **Char"** as the second parm. That is because the two sets of double quotes collapse into a single double quote, but do not do anything to mark it all as one parameter.
"""Quote Char"""	**"Quote Char"** *(All one parm)*. Note the usage of the triple quotes.

Command line editing

If you are using FTP interactively you can take advantage of the command line editor. It does all that you would expect it to do: backspace, delete, insert, home, end, etc. It also includes a ten line command buffer so that you can recall older commands easily without retyping them.

Reading from STDIN

FTP will take input from a file redirected to STDIN. This allows you to automate some simple tasks, like logging in to fetch a file.

For example, start with a simple text file that has the input you would normally type. This should include your name and password followed by FTP commands:

```
brutman
badpassword
BIN
CD DOWNLOADS
GET NEWFILE.ZIP
QUIT
```

We will call that file “script.txt” for this example.

Then when you start FTP use standard file redirection to tell FTP to take its input from script.txt instead of the keyboard:

```
ftp ftp.yourserver.com < script.txt
```

Here are some notes for using this:

- Because you have to start FTP with STDIN redirected from a file, you have to include your login information as the first two lines of the script file. This is obviously not a good security practice, so don't do it on public servers on the Internet and never reuse passwords.
- FTP will interpret "end of file" as "quit". So when the script ends, so does FTP.
- If you don't want FTP to quit when it hits the end of the script file, add the "interactive" command to the end of the file. This will tell FTP to start reading from the keyboard again.

FTP Performance

Good FTP performance requires good flow control. This FTP client provides two configuration settings for helping achieve the best flow control for your machine.

FTP_TCP_BUFFER controls how much memory is used to hold data from incoming TCP packets. In general larger buffer sizes improve performance by letting the sending side send more data without having to wait for an acknowledgement.

FTP_FILE_BUFFER controls the size of disk writes. In general larger disk writes are more efficient, as they amortize the cost of calling DOS over more data. But larger disk writes can also take longer, causing hiccups in the flow of data.

In general, use a large FTP_TCP_BUFFER because that should always work. But run some experiments on your machine with FTP_FILE_BUFFER at different sizes. FTP_FILE_BUFFER should be the size of your disk clusters, or a multiple of that.

Configuration Parameters

The following are optional configuration settings that you can put in your mTCP configuration file.

`FTP_CONNECT_TIMEOUT <n>`

Set the timeout for a socket connection to n seconds. The default timeout is 10 seconds.

`FTP_TCP_BUFFER <n>`

Set the size of buffer used by TCP for receiving data to n bytes. The default is 8192 bytes. For best performance this should be set to a multiple of 1024 (1KB). The default is reasonable compromise between size and performance. You can go lower to conserve memory or up to 16384 bytes to improve performance.

`FTP_FILE_BUFFER <n>`

This controls the size of the buffer used to store data when receiving a file. The default is 8192 bytes. For best performance this should be set to a multiple of 1024 bytes. The default is a reasonable compromise between size and performance. The maximum size is 32768 bytes.

`FTP_SEND_FILE_BUFFER <n>`

As of this release this configuration option is ignored. The value set by FTP_FILE_BUFFER is used for both sends and receives.

`FTP_MLIST_BUFFER <n>`

When you use MGET, MPUT or MDELETE FTP has to fetch a list of files from the server or from your local machine. This buffer holds that list of files. If the list of files is too long the MGET, MPUT or MDELETE command will fail.

One solution is to change your filespec so that the list of be transferred is smaller, and use multiple MGET or MPUTS to do what you need. The other solution is to increase this buffer size. The default size is 4096 bytes, which is plenty for most usage. The maximum size is 16384 bytes.

`FTP_BIN_CMD_STUFF false`

By default the FTP client will try to send a BIN command to the server right after you log in. This is an attempt to set the default to BINARY file transfers which is what most people want. If you are a power user and this behavior annoys you then use this configuration parameter to disable this behavior.

FTP server

Introduction

This is an FTP server for DOS. Some of the features include support for passive and active connections, support for multiple users at the same time, a sandbox mode which is useful for setting up anonymous FTP, and good compatibility with a wide variety of FTP clients.

Special hardware or software requirements

The FTP server requires as little as 200KB if you just have one FTP client connecting at a time. If you want to support multiple connected clients more memory will be needed.

Before you fire it up ...

Setting up and running an FTP server is a little more involved than running the other mTCP applications. You need to consider your security and performance requirements. Security is especially important; an FTP user with full privileges to your system can do serious damage, accidental or otherwise.

Quick Start Instructions

1. Configure mTCP and ensure that you have basic network connectivity. You should be able to run the other mTCP applications first.
2. Add the FTPSRV_PASSWORD_FILE and FTPSRV_LOG_FILE configuration parameters to the mTCP configuration file. (FTPSRV_LOG_FILE is optional but recommended.)
3. Create a minimal password file. (See the included sample file, ftppass.txt)
4. Run. You should be able to use the FTP server from machines on your home network with no further configuration.
5. If you are planning on offering anonymous FTP on the internet or using it with FTP clients on the other side of a firewall, read the rest of these instructions.

FTP server background information

Active connections vs. Passive connections

FTP clients connect to FTP servers on a socket connection known as the control connection. The client sends commands and gets return codes on the control connection. FTP servers usually listen for new control connections on port 21, although other non-standard ports are possible.

When an FTP client and server wish to exchange data they will open a new socket connection that will be used for that one data transfer. A data transfer is a directory listing, a file send or a file receive. (Yes, even a directory listing requires a new socket connection.) After the data transfer is done the new socket is closed, leaving only the original control connection. Data connections are not used for more than one data transfer.

There are two ways the new socket connection for data can be created. An "active" connection is when the client tells the server where to connect to, listens on a port, and waits for the server to open the new socket

connection. A "passive" connection is when the client asks the server to listen on a port and the client opens the new socket connection. Note that "active" and "passive" are from the point of view of the server.

In the real world active connections are more difficult to support because most home users are behind firewalls. Firewalls cause problems because they usually are permissive about letting machines make outgoing connections to the rest of the world but they do not readily let incoming connections in. This is a problem because when in active mode, an FTP server will have to make a new incoming data connection back to the client. Firewalls have special code to allow this incoming connection but it often doesn't work correctly for a variety of reasons.

On the other hand, passive mode connections almost always work. If your client was able to create the control connection to the FTP server in the first place, then it probably can create other connections to the FTP server. Passive mode is generally preferred by FTP clients because it works with a wider variety of firewalls.

If you have a firewall it is probably optimized to allow FTP clients make connections to the outside world. Running an FTP server behind a firewall becomes a problem because you need to allow incoming connections from the outside world. (We'll talk about that more later.)

UserIDs, Anonymous FTP, and permissions

Anonymous FTP allows anybody to connect to and use your FTP server. There is no authentication required. While useful for distributing data, anonymous FTP can be dangerous when not configured correctly.

In addition to anonymous FTP the mTCP FTP server supports named FTP accounts that require a password be entered. Named accounts provide more security but they are not foolproof either - FTP transmits passwords in the clear and a named user can still do accidental damage to your system.

To give you some control and security the mTCP FTP server has the following features:

- The ability to lock a user into a specific part of the directory hierarchy. This ability can be used to create "sandboxes" that can limit what users can see and do.
- The ability to restrict uploads to a particular directory whether the user is in a sandbox or not.
- The ability to set permissions for individual operations that might be destructive in some way, such as delete, make directory, remove directory, rename, and store (store, store unique, or append) files.
- mTCP FTP server does not let you overwrite files when uploading. You must explicitly delete a file before replacing it.

It isn't difficult to set up the server in a secure way. But it does require a little planning.

Configuration parameters

There are no command line options that are used when starting the FTP server; all configuration parameters are set in the mTCP configuration file. The following configuration parameters are recognized:

Required parameters

`FTPSRV_PASSWORD_FILE <path_and_filename>`

This should be set to a full path and filename of a password file. The password file format is explained in

"Setting up the password file."

Optional parameters

`FTPSRV_LOG_FILE <path_and_filename>`

This should be set to a full path and filename of a file that will be used to log user connections and file transfers. The file will be appended to each time it is opened.

If you do not specify this option then no logging of user activity will be done.

`FTPSRV_SESSION_TIMEOUT <seconds>`

This configuration parameter allows you to set the inactivity timeout for FTP sessions. The default is 180 seconds (three minutes). The maximum is 7200 seconds (two hours). This configuration parameter is optional.

`FTPSRV_CONTROL_PORT <port number>`

The standard port for an FTP server to listen to for new control connections is port 21. You may specify an alternate port if 21 is not available or usable. A lot of ISPs will block incoming access to port 21, so to accept connections from the outside world you may need to use a high numbered port such as 2021. This is not standard so you will need to make it clear to any possible client what the new port number is.

`FTPSRV_EXT_IPADDR <numerical IP address or hostname>`

If you are behind a firewall and you are accepting passive connections from the outside world you will need to use this option to tell the FTP server what your external (public) IP address is. This allows the FTP server to put the correct connection information in the response to the PASV command sent from the client.

See the section on firewall considerations for more information.

`FTPSRV_PASV_BASE <port number>`

If you are accepting passive connections you can use this option to specify a starting port number for passive data connections to use. The default starting port number is 2048.

Being able to set this value might be helpful for configuring your firewall.

`FTPSRV_PASV_PORTS <number of ports to use>`

This option can be used to set the number of ports used for passive data connections. The default is to use 1024 different ports.

Being able to set this value might be helpful for configuring your firewall.

`FTPSRV_CLIENTS <number of clients>`

This option is used to set the maximum number of clients that can be connected at one time. The default is three clients. The maximum number of connected clients is 10.

`FTPSRV_MOTD_FILE <path_and_filename>`

This should be set to a full path and filename of a file that will be used to great new users when they first login to the FTP server. The file should consist of plain ASCII text with lines no longer than 70 characters. The total file size should be no more than 600 bytes.

`FTPSRV_FILEBUFFER_SIZE <n>`

Each connected client has a buffer used for reading and writing files. The default size is 8KB. You can set it between 4KB and 16KB in 1KB increments. Larger buffer sizes perform slightly better, but on small systems the total amount of buffer space can add up.

`FTPSRV_TCPBUFFER_SIZE <n>`

Each connected client has a buffer used for receiving data on the TCP/IP data socket. The default size is 8KB. You can set it between 4KB and 16KB in 1KB increments. Larger buffer sizes perform slightly better, but on small systems the total amount of buffer space can add up.

Buffer sizes below 8KB might cause flow control problems with faster systems.

`FTPSRV_PACKETS_PER_POLL <n>`

This is an advanced option that only makes sense on smaller systems.

The FTP server checks the packet driver for new incoming packets to process each time through the main program loop. The most efficient way to process packets is to scoop up any available packets that might be available and process them all at the same time. While efficient, this can delay the sending of outgoing packets for a short time. On slow systems that delay can cause flow control problems with faster systems.

This setting tells the FTP server how many packets to process each time through the main program loop. The default is 1, which is safe for any system. If you want to experiment with larger values for better performance you can set it as high as 10.

`FTPSRV_EXCLUDE_DRIVES <drive_letters>`

This option allows you to exclude specific drive letters from the scan for valid drive letters that is done when FTPSRV starts. This is useful if you want FTPSRV things like network drives or floppy drives. For example:

```
FTPSRV_EXCLUDE_DRIVES AB
```

will tell FTPSRV to ignore drive letters A and B, making them unavailable to users even if they have diskettes in them.

`FTPSRV_CLOBBER_FILES <true|false>`

This option allows you to tell the FTP server that a user may overwrite existing files. By default trying to send a file that already exists will be rejected, as you would be overwriting the existing file. (You are able to delete the file first, but this is an extra step.) If you set this option to true existing files will be overwritten

without warning.

This option has no effect on anonymous FTP users; they may not overwrite existing files and this option will not enable them to.

Choosing the number of clients

It doesn't take a lot of CPU to have a connected client. In theory a low end machine can handle many connected clients. The problems start when those clients try to transfer large amounts of data.

If you have multiple clients transferring files at the same time the clients will be competing with each other for the following:

- CPU time
- Disk I/O
- Network I/O

Clients that are browsing directory listings are not a burden on the system. Clients that are actively transferring files will split these resources. Depending on the speed of your system multiple active clients may reduce transfer speeds to an unacceptable level.

At the lowest end of the spectrum, an 8088 based PC XT with the original MFM hard drive and a reasonable Ethernet card should be able to service one file transfer at a rate of 30KB/sec. Multiple clients will split that bandwidth.

A higher end system like a 386-40 with an IDE hard drive can service one file transfer at a rate around 400KB/sec. Multiple clients splitting this bandwidth is obviously not as much of an issue.

Besides the available bandwidth memory requirements are another consideration. Each connected client requires around 11KB. If a file transfer is in progress the client will require another 8KB for the length of time the transfer is in progress. Since it is unlikely that every client will try to transfer a file at the same time, you can probably get by with less than the worst case memory scenario. I haven't measured it precisely, but you should not need more than 400KB in the worst case.

Setting up the password file

The location of the password file is set by the FTPSRV_PASSWORD_FILE configuration parameter described in "Configuring the mTCP FTP server."

The password file is used to define which users are allowed to connect to the FTP server and what directories and functions they can use. There is one entry per line. Entries are in the following format:

```
userid password sandbox_directory incoming_directory permissions
```

`userid`

This can be up to 10 characters in length.

`password`

Passwords can be up to 10 characters in length. You may choose any reasonable mix of characters for a password; one notable exception is that spaces are not allowed.

If you use the keyword "[email]" (without the quotes) the FTP server will prompt for an email address instead of a password and will permit any user input except a blank line. The "[email]" keyword is used to implement anonymous FTP; be sure to read "Setting up Anonymous FTP" before setting up anonymous FTP.

FTP passwords are always transmitted "in the clear" on the wire, so they are not terribly secure.

`sandbox_directory`

If you wish to restrict a user to a specific directory (and all of its subdirectories) set this field to be the fully qualified drive letter and path of the directory. If you do not want a user restricted use the keyword "[none]" without the quotes but with the brackets.

The drive and path format is not in the standard DOS format. The format used by the program is:

/drive_X/path

where X is a DOS drive letter. For example, D:\FTP\INCOMING is expressed as /drive_d/ftp/incoming.

When a sandbox is in use the user will be unaware that they are being restricted to a sandbox. They will just see the designated directory as the top of the directory hierarchy. For example, if you specify a sandbox directory of "/drive_c/ftpdata" the user will only see that directory as "/". The FTP server automatically puts "/drive_c/ftpdata" in front of any requests, and it prevents the user from trying to escape the sandbox.

The sandbox must be specified using a full path, including drive letter, in the format described above.

If a user is not restricted to a sandbox they are free to move to any directory and drive letter on the system. To switch drives the user must use a path in the format described above. For example, to switch to drive E: the user would enter "cd /drive_e/". This syntax looks like a standard Unix path, which provides the best compatibility with the widest range of clients.

(Previous versions of the FTP server used a more DOS-like syntax. But compatibility with clients was a nightmare. This works much better now ... -Mike)

`incoming_directory`

If you want to restrict uploads to a specific directory enter the directory name here. If the user tries to upload a file in a different directory the request will be denied.

If the user is also in a sandbox then the incoming directory specified here should be a relative pathname within the sandbox. For example, if the sandbox is "/drive_c/ftpdata" and this parameter is "/incoming", then all uploaded files will be deposited in "/drive_c/ftpdata/incoming". (The user will see it as "/incoming" due to the way the sandbox works.)

If the user is not in a sandbox then the incoming directory should be an absolute pathname including a drive

letter in the same format described in "Sandbox directory".

If the user is not to be restricted then use the keyword "[any]" (without the quotes).

`permissions`

This field is used to control permission to commands that might alter the filesystem. You can choose to grant permission to all commands by using the "all" keyword (without the quotes). If you don't use "all" you must grant permissions to each command individually.

The commands you can grant permission to are:

DELE	delete a file
MKD	create a directory
RMD	remove a directory
RNFR	rename from (this is used to start a file rename)
STOR	store (receive a file from the client)
APPE	append (create a new file or append to an existing file)
STOU	store unique (receive a file form the client, but the server picks a unique filename.)
SITEQUIT	Use a site command (quit) to end the server.

If you do not grant permission to a command the user can not use it. The other FTP commands for changing directories, getting directory lists, changing modes, etc. are always available.

Here is a sample password file that demonstrates a named user with no restrictions, a named user with restrictions on uploading, an anonymous account with uploading, and a read-only user with no uploading:

```
brutman    passw0rd  [none]        [any]                  all
tester     b0guspw   [none]        /drive_e/ftp/incoming  stor stou
anonymous  [email]   /drive_e/ftp  /incoming              stor stou
readonly   [email]   /drive_e/ftp  [any]
```

The first user has no restrictions on directories and permission to use all of the filesystem altering commands. The second user can see the entire system but is restricted to only uploading in one directory, and they can only put files - they can't delete, create or remove directories, or append to existing files. The third user is restricted to one directory and its subdirectories (a sandbox) and may only upload to one directory within that sandbox. The last user can only look in one directory and its subdirectories (a sandbox) and may not upload. (Ignore the [any] field, it's just holding a place.)

Setting up Anonymous FTP

Setting up anonymous FTP is fairly easy but you should be aware of the potential security problems if you do it wrong.

First, as an operating system DOS has no real user permission system. The FTP server can do anything, and by extension your users can do anything that the FTP server allows them to. The operating system does nothing to help keep things secure.

An anonymous user account is created when you set the password field of a user account to "[email]". The

userid is traditionally named "anonymous" or "ftp", but it can be any name if the password field is set to "[email]". Setting the password field to "[email]" turns off all password checking for that user and simply records what the user enters for the password as their email address.

It is absolutely necessary to lock anonymous users into a sandbox. If you do not, anonymous users will be able to look in any drive and any subdirectory on your system. Depending on how you set the permissions to the FTP commands, it would be possible for an anonymous user to trash your system.

If you plan on letting anonymous users upload you should designate a directory within the sandbox that will be used to hold all uploaded files. That allows you to easily identify and sort new files in a single place instead of having them spread through the sandbox.

The recommended user permissions for anonymous users are STOR and STOU. This lets them upload files with a unique name that they choose, or with a unique name that the server chooses. mTCP FTP server will not overwrite an existing file so both STOR and STOU are safe for existing data. APPE (append) is not safe; it can alter existing files, so it is not recommended for anonymous FTP users.

If you do not plan on accepting uploads from anonymous users then do not even give them STOR and STOU. They will still be able to retrieve files from anywhere in the sandbox.

Below are sample lines that can be used in the password file to implement anonymous FTP.

```
anonymous [email]   /drive_e/ftpdata  /incoming stor stou
ftp       [email]   /drive_e/ftpdata  /incoming stor stou
```

These lines assume that the e:/ftpdata is the sandbox that the FTP users will be using. Store (STOR) and Store Unique (STOU) are granted, and the user has to be in the /incoming directory to upload files. (With the sandbox taken into account the full path for the incoming directory is /drive_e/ftpdata/incoming.)

Two different users with the same sandbox, incoming directory and permissions are used because anonymous FTP usually honors both "anonymous" and "ftp" as usernames.

It is possible to assign different users to different sandboxes; there is no limitation on the number of sandboxes on a system. (The sandbox and upload directory are a per user setting.) You can use this to have different classes of anonymous users; perhaps one class is not allowed to upload, while a second and slightly more privileged class is allowed to upload.

DOS filesystem limitations

For those of you used to Unix or modern Windows systems this is going to be a shock. Here is the brief run-down on the DOS filesystem and what the limits are:

A fully qualified filename in DOS looks like this:

```
d:\path\filename.ext
```

The maximum length for a fully qualified filename is 78 characters. The components are:

drive letter	2 characters (letter plus a colon)
path	63 characters including a beginning and ending "\"
filename	13 characters in 8.3 format

The official directory delimiter is always a "\" (backslash).

Valid characters for a filename are all letters, numbers, the following punctuation characters !@#$%^&()-_{} `'~*? and all "high bit" characters numbered ASCII 128 and above. The rendering of high bit characters depends on your active DOS code page - mTCP FTP server assigns no meaning to them at all.

If your user tries to use an invalid path or filename their request will be rejected. Gently remind them that they are on DOS, and that 8.3 isn't just a pain in the rear, it's the law.

To keep compatibility with the largest number of FTP clients the FTP server doesn't show DOS style drive letters and paths to connected users. It rewrites the drive letter part to look like /drive_X/ where X is the drive letter and the path delimiters use forward slashes, not backslashes. The FTP server also presents the image of a root file directory '/' where the subdirectories are the available drive letters.

A connected user must use a full path (in /drive_x/path/ form) to change directories. If a relative path is used it is relative to the current drive letter and path for that user.

Using mTCP FTP server with a firewall

If you plan on allowing users to use your FTP server through a firewall, read on. If you are not allowing access from people outside of your LAN you can skip this section.

Most firewalls are designed with the needs of FTP clients in mind. Now you will be running a server, where things don't work so well.

Choosing a control port to listen on

If your ISP allows it use the standard FTP port (21).

If your ISP does not allow it or you want to improve your security slightly choose an alternate port. ISPs generally don't block ports numbered 1024 and higher. While it works, non-standard ports are not as convenient for your clients to use. (Your users will have to find a way to specify the new port number in their client.)

Set up your firewall to forward the control port

Most firewalls block incoming connections from the internet. This makes sense, as your personal machines behind the firewall will make whatever outgoing connections they need. This doesn't work when you run a server though. You have to be able to allow incoming connections from the outside world to get to your FTP server.

Most firewalls allow you to specify where incoming connections from the outside world should be forwarded to. The forwarding is done based on port numbers. Set up your firewall to forward incoming connections on your chosen control port to the machine running the FTP server.

Set up your firewall to forward data ports

Clients connecting to your FTP server from the outside world are probably behind firewalls that also do not allow for incoming connections. They will probably need to use passive mode for data connections, which means that they expect to make incoming connections to your FTP server.

By default mTCP FTP server will use ports 2048 through 3071 for data connections. (The starting port number and the number of ports to use can be configured if the defaults are not suitable.) You will need to configure your firewall to pass all of these port numbers through to the FTP server.

Tell mTCP FTP server your public IP address

If you are behind a firewall you are probably using a private, non-routable address assigned by your firewall. This address was obtained using DHCP or configured by hand. The firewall is responsible for translating incoming and outgoing IP packets between the private non-routable addresses and the public address assigned by your ISP.

This mapping is called "Network Address Translation" and it is normally transparent to you. There is a snag when running an FTP server; when a client requests a passive connection for a data transfer the server will respond with an IP address and port that the client should connect to. The FTP server knows its IP address, which is the local non-routable address on your home network. Without any intervention that is the address it will give to the connected client requesting a passive mode connection. Clearly that address is wrong for an external client.

To get around this you can use the FTPSRV_EXT_IPADDR configuration parameter to tell mTCP FTP server what address it should send to clients instead of the private, non-routable address that it knows. The address to use is the public address that your ISP has assigned to your firewall/cable modem/DSL modem.

A commercial grade firewall will modify the FTP server passive response automatically, but home grade firewalls will not. When a server is running on a non-standard port almost no firewalls will be able to fix the problem. FTPSRV_EXT_IPADDR gives you a work around.

mTCP FTP server is smart enough to know if a client is on the same internal LAN as it is, or if the client is coming from a different LAN. If the client is on the same LAN as the server then the passive command response will use the local non-routable private address. That way your local machines can use passive mode connections too.

There are two ways to set FTPSRV_EXT_IPADDR:

- Setting it to a numerical IP address: This works well if your public IP address assigned by your ISP does not change. Example: FTPSRV_EXT_IPADDR 50.22.19.21
- Setting it to a fully qualified host name. Use this option if yo u are using Dynamic DNS. When you do this the FTP server will do a DNS query every 5 minutes to resolve the name and convert it to a numerical IP address. Example: FTPSRV_EXT_IPADDR yourhomedomainname.com

Starting, running and ending the FTP server

This assumes that your packet driver is loaded and that the basic work to set up mTCP is done. (You should have the MTCPCFG environment variable pointing at a valid mTCP configuration file, and you should have tested your network connectivity.)

Starting

To run ftpsrv, just go to the directory where it is located and run it. Ftpsrv will read the mTCP configuration file, and then scan the password file. If all looks good and it can get the memory that it needs, it will start running. If anything fails you should get an error message that tells you what is wrong.

The three most common reasons for it to fail are:

- The mTCP configuration file can not be found or is not correct.
- The password file has an error in it. Not all errors will be spotted, but it has to look like a reasonable password file.
- You don't have enough memory available.

If mTCP starts running it will give you a set of messages, and then it will start to use the screen to log events. Events will also be written to a log file on disk specified by the configuration file, if desired.

Successful start-up will look something like this:

```
mTCP FtpSrv by M Brutman (mbbrutman@gmail.com) (C)opyright 2009-2022
  Version: May 27 2022

Opening password file at c:\mtcp\ftppass.txt
  Password file looks reasonable.

mTCP FtpSrv version (May 27 2022) starting

Clients: 10, Client file buffer size: 8192, TCP buffer size: 8192
Packets per poll: 1, TCP sockets: 21, Send buffs: 35, Recv buffs: 40
Client session timeout: 120 seconds
Control port: 21, Pasv ports: 2048-3071
mTCP IP address: 192.168.2.100, Pasv response IP addr: 192.168.2.100

Press [Ctrl-C] or [Alt-X] to end the server
```

The Pasv response IP address and other parameters will be different for your machine and configuration.

The top line of the screen has a small status line that shows you the total number of connections and the number of active sessions. The rest of the screen is used for log messages.

At this point your server is ready for incoming connections.

Running

The following keys are recognized on the console:

Alt-B	Toggle the beeper that signals when users sign on
Alt-H	Help
Alt-S	Show statistics
Alt-X	Exit

There is not much to do while the server is running. I would turn off the monitor and check on it once in a while by connecting to it with an FTP command and using the SITE STATS or SITE WHO command.

If you are accepting uploaded files you should check for new files and process them once in a while. You can do this by downloading the files to another machine, examining them, and then either moving them using the rename command or deleting them from the server. You should have a rough idea of how much free space you started with; file uploads will fail if you run out of disk space. You can check the amount of available disk space with the SITE DISKFREE command.

With five or less clients configured the mTCP FTP allocated enough TCP/IP sockets to cover every possible combination of connections. After five users it is assumed that not every client will try to do a data transfer at the same time, so the worst case number of TCP/IP sockets are not created. If by some freak accident every connected client does try to do a passive mode data connection at the same time some of them might fail - the failure will be temporary and they will be able to try again.

If more people try to connect than there are configured clients each new connection will get a "421 - please try again later" message.

The default timeout period for an inactive session is three minutes. You can set it as low as one minute or as high as 10 minutes. An inactive user should be disconnected within 10 seconds after exceeding the limit.

Ending

To end the server hit Ctrl-C at the keyboard. Ctrl-C will be recognized within a second, and it will take up to 10 seconds to close all of the current data connections.

Notes on specific FTP clients

The FTP server presents the illusion of a Unix filesystem to connected clients, even while allowing users to switch drive letters. Hiding the fact that it is running under DOS keeps the clients from getting confused. As a result, there should be very few client specific problems.

Some clients send a LIST command assuming that there is a Unix system running /bin/ls on the other end. They send options with the LIST command that are supposed to be interpreted by /bin/ls, but look like a valid DOS filename. (Hint: A '-' character is valid as the first letter of a DOS filename.) The FTP standard does not have a mechanism for sending options on a LIST command so these clients are technically broken.

Previous versions of FTP server did not handle these options and thus it looked like clients could not get directory listings. This version ignores those options, which should improve compatibility with those clients.

When running on a slow machine it is possible that faster clients will flood the TCP/IP socket, causing the TCP/IP flow control mechanisms to kick in. Sometimes they kick in too hard and make file uploads very unpredictable. If this happens you can increase the file buffer and TCP receive buffer sizes used, and if your client supports it throttle the rate at which it sends data and reduce the number of concurrent connections. (FileZilla is one client that can do this.)

The standard command line Windows XP FTP client does not support passive mode data connections. On an internal network this is not a problem, but if there is a firewall between you and the Windows XP machine it will be nearly worthless. Get a better client.

Internet Explorer, Firefox and Chrome work well against mTCP FTP server when an FTP URL is used to access

the server.

Command line Linux clients have no known problems. And of course the mTCP FTP client has no known problems. ;-0

Implemented RAW FTP commands

USER	Set user (only valid during initial login sequence)
PASS	Send password (only valid during initial login sequence)
RNFR	rename from (original filename)
RNTO	rename to (new filename)
DELE	delete file
CWD/XCWD	change directory
CDUP/XCUP	move "up" one directory in the hierarchy
PWD/XPWD	print working directory
MKD/XPWD	make directory
RMD/XRMD	remove directory
PASV	request passive mode
PORT	use PORT mode
ABOR	abort current data transfer
LIST	directory list
NLST	name list
RETR	retrieve (send file to client)
STOR	store (receive file from client)
STOU	store unique (receive file, but server chooses a unique name)
APPE	append to existing file (or create a new file)
MODE	set transfer mode (only stream is supported)
STRU	structure (only file is supported)
TYPE	structure (only file is supported)
SIZE	get the file size of a file (BINARY mode only)
HELP	help
ALLO	allocate (obsolete, but not an error)
FEAT	feature negotiation; will respond with MDTM
MDTM	modification type
NOOP	do nothing successfully
STAT	give server status or do a dir listing over the control socket
SYST	identify system type (responds with "215 DOS Type: L8")
SITE	do SITE specific commands
REIN	re-initialize (not supported)
ACCT	account (not supported)
REST	restart position (not supported)

SITE commands:

SITE DISKFREE	show free disk space for a given drive letter
SITE HELP	give help for the SITE command
SITE QUIT	Tell the server to exit to DOS
SITE STATS	respond with some fun statistics
SITE WHO	show who is online

HTGet

Introduction

HTGet is a small utility that can fetch files served from HTTP servers. It uses HTTP 1.1 which is the current HTTP protocol that enables use with virtual hosts. HTTP 0.9 and HTTP 1.0 are supported for compatibility with older HTTP servers.

HTGet was inspired by a version written by Ken Yap (ken@syd.dit.csiro.au) in 1997 for the WATTCP stack. This version is a complete rewrite using some of the concepts from the original.

Using HTGet

Usage is like this:

```
htget [options] <URL>
```

Options are:

`-h`	Shows help text
`-v`	Print extra status messages
`-quiet`	Quiet mode: suppress status messages
`-headers`	Fetch only the HTTP headers
`-showheaders`	Fetch content and show headers too
`-m <file>`	Fetch file only if modified after <file>'s date
`-o <file>`	Write content to <file> instead of STDOUT
`-pass <user:password>`	Send authorization
`-09`	Use HTTP 0.9 protocol
`-10`	Use HTTP 1.0 protocol
`-11`	Use HTTP 1.1 protocol (default)

The URL format is:

```
http://yourservernamehere[:port]/path
```

The port parameter is optional - if it is not specified the default is port 80. Authentication information is not included in the URL syntax; this is slightly nonstandard but it makes the code simpler. (See the description of the -pass option below for details on how authentication is supported.)

When you run HTGet it will attempt to connect to the HTTP server, send headers, and receive content. If you do not give any options the downloaded content is sent to STDOUT. You can redirect STDOUT to a file or use the -o option to specify a filename for HTGet to write to directly.

Using -v gives you extra messages that tell you how the data transfer is progressing. The extra output is sent to STDERR to avoid interfering with the content that is being set to STDOUT.

The -m option checks the modification date of the file that was specified with the -o option and sends that date to the server during the request phase. If the content on the server has not been updated since the timestamp on

the local file the server will return a "304" status code and no content. If the content on the server has been updated you will get the new content instead and the file timestamp should match the timestamp from the server. This option is useful for when you only want to update a file based on if it has been updated on the server. To use it you should have the TZ environment variable set; see Setting the TZ environment variable for instructions on how to do that.

If the content you are requesting is password protected try using the -pass option. The -pass option takes a string in the form of username:password and encodes it using BASE64 before sending it to the HTTP server. The authentication type is known as "BASIC" - it is not secure but it is widely supported. If the HTTP server requires something more secure then this option probably will not work.

The current HTTP protocol version number is 1.1 which is what this program is using. The big improvement that version 1.1 adds is support for virtual hosts, which are HTTP servers that share the same IP address. Since this program is designed to run on vintage computers it is quite possible that you might run into a vintage web server using HTTP 1.0 or even HTTP 0.9. These older servers might not understand the headers used by HTTP 1.1; if that happens the request will fail. Try using the -10 option to force the program to fall back to HTTP 1.0 if you suspect you have an ancient server that is confused by HTTP 1.1. Use the -09 option to specify HTTP 0.9.

Examples:

```
htget -o google.htm http://www.google.com/
```

This downloads the Google search page and writes it to google.htm.

```
htget -m -o google.htm http://www.google.com/
```

This does the same thing as above, but if the copy on the server has not changed since the local copy was written it will not be downloaded and rewritten.

```
htget -pass john:secret http://someserver.com/protected/secret.file
```

This grabs the password protected file "secret.file" using "john" as the userid and "secret" as the password. The userid and password get converted to BASE64 before being sent to the server. If the server supports HTTP BASIC authentication you will get the file.

```
htget -10 http://yeancient.server.net/readme.txt
```

Force HTTP 1.0 mode for a prehistoric server that does not use HTTP 1.1.

```
htget -headers http://www.google.com/
```

Don't get the content - just look at the headers that would have been sent.

When the server sends content ...

The server can send content even when the server reports a bad return code. The content might be a page that describes the error and tells you what to do.

HTGet will not write output if there is a socket error or if the file has not been modified and you use the -m switch. But due to the way HTTP works, it will write data if it is given content, even if the content is an error message. So don't overwrite anything important assuming that the server will give you a fresh copy - it might get replaced with an error message.

Return codes

The HTTP server always returns a result code which can be used to tell you if the request was successful. Unfortunately, the range of possible return codes from the HTTP server is greater than what can be expressed in a program exit code so HTTP return codes can not be used as exit codes directly.

It is valuable to be able to determine if a request was successful or not using a program instead of having a human make the decision. To make this easier HTGet maps HTTP response codes to a smaller set of return codes. Here are some of the mappings:

0	No socket errors, but unknown HTTP code (should never happen)
1	Socket or protocol error
10	HTTP response code from 100 to 199
20	HTTP response code from 200 to 299 if not shown below
21	200 OK
24	203 Non-Authoritative Information
30	HTTP response code from 300 to 399 if not shown below
32	301 Moved Permanently
35	304 Not Modified
37	307 Temporary Redirect
40	HTTP response code from 400 to 499 if not shown below
41	400 Bad Request
42	401 Unauthorized
43	403 Forbidden
44	404 Not Found
45	410 Gone
50	HTTP response code from 500 to 599 if not shown below
51	500 Internal Server Error
53	503 Service Unavailable
54	505 HTTP Version Not Supported
55	509 Bandwidth Limit Exceeded

The general idea is to give the common response codes a unique return code and compress the rest into a generic return codes. For a complete list of the mappings refer to the source code.

Differences with the original WATTCP HTGet

As mentioned earlier, this is a complete rewrite of the original HTGet. A few variable names survived - not much else. Here are the major differences:

- HTTP 1.1 request support with a switch to use HTTP 1.0 if necessary
- The ability to hit Ctrl-Break, Ctrl-C or ESC to abort a transfer
- Far better performance due to better buffer handling.
- Return codes that try to give a better indication of what happened.

HTTPServ

Introduction

HTTPServ is a small HTTP (web) server designed for early IBM PCs and similar machines. Features include:

- HTTP/0.9, HTTP/1.0 and HTTP/1.1 support
- HTTP/1.0 and HTTP/1.1 persistent connections
- HTTP/1.1 request pipelining support
- Up to 8 simultaneous requests can be serviced with up to 8 more queued
- Ability to improve performance by serving pre-compressed content
- Easy mapping of incoming long URL path components to DOS 8.3 filenames
- HTTP Basic Authentication
- Improved performance through caching of directories and HTACCESS files
- Optional server status page
- Caching of selected files in RAM
- Optional logging
- Optional time synchronization with an SNTP server

HTTPServ is not a replacement for a modern HTTP server but if you are looking for an easy way to make the files on an old machine browsable from modern machine then it should work well for you. And if you are a hardcore retro-computing enthusiast, you can serve your static HTML pages directly from an old machine.

Special hardware or software requirements

- More memory: At least 384KB available memory, 512KB or more is recommended
- A mass storage device; serving from a floppy drive will work but it will be slower and limited.

HTTPServ needs more memory than most mTCP applications because it is assumed that it should try to handle multiple requests from connected clients at the same time instead of just queuing the requests and making the connected clients wait until it is their turn. You can use configuration options to reduce memory requirements at the cost of some performance and reduced function.

Unless you have very little content to server you will need a mass storage device such as a hard drive or a Zip drive. The faster the device is, the better – it is assumed that most of the objects being served will be served from disk storage and will not be cached in memory.

On older 8088 class machines some versions of DOS are extremely slow to calculate the free space on the hard drive. This can lead to long pauses where HTTPServ is unable to do any work, possibly causing TCP/IP errors. DOS 5 and later versions seem to be affected the most. For these slower machines DOS 3.3 is recommended.

Using HTTPServ

The HTTPServ command line looks like this:

```
httpserv [options]
```

Options are:

`-help`	Shows help for the command line
`-clients <n>`	Set the number of concurrent clients supported (default=8)
`-dir_indexes`	Enable generating directory listings (default is not to)
`-doc_root <dir>`	Set dir as the document root
`-port <n>`	Use port n for incoming requests (default is port 80)
`-shutdown_at <xx:xx>`	Shutdown at the specified time (24 hour format)
`-udp_logging <x.x.x.x:p>`	Enable logging over UDP

The -clients option allows you to set how many concurrent clients are supported. By default HTTPServ will try to serve up to 8 concurrent clients. You can select a lower number to reduce the memory required. (Each possible client takes about 10KB of RAM.)

The -dir_indexes option allows you to specify whether HTTPServ will automatically generate indexes for directories or not.

The -doc_root option tells HTTPServ the directory to use as the HTML document root. The content that you are going to serve must be in this directory or a subdirectory of it.

The -port option allows you to tell HTTPServ to listen on a port other than the standard HTTP port 80.

The -shutdown_at option is useful for having the server exit to DOS at a predetermined time each day. This allows you to sync the system time, rotate the logs, or do other maintenance tasks. The server must be running for at least two minutes before it will check for the shutdown time; this is to ensure that it does immediately trip over itself if it shuts down and restarts in the same minute. The shutdown time is also imprecise; it is only checked every 5 seconds and the server must be idle.

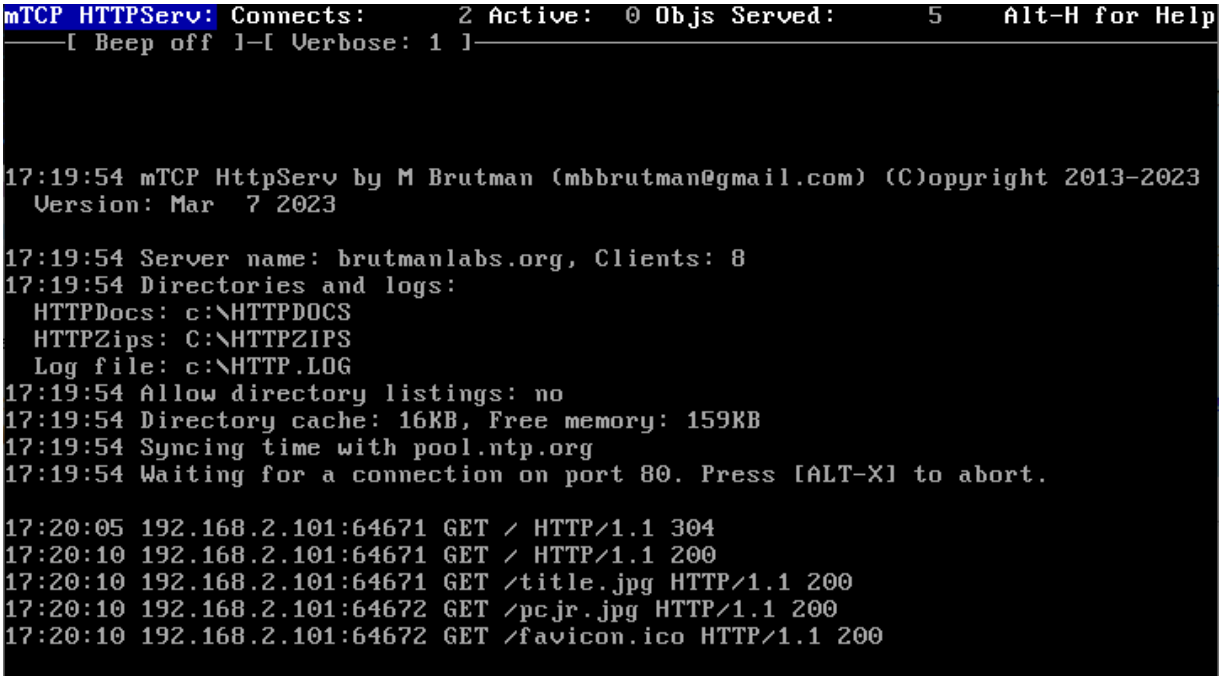
The -udp_logging option is useful for being able to check a running server without requiring the screen to be turned on. You can send log messages to a specific machine and use that as a remote logging device, or broadcast to your subnet allowing you to check the machine from any another machine on your network. (The Netcat utility can be used to listen to those messages.) Remember that remote logging to any machine is possible, while broadcast UDP will only work with machines on your local network. Also remember that UDP does not protect against lost packets, so you may miss some log entries.

The command line options are not not complete; many more configuration type options are set using the mTCP configuration file. Command line options are treated as a temporary override for anything found in the configuration file. This allows you to do quick tests or run HTTPServ for quick file transfers without altering your normal configuration options in the configuration file.

For a more detailed discussion of these options see their corresponding configuration file options in the "Configuration options" section.

Console messages and controls

The HTTPServ console is designed to show you what is going on at the moment. At a glance you can see the total number of connections served, the total number of objects served, and the current number of active connections. Recent requests and error messages are shown on the screen.



The following key combinations can be used:

Alt-H	Show help text
Alt-B	Toggle beeping on and off
Alt-V	Set the verbosity level
Alt-X	Shutdown the server

The state of the beeper and the verbosity level are shown on the second line of the screen.

The default verbosity level is 1, which gives you an overview of what each request is and what the return code was. At verbosity level two you can see some of the additional fields from the request, including the user-agent and referrer. Verbosity level 3 includes debug messages which are probably not very interesting.

If a log file was provided in the configuration options then all messages from the server are logged to the log file; the verbosity level just determines what gets shown on the screen. (Well, almost. Debug messages do not get logged unless you are also showing them on the screen.)

Serving static content

Like on other servers, the content that you want to serve is located under a directory called the document root. On HTTPServ you can set the document root using the command line or in the configuration file. All URLs served by HTTPServ are relative to the document root.

Remember, you are on a DOS machine so there are some limits you have to observe:

- DOS filenames and subdirectories have to use the 8.3 filename.ext format.
- DOS filenames are not case sensitive.
- DOS does not like filenames that contain special characters, like spaces

Serving content is a fairly straightforward process. The URL should identify a directory or filename relative to the document root. Just start at the document root, open the directory or filename specified by the URL, and send the bytes over the socket. If the URL does not exist send back an error message instead. In reality serving is more complex, but this is the basic idea.

If you are designing a site from scratch to be served by HTTPServ then all of the links have to follow the DOS filename conventions. This can make the links look strange, as they will look terse compared to other web sites.

If you are moving existing content or you want to use URLs with longer names than DOS allows there is a mapping function you can use to map the longer names to the actual DOS 8.3 filenames. That allows you use use links like "http://servername/ThisIsALongSubdir/MyContent.html" instead of "http://servername/long.sub/mcontent.htm". The magic happens each time a request comes in; the parts of the URL are parsed and optional mappings from long names to short DOS 8.3 filenames are done. If a mapping exists the server substitutes in the proper DOS 8.3 filename automatically.

To use this feature please see the "Map" directive in the HTACCESS files section. Using HTACCESS files can slow things down and they require memory but when used correctly the performance and memory costs are minimal. This feature also makes it easier to move existing content to a DOS machine for serving without having to rewrite all of the links within the content.

Serving pre-compressed static content

Sending data over a network requires both network bandwidth and CPU usage. The CPU usage to send and receive data is roughly proportional to the amount of data being sent.

Compressing data before sending it requires additional CPU usage for the compression step but it can result in less CPU used during the transmit stage and less network bandwidth required. Depending on your network connection and CPU speed, compressing content before sending it can make sense. Modern HTTP servers usually choose to compress content because their CPUs are so powerful and preserving network bandwidth is more important.

HTTPServ can not compress content on-the-fly but it can serve compressed content if you prepare it ahead of time. Any content that you choose to compress ahead of time is stored on disk in a directory tree that parallels the HTTPDocs directory hierarchy. This parallel directory tree is known as HTTPZips.

At run-time when a client makes a request HTTPServ will search the normal document root (HTTPDocs) for the file the user requested. If the client indicates that it can accept gzip'ed content HTTPServ will then check HTTPZips for a compressed version of the file. If a compressed version of the file is found it will be sent instead of the standard version of the file. If no compressed version is found, the standard version of the file is sent instead.

HTTPServ expects the directory hierarchy in HTTPZips to be the same as it is in HTTPDocs. It will look for the compressed version of the file using the same path, substituting your HTTPZips directory for your

HTTPDocs directory when looking for compressed content. If the compressed file is not in the same relative place it will not be found. (Note: HTACCESS files can be used to alter how URLs are mapped to directories. Serving compressed content does not change this; the HTACCESS files in the the normal document root are still used. The only change that happens is that instead of reading the final target out of HTTPDocs HTTPZips will be tried instead.)

HTTPServ also has no idea if the compressed version of a file is equivalent to the standard version of the file. You need to make sure that the two stay in sync; if you update a file in HTTPDocs then you should update the compressed version in HTTPZips or delete it until you can update it.

To make compressed content available follow these high-level steps:

- Create an HTTPZips directory and subdirectories under it that mirrors your HTTPDocs directory hierarchy.
- Gzip files that will compress easily; human readable text files are good candidates. JPG, ZIP archives, and other files that are already compressed will not benefit from further compression.
- DOS filenames are limited to the standard 8.3 format. Convention states that gzip'ed files should have a ".gz" suffix on them, but that is not possible in DOS. That is why HTTPZips is a separate parallel directory hierarchy - it is assumed that everything located in HTTPZips has a ".gz" suffix on it. As a result, the compressed files will have the same name as the original files but in a different directory to make it clear that they are compressed without a .gz suffix.

Pre-compressing the content requires additional disk space and some extra file management but if you have the space I recommend doing it. It cuts down on the amount of CPU time required for requests that can be served from the compressed copies, giving the clients faster response time and letting you serve more potential clients.

The use of HTTPZIPS is completely optional. If you are only an occasional user of HTTPServ then it probably is not worth the extra work.

Protecting files and directories

HTTPServ supports HTTP Basic Access Authentication. With "basic auth" there are no cookies, sessions or login pages. But there is minimal security; the userID and password is transmitted by the web browser "in the clear." This means that basic auth is best suited to things that you want to lightly restrict access too, but nothing that would be harmful if it were more widely available.

To protect a file using basic auth you mark the file as belonging to a realm. A realm is a tag (often a description) that allows you to logically group files that need protection together. Each realm has a password and a list of users associated with it.

When a user requests access to a protected file a challenge is sent back asking for a userID and password. The challenge generally appears as a dialog box on the web browser with a message that includes the realm the file is part of. The user then enters a userID and password and the browser retries the request, sending the information along in an additional HTTP header. If the entered userID and password are correct for the realm the user is granted permission to view the content.

Web browsers often cache the entered userID and password for a few minutes. If the user tries to access another file protected by the same realm the caching allows the web browser to retry the request without prompting the user for the userID and password again. (If the realm different or not in the cache then the user will get the

dialog box.)

Files and directories marked with a particular realm do not need to be located near each other; realm is a tag that can be applied wherever it is needed.

HTTPServ lets you mark a file or directory as access protected using the FileAuth directive in the HTACCESS file where the file or directory resides. One very important thing to be aware of is that if you mark a directory as being protected, it is assumed that if the user gets the userID and password correct that everything under that directory (files and subdirectories) is also available. This is true even if those files and directories have different protection specified - that protection is ignored because the user has been granted access at a higher level. This is a shortcoming of basic auth - only one authorization header can be sent for a request so it is not possible to provide layers of protection for files and directories along a URI path.

Realms are defined in a text file that you provide. The HTTPSERV_REALMS configuration directive in the mTCP configuration file tells HTTPServ where to find the file. Each line in the realms file defines a realm, the password and a list of users authorized to the realm. Below is a sample REALMS.TXT file:

In your mTCP configuration:

```
HTTPSERV_REALMS e:\realms.txt
```

In e:\realms.txt:

```
"Server functions" s3cr3t   root
"Restricted Files" passw0rd brutman eddie
"Restricted Dirs"  pfft123  brutman warez anon
```

In this example three different realms are declared. The realms have spaces in their names so they are surrounded by quotes to indicate that both words are part of the realm name. Next comes the password and after comes a list of users that are valid for the realm. All fields are delimited by spaces and have the following length restrictions:

- realm name: Up to 50 characters (not including quotes)
- passwords: Up to 12 characters
- userids: Up to 16 characters

The length of each line is limited; as of this writing the maximum line length is 255 bytes.

Do not locate your realms file under HTTPDOCS where it can be served. That's just an accident waiting to happen.

Serving to the public Internet

Teaching you how to run a public HTTP server is beyond the scope of this document. I assume a basic knowledge of networking and security. These are just a few reminders and things you should be aware of.

Security, Security, Security ...

If you connect a machine to a network you open it up to a wide array of possible problems. Viruses or malware may try to infect your machine. People might try to find sensitive information on your machine. People might

try to crash your machine just for fun.

I have made great efforts to protect against these things but nothing is perfect. You are responsible for not exposing sensitive information; if you don't want it accidentally published don't put a web server on the machine. I have tried to protect against buffer overruns, malicious URLs, etc. but this work pretty much never ends which is why it needs to be tested and bugs need to be reported. And lastly, if somebody decides to DoS (Denial of Service) your DOS machine with a cell phone or a faster machine, there is not much that can be done.

You take responsibility for running a web server. It can be fun, but it takes some care. Here are some things I've done to make HTTPServ safer:

- If something is protected with HTTP Basic Authentication then the default is to fail to serve the file unless the userid and password for the realm match up.
- HTACCESS.--- files are never served.
- The server aggressively defends against buffer overflow attacks.
- The server does not allow access to special filenames, like LPT1:, $CLOCK, etc. If it is not a file or a directory, it will not be served.
- The server times idle connections out automatically, and more quickly if connections are limited.
- A fair amount of input checking is done to weed out malformed requests.

There are probably things that will break and I will fix them as they come up.

Your ISP probably hates private servers

Most residential grade internet service providers prohibit running servers on your network connection. The usual method of enforcement is to block incoming requests to "well known ports" that correspond to services that you might want to run.

The usual way around this restriction is to run your service on an alternative port. For example, the standard port for HTTP traffic is port 80. Incoming requests to port 80 will be blocked by your ISP but if you choose to use port 8080 instead the traffic will probably not be blocked. (And for real vintage computing street cred, use port 8088.)

The HTTPSERV_PORT configuration option can be used to tell HTTPServ what port to use for incoming requests. If you use a non-standard port the port must be included in any URLs that point to your server. For example, while http://yourserver.net/ works for a server running on the standard HTTP port the same server using port 8080 would need a URL that looks like http://yourserver.net:8080/.

Configuring your home firewall

Most people have a home firewall of some sort that allows you to have multiple machines in your house sharing the same connection from your ISP. While this is a valuable function the firewall also serves another valuable purpose - it generally blocks incoming requests from arbitrary machines on the Internet. One reason for doing this is that a request from an unknown machine might be hostile. The other reason for doing this is that the firewall really does not know which machine on your network should handle a particular incoming request. Dropping those incoming requests helps improve security but it becomes a problem when you are trying to run a web server; all of the traffic to the web server looks like incoming requests from possibly hostile machines on the Internet.

Most home firewalls can be configured to allow incoming requests on the public IP address on specific ports to be sent to a specific machine and port number on your internal network. This solves the technical problem of where to route incoming requests to. The forwarding is done on a port by port basis so you generally tell the firewall to only forward anything coming in on port 80 to the specific machine in your network running the web server. Requests on other ports still get dropped.

Keep in mind that your ISP is probably blocking incoming requests to port 80 and that you probably have to use an alternate port, such as 8080. If that is the case then your port forwarding should also reflect that; incoming requests to your cable/DSL modem on port 8080 will need to be forwarded to your HTTP server. The HTTP server can be running on any port; they do not have to match. This allows you to set up HTTPServ to run on port 80 on your private LAN and only people connecting from outside will have to use the non-standard port that you chose to get around your ISP. (The firewall handles the forwarding of the traffic including the mapping of ports automatically.)

Telling people your server name

HTTP servers sometimes have to give a response code that includes follow-up information, such as a new URL to use. A good example of this happens when somebody sends a request for a directory but does not include the trailing slash ('/') - the HTTP server will redirect the client to use the correct name for the directory which has the trailing slash.

For that mechanism to work the server has to tell the client a full URL, including the server name. The HTTPSERV_SERVERNAME configuration option is used for this purpose. If it is not set the HTTP server will use its assigned internal IP address which will work on your internal network but it is not valid for outside connections.

The HTTPSERV_SERVERNAME configuration option should be set to a fully qualified host name that can be resolved using DNS. For machines on a home internet connection dynamic DNS can be used to keep the DNS address up to date, isolating you from changes to the IP address assigned to you by your ISP. If you do not have a domain name you can use the numerical IP address assigned by your ISP, but if that changes you will need to detect the change, stop the server, update HTTPSERV_SERVERNAME, and restart the server. (As you can see a real domain name with dynamic DNS is far easier to work with.)

Example:

- If the internal IP address of your HTTP server is 192.168.1.124 and it is listening on port 80 and …
- Your cable/DSL modem (firewall) has a public address of 67.43.212.82 and …
- Your ISP blocks port 80 so you are using port 8080 for connections from the public Internet and you have set your router up to forward them to your HTTP server, then ...

- Your HTTPSERV_SERVERNAME can be set to 67.43.212.82:8080.

However, if your ISP also doesn't give you a static IP address that might change, and the HTTP server has no way to know when that happens. If you want to guard against that you can use Dynamic DNS:

- Prerequisite: your router supports Dynamic DNS.
- Your internal IP address and port and the router forwarding remain the same.
- You register a new domain name and set it up for Dynamic DNS.
- You set HTTPSERV_SERVERNAME to the registered name and port: (yourserver.org:8080)
- When your home router detects the assigned IP address from the ISP changes, it updates the DNS

system automatically.
- HTTPServ always refers to itself by the DNS name, so it never goes out of date.

In summary, try not use a numerical IP address with HTTPSERV_SERVERNAME unless you know your IP address doesn't change. If it is subject to change then using a hostname resolvable with DNS and keeping it updated with dynamic DNS is the best option.

Time synchronization

DOS was not designed to keep accurate time over days or weeks. Without some form of time synchronization the clock will drift, which can affect clients that use the timestamp to know if content has been updated.

If you are going to be running HTTPServ for days or weeks then you probably want to configure it to synchronize time with an SNTP server. Do this using the HTTPSERV_SNTP_SERVER configuration parameter. When that is set to an SNTP server HTTPServ will check the time every 10 minutes. If there is more than a five second difference the system time will be updated.

To ensure that the time server address does not become stale use a DNS address (human readable host name), not an IP address. The DNS address will be re-resolved periodically to an IP address, allowing HTTPServ to keep up with address changes. You can use a numerical IP address instead but that is discouraged as it breaks this feature and SNTP servers are often organized as pools of machines so you are potentially circumventing their load balancing scheme.

Enabling HTTPS connections

The mTCP HTTP server does not support HTTPS, but it may be desirable to allow for HTTPS connections.

This can be done by using a reverse proxy, which takes the SSL connection and forwards it to the mTCP HTTP server without the encryption. Apache and other web servers can be used to implement a reverse proxy.

Miscellaneous tips

Choosing files to cache in memory

Since most clients support gzip'ed content you should cache only gzip'ed versions of your content to stretch your memory; clients that do not support gzip'ed content are rare.

Cached files need to be under 64KB in size; if a file is above that size the request will be rejected. You get more bang-for-the-buck by caching several small files instead of fewer large ones. The reason for this is that there is overhead to finding and opening a file and that overhead is the same whether the file is large or small. As a fraction of total transfer time that overhead is larger for small files than it is for large files. So it makes sense to try to eliminate that overhead for small files.

Files that are used across multiple pages are good candidates for caching. This includes CSS files, small logos and icons, etc. Your site's home page and any pages that you think will be visited often are also good candidates for caching.

Log file considerations

If you are running a server for your own personal use on a private LAN then you probably do not need to have the log file enabled. Not having a log file may provide some modest speed improvement. It will also cut down on the required memory.

If you are running a public server then you really should have logging turned on.

The data written to the log file is pretty verbose. Make sure that the drive the log file resides on has enough space; a few megabytes is a safe number to start with. During my large scale test HTTPServ served 7175 requests and generated a logfile that was 6.4MB in size. That is an average of 900 bytes for each request.

In the future logs will either be written in a compressed format at runtime or there will be a "log rolling" function added to periodically close the log file, compress it, and start another one.

Do not write your log file to a RAM disk. If the machine crashes you will lose the log file.

Using disk caching software

If you are on a machine that has expanded or extended memory that can be used to cache disk operations then you should use it. Such caching software will generally cache the most used disk blocks, automatically removing blocks from the cache that are no longer being used. This makes more sense than using conventional memory to cache specific files using the cache function of HTTPServ; you have much more memory available to you if you have expanded or extended memory.

The only danger might be that the log file writes would also be cached; if something crashes or you lose power you might lose some of your logs.

HTTPServ Configuration options

HTTPServ can be run using command line arguments only; no further configuration is needed unless you want to use advanced options. If you want to set configuration options they can be placed in the mTCP configuration file pointed at by the MTCPCFG environment variable.

The following options can be used to configure HTTPServ behavior:

`HTTPSERV_HTTPDOCS`	Set the document root
`HTTPSERV_HTTPZIPS`	Set the document root for pre-compressed files
`HTTPSERV_LOG_FILE`	Set the location of the log file
`HTTPSERV_REALMS`	Set the location of the realms file for HTTP Authentication
`HTTPSERV_SERVERNAME`	Set the name used when sending responses to clients
`HTTPSERV_PORT`	Set the port used for incoming HTTP requests
`HTTPSERV_VERBOSITY`	Control how verbose the on-screen messages are
`HTTPSERV_DIR_INDEXES`	Set the default for generating directory indexes
`HTTPSERV_DIR_CACHE`	Set how much memory is used for caching HTACCESS files
`HTTPSERV_CLIENTS`	Set the number of concurrent clients to support
`HTTPSERV_LONG_EXPIRE_HOURS`	Expiration time for static content
`HTTPSERV_SNTP_SERVER`	The SNTP server to use for time synchronization

Each configuration option is explained in more detail below:

`HTTPSERV_HTTPDOCS`

Equivalent command line option: -doc_root

Format:

```
HTTPSERV_HTTPDOCS <fully qualified dos directory name>
```

Every web server has a "document root" that is a directory that serves as the root of the directory tree for files being served by the server. This configuration option is used to tell HTTPServ which directory to use as the document root.

The document root should be provided as a fully qualified path including a drive letter. Files in the directory and subdirectories within the doc root can be seen by people browsing the server; files outside of the doc root are hidden and not available.

Examples:

```
HTTPSERV_HTTPDOCS d:\httpdocs
HTTPSERV_HTTPDOCS c:\
HTTPSERV_HTTPDOCS e:\public\toshare
```

If you do not set this configuration option then you can set it on the HTTPServ command line using the "-doc_root" option. If both the -doc_root option and HTTPSERV_HTTPDOCS are specified then the -doc_root will have priority and be treated as though it was overriding HTTPSERV_HTTPDOCS.

`HTTPSERV_HTTPZIPS`

Equivalent command line option: Not applicable (see below)

Format:

```
HTTPSERV_HTTPZIPS <fully qualified dos directory name>
```

HTTPSERV_HTTPZIPS tells HTTPServ where your pre-compressed files are. Pre-compressed files allow for better performance at the cost of additional disk space. See Serving pre-compressed static content for the details on how it works.

Example:

```
HTTPSERV_HTTPZIPS d:\httpzips
```

The use of HTTPZIPS is completely optional.

HTTPSERV_HTTPZIPS is only valid if the document root is set using HTTPSERV_HTTPDOCS in the configuration file. If you use the -doc_root command line option to temporarily change the document root then serving pre-compressed content will be disabled. No equivalent command line option for

HTTPSERV_HTTPZIPS is provided.

HTTPSERV_LOG_FILE

Equivalent command line option: Not applicable

Format:

```
HTTPSERV_LOG_FILE <fully qualified dos file name>
```

If you are going to run any sort of public server you probably want a log file so you can keep track of which IP addresses are connecting to your server and what they are doing. This configuration option is used to tell HTTPServ which file to use as the log file.

HTTPServ always appends to this file. As of this writing HTTPServ does not compress log files or rotate them so you should locate your log file on a drive letter that has plenty of available space and you should periodically check to see if the log files needs to be compressed and moved offline. Do not locate your log file in an area of the disk being served by HTTPServ - that is a potential security problem.

HTTPSERV_LOG_FILE is optional. Using it can slow the server down a little bit but if you are allowing traffic from outside of your network you really should use it. Please see Log file considerations for more discussion on the log file and logging techniques.

HTTPSERV_REALMS

Equivalent command line option: Not applicable

Format:

```
HTTPSERV_REALMS <fully qualified dos file name>
```

If you are using HTTP Basic Authentication to protect any of your files or directories then you need a realms file to define the password for each realm you are using in the HTACCESS files. Please see Protecting files and directories for the details on how to use this file.

HTTPSERV_SERVERNAME

Equivalent command line option: Not applicable

Format:

```
HTTPSERV_SERVERNAME <alternative_server_name>
```

HTTPServ knows the IP address of the machine that it is running on but it does not really know what the hostname of the machine is. Hostname can be ambiguous, especially for machines that are DHCP clients or that are behind firewalls. The hostname that you tell clients to use may have nothing to do with the hostname the machine is configured for, especially if port forwarding is involved.

If HTTPServ is going to serve traffic to machines outside of your immediate LAN or to machines that are on the other side of a firewall that changes network addresses then you probably did some configuration to get the requests routed to HTTPServ. Whatever host name you are telling the clients to use to connect to HTTPServ should be set on this configuration option.

One way in which this is used is when the client requests a directory but does not include a trailing slash ('/'); that requires HTTPServ to send a redirect (301 HTTP header) to let the client know that the correct URL includes a trailing slash. The redirect message needs to have a server name that the client can use.

Example:

(Assume your machine is behind a firewall that uses NAT and the address is 192.168.1.124)

```
HTTPSERV_SERVERNAME bogus.dyndns.org:8080
```

In this example you are using a dynamic DNS service and port forwarding to provide connectivity to your machine from the outside world. The real IP address does not make sense to clients outside of your network so if your server needs to send a response that includes the server name, it will use "bogus.dyndns.org:8080" instead of 192.168.1.124.

See the section entitled Serving to the public internet for a more detailed explanation of this and other firewall related issues.

HTTPSERV_PORT

Equivalent command line option: -port

Format:

```
HTTPSERV_PORT <n>
```

By default HTTPServ listens on the standard WWW port, which is port 80. If you need to use an alternative port use this option or the "-port" option on the command line.

HTTPSERV_VERBOSITY

Equivalent command line option: Not applicable

By default HTTPServ prints a minimal amount of information on the console about each request it is serving. If you would like to see more detail for each request you can set the verbosity level higher. Supported levels are "low" (the default), "medium" and "high".

Example:

```
HTTPSERV_VERBOSITY medium
```

"Medium" is actually pretty verbose. "High" is used for debug messages.

`HTTPSERV_DIR_INDEXES`

Equivalent command line option: -dir_indexes

By default HTTPServ will not generate directory listings for directories; this is for security reasons. If you would like to change the default to enable directory listings then use this option.

Enabling directory listings by default is safe and harmless if you are the only person using the server. But if you are going to allow other people to browse your machine you should disable directory listings and provide a static, per directory file that links to what you want users to be able to see. That gives you an opportunity to provide descriptions for each file next to the link to it. For some cases like a large shareware repository that might not be practical. Use your best judgment.

Also keep in mind that just because a file does not have a link to it does not mean that an external user won't be able to see it. They can always guess at links. If you don't want something served, don't put it on your server.

Legal values for this option are "yes" and "no." A setting for a specific directory in its HTACCESS file overrides this global setting.

Example:

```
HTTPSERV_DIR_INDEXES yes
```

That will enable directory indexes by default across all directories.

`HTTPSERV_DIR_CACHE`

Equivalent command line option: Not applicable

Serving files using HTTP is a fairly expensive operation for an old machine, especially when HTACCESS files are used. There may be an HTACCESS file in any directory and each HTACCESS file may have a lot of directives. Serving a file requires walking the entire path in the URL and examining each directory and possible HTACCESS file along the way.

To make this reasonable on slower machines the directory structure under HTTPDocs and the contents of HTACCESS files are cached in memory. The cache can be set from 8KB to 63KB in size.

Example:

```
HTTPSERV_DIR_CACHE 32
```

That will set the directory cache to 32KB in size.

If you make extensive use of HTACCESS files you should increase this value to the maximum setting. You can run with less but it might slow things down. The cache is fixed in size and not very sophisticated; to

keep the programming and memory management reasonable some compromises were made. If the cache overflows it gets erased and rebuilt starting with the data from the current request. If the cache is too small it will keep getting erased, which is almost like having no cache at all.

HTTPSERV_CLIENTS

Equivalent command line option: -clients

Format:

```
HTTPSERV_CLIENTS <n>
```

Modern web browsers open multiple connections to a web server at the same time to speed the processing of requests. To the web server each connection is an independent transaction so one clients with multiple connections is treated the same as many clients each with one connection.

HTTPServ allows you to set the number of simultaneously connected clients that it can serve requests to. The default is eight, meaning that eight different clients (or one client being a pig) can connect at the same time, issue a request, and have the request processed. HTTPServ will do its best to service the requests; interleaving the processing. This is better than just simply queuing the requests and processing them one at a time. However, it is not magic - when interleaving the processing everything takes longer.

If more than <n> requests arrive at the same time the additional requests will be queued and made to wait.

HTTPSERV_LONG_EXPIRE_HOURS

This option sets the expiration time for served objects. By default the server will tell the client that content is cachable for up to 24 hours. This allows clients to use their cached version of the content or do a lightweight expiration check on content if they request the object again in the same day.

Example:

```
HTTPSERV_LONG_EXPIRE_HOURS 168
```

This instructs HTTPServ to tell the client that content is cachable for up to 168 hours (1 week).

HTTPSERV_SNTP_SERVER

This option sets the address of the SNTP server to synchronize time with. When used HTTPServ will check the time every ten minutes, and if more than a five second difference is detected the system time will be updated. If a human readable host name is used the DNS address will be re-resolved periodically to ensure that a stale address for the server is not used.

Example:

```
HTTPSERV_SNTP_SERVER pool.ntp.org
```

This instructs HTTPServ to use pool.ntp.org to synchronize the time.

HTACCESS Files

(Note: The actual filename used for an HTACCESS file is "HTACCESS.---" . That should be obscure enough to avoid collisions with real filenames.)

In any directory that has content being served by HTTPServ you can create a special file called HTACCESS that can be used to tell HTTPServ to perform special functions. These special functions include marking files and directories as protected using HTTP Basic Authentication, mapping parts of a URL to the standard DOS 8.3 filename format, providing a default file to display when the user navigates to a directory (instead of providing a directory list), etc.

An HTACCESS file in a directory only affects the handling of names in that directory. Names includes names for real files and names for subdirectories in that directory. If you navigate away from that directory the HTACCESS file is no longer relevant; nothing is inherited from it. (Well, almost nothing ... read the description of the FileAuth directive for details.)

Each line of the file has a directive and some options. All fields on the line are delimited by a single space. Below is a description of the directives.

Map Directive

HTTP URLs are supposed to loosely map to subdirectories and files on the HTTP server. Exceptions to this include parameters in the URL and paths that refer to generated content. As HTTPServ runs on DOS it expects to serve URLs that are composed of legal DOS file and directory names. This becomes a problem when you consider that DOS file and directory names are at most 8 bytes long with an optional 3 byte extension - the dreaded 8.3 format. DOS file and directory names are also not case sensitive.

Map lets you get around these limits by allowing you to define long, case sensitive names that are aliases for DOS 8.3 names. This is useful if you want to host an existing web site where the existing HTTP URLs do not conform to DOS 8.3 naming conventions. Of course the files themselves have to conform to DOS 8.3 naming conventions but you do not have to go into the files and change all of the links to match or update all of the existing URLs.

The Map directive has the following format:

```
Map NiceLongName filename.ext
```

where:

- NiceLongName is a case sensitive name up to 51 bytes in length. If this name contains spaces please wrap it in quotes.
- filename.ext is the DOS 8.3 filename that should be used in place of NiceLongName. This name is not case sensitive.

For example:

```
Map Downloads.html dls.htm
```

That example says that if the URL from the user has "Downloads.html" as the filename in it that HTTPServ should convert that to "dls.htm" instead.

When HTTPServ is processing a URL it "walks" through each part of the URL and considers it on a component by component basis. The mapping if long names to short names is done at each stage of the walk. Consider the following URL:

```
http://yourserver.com/Reference/Downloads/Downloads.html
```

In the path part of the URL we find three parts- "Reference", "Downloads", and "Downloads.html". Reference and Downloads refer to subdirectories and Downloads.html refers to a file. But none of these are valid DOS 8.3 filenames that HTTPServ can read.

To handle this URL we can do the following. Assume that your server is configured with "E:\HTTPDOCS" as your HTTP document root and that E:\HTTPDOCS\REF\DLS\DLS.HTM is the full path that should be served for that request. "E:\HTTPDOCS" is the root of your server and it has an HTACCESS file that includes a Map directive:

```
Map Reference REF
```

That allows HTTPServ to move into the real REF subdirectory when it sees "Reference" at the first level of the URL. Inside of the REF subdirectory there is another HTACCESS file that includes another Map directive:

```
Map Downloads DLS
```

That allows HTTPServ to move into the DLS subdirectory when it sees "Downloads" at the next level of the URL. Inside of the DLS subdirectory we need one more HTACCESS file with one more Map directive:

```
Map Downloads.html DLS.HTM
```

That allows HTTPServ to substitute DLS.HTM as the file to be sent to the web browser.

Mapping only happens once per URL component. You can not have one mapping map to another long name, which in turn maps to something else; if that were allowed a user mistake could lead to an endless loop.

FileAuth Directive

The FileAuth directive lets you mark a file or directory as being protected by an HTTP Basic Authentication realm. (See Protecting files and directories above for an overview.) The format of the FileAuth directive is:

```
FileAuth filename.ext realm
```

where:

- filename.ext is the DOS file or directory that you are protecting in DOS 8.3 format. Only DOS 8.3 format file and directory names may be protected; if you are using long names in your URLs you need to use MAP directives to convert them to DOS 8.3 format to be able to serve them.
- realm is a string that describes what you are protecting. Most web browsers show this string to the user when prompting for a userid and password. The realm must be defined in a "realms" file that your

mTCP configuration tells HTTPServ to use. If the realm is not found in that file nobody will be able to authenticate and see the file or directory.

Example:

```
FileAuth secret.txt "Sensitive Files Area"
```

The example shows that "secret.txt" is marked as belonging to realm "Sensitive Files Area" . Only users listed in the realms file who know the password will be able to see secret.txt.

GeneratedPage Directive

mTCP HTTPServ has two special pages and two special functions that can be accessed using a URL. By default these pages/functions are not available. You can use the GeneratedPage directive to make them available as a page in the directory where the HTACCESS file is located.

The format is:

```
GeneratedPage filename.ext special_function
```

where

- filename.exe is a DOS 8.3 format filename that will be used as the name of the page/function. Yes, it is a little strange to require a DOS 8.3 format name for something that does not actually exist but it makes the code simpler.
- special_function names the special page or function. The current pages/functions are:

serverstatus	generate a page of statistics for the server
clients	generate a page that shows the connected clients
cachecmds	server cache control commands
serverquit	end the server immediately

Example:

(Assume these directives are in E:\HTTPDOCS\SPECIAL\HTACCESS.---)

```
Map Status status
Map Clients clients
Map QuitQuitQuit quit
GeneratedPage status serverstatus
GeneratedPage clients clients
GeneratedPage quit serverquit
FileAuth clients "Slighty sensitive"
FileAuth quit "Admin operations"
```

In this example all three special pages/functions are active. The server status page can be obtained by going to http://yourserver.com/Status or http://yourserver.com/status because we have made the page available as "status" and providing a mapping for "Status". The clients page is similar - it has a DOS 8.3 name and a mapping for "Clients". Clients is also marked as being in the "Slightly sensitive" realm so it requires a userID and password to view. The quit function is also available through two names and it is extremely dangerous so it is marked as belonging to the "Admin operations" real. Presumably your realms file has less users authorized to

"Admin operations" than to "Slightly sensitive".

See Special server functions for more details on how these functions work.

DirectoryIndex Directive

If a user uses a URL that turns out to be a directory the server may either generate a list of files in the directory, it may refuse access to the list of files in the directory, or it may serve a default file that exists in the directory. If INDEX.HTM exists in the directory it will be served by default. If INDEX.HTM does not exist you can specify what file should be served by default using this directive. The format is:

```
DirectoryIndex filename.ext
```

If the user requests a URL that ends in the directory where this is set then whatever file specified on this directive will be served to them. Use it if you want to serve a file when a directory is specified but you are not using INDEX.HTM.

If you use this directive the user will not be able to generate a directory list for the directory, even if that is explicitly allowed.

IndexAllowed Directive

The IndexAllowed directive allows you to specify if HTTPServ should generate a file list or not if a user specifies a URL that resolves to a directory. The format is

```
IndexAllowed yes|no
```

Yes and no are used to explicitly allow or deny a directory list to be generated for a given directory. If no directive is provided then the default is "no" unless that is overridden in the mTCP configuration file using the HTTPSERV_DIR_INDEXES configuration option.

Special server functions

Besides serving static files HTTPServ has some special functions that it can perform. The special functions can be used to show status pages or make the server perform operations such as caching content in memory or performing a shutdown.

By default all of the functions are enabled but there is no way to get to them. There is no hard-coded URL that you can use and find the functions. To enable a function use the GeneratedPage directive in an HTACCESS file to create an pseudo filename entry for it in a directory that corresponds to the special function. This seems complicated but it lets you define what the URL path to use to get to function and allows you to protect the function using the existing HTTP Basic Authentication method.

See the GeneratedPage directive in the HTACCESS Files section for the details on how to enable each of these special functions.

HTTPServ Status Page

This is a fun page that gives you statistics such as how long the server has been running, how much memory is available, how many objects have been served, TCP/IP statistics, etc.

Sample lines to enable the status page at /ServerStatus:

(Added to HTACCESS.--- in your document root.)

```
Map ServerStatus status.htm
GeneratedPage status.htm serverstatus
```

Line 1 creates a mapping from a long name ("ServerStatus") to the DOS 8.3 filename status.htm. Line 2 tells the HTTP server that whenever it sees a request for status.htm that it should generate the server status page. (As a side effect of the mapping the server status page is available as both /ServerStatus and /status.htm.)

HTTPServ Clients Page

This page shows you the IP address of currently active connections, and if possible what request is being served on each of them. You can only see active requests which limits the usefulness of the page.

Sample line to enable the clients page at /Admin/Clients:

(Added to HTACCESS.--- in the /Admin directory of your document root.)

```
GeneratedPage Clients clients
```

The first parameter of the line tells the HTTP server that whenever it sees a request for Clients that it should generate the clients list page. We're taking advantage of the fact that DOS filenames are not case sensitive, and that our desired URL ("Clients") is also a legal DOS filename.

File cache functions

HTTPServ lets you use RAM to cache files. If you choose the files to cache carefully this can really speed things up, especially on older systems with slow disk access.

Three cache functions are provided - one to read a file into memory, one to tell you what files are in memory, and one to clear all of the files from the cache. The cache functions are implemented with one handler that uses URL parameters to pass commands to the server.

To enable the caching function map it to a URL by using the GeneragedPage directive:

(Added to HTACCESS.--- in a directory under your document root where you want it to appear.)

```
Map CacheControl cache
GeneratedPage cache cachecmds
FileAuth cache "Admin Operations"
```

Line 1 creates a mapping from a long name ("CacheControl") to the DOS 8.3 filename cache. Line 2 tells the HTTP server that whenever it sees a request for cache that it should process file cache commands. (As a side effect of the mapping the file cache commands are also available as both CacheControl and /cache in the directory where the HTACCESS file was edited.) Line 3 specifies that the cache URL should be protected by basic authentication. (See Protecting files and directories for how to use basic authentication.)

The file cache function does not generate a page by default. It does accept the following URL parameters :

- add: add a file to the cache
- show: generate a page to show the files that are in the cache
- clear: wipe out the entire cache

To add a file to the cache send a parameter with the name "add" and the URL of the file to cache as the value of the parameter. HTTPServ will process the URL and resolve it to a DOS file, including doing any long name to DOS filename mapping specified in the HTACCESS files. For example:

```
http://192.168.2.148/proc/cache?add=%2Fpcjr.css
```

In this example I have a subdirectory called /proc hanging off my HTTPDOCS directory. The HTACCESS file for /proc defines a pseudo-filename called "cache" to enable the caching function. The "/pcjr.css" object (relative from the document root) is to be added to the cache. The string looks awkward because the special characters need to be escaped in the URL. An easier way to do this is to provide a simple HTML file in the directory with a form that looks like this:

```
<form action="cache" method="get">
  <INPUT type="text" name="add" />
  <INPUT type="submit" />
</form>
```

The form version is much easier to use and the browser handles escaping the special characters for you. The form method must be "get" as HTTPServ does not handle the post method yet.

When you add a file to the cache HTTPServ will check the document root to see if it exists. If the file does exist it will be read into memory and served from memory from that point forward. If a gzip'ed version is available in HTTPZips that version will be cached instead; most modern clients will request compressed content and this saves memory in RAM.

Remember, even though DOS is not case sensitive URL paths are. Make sure you use the proper combination of upper and lower case when specifying the URL paths of files to cache.

To see the list of files in the cache send a parameter with the name "show" and no value. For example:

```
http://192.168.2.148/proc/cache?show=
```

A form can also be used to simplify this:

```
<form action="cache" method="get">
  <INPUT type="hidden" name="show" value="">
  Show cached files: <INPUT type="submit"/>
</form>
```

The cache clear command is similar; the parameter name to use is "clear". Individual files can not be cleared from the cache; the entire cache must be cleared at the same time. This is done to prevent memory fragmentation in the DOS heap.

Remote shutdown function

The remote shutdown (quit) function does what you think that it does. Be careful; if you enable this you must

protect it with HTTP Basic Authentication or anybody who finds the link will be able to shut your server down remotely. And there is no "Are you sure?" prompt, so if you do invoke this function the server will shut down and return to DOS.

To enable the quit function map it to a URL by using the GeneragedPage directive:

(Added to HTACCESS.--- in a directory under your document root where you want it to appear.)

```
GeneratedPage quit serverquit
FileAuth quit "Admin Operations"
```

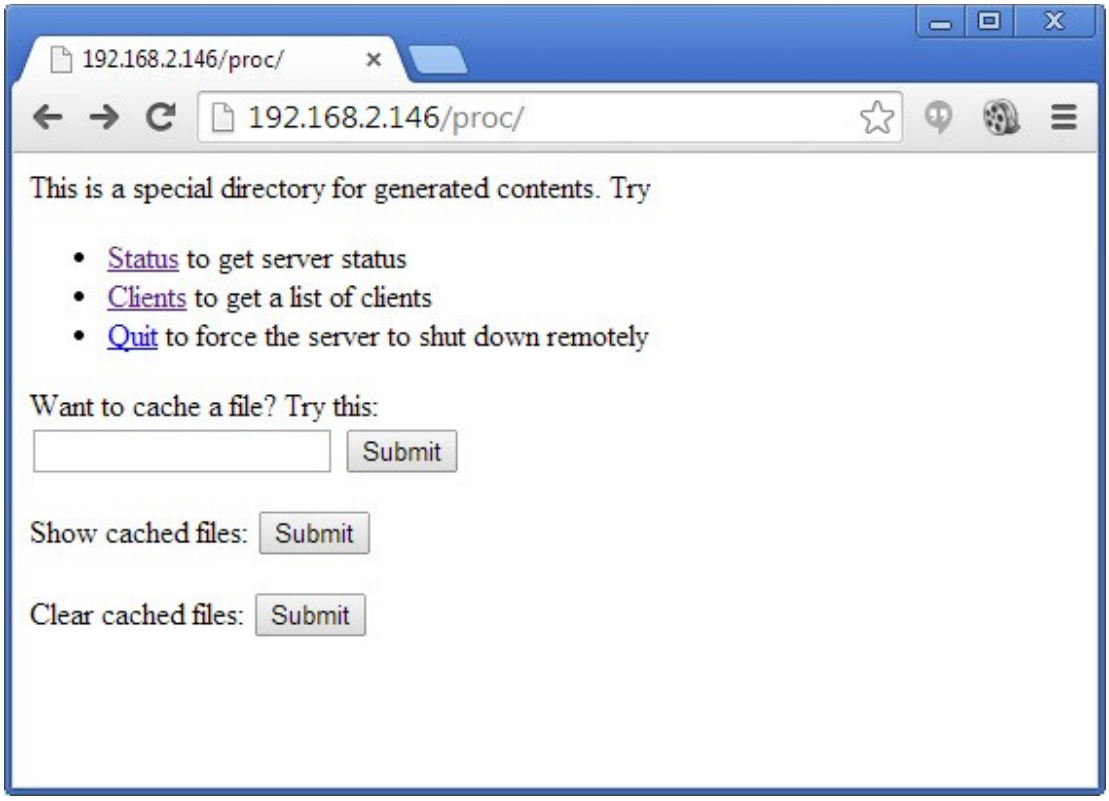
In this example the "quit" function is active in the directory tree where you edited the HTACCESS file. A realm called "Admin Operations" should be defined in the realms file, and it will include a password and the list of users who can see content or used functions tagged with "Admin Operations." (See Protecting files and directories for how to use basic authentication.)

Sample files for exposing the special server functions

Let's assume the following:

- Special functions live in a directory called "proc" that is a subdirectory directly under your HTTPDOCS directory.
- All functions are enabled.
- The clients page, cache and quit functions are protected by HTTP Basic Authentication.
- When a user goes to that directory a default file will be served to them to make things easier.

When a user navigates to the proc directory we want them to see something like this:



Here is a simple HTML file we will call default.htm that will be served to the user when they go to /proc on your server:

```
<html>
  <head></head>
  <body>
    This is a special directory for generated contents.  Try
    <ul>
      <li><a href="Status">Status</a> to get server status
      <li><a href="Clients">Clients</a> to get a list of clients
      <li><a href="Quit">Quit</a> to force the server to shut down remotely
    </ul>
    Want to cache a file?  Try this:
    <form action="cache" method="get">
      <INPUT type="text" name="add" />
      <INPUT type="submit" />
    </form>
    <p>
    <form action="cache" method="get">
      <INPUT type="hidden" name="show" value="">
      Show cached files: <INPUT type="submit"/>
    </form>
    <p>
    <form action="cache" method="get">
      <INPUT type="hidden" name="clear" value="">
      Clear cached files: <INPUT type="submit"/>
    </form>
  </body>
</html>
```

And here is the HTACCESS file for the /proc directory. It does the following:

- Tells HTTPServ to serve default.htm by default when no filename in the directory is specified.
- Defines the four "pseudo filenames" that represent the four server functions.
- Declares three of the functions to be protected by the "Server Functions" realm for HTTP Basic Authentication.

```
directoryIndex default.htm
indexAllowed   no
GeneratedPage  status  serverstatus
GeneratedPage  quit    serverquit
GeneratedPage  clients clients
GeneratedPage  cache   cachecmds
FileAuth       quit         "Server Functions"
FileAuth       clients      "Server Functions"
FileAuth       cache        "Server Functions"
```

Remember, DOS filenames are not case sensitive. The HTML file is using "Status" for a reference, which refers to "status" on the GeneratedPage directive in the HTACCESS file. And there really is no file in /proc called "status" - it is a pseudo filename that results in a status page being generated.

IRCjr

Introduction

This is a fairly simple IRC client that I have been working on since the summer of 2008. Originally it started as a bare bones program that let you connect to one channel on one server. All messages came out on the same screen and sending private messages to other users was a hassle. But, it worked!

Things have changed quite a bit since then. The program has grown in size since those early versions but there are a lot of improvements that make the program much easier to use. IRCjr now supports multiple open channels and private conversations, Client to Client Protocol (CTCP) messages, etc.

Special hardware or software requirements

None

Note for genuine IBM CGA Users: see the section on "Removing CGA Snow" if you are bothered by glitches/noise caused by direct screen writes.

Set up

IRC networks require you to specify a nickname that other people will know you by, a user name which is the userID that you use on your local machine, and a real name. The nickname is the most important of these, as your nickname will be used the most. The user name and real name can not be verified or trusted and most people probably don't provide accurate information for them anyway.

The nickname, user name, and real name are set by adding lines to the mTCP configuration file. The parameter names are:

IRCJR_NICK	Your nickname for IRC
IRCJR_USER	Your username on this machine
IRCJR_NAME	Your name in real life

Here is an example that shows how to use these:

```
# IRCjr parameters for nickname, user name, and real name
#
ircjr_nick Zoidberg
ircjr_user Jzoidberg
ircjr_name John Zoidberg
```

The official IRC protocol specification states that nicknames are a maximum of 9 characters long. That is routinely ignored and you will find that your favorite IRC servers support much longer nicknames. IRCjr allows you to use up to 50 characters for a nickname but if your server does not support nicknames that long it may truncate the nickname or give you an error message.

Some IRC servers are picky about the user name and real name and want you to use something plausible or they

will reject your connection. For example, if you pick a single word for your real name it is possible that the server will assume that you are a "bot" and reject your connection. If you are having trouble connecting to a server check the error messages carefully and see if it is complaining about these fields.

Starting IRCjr

The command line arguments are:

```
ircjr [-port <n>] <server_name> [<channel>]
```

The server name is always required. The channel is optional - you can use the /join command to join a channel once you are connected to the IRC server.

Be sure that the packet driver is loaded first. When the program first starts it is going to initialize the TCP/IP stack and try to find the IP address of your IRC server. It is going to use the DNS system to do this. If you have not set up your TCP/IP environment variables correctly you will not be able to connect to your IRC server. (If you know the numerical IP address of your server you can use it directly to avoid the DNS lookup.)

If all goes well you will get a connection to the IRC server. If you specified a channel on the command line you will be signed onto that channel automatically by the program. At any time during the connection process you can hit the [ESC] key to quit.

The optional -port <n> can be used to connect to IRC servers running on non-standard ports.

Here are some examples:

```
ircjr irc.slashnet.org #vc
```

connects to my favorite channel on my favorite IRC network.

```
ircjr -port 2000 irc.privateserver.net
```

shows how to use the -port option to connect to a server with a non-standard IRC port.

Using IRCjr

Basic knowledge of IRC is assumed. But just in case ...

IRC stands for "Internet Relay Chat" and it is a system that allows clients to connect to a common server and chat with each other. Servers allow for the creation and management of different "channels" that clients can join. The channels are usually named after a topic of interest. Users may also chat directly with each other out of the view of others.

Servers can be connected to other servers, allowing for IRC networks that span the globe. Each network has its own policies regarding user behavior, and even within a channel there might be "operators" who enforce policies on that channel. Operators have the ability to remove people from channels, so be polite to them. ;-0

Screen layout

The IRCjr screen is split into two areas - the user input area and the dialog area. The dialog area shows the running conversation in a channel or a private chat session. The user input area is where you can compose your comments before sending them to the other users or where you can enter IRC commands.

A status line separates the two areas of the screen. The status line tells you which session you are viewing, which sessions are active, and also has indicators for the various toggle settings.

During normal usage the text that you enter will be sent to the other users as part of the conversation. The text that you enter will be interleaved in the conversation with text from other users in the order that the server receives it.

Sending IRC commands

It is possible to send IRC commands to the IRC server. IRC commands meant for the server start with a '/' character at the start of the line. IRC commands are used for many different reasons; here are some of the common ones:

/join	Join an IRC channel (e.g.: /join #vc)
/msg	Send a private msg to another user (eg: /msg Leela Hello!)
/query	Same as /msg, but open a new session to do it
/names	See who is in the channel (eg: /names #vc)
/part	Leave a channel (eg: /part #boring)
/nick	Change your nickname (eg: /nick newnickname)
/list	List the channels on a server
/quit	Disconnect from the server and end IRCjr

A hint on the /list command - on a large server or server network the list of open channels can be hundreds or thousands of lines long. You can't see all of the output on the screen or in your backscroll buffer, but you can use the LOG toggle to save the list to a file and read it later.

Some less common but still fun commands are:

/info	Get server information
/motd	Read the message of the day on the server
/whois	Get information about a particular nickname
/away	Mark yourself as away (or back)

IRCjr processes the commands that it understands and passes unfamiliar ones straight to the server. This lets you try to use any IRC command even if IRCjr does not recognize it.

Sending CTCP commands

Besides commands for the server there are "Client To Client Protocol" (CTCP) commands you can use to interact directly with other IRC clients. They are:

/me	Send an "Action" command
/ctcp version	Find out what client a user is running
/ctcp ping	Ping a client to see if they are active
/ctcp time	Find the local time for another user
/ctcp userinfo	Get the user name for a nickname

The /me CTCP command is pretty simple to use:

`/me ducks and runs!`	*Sends a message that you are ducking and running*

The other CTCP commands require you to provide a target (a nickname or even a channel name). For example:

`/ctcp version twinkie`	*(find what IRC client user "twinkie" has)*
`/ctcp time otherguy`	*(find out what time otherguy thinks it is)*

IRCjr Hotkeys

Besides the IRC commands there are special key combinations that make IRCjr perform tricks:

ALT-B	Toggle the new message beeper
ALT-H	Display the help window
ALT-C	Close current session/window
ALT-L	Toggle session logging on and off
ALT-S	Show TCP/IP statistics
ALT-T	Toggle timestamps on incoming messages
ALT-X	Exit IRCjr
PgUp	Go back in time through the backscroll buffer
PgDn	Go forward in time through the backscroll buffer
Alt 0-9	Switch to sessions 0 thru 9 (if applicable). (Session 0 is always the "Server" session. Channels and private chats use sessions 1-9.)

Session/Window handling

Each open channel and private chat has its own virtual session. The main window only displays one session at a time, so to flip between the virtual sessions you use Alt and a number key.

When you first start you are in a session reserved for server messages. This window is known as the "Server" session. You can get to it at any time by hitting Alt-0. It can not be closed.

When you join a new channel a new virtual session is created and you are automatically switched to that new session. The same thing happens if somebody sends you a private message or if you send somebody a private message.

The name of the current session is shown on the status indicator line. Also on the status indicator line there is a sequence of digits that show you how many sessions are open and what state they are in:

- Normal digit: a session with no recent activity

- Bright digit: the session currently being displayed
- Reverse digit: a session that has new activity in it

You can flip directly to any session by hitting Alt and the number that represents the session.

You can have up to 10 sessions including the server session. After that trying to create a new session will fail and whatever messages that are sent to that channel will be put on the Server session. If a new session is required and it can't be created because you have run out of memory those messages will go to the Server session too. You should probably close a session to make room for the new session; sending messages will be difficult from the server session.

Sessions can be closed by hitting Alt-C while in the session. If it is a session for a channel an IRC /PART command will be sent automatically.

Logging sessions

Previous versions of IRCjr logged all messages to a single file and you could not control logging on a per-session basis. Newer versions do logging on a per-session basis.

If you turn logging on in a session a new file with the name of the session and the extension "irc" will be created. If the file already exists it will be appended to so that you do not lose your previous log. If the session name is longer than eight characters the filename will only use the first eight characters - this is a limitation of DOS.

By default the current directory is where the log file will be created or opened. You can specify a directory to use for logs by adding a configuration parameter to the mTCP configuration file. (How to do that is described later on.)

When you turn logging off the file is closed. You may resume logging again at any time by using Alt-L. An indicator on the status line tells you if logging is on or off for the session that is currently open.

IRCjr does not try to detect when you connect to channels with the same name on different servers and turn logging on. The same log file name will be used in that case. IRCjr does write the server name and full channel name each time it starts logging so this should not be a large problem.

Using attributes (bold, reverse, etc) and mIRC color codes

IRC is generally a plain text service with not too many extras. You can highlight your messages in a few different ways by using a few special keys:

Bold:

You can select parts of your messages to appear in bold by turning bold on where you want it to start and turning bold off when you are done. To turn bold on press Ctrl-B - a special "block" character will appear to signal that you have inserted a special attribute command. After turning bold on, type what you want to appear in bold and then press Ctrl-B to turn bolding off. (The Ctrl-B behaves like an on/off toggle.)

Your IRCjr text display really is not set up to display a true boldface font. IRCjr simulates bolding by making the text brighter. (True bolding would show a thicker font.)

Reverse video:

You can select parts of your messages to appear in reverse video by turning it on and off in the same way that bold is used. To turn reverse video on press Ctrl-R and to turn it off after you are done using it press Ctrl-R again. (Ctrl-R works like an on/off toggle.)

Italics:

You can select parts of your messages to appear in italics by turning italics on and off in the same way that bold is used. To turn italics on press Ctrl-I and to turn it off after you are done press Ctrl-I again. (Ctrl-I works like an on/off toggle.)

IRCjr can send the italics code, but it can not display italics. When somebody sends italics to IRCjr it inserts the string "`<italics on>`" where the italics attribute starts and "`<italics off>`" when it stops.

Underline:

You can select parts of your messages to appear underlined by turning underlining on and off around the parts of the text that you want underlined. To turn underlining on press Ctrl-U and to turn it off when you are done with it press Ctrl-U again. (Ctrl-U works like an on/off toggle.)

IRCjr can send the underline code but unless you are running on an IBM Monochrome Display or a Hercules adapter it can not show underlined text. (This is a hardware limitation - CGA, EGA and VGA cards can not show underlined characters in text mode.) When something sends the underline attribute to IRCjr it inserts the string "<underline on>" where the attribute starts and "<underline off>" when the attribute stops.

mIRC color codes

If you really want to get fancy you can insert color codes into your text.

There are 16 different colors you can use and you can specify both the foreground and background color. To specify a color code start by pressing Ctrl-K. When you do a pop-up color chart with the color code numbers will show up. To specify the foreground color enter the number of the color on the chart. If you want to specify a background color then add a comma and the number of the background color. When you are done, just start typing your message and the pop-up color chart will disappear. The background color is optional and does not have to be used.

To turn color off just enter a Ctrl-K by itself.

For example, the following sequence of keys will give you the word "Hello" written in red on a yellow background:

```
<Ctrl-K>4,8Hello<Ctrl-K>
```

While this sequence will just give you the word "Red" with the default background color:

```
<Ctrl-K>4Red<Ctrl-K>
```

Some people fine color codes annoying, so don't go overboard.

Reset (turn all off)

To disable any attributes that might have been set use Ctrl-O. The special "box" character will be shown where the Ctrl-O was inserted. After that point the text you send should be normal with no bold, reverse, italics, underlining or color.

Selecting your display mode and size

IRCjr uses whatever the current display mode is on your active screen. So if you want to use a particular number of rows and columns, set up the display mode first using your favorite utility and then start IRCjr.

CGA and MDA users do not have any choices - only 80x25 is available.

EGA users can use 80x25 or 80x43 text mode.

VGA and SVGA users can use a variety of text modes depending on what their card supports. I've tested it with 80x25, 80x43, 80x50, 90x60, and 132x43.

If you have a monochrome display and a color display on the same machine choose the display you want to use using the DOS MODE command first. IRCjr will detect the current active display and use it.

If you have an oddball display adapter like the MDSI Genius Display, IRCjr can support that too! IRCjr doesn't care what card you have, as long as it has a video buffer that is addressed linearly the same way that MDA and CGA do in their text modes. The MDSI in particular reports as an MDA display, but it is capable of 80x66. To override the number of rows IRCjr detects use the MTCP_SCREEN_ROWS environment variable. For example:

```
SET MTCP_SCREEN_ROWS=66
```

will instruct IRCjr to use 66 rows of text, even if it detects something different. This works for the MDSI Genius because it reports as MDA (25 rows only), but has enough video memory laid out linearly to support 66 rows of text. Other oddball graphics cards might be able to use this trick too.

Removing CGA Snow

CGA "Snow" is static or noise that is caused by writing directly to the video buffer memory on CGA cards that do not have dual ported memory. Only the original IBM CGA adapter and close clones are affected by this, and it can be very distracting.

If your CGA card is displaying CGA Snow you can use an environment variable to tell IRCjr to change the way it writes to the screen, thus eliminating the CGA Snow. It will cause screen updates to be slower but it removes the noise.

To enable this feature set the MTCP_NO_SNOW environment variable to any value:

```
SET MTCP_NO_SNOW=1
```

Unicode support

IRCjr supports Unicode using the UTF-8 encoding which is the preferred encoding on modern IRC servers. In general Unicode is allowed in messages and channel topics. Unicode support in other places suck as nicks depends on the server being used.

If enabled Unicode characters will be mapped to characters in your local code page. Characters that can not be displayed locally will be represented by the “tofu” character (a white block that looks like “■”).

Characters generated on your keyboard that are not 7 bit ASCII will be mapped to Unicode characters before being sent to the server. These include accented Latin alphabet characters, math symbols, characters generated by holding the Alt key and entering a three digit character code on the numeric keypad, etc. Direct entry of Unicode code points via a “compose” sequence is not supported yet.

See Unicode support for the details on how to enable Unicode support.

Optional Configuration Parameters

You only need the three configuration parameters explained earlier to start using IRCjr. But there are other configuration parameters you can use to alter IRCjr's behavior, customizing it to your needs. Here is the list of optional configuration parameters you can use:

IRCJR_PASS	Specify a connection password
IRCJR_CONNECT_TIMEOUT	Set the timeout period for the socket connect
IRCJR_REGISTER_TIMEOUT	Set the timeout for the registration process
IRCJR_BACKSCROLL	Set # of backscroll lines for each channel
IRCJR_BACKSCROLL_CHAT	Set # of backscroll lines for each chat session
IRCJR_BACKSCROLL_SERVER	Set # of backscroll lines for the server session
IRCJR_COLOR_SCHEME	Override the default color scheme
IRCJR_TIMESTAMPS	Turn timestamps on at program start
IRCJR_LOGGING_DEFAULT	Turn on logging at program start
IRCJR_LOG_DIR	Specify a directory for log files
IRCJR_NICK_UPDATES	Specify where nickname updates go to
IRCJR_QUIT_UPDATES	Specify where quit notifications go to

And now for the details on how to use these ...

`IRCJR_PASS`

> IRCJR_PASS allows you to send a "connection password" during the connection and registration process. Connection passwords can be used to authenticate registered nicknames on many servers. Keep in mind that connection passwords are not very secure; they are transmitted in the clear.

`IRCJR_CONNECT_TIMEOUT`

IRCJR_CONNECT_TIMEOUT controls how long IRCjr will wait for a TCP/IP socket connection to an IRC server. The default is 30 seconds, which should be long enough for almost anything. But if your connection is poor and you need more time you can set it to something longer.

`IRCJR_REGISTER_TIMEOUT`

IRCJR_REGISTER_TIMEOUT controls how long IRCjr will wait for a server to recognize the client after the TCP/IP socket is created. The default is 30 seconds which is usually long enough, but on a busy server you might need to allow for more time.

`IRCJR_BACKSCROLL, IRCJR_BACKSCROLL_CHAT, IRCJR_BACKSCROLL_SERVER`

The IRCJR_BACKSCROLL_* settings are used to tell IRCjr how much memory to reserve for backscroll buffers. The backscroll buffer allows you to see lines that have scrolled off of the current screen. The backscroll buffers are a great feature but they require more memory.

The default backscroll buffer settings are:

Server session:	50 lines
Channels:	150 lines
Private chats:	75 lines

These are reasonable for most users. With these settings:

- Just connecting to a server requires ~164KB
- Connecting and joining a channel requires ~191KB
- Connecting, joining a channel and having one private chat requires 206KB

This is a bit more than in the original IRCjr but still manageable. On an 512KB system you can be on nine different channels with a generous amount of backscroll in every channel.

You can set the backscroll buffers to be larger, but you are limited by how much RAM you have available. When IRCjr runs out of memory you will not be able to open new sessions to chat in new channels or receive messages from private users. It will put private messages in the server window, which is ugly but will work. You will probably be happier if you close some idle sessions to get some memory back, or reduce the number of backscroll lines by using these settings. (If you are really tight on memory you can disable a class of backscroll buffers by setting them to 0.)

`IRCJR_COLOR_SCHEME`

IRCJR_COLOR_SCHEME only has one setting at the moment - CGA_MONO. This is useful for machines that have CGA cards but use an LCD or monochrome monitor where the different shades of gray, green or amber might be difficult to distinguish. Setting this will give you a high contrast color scheme.

`IRC_TIMESTAMPS`

IRC_TIMESTAMPS is used to turn timestamps on when the program first starts. The default is to start without timestamps turned on. While the program is running you can always turn the timestamps on by using ALT-T (a toggle setting), but if you usually do that then you can use this setting to make that the

default.

IRCJR_LOGGING_DEFAULT

IRCJR_LOGGING_DEFAULT allows you to specify the default for the logging state. Normally the program does not log the contents of sessions unless you turn logging on using Alt-L. If you see this "on" the logging will be started by default as soon as the program starts.

IRCJR_LOG_DIR

IRCJR_LOG_DIR allows you to specify a directory path where log files will be written to. By default log files are written to whatever the current directory is when IRCjr starts. If you want to specify a directory enter it here in standard DOS format with a drive letter and backslashes as the path delimiter. (A full or relative path can be specified. The path must end in a backslash.)

Log files are always appended to and never deleted or overwritten.

IRCJR_NICK_UPDATES

IRCJR_NICK_UPDATES allows you to tell IRCjr where to send nickname updates generated by other users. Four options are available:

none	No nickname updates will be shown anywhere
server	Nickname updates will be shown on the server session
current	Nickname updates will be shown on the current open session
all	Nickname updates will be shown on all sessions

Previous versions of IRCjr used the equivalent of "all" which could get annoying on busy channels. The current default is "none".

IRCJR_QUIT_UPDATES

IRCJR_QUIT_UPDATES allows you to tell IRCjr where to send the notifications that are generated when other users disconnect from the server you are using. The same four options as used by IRC_NICK_UDPATES are available.

Previous versions of IRCjr used the equivalent of "all" which could get annoying on busy channels. The current default is "none".

Here is an example of an mTCP configuration file with these parameters set:

```
DHCPVER DHCP Client version Apr 26 2009
TIMESTAMP Sun Apr 26 17:59:54 2009

# Parameters for my machine; your machine will be different
#
packetint 0x60
mt 576
hostname DOSBox

# IRCjr parms
ircjr_nick Zoidberg
ircjr_user Jzoidberg
ircjr_name John F Zoidberg

# All parameters after this are optional

# IRC connection password
ircjr_pass secretgoeshere

# Use these for really slow servers
ircjr_connect_timeout 45
ircjr_register_timeout 60

# Setup for large backscroll buffers

ircjr_backscroll 375
ircjr_backscroll_chat 150
ircjr_backscroll_server 100

# Use this if you need a high contrast screen
ircjr_color_scheme cga_mono

# Turn timestamps on at the start
ircjr_timestamps on

# Turn on logging by default and log to a specific directory
ircjr_logging_default on
ircjr_log_dir e:\data\irclogs\

# Send nickname updates and quit notifications to the server session
ircjr_nick_updates server
ircjr_quit_updates server

# DHCP generated settings will appear here ...
```

Obviously, substitute in values that make sense for your machine. "Zoidberg" is a great name but it is probably already taken.

Netcat

Introduction

Netcat (nc) is a utility that can send and receive data using a TCP/IP socket. It has the ability to open a connection to another machine or to listen for incoming connections. Input can be entered interactively through the keyboard or redirected in from a file. Output can appear on the screen or redirected to a file. Nc can be used to send large files or short messages, and it can be called from within batch files too.

Any status messages that netcat creates that are not part of the data received from the socket are sent to stderr. This ensures that if you redirect the socket output to a file that only the data that came across the socket will be in that file.

Using Netcat

Netcat uses the following syntax:

```
nc -target <ipaddr> <port> [options]
nc -listen <port> [options]
```

The first form of the command is used to create a socket connection to another machine, as a client program normally would. The second form of the command is used to make nc wait for an incoming connection from another machine, as a server program normally would.

Interested in getting started quickly? Skip down to the section entitled 'Examples'. The discussion of options is fairly long.

Netcat Options

Options are:

`-help`	Show basic help text
`-verbose`	Print extra status messages
`-bin`	Treat STDIN and STDOUT as binary streams
`-telnet_nl`	Send and receive newline (NL) as Telnet newline (CR/LF)
`-echo`	Turn on local echoing
`-w <n>`	Wait <n> seconds for traffic after STDIN closes
`-nocorc`	Do not Close on remote close
`-srcport <port>`	Use <port> as the source port when connecting

Usage

First, some background definitions:

Newline character (NL):

This varies depending on the computer and operating system being used.

- For MS-DOS: A newline is composed of a CR/LF pair
- For Unix: A newline character is an LF
- For internet programs using the Telnet NVT standard: A newline is composed of a CR/LF pair.

Socket 'closes'

In this context a socket is a connection to another machine using TCP/IP. The socket is a full duplex connection - each side can send at the same time, and each side is treated as a peer.

When a socket connection is "closed" it means that one one side of the connection has signaled that it is done sending data. The socket still exists, and the side that did the close is still free to receive data from the socket - it just can't send any more data. The other side of the connection is free to send more data to the half that did the close, but it should not expect to be able to read any more data from the closed half.

It is only when both sides have closed their half of the socket that the socket connection is terminated and cleaned up.

Most programs communicating via a socket use a protocol over the socket that warns both sides when the connection is going to be closed. That allows them to close their half of the socket and end the application while being assured that the other side is going to close and exit too. Netcat's default behavior is to close and exit when the other side closes, but this can be overridden.

Netcat default operation

With no extra options netcat will behave as follows after a connection is made:

- Input will come from the keyboard and be sent over the socket
- Data received on the socket will be written to the screen
- Pressing Ctrl-Z at the keyboard will specify the end of input and close your half of the socket connection. It will also terminate the socket connection and end netcat even if the other side of the connection had not closed their half of the socket. (It is expected that the other side will close when they see you close.)
- Pressing the Enter key will generate a Unix newline char (LF)
- Unix newline characters received on the socket will be converted to MS-DOS newline characters (CR/LF)
- If the remote side does a close on their half of the socket netcat will close your half of the socket and exit. (This can be overridden.)
- Each keystroke is sent in a separate packet. There is no line editing capability because that would block the receipt of incoming data from the socket.

There are no surprises here, unless you come from an entirely DOS universe where newline characters are always a CR/LF pair instead of the Unix style newline which is an ASCII LF.

Redirecting STDIN and STDOUT from/to files

You can specify that STDIN should be read from a file or that STDOUT should be written to a file. Just use the

normal DOS redirection characters to do this. Netcat behavior when STDIN/STDOUT are redirected is similar to its default behavior, except for these changes:

- If netcat detects that STDIN is redirected from a file it will try to read large chunks of the file in a single read operation and then send those chunks in large packets. This is more efficient than the normal behavior when netcat is reading from the keyboard.
- If netcat detects that STDIN is redirected from a file it will close its half of the socket when EOF is hit on STDIN. This will cause netcat to terminate the socket connection and exit whether the other side had closed their half of the socket or not. (There is an option to change this behavior - see "-w".)

Netcat options (details)

`-verbose`

Verbose mode gives you extra status message that tell you when the local side or the remote side closed their connection. These messages are written to stderr so they will not affect data that you redirect to a file on STDOUT.

`-bin`

MS-DOS can treat files as either BINARY or TEXT.

In BINARY mode:

- All characters are read and written exactly with no changes.
- The End of File (EOF) is determined by the file size, not a Ctrl-Z char. If you are reading from the keyboard (and not a redirected file) then you need to use Alt-X to close the connection, as Ctrl-Z will not be interpreted as EOF.

In TEXT mode:

- When writing on STDOUT Unix newline characters (LF) will be converted to MS-DOS newline characters (CR/LF).
- When reading from STDIN MS-DOS newline characters (CR/LF) will be converted to the Unix newline character (LF).
- A Ctrl-Z will be interpreted as End of File

If you are transferring binary data, use -bin to ensure that netcat and MS-DOS do not alter your data in any way. If you are working with normal ASCII text then use the default, although binary mode might work too.

Binary mode is going to perform much better than the normal text mode because in binary mode no attempt is made to process the incoming and outgoing characters, saving on processing time.

Note: You can use -bin even if you are not redirecting STDIN and STDOUT to a file. The behavior change is the same whether you are redirecting or not.

`-telnet_nl`

The standard for sending newline characters across the internet using the Telnet "network virtual terminal"

(NVT) is to send a CR/LF pair. This is because some machines use CR, some use LF, and some use other oddball characters.

If you are connecting to a service that uses the Telnet NVT standard then turn this option on. This will convert your outgoing newline characters to the Telnet standard of CR/LF. Services like HTTP, FTP, SMTP, etc. all use the Telnet NVT standard so they will probably work better with this option. Incoming CR/LF pairs will be translated to normal newlines.

This option can not be mixed with -bin. If you need binary mode, then turning on -telnet_nl makes not sense.

(Slightly fun trick .. MS-DOS uses the same newline characters as the Telnet NVT standard. So you can probably use -bin in place of -telnet_nl. But -bin changes other things, so it is probably best not to mess with it.)

This option does not turn netcat into a Telnet client - it only changes the way that newline characters are sent and received. It will not do Telnet option negotiation, which is an important part of the Telnet protocol that a real Telnet server will need. But this option will help you connect to FTP, SMTP and other servers that use the Telnet NVT standard but not do Telnet option negotiation.

`-echo`

By default no echoing of characters is provided. If you turn this on any characters from STDIN will be echoed to STDOUT.

If the other side is echoing your characters and you turn this on you will see things twice.

Keep in mind that because we are echoing to STDOUT with this option, your local input will show up in STDOUT even when you redirect STDOUT to a file.

`-w <n>`

The default is to close the socket connection and exit from netcat when EOF is detected on STDIN. This can be a problem if there was data in flight from the other side at the time the socket was closed. (You will miss the data that was in flight. This is more likely to happen when STDIN is redirected from a file, as netcat will send it very quickly and hit EOF possibly before the other side has had a chance to do everything.)

To get around this the -w option tells netcat how long to wait for data to come from the other side after your side of the socket hits EOF on STDIN. This can give netcat a chance to receive the last of the data from the other side before closing the connection. <n> is the number of seconds to wait.

`-nocorc`

Corc stands for "Close On Remote Close", so -nocorc means don't do that.

By default if netcat detects that the other side closed its half of the socket connection it will close the socket and exit. This will happen even if EOF was not yet detected on STDIN.

With this option netcat will ignore the close from the other side until EOF is detected on STDIN.

This option probably is not very useful for most people, but if you are experimenting with 'half closed' sockets it can be very useful.

`-srcport <n>`

By default if you are connecting to another machine netcat will use a random port number as the source port on your machine. If for some reason you need to specify a source port for your machine you can use this option.

Keyboard Handling

By default netcat reads interactively from the keyboard. This is normally straightforward, but there are some things you should know.

In binary mode:

- Ctrl-Z is sent as ASCII 26 and does not mark EOF. (You must signal 'EOF' using Alt-X.)
- Enter is sent as ASCII 13 (CR).
- Ctrl-Enter is sent as ASCII 10 (LF).
- All other keys are sent exactly as they are read.

In text mode:

- Ctrl-Z means EOF for STDIN, and will close your half of the socket.
- Enter is sent as a Unix newline (LF)
- Ctrl-Enter is also sent as ASCII 10 (LF). (No change from above)
- To send a raw CR use Ctrl-M.

In text mode with -telnet_nl turned on the Enter key changes slightly:

- Enter is sent as a telnet NVT newline (CR/LF)
- To send a raw CR use Ctrl-M. To send a raw LF use Ctrl-J Or Ctrl-Enter.

Special keys:

Whether you are reading from the keyboard or not the following keys will always be honored:

Alt-E	Toggle local echoing on and off
Alt-H	Help
Alt-X	Close your half of the socket connection
Alt-S	Show Status
Ctrl-Break	Unconditional exit

Alt-S is really handy, as it will show you if the other side has closed their half of the connection or not. If you are doing a large file transfer it can also show you how many bytes have been sent and received.

Environment variables

Two environment variables are provided to help tailor netcat to your memory and performance requirements.

`READBUF`

Default 8192, Range 512 to 32768

Use: Sets the size of the buffer used to read from STDIN when STDIN is redirected from a file on your local disk. Bigger numbers generally result in higher performance. Use a smaller number if you are tight on memory and need to cut back on what netcat whats to use.

If you are low on memory 4KB is a reasonable setting that will not hurt performance too much.

This environment variable has no effect when you are reading from the keyboard.

`WRITEBUF`

Default 8192, Range 512 to 32768

Use: Sets the size of the buffer used to write to STDOUT when STDOUT is redirected to a file on your local disk. Bigger numbers generally result in higher performance. Use a smaller number if you are tight on memory and need to cut back on what netcat whats to use.

If you are low on memory 4KB is a reasonable setting that will not hurt performance too much.

For both variables DOS likes it best if you use a power of 2, or at least a multiple of 512 bytes. Other numbers will cause extra work for DOS and limit performance.

Examples (and Quick start!)

```
nc -target 192.168.2.10 7
```

Opens a socket connection to 192.168.2.10 on port 7. Your keystrokes will be sent and data coming back will be displayed on the screen. (On a Unix machine port 7 is the echo service, so you should see everything that you send.)

```
nc -target 192.168.2.5 1234 < input.txt
```

Opens a TCP/IP socket to 192.168.2.5 port 1234 and sends the file 'input.txt' to the other machine instead of sending your keyboard input. The program automatically ends after input.txt has been transferred.

```
nc -target 192.168.2.5 1234 -w 5 < input.txt
```

Same as above, but wait 5 more seconds after EOF is detected on STDIN so that any response from the other side can be read.

```
nc -listen 2000 > newfile.txt
```

Listens on port 2000 for a connection from a remote machine. When the connection is established any incoming data is sent to newfile.txt. When the remote machine closes the connection netcat will exit.

```
nc -listen 3000 -nocorc > newfile.txt
```

Similar to above, but netcat does not automatically exit until it is told to. (This is not terribly useful when redirecting files, but is useful when using netcat interactively.)

```
nc -target www.cnn.com 80 -telnet_nl -verbose -echo
```

Connect to the HTTP server at www.cnn.com on port 80. Automatically convert newlines to telnet NVT newlines, enable verbose status messages, and start with local echoing turned on. (If you send "GET /" without the quotes and press Enter, you will get a web page back!)

Network printing

If you have a network attached printer that supports “raw” printing on TCP/IP port 9100 then it is possible to send files from your DOS machine to the printer for printing using the following steps:

- Send your output to a DOS file instead of a printer. The output file needs to be in a format the printer understands, such as PCL for printers that use Printer Control Language or PostScript for printers that use that. (This implies that your DOS program has a suitable printer driver for your printer.)

- Send the output file to the printer using the netcat command:

  ```
  nc -target <printer address> 9100 -bin < <filename>
  ```

Substitute <printer address> with the IP address of your printer and <filename> with the filename your program generated when you said "print to file." For example, on my printer:

```
nc -target 192.168.2.20 9100 -bin < testfile.txt
```

will print testfile.txt over the network to the printer. (testfile.txt gets redirected on STDIN to netcat.)

If you are unsure if your printer listens on port 9100 then just try the steps above anyway; at worst case nothing will happen.

Limitations

This version of Netcat does not support UDP. (Most Unix versions do.)

NetDrive

Introduction

mTCP NetDrive is a DOS device driver that allows you to access a remote disk image hosted by another machine (the server) as though it was a local device with an assigned drive letter. The remote disk image can be a floppy disk image or a hard drive image. When your DOS machine boots the device driver reserves one or more drive letters for remote disk images. When you are ready to use a remote disk image you run a command-line utility to connect a reserved drive letter to the remote disk image. From that point forward DOS treats it like any other device with a drive letter. When you are done with the disk image you use the command-line utility to disconnect it.

Other programs exist for accessing remote data. Here is what makes NetDrive interesting:

- It works with all versions of DOS starting with version 2.0 (the first version to use device drivers).
- The device driver uses just 9KB of RAM.
- It uses UDP so that it can be used easily on a private network or across the public Internet.
- It works with floppy images and hard drive images up to 2GB in size.
- Multiple remote drives from different servers can be connected at the same time.
- You can connect and disconnect quickly, right from the DOS command line.
- It correctly supports the DOS Verify flag for improved data protection when writing data.
- The remote disk images are raw sector dumps that can be mounted and used directly by Linux. Tools for Windows for manipulating disk images are also available.
- Compatible with class DOS commands such as ASSIGN, JOIN, LABEL, and SUBST that are not usable on other network drives.
- Disk images can be journaled on the server side, allowing for an advanced “go back/undo” feature that can be controlled from DOS. (No other remote storage server for DOS has this ability.)
- Performance can be quite good, even on servers that are distant.

NetDrive requires a packet driver to be loaded so that it can send and receive packets using your Ethernet or emulated Ethernet device. A packet driver usually requires 5 to 30 KB of RAM in addition to the RAM for the device driver.

And of course, there are limitations too:

- This is a block device driver, not a remote filesystem; everything happens at the block level. The server sees the disk images as large files. (The server will need to use tools to interact with the files within the disk images.) Think of a NetDrive drive letter like a portable hard drive – one system has read-write access to it at a time.
- The performance of a remote disk image is somewhat dependent on the distance to the server, especially for small reads and writes.
- It is possible for multiple clients to use the same image at the same time, but only if the image is marked read-only. (Two active writers to the same image should never be allowed because they will probably corrupt the image.)
- There is no encryption or authentication between the client and the server. If you set up a server you can use firewall rules to control the traffic allowed to the server. (Better authentication will be added if there is a demand for it. Encryption can be provided by setting up a VPN.)

The mTCP NetDrive server is a reasonably small program (under 20 MB working size) that runs on Windows 10, Windows 11, Linux (x86 64 bit), Linux (Arm 64 bit), BSD variants and other operating systems where a Go runtime is supported. No special permissions are required to run it.

DOS installation and usage

NetDrive is part of mTCP so you must have mTCP set up and working before using NetDrive. If you can ping to other machines on the network and your MTU is set to at least 1200 then you are ready to set up NetDrive.

The DOS side of NetDrive consists of two files:

- NETDRIVE.SYS DOS device driver
- NETDRIVE.EXE DOS utility program

Add NETDRIVE.SYS to your CONFIG.SYS file using the following line:

```
DEVICE=NETDRIVE.SYS
```

The default is still to reserve one drive letter; use the -d option to reserve more drive letters. For example:

```
DEVICE=NETDRIVE.SYS -d:4
```

will reserve four drive letters for NetDrive. Each drive letter requires 96 bytes of device driver memory and 48 bytes of DOS memory.

Note: Do not use tricks like DEVLOAD. There is nothing in the device driver that cares about where it is loaded so using DEVICEHIGH is fine but DEVLOAD seems to break things in strange ways.

When NETDRIVE.SYS installs during the boot process it reserves its drive letters. When used with DOS version 3.0 or newer the device driver will tell you the first drive letter that was assigned to it.

- DOS assigns drive letters based on the ordering of device drivers in CONFIG.SYS so you may need to adjust where you put this line to get a specific DOS drive letter.
- You may also need to use the LASTDRIVE directive in CONFIG.SYS to tell DOS to reserve sufficient drive letters.
- DOS 2.x does not provide drive letter information to the device driver, but you should be able to find the assigned drive letter fairly easily.

When NetDrive is connected to a remote image it will be using the packet driver provided by the network card, making that packet driver unavailable for other programs. If you want to be able to use other mTCP programs while NetDrive is active you need to add a secondary software interrupt to the PACKETINT configuration parameter in the mTCP configuration file. For example:

```
PACKETINT 0x63, 0x65
```

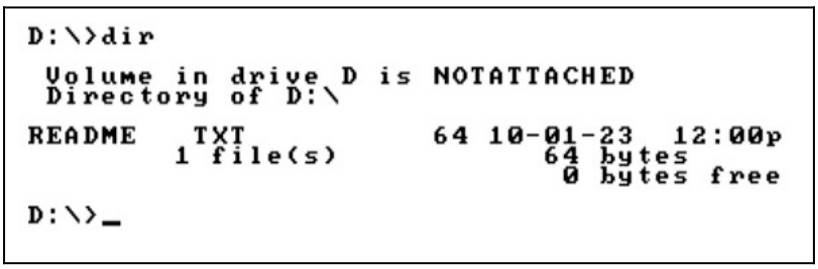
says that your network card packet driver is loaded at software interrupt 0x63 and that NetDrive should provide a secondary passthrough packet driver at software interrupt 0x65.

See Passthrough: Running other networking code with NetDrive later in this chapter.

Primary commands

CONNECT

Before being connected to a remote disk image the drive letter shows a minimal, read-only RAM disk that signals that a remote disk image has not been connected yet. The label on the RAM disk is “NOTATTACHED” and there is a README.TXT file:



To connect to a remote disk image:

- Load your packet driver and set up mTCP as you would for any other mTCP program. (If you can ping to other systems you are ready to run NetDrive.)
- Run the NETDRIVE CONNECT command and tell it what to connect to.

The format for the NETDRIVE CONNECT command is:

```
NETDRIVE CONNECT <machine:udp_port> <image_name> <x:> [options]
```

The options are:

`-ro`	Force “read-only” mode on the connected remote image
`-mra <n>`	Performance option: set maximum read-ahead to n KB for this connection
`-mws <n>`	Performance option: set maximum write size to n KB for this connection
`-sd <n>`	Performance option: delay <n> milliseconds when sending write packets

Example:

```
NETDRIVE CONNECT CALCULON:2002 BIGDISK.DSK D:
```

“CONNECT” can be shortened to “C” to save on typing:

```
NETDRIVE C CALCULON:2002 BIGDISK.DSK D:
```

When NETDRIVE.EXE runs it finds the device driver in memory, contacts the server to ensure the remote disk image is available, connects the device driver to the packet driver and hands over network processing to the device driver. On first touch of the drive letter DOS detects that the device has changed and loads a new BIOS Parameter Block (BPB) from the device. From that point forward the remote disk image is available for reads and writes as though it were a local device.

```
D:\>dir

 Volume in drive D is NOTATTACHED
 Directory of D:\

README   TXT            64 10-01-23  12:00p
       1 file(s)             64 bytes
                              0 bytes free

D:\>netdrive c brutman.com:2002 bigdisk.dsk d:
mTCP NetDrive by M Brutman (mbbrutman@gmail.com) (C)opyright 2008-2026
Version: Sep 30 2026

Resolving brutman.com, press [ESC] to abort.

Session (49780) started for image "bigdisk.dsk"
Drive size: 209715200, Media descriptor byte from FAT: F8
Session write overlay is turned on; writes will not persist.

NetDrive uses the packet driver at 0x60, pass-through driver available at 0x61

Max read-ahead: 16KB, Max write size: 32KB, Send Delay: 0.5ms
One of those values was auto selected or is a default.  Use the -mra, -mws
and -sd options to set better values for this connected image.

D:\>dir

 Volume in drive D is BIGDISK
 Directory of D:\

APPS         <DIR>          11-27-23  12:08a
DRIVERS      <DIR>          11-27-23   8:06p
GAMES        <DIR>          11-25-23   8:09p
IOTEST       <DIR>          11-25-23   8:58p
MODEM        <DIR>          11-27-23  12:20a
MTCP2023     <DIR>          11-26-23  11:57p
PACKET       <DIR>          11-25-23   8:46p
PCJR         <DIR>          11-27-23  12:05a
READMAC      <DIR>          12-03-23   8:05p
RLE          <DIR>          12-03-23   8:04p
UTILS        <DIR>          11-25-23   7:05p
00README TXT          1,029 04-21-26   7:05p
       12 file(s)          1,029 bytes
                     170,475,520 bytes free

D:\>
```

STATUS

The STATUS command (“s” for short) can be used to check the state of NetDrive drive letters.

```
NETDRIVE STATUS D:
```

A drive letter is required because NetDrive has to do an IOCTL call against the drive letter to get to the data structure that it needs. Regardless of the drive letter provided, NetDrive will display its status and the status of all drives it is responsible for.

```
D:\>netdrive s d:
mTCP NetDrive by M Brutman (mbbrutman@gmail.com) (C)opyright 2008-2026
Version: Sep 26 2026

NetDrive units: 4, First drive letter: D:

Our IP addr: 192.168.2.240, UDP Port:7758, MAC: 08:00:27:7C:50:D1
Packet driver int: 0x60, Netdrive shim int: 0x61
Received packets: All: 636, To our MAC addr: 203, To our IP/UDP: 202
Arp replies sent: 2, Packet send errors: 0

#: Server address, Session number, Remote Image, Blks In, Blks Out, Retries

D: 35.206.78.10:2002, 50454, bigdisk.dsk (RW), 202, 0, 0
E: (Not connected)
F: (Not connected)
G: (Not connected)

D:\>_
```

DISCONNECT

Use the DISCONNECT ("d" for short) command to cleanly end your session:

```
NETDRIVE DISCONNECT D:
```

This will tell the device driver to stop processing packets from the packet driver and send a disconnect message to the server. The minimal RAM disk will then appear again at the drive letter.

The DISCONNECT command is guaranteed to work correctly on the DOS side even if the server can not be reached. This is a fail-safe in case there are networking problems.

Checkpoint commands

A standard NetDrive image mimics a block device. NetDrive has a feature that adds a write journal to a standard image, making it possible to create checkpoints and then revert back to them, effectively adding an "undo" feature to changes that you make to the image. (Being able to undo a mistake or an experiment, such as accidentally deleting files right from the DOS command line is an incredibly powerful feature.)

The feature must be enabled on the image on the server before it is available to DOS. See Advanced disk images: Session scoped images and Journaled images for instructions on how to do that.

In short, the feature is a simple "checkpoint, possibly make changes, and then be able to jump back to the checkpoint if you don't like what you did" type of undo feature. Changes refer to all writes that are made to the remote disk image including changing file contents, adding new files, deleting files, etc.

The first step is to write a checkpoint marker into the disk image. Checkpoint markers are time stamped and you provide a text string to describe the purpose of each checkpoint marker. After writing a checkpoint marker you can then continue working, possibly writing and altering the remote disk image. If you want to undo the changes since the time a checkpoint marker was written you can use the goto checkpoint marker command to undo the changes that were made after the checkpoint marker was written.

A new checkpoint marker is written automatically every time a client connects to a journaled disk image. If you forget to write a checkpoint marker before doing something risky, at worst you'll just have to jump back to the

state at the start of the connected session, limiting any damage you might have done.

LIST_CP

LIST_CP will fetch and display a list of checkpoint markers on the remote disk image connected to a drive letter. Example:

```
NETDRIVE LIST_CP E:
```

Checkpoint markers will be shown in reverse age order, as the newest checkpoint markers are probably the most interesting while connected. If you have many checkpoint markers you may need to pipe the output through the DOS MORE command or redirect the output to a file to view it.

Note: When NETDRIVE.EXE runs it disconnects the packet driver from the device driver temporarily so that it can use the network connection.

- Do not redirect the output of NETDRIVE.EXE to a connected NetDrive; it will fail because connected images are not available while the NETDRIVE.EXE is running. Redirect the output to a locally connected drive.
- If you pipe the output through the MORE command be sure to set the DOS TEMP environment variable first and to have the temporary output go to a local drive. The pipe to MORE creates a temporary file and that temporary file can't be on a NetDrive image for the same reason explained above. If you are on an older version of DOS that doesn't use the TEMP environment variable then change the logged drive to a local hard drive first before using MORE with the output of NETDRIVE.EXE

Nothing terrible will happen if you try to write to a connected image while NETDRIVE.EXE is active, but you will have to wait for the DOS error messages and chose the "Fail" option a few times. (Never chose "Abort" if you get "Abort, Retry, Ignore, Fail" because Abort kills the program immediately without giving it a chance to clean up interrupts that it hooked; you system will probably crash shortly after using the Abort option.)

MARK_CP

MARK_CP writes a new checkpoint mark in the remote disk image connected to drive x: The tag is a text string that can be used to describe why you are writing the checkpoint marker. (Use quotes around the tag if it has spaces in it.) Example:

```
NETDRIVE MARK_CP E: "Before deleting a bunch of files"
```

GOTO_CP

GOTO_CP lets you undo any changes made after a specified checkpoint marker was created. The target checkpoint marker can be specified with a number (obtained using LIST_CP) or by matching the text supplied when the marker was created.

Format:

```
NETDRIVE GOTO_CP <x:> -tag <tag> [-del_marker]
NETDRIVE GOTO_CP <x:> -num <checkpoint_number> [-del_marker]
```

The first form of the command jumps back to the newest checkpoint marker that used the specified tag. The second form of the command allows you to specify the checkpoint marker number, as shown by the LIST_CP command. (Always run the LIST_CP command first to get the current set of checkpoint markers and numbers.)

When you use GOTO_CP all changes made after the specified checkpoint marker are deleted, effectively reverting the remote disk image back to what it was at the time the checkpoint marker was written. The checkpoint marker itself is left in place so that you don't have to explicitly write a new checkpoint marker again. The -del_marker option allows you to delete the checkpoint marker too.

These commands will not work on an image that is not journaled. See Enabling journaling on a disk image for instructions on how to enable journaling on a disk image.

To use this feature effectively you need to get into the habit of creating a checkpoint mark before doing something potentially risky. Checkpoint marks take very little storage and represent safe points in time that you can restore the drive back to using the GOTO_CP command.

Suggested workflow:

1. Connect to a remote disk image.
2. Before doing something that you might want to undo, use MARK_CP to create a checkpoint mark.
3. Continue working.
4. If you decide you need to undo the changes, use GOTO_CP to go back to the checkpoint mark you created, or an even earlier one.
5. If you decide to keep the changes there is no cost to having the checkpoint mark in place. Eventually you will need to commit the writes to the image because the journal file will get too large - a command line function on the server program handles this.

Passthrough: Running other networking code with NetDrive

When NetDrive is connected to a remote image it will be using the packet driver provided by the network card, making that packet driver unavailable for other programs. This is a problem if you want to run other networking programs while NetDrive is in use. Fortunately NetDrive has a "passthrough" feature that allows other networking programs to run even while NetDrive is using the network card packet driver. To enable this feature add a second software interrupt number to the PACKETINT directive in the mTCP configuration file. The second software interrupt number needs to be unused and available. Example:

```
PACKETINT 0x63, 0x65
```

That says that your network card packet driver is loaded at software interrupt 0x63 and that NetDrive should provide a secondary passthrough packet driver at software interrupt 0x65.

- If NetDrive is not connected to any remote drive images it will not use a packet driver, making 0x63 (the network card packet driver) available for mTCP or other programs to use.
- When NetDrive is connected to a remote drive image it takes over the network card packet driver and provides a secondary passthrough packet driver at 0x65. mTCP programs will detect that the network card packet driver is in use and try to use the passthrough packet driver at 0x65.

See Networking related configuration parameters for more information about PACKETINT.

This feature works with mTCP programs and with other programs that scan the software interrupts for an available packet driver to use. (mTCP NetDrive will "hide" the network card packet driver so that it is not detected by software that scans for loaded packet drivers, allowing the search to find the passthrough driver.) If your other network code has to be explicitly told which software interrupt to use then you will need to deal with giving it one value (the network card packet driver software interrupt) when NetDrive is not connected to a remote image and a different value (the passthrough software interrupt) when NetDrive is connected to a remote image.

Limitations:

- NetDrive will filter traffic destined for it (IP address and UDP port.) Programs using the passthrough packet driver can not use the same IP address and UDP port.
- The passthrough packet driver implements the minimum required feature set for a packet driver. If you need extended or high performance features it will not provide those.

Setting the DOS Verify mode

The DOS VERIFY command can be used to toggle verify mode, which when enabled tells the code controlling devices to re-read the data after a write operation to verify that the data was written correctly, if feasible. By default this mode is not enabled as most computer hardware is reliable enough and the extra read-back operation slows things down. (Floppy disks are generally not as reliable and they should probably always use verify mode, but that ship sailed a long time ago …)

Network drives in DOS generally use the DOS Network Redirector interface, and thus the verify mode has no effect on them because only block mode device drivers implement the write-verify operation. Even though mTCP NetDrive operates over the network, it is a block mode device to DOS and therefore it correctly handles write verify operations.

When verify mode is on mTCP NetDrive will:

- Send a write-verify command to the server.
- Compute a CRC16 or Fletcher-32 checksum ("the check data") over the original data.
- Wait for the server to respond.
- Compare the check data computed by the server with its locally computed check data.
- If the received check data matches, the write is marked as successful.

On the server side the server will:

- Receive and execute the write-very command.
- Re-read the same data back from the image file.
- Compute a CRC16 or Fletcher-32 checksum ("the check data") over the re-read data.
- Send the result back to the client.

While this does not guarantee that the server operating system wrote the data to media correctly, it does provide a high level of assurance that it happened. Using the check data instead of sending back the actual bytes for a byte-by-byte comparison is used to reduce network traffic, enable support for larger writes, and also improves performance by allowing the computation of the client check data to overlap the network transit time. (The performance of write and write-verify can be nearly identical if the client is fast enough or if there is a small

distance between the client and the server.)

Error handling

When DOS has a problem reading or writing to a block device it shows the famous "Abort, Retry, Fail" prompt. The prompt comes from a part of DOS called the "critical error handler." As the NetDrive device driver is called by DOS, if it returns an error during a read or a write operation it will cause the critical error handler prompt to appear.

mTCP NetDrive uses UDP to communicate with the server. UDP does not guarantee the successful delivery of packets. To compensate for this the mTCP NetDrive protocol has its own timeout and retry logic.

When NetDrive makes a request to the server it starts a timer. If the timer expires without a response that means that either the original request packet or the response packet from the server was lost. The device driver will automatically retry the request up to two times. If the request continues to fail the device driver returns an error, causing DOS to present the "Abort, Retry, Fail?" prompt.

It is unusual (but not uncommon) to lose a packet on a modern network and most packet loss problems are resolved on the first automatic retry. If you see an "Abort, Retry, Fail?" prompt on the network drive you might have a very bad network connection or a problem on the server side. If the network problems have passed then answering "Retry" will work and everything will resume correctly. This is similar to what happens if you interrupt reading or writing to a floppy disk by opening the floppy drive door.

(For testing purposes I physically disconnect the Ethernet cable from the machine to simulate a bad network, wait for the "Abort, Retry, Fail?" message, reconnect the Ethernet cable, and then use "Retry" to get things moving again. I also have a test mode on the server that artificially injects errors on good commands to test the DOS client's error handling.)

Avoid using "Abort"

In general, use the Retry or Fail options and avoid using Abort at all costs. Abort basically ends the program immediately and does not give the program a chance to restore interrupt handlers or do any other cleanup that is needed. This can cause corruption and system crashes, and that is true even when NetDrive is not in use – it is a "feature" of DOS. Retry works better, assuming the error is transient, and Fail gives the calling program a chance to figure out what to do instead.

CHKDSK errors

Some programs replace the DOS critical error handler with their own critical error handler so you won't see those error messages. CHKDSK and older versions of the Norton Utilities are two examples of programs that do this. This is potentially bad because CHKDSK might interpret a timeout improperly and report that data is corrupted when it really is not. This can lead you to allow CHKDSK to fix the errors, making things worse and leading to data loss.

To prevent this from happening:

- Never run CHKDSK if you are seeing lost packets on a connection. (The NETDRIVE STATUS command will tell you how many retries have happened on a connection.)
- If CHKDSK does report a problem, do not allow it to fix the problem the first time. Use CTRL-Break

and then run it again to see if the same exact error is reported again.
- If the error changes or goes away use NETDRIVE STATUS to see if you are having packet loss. Whether CHKDSK reports an error or the type of error it reports will vary depending on the degree of packet loss.
- Some errors might be legitimate and there will be no packet loss. (Writing bad data to the FAT can happen due to software bugs. If the same error is reported several times and there is no packet loss then make a backup of the disk image first before allowing CHKDSK to fix the problem.

If the server goes away

If the server ends or the session times out while a remote disk image is connected the retry logic won't be able to get around the error as there is no longer a server to talk to. If this should happen choose "Fail" when you see the critical error handler prompt to get back to the DOS prompt. (You may need to tell the current operation to do this several times; the number of times depends on the error handling in the program that was interrupted.) Then switch to a valid drive letter and run the NETDRIVE DISCONNECT command. That will clean up the DOS side and restore the tiny read-only RAM disk to the reserved drive letter. As noted above, the NETDRIVE DISCONNECT command does the right thing even if the server can't be reached.

Having the server terminate while a remote disk image is connected is like pulling the power plug on a local hard drive. You will want to run CHKDSK on the network drive once things are working again.

Changing the error handling behavior

If you need to change the timeout or retry behavior two environment variables can be used:

```
SET NETDRIVE_TIMEOUT=x
```

where x is the number of seconds before timing out a command and

```
SET NETDRIVE_RETRIES=y
```

where y is the number of automatic retries to do before letting DOS report an error.

Set these environment variables, then run NETDRIVE STATUS to update the device driver to use the new values.

Performance tuning: Read-ahead and Large Writes

NetDrive implements two features to improve performance. While primarily designed to hide the latency of distant servers, these features also provide some benefit on local networks. Depending on your hardware you may need to adjust the default settings to improve your performance.

Read-ahead

A program running in DOS calls DOS to read data from a disk. The size of the read varies with the program; accessing the FAT is generally done one block at a time while copying files or loading a program uses larger reads. DOS routes the request to the code that owns the drive letter.

When the request is delivered to NetDrive it is told the size of the read to perform. Small reads (one or two blocks) fit in a network packet while larger reads must be broken up into several network packets. Each packet

requires a full round trip to the server; when the server is close that time is negligible but when the server is distant the round time becomes larger than the processing time itself and it reduces how fast data can be moved.

To get around the round-trip latency problem the NetDrive client will send a hint to the server telling it how many consecutive blocks it is going to request as part of the high-level DOS read it is executing. The hint enables the server to send additional blocks that it knows the client will be requesting shortly. If those blocks arrive before the client formally requests them that round-trip time is eliminated, greatly improving performance.

If the high level DOS read is only reading one or two blocks at a time then there is no opportunity for improvement. This commonly happens when DOS is manipulating the FAT. However, if DOS is loading a program or reading a file it will often read multiple blocks at a time. This is more likely to happen on hard drive images, where the cluster size is a multiple of disk blocks. (Larger images generally have more blocks per cluster, increasing the opportunity for performance improvements.)

In general, more read-ahead means better performance. But it is possible to have too much of a good thing; a fast server can easily flood a slower DOS machine with read-ahead packets. To prevent that the NetDrive client probes the CPU class and limits the number of read-ahead packets that the server will send based on what it detects.

CPU class detected	Read-ahead
8088	4KB
80286	4KB
80386	8KB
80486 or better	16KB

Slower machines limit read-ahead to 4KB based on CPU ability and network card limitations – most network cards designed for 8 bit buses have a fairly small onboard memory buffer. Faster machines have higher limits based on their speed and generally better network cards.

You may need to override the read-ahead limit for your specific hardware. For example, if you use an 8 bit 3Com 3C503 network card on an 80386 system you are going to have a bad time because that card is fairly slow and only has about onboard buffers for about 4 packets. You will see this in the number of retries for the connection when you use the STATUS command.

There are two ways to set the read-ahead value for a specific connection:

- Set it at connect time using the “-mra” option of the connect command.
- Use the “set_max_ra” command in NETDRIVE.EXE or ND_CACHE.EXE programs.

The range of values for read-ahead is 0 (no read-ahead) to 64KB. Larger values generally offer greater performance, limited by the speed of your computer and the buffer on the network card. If you are experiencing slow performance due to a high number of retries, try lowering the value for read-ahead or using the ND_CACHE program to find a good value for a specific connection.

Finding a good read-ahead value with ND_CACHE

The nd_cache.exe program can be used to test your specific machine against a specific server to find the optimal read-ahead value. It works by doing a series of timed tests with different read-ahead values, measuring the actual throughput of your connection. Once you have found a good read-ahead value for the connection nd_cache.exe or netdrive.exe can be used to set that value for the connection. That read-ahead value can also be set by the netdrive.exe program when you first make a connection using the -mra command line flag.

To run a test give nd_cache the drive to test against, the number of seconds to test each variation, and optionally provide a time source. For example:

```
ND_CACHE TEST E: 5 -TS CMOS
```

says to run a test against drive E:, use the CMOS clock for timing, and each test each variation for 5 seconds. The output will look something like this:

```
C:\MTCP>nd_cache test d: 5 -ts cmos
mTCP nd_cache by M Brutman (mbbrutman@gmail.com) (C)opyright 2025-2026
Version: Sep 30 2026

NetDrive device opened, current read-ahead cache size is 4 KB

Running trials for 5 seconds, use Ctrl-Break to stop early.
A ** after a measurement indicates a new best speed.

Cache:   0 KB,      90 KB read in  5 seconds,     18.0 KB/sec, 0 retries   **
Cache:   1 KB,     176 KB read in  5 seconds,     35.2 KB/sec, 0 retries   **
Cache:   2 KB,     201 KB read in  5 seconds,     40.2 KB/sec, 1 retries   **
Cache:   3 KB,     312 KB read in  5 seconds,     62.4 KB/sec, 0 retries   **
Cache:   4 KB,     375 KB read in  5 seconds,     75.0 KB/sec, 0 retries   **
Cache:   5 KB,     426 KB read in  5 seconds,     85.2 KB/sec, 0 retries   **
Cache:   6 KB,     476 KB read in  5 seconds,     95.2 KB/sec, 0 retries   **
Cache:   7 KB,     528 KB read in  5 seconds,    105.6 KB/sec, 0 retries   **
Cache:   8 KB,     567 KB read in  5 seconds,    113.4 KB/sec, 0 retries   **
Cache:  10 KB,     638 KB read in  5 seconds,    127.6 KB/sec, 0 retries   **
Cache:  12 KB,     702 KB read in  5 seconds,    140.4 KB/sec, 0 retries   **
Cache:  14 KB,     750 KB read in  5 seconds,    150.0 KB/sec, 0 retries   **
Cache:  16 KB,     289 KB read in  5 seconds,     57.8 KB/sec, 4 retries
Cache:  18 KB,     114 KB read in  5 seconds,     22.8 KB/sec, 5 retries
Cache:  20 KB,     168 KB read in  5 seconds,     33.6 KB/sec, 5 retries
Cache:  22 KB,     161 KB read in  5 seconds,     32.2 KB/sec, 5 retries
Cache:  24 KB,     200 KB read in  5 seconds,     40.0 KB/sec, 5 retries
Cache:  26 KB,     270 KB read in  5 seconds,     54.0 KB/sec, 4 retries
Cache:  28 KB,     145 KB read in  5 seconds,     29.0 KB/sec, 5 retries
Cache:  30 KB,     217 KB read in  6 seconds,     36.1 KB/sec, 5 retries
Cache:  32 KB,     132 KB read in  6 seconds,     22.0 KB/sec, 6 retries

Read-ahead cache restored to 4 KB, use the set_max_ra command
to set it to a new value.

C:\MTCP>_
```

Each iteration tests a different read-ahead size and reports how much data was read and how many retries happened during the iteration. This test was conducted against a remote machine so read-ahead have a fairly large impact. This class of machine (80286) defaults to a 4KB read-ahead, but it can clearly benefit with 8 or 12KB of read-ahead. At 16KB of read-ahead you can see the drop in performance and the increase in retries, which indicates the system is saturated and can't handle any more read-ahead packets. The next step would be to set the read-ahead to 12KB for this connection and see how it behaves under real-world usage.

nd_cache.exe should be used with a CMOS clock as the time source, as the regular DOS time of day function

can be unreliable when a machine is under a high network load. If you are on an XT class machine that does not have a CMOS RTC you will need to use a stopwatch to determine if DOS is keeping accurate time, especially at larger read-ahead values.

Large Writes

Writes on older versions of NetDrive have the same round-trip latency problem that reads do. Any write larger than two blocks in size (the size of a network packet) requires multiple round trips, reducing performance. The larger the round-trip latency, the greater the performance loss.

This version of NetDrive gets around the problem by using IP fragments. Instead of breaking a large write into several independent write packets that each have to be answered, NetDrive will break the large write into several several IP fragments that represent a single large packet. Logically only one write operation is being requested to the server, but now the maximum write size increases from 1KB to 32KB. This works because the client can blast the IP fragment packets at the server as quickly as possible and it only waits for the return code after all of the fragment packets are sent. The server does the heavy lifting of reassembling the packets and sending a response once it has all of the data and all of the processing is done.

The only downsize to large writes is that a lost network packet causes the entire sequence of packets that make up the large write to need to be resent. (If the server does not receive all of the fragments in a reasonable time it will give up and not send a return code to the client. This will cause a retry on the client side.) These days on most networks dropped packets are unusual so this should not be a problem.

If you need to reduce the size of writes or disable the large writes feature for a specific connection there are two ways to do it:

- Set it at connect time using the "-mws" option of the connect command.
- Use the "SET_MAX_WS" command in NETDRIVE.EXE program.

Setting the maximum write size to 1KB disables the large writes feature, making it look like older versions of NetDrive.

On very fast systems it is possible for the large writes feature to send packets to the network card faster than the network card can handle them. NetDrive tries to detect this and when it happens NetDrive will automatically insert a delay between back-to-back packet sends until the network card is happy.

Some packet drivers do not report errors when this is happening so NetDrive does not detect the problem. In that case writes will fail. If this happens to you first disable large writes by using "SET_MAX_WS" command with a value of 1KB in NetDrive. If that works you can turn large writes back on and force a delay value using the "SET_SEND_DELAY" command. For example:

```
NETDRIVE SET_SEND_DELAY D: 0.5
```

will tell NetDrive to add a 0.5 millisecond delay between packets when sending packets for large writes. The range of values for the delay is 0.0 (force no delay) to 15.9 ms.

If NetDrive detects a fast machine (486 or better) it will automatically set the send delay to 0.5 ms. If your machine does not need the delay and you want extra performance on writes use the "-sd" option on the connect command to force it to no delay (0.0ms) when the connection is made.

Known issues

DOS issues

NetDrive can not be ask for specific drive letters

> NetDrive can not be told to install at a specific driver letter; the drive letters are assigned to it by DOS. This is a limitation of DOS and it can not be fixed by NetDrive. (To make it work NetDrive would have to be installed as a TSR and it would have to be aware of the internals of all of the different DOS versions. That would make it much larger and far more complicated.)

DISKCOPY and DISKCOMP do not work with NetDrive

> DISKCOPY and DISKCOMP do not work on NetDrive drive letters, as they do not look like removable storage. (This is true even when floppy images are connected.)

FORMAT.COM and SYS.COM should not be used on a NetDrive image

> DOS version 3.2 and later expect a device driver to implement the “Generic IOCTL” call, which is used to do things like interrogate the geometry of a block device, read tracks, format and verify tracks, etc. NetDrive does not implement Generic IOCTL because it is generally not needed. However, this means you can not format or add the boot system files to a NetDrive image using these commands.
>
> Use the CREATE command in the NetDrive server to create formatted disk images for use with NetDrive. Other tools that create standard FAT disk images may also be used.
>
> If you need to have the boot system files on a NetDrive image, start with an image created by something like QEMU and move the image over to NetDrive.

PCjr BIOS keyboard handling

The device driver works on the PCjr but the non-standard keyboard handling in the BIOS can cause the machine to crash if the keyboard buffer is filled while the device driver is active.

The root of the problem is the PCjr design, which requires the keyboard to use the Non-Maskable Interrupt to handle the keyboard decoding. Bugs in the BIOS can cause interrupts to be re-enabled when the keyboard handlers run, causing problems for code that requires interrupts to be disabled.

The easiest thing to do is not to “lay on the keyboard” while the machine is performing disk I/O through the device driver. You can, and the device driver will be fine. But your Ethernet card might not be as tolerant due to timing requirements, and that might cause the Ethernet card to stop responding. (The WD 8003 series seems to have this problem, but a reboot clears it up.)

If you need a more comprehensive fix try using the NOBEEP.SYS device driver. NOBEEP.COM is similar, but is packaged as a TSR. These two programs work around the BIOS bugs, allowing you to use NetDrive without worrying about the keyboard crashing the system.

Ping

Introduction

This is a DOS version of the 'Ping' utility commonly found on systems that support TCP/IP. Ping allows you to do a basic check to see if a remote system is alive and responding to TCP/IP even if it is not accepting connections on specific ports. Ping is also useful for measuring how far away (also known as the ping latency) of a remote system - lower ping times mean better performance.

Using Ping

Usage is like this:

```
ping [options] <your.machinename.here>
```

Options are:

`-a`	Audible ping
`-help`	Show a help message
`-count`	Tell ping how many ping packets to send
`-size`	Set the size of the test data in each packet
`-timeout <n>`	Number of seconds to wait for response packets.

By default ping will send four packets and each packet will have a 32 byte test message in them. The total packet size will be 60 bytes, including the TCP/IP headers.

You can tell ping to send more packets or to send bigger packets. The test data portion of a packet can be up to 1024 bytes in size. Ping sends one packet per second by default.

Here is some sample output:

```
mTCP Ping by M Brutman (mbbrutman@gmail.com) (C)opyright 2009-2022
Version: May 27 2022

Sending ICMP packets to 35.206.78.10
ICMP Packet payload is 256 bytes.

Packet sequence number 0 received from 35.206.78.10 in 54.40 ms, ttl=55
Packet sequence number 1 received from 35.206.78.10 in 56.10 ms, ttl=55
Packet sequence number 2 received from 35.206.78.10 in 55.25 ms, ttl=55
Packet sequence number 3 received from 35.206.78.10 in 56.10 ms, ttl=55

Packets sent: 4, Replies received: 4, Replies lost: 0
Average time for a reply: 55.46 ms (not counting lost packets)
```

You can press Ctrl-C or Ctrl-Break to stop. If you ping a non-existent machine you will either get a DNS error message or an ARP failure.

Return codes

Ping uses the following return codes:

0	All good – every ping packet received a response
1	Usage error
2	DNS error
3	No ping responses received
4	Some ping responses received

Notes

Broadcast ping

Ping now supports sending pings to IP broadcasts addresses. This can be useful for discovering what machines are active on a network without trying to ping each address directly.

When using a broadcast address:

- Not all machines will answer a broadcast ping. (Some consider broadcast ping to be a security issue.)
- Each machine that chooses to respond will send a packet back, so each sent packet results in many more received packets.
- It is not possible to detect lost requests or responses when using broadcast ping.

Ping timer resolution

Normal timer resolution under DOS is limited to 55 milliseconds. This is how fast the BIOS timer clicks, and it works out to about 18 times a second.

Ping can measure as little as 0.1ms of elapsed time. It does this by reading the 8253/8254 Programmable Interrupt Timer chip directly. On the oldest and slowest systems the measured ping round-trip-time will report higher than on faster systems, as the time to receive the packet and process it is also part of the elapsed time. The difference should never be more than 1 or 2 milliseconds unless the system is under a high load, such as when it is dealing with many unrelated network packets.

Ping responses from other mTCP programs

mTCP programs will generally respond to ping requests when they are actively running. The size of the incoming ping requests should be 256 bytes or less, as that is all mTCP allocates for ping responses.

The Ping program itself is capable of sending and responding to Ping requests as large as 1024 bytes.

PktTool

Introduction

Pkttool is a small utility that can:

- scan memory for loaded packet drivers
- read statistics from a loaded packet driver
- capture and display packets flowing on your network (packet sniffer)

The first two functions are useful when trying to diagnose problems with an Ethernet card on your system. The packet sniffer function is an advanced function that can help you diagnose problems on your network.

Using PktTool

No set up is required. Unlike most of the mTCP programs, pkttool does not use the mTCP configuration file. All you need is to have your packet driver loaded.

Pkttool has three primary commands, each of which may have options. The commands are:

scan	scan memory for loaded packet drivers
stats	read statistics from a loaded packet driver
listen	capture and display packets flowing on your network

Scanning for packet drivers

To scan for loaded packet drivers use the scan command:

```
pkttool scan

mTCP pkttool by M Brutman (mbbrutman@gmail.com) (C)opyright 2008-2015
Version: Apr 12 2015

Scanning:

Details for driver at software interrupt: 0x60
  Name: NE2000
  Entry point: 0192:03CE
  Version: 11   Class: 1   Type: 54  Interface Number: 0
  Function flag: 2  (basic and extended functions)
  Current receive mode: packets for this MAC and broadcast packets
  MAC address: AC:DE:48:88:99:AA
```

If any packet drivers are detected information about them will be displayed.

If the Trumpet DOS TCP/IP TSR is found in memory you will be warned about it. mTCP is not compatible with Trumpet but both can be used if you have each one assigned to a different packet driver.

After the pkttool program ends the DOS ERRORLEVEL variable will be set to the number of packet drivers

found, or 0 if none were found.

Displaying packet driver statistics

To display the statistics tracked by each packet driver use the stats command:

```
pkttool stats <x>
```

where <x> is the software interrupt that the packet driver was loaded at. It is a hexadecimal value ranging from 0x60 to 0x80.

And here is some sample output:

```
pkttool stats 0x60

mTCP pkttool by M Brutman (mbbrutman@gmail.com) (C)opyright 2008-2015
Version: Apr 12 2015

Details for driver at software interrupt: 0x60
  Name: NE2000
  Entry point: 0192:03CE
  Version: 11   Class: 1   Type: 54  Interface Number: 0
  Function flag: 2  (basic and extended functions)
  Current receive mode: packets for this MAC and broadcast packets
  MAC address: AC:DE:48:88:99:AA

All statistics are from when the packet driver was loaded.

Packets in:           10  Packets received
Packets out:           4  Packets sent
Bytes in:           2700  Bytes received (includes Ethernet headers)
Bytes out:          1210  Bytes sent (includes Ethernet headers)
Errors in:             0  Receive errors
Errors out:            0  Sending errors
Packets lost:        135  Lost due to no handler, out of buffers, etc.
```

Listening for packets on the network

The most powerful feature of pkttool is the ability to listen for packets being received and to give you a quick idea of what they are.

Under normal circumstances an Ethernet card is set up to listen only for packets that are addressed directly to it or packets that are sent for broadcast to the entire network. A special mode known as "promiscuous mode" can be used to let the packet driver receive all packets on the wire, regardless of the destination. Pkttool uses promiscuous mode to let you see those packets, making it a valuable debugging tool.

To run pkttool in listen mode use the following syntax:

```
pkttool listen <x> [-v] [-raw] [-s <n>] [-outputfile <filename>]
```

where:

<x> is the software interrupt that the packet driver was loaded at. It is a hexadecimal value ranging from

0x60 to 0x80.

-v means you would like to see some basic output that shows each incoming packet. (Think of it as "verbose mode".)

-raw means that you want to see the raw hexadecimal data for each incoming packet.

-s <n> says to capture only n bytes when using -raw mode. So if you have large packets and you only care about the first 128 bytes, this option will let you only capture the first 128 bytes of those larger packets.

-outputfile <filename> instructs pkttool to write its output to the specified file.

When started in listen mode, pkttool will switch your Ethernet card into promiscuous mode and start listening for traffic. Some of the traffic on your network that you might see includes:

- ARP requests
- IPv4 traffic
- IPv6 traffic
- Network switch management traffic

Not all possible traffic will be visible; what you will be able to see will depend on your network switches and routers.

Using the -v option:

If you use the -v option you will see output like this for each incoming packet:

```
16:45:34.06 To: 30:85:A9:68:E6:60 From: 00:21:9B:29:79:BE Len: 80 T:0800
IPv4: From: 192.168.2.100 To: 216.58.216.132 Len: 66 Type: (17) UDP

16:45:34.11 To: 00:21:9B:29:79:BE From: 30:85:A9:68:E6:60 Len: 131 T:0800
IPv4: From: 216.58.216.132 To: 192.168.2.100 Len: 117 Type: (17) UDP

16:45:34.72 To: 33:33:00:00:00:0C From: 00:21:9B:29:79:BE Len: 208 T:86DD
IPv6: From: FE80:0000:0000:0000:FD79:90D5:2FD7:832E
       To: FD79:90D5:2FD7:832E:FF02:0000:0000:0000
      Len: 154 Type: (17) UDP

16:45:35.21 To: FF:FF:FF:FF:FF:FF From: 00:21:9B:29:79:BE Len: 60 T:0806
ARP: Request: 192.168.2.100 is looking for 192.168.2.129
```

For each packet you will see a timestamp, the destination and source MAC address, and the Ethernet packet type. (If you see an Ethernet packet type that is less than 0x0600 that is an older style Ethernet 802.3 frame as opposed to a newer style Ethernet II style frames.)

If the packet type is ARP, IPv4, or IPv6 you will get some additional information from the packet.

Using the -raw option:

If you use the -raw option you will be able to see the data in the packets. That data might contain user IDs and passwords that are not encrypted. Be careful and use this responsibly! You do not want to leak a password because you shared a cool packet trace ...

Each packet will be displayed as a raw hexadecimal dump:

```
0000  30 85 A9 68 E6 60 00 21 9B 29 79 BE 08 00 45 00   0..h.`.!.)y...E.
0010  00 42 5D 06 00 00 80 11 69 D9 C0 A8 02 64 D8 3A   .B].....i....d.:
0020  D8 84 DE FA 01 BB 00 2E 38 D6 0C F0 0C 29 5A 1F   ........8....)Z.
0030  DE 33 EF C2 7A AB 07 81 46 03 21 43 40 D5 69 C4   .3..z...F.!C@.i.
0040  9C 72 D8 5E 08 E3 C2 9E A6 E4 09 48 B8 3C F4 41   .r.^.......H.<.A

0000  FF FF FF FF FF FF 00 21 9B 29 79 BE 08 06 00 01   .......!.)y.....
0010  08 00 06 04 00 01 00 21 9B 29 79 BE C0 A8 02 64   .......!.)y....d
0020  00 00 00 00 00 00 C0 A8 02 81 0C F0 0C 29 5A 1F   .............)Z.
0030  DE 33 EF C2 7A AB 07 81 46 03 21 43               .3..z...F.!C
```

In this example the first packet is an IPv4 packet. The next packet is an ARP packet.

Dumping the hexadecimal for each packet like this takes time and your computer might not be able to keep up. If you see warnings about packets being dropped then consider using the -s option to not dump all of the contents of each packet, or just don't use the -raw option at all. If you are writing to an output file you can improve your speed by writing to a file that is on a RAM disk.

Pkttool will continue to listen for traffic until until you press [ESC] to end the program. At the end of the program you will get some output that summarizes the traffic that was seen. Here is a sample of that output:

```
Tracing active for 2 minutes, 8 seconds

Packets received:            1018
Packets dropped:                0
Ethernet broadcast packets:    73
Unique MACs detected:           4

MAC addr: A4:70:D6:FD:F1:22  Packets sent: 10  Bytes sent: 910
  Packets:     10  Bytes:     910  Type: 0800 Ipv4

MAC addr: 00:1B:A9:7D:8A:8A  Packets sent: 35  Bytes sent: 6330
  Packets:     34  Bytes:    6270  Type: 0800 Ipv4
  Packets:      1  Bytes:      60  Type: 0806 ARP

MAC addr: 00:21:9B:29:79:BE  Packets sent: 545  Bytes sent: 109741
  Packets:    381  Bytes:   92317  Type: 0800 Ipv4
  Packets:     88  Bytes:   12864  Type: 86DD Ipv6
  Packets:     76  Bytes:    4560  Type: 0806 ARP

MAC addr: 30:85:A9:68:E6:60  Packets sent: 428  Bytes sent: 149225
  Packets:    356  Bytes:  144565  Type: 0800 Ipv4
  Packets:     63  Bytes:    3780  Type: ---- (IEEE 802.3 packet)
  Packets:      5  Bytes:     300  Type: 0806 ARP
  Packets:      4  Bytes:     580  Type: 888E EAP over LAN


All hosts combined: Packets sent: 1018  Bytes sent: 266206
  Packets:    781  Bytes:  244062  Type: 0800 Ipv4
  Packets:     88  Bytes:   12864  Type: 86DD Ipv6
  Packets:     82  Bytes:    4920  Type: 0806 ARP
  Packets:     63  Bytes:    3780  Type: ---- (IEEE 802.3 packet)
  Packets:      4  Bytes:     580  Type: 888E EAP over LAN


MAC addresses and their possible IP addresses (based on ARP traffic):

00:1B:A9:7D:8A:8A  ->  192.168.2.20
00:21:9B:29:79:BE  ->  192.168.2.100
30:85:A9:68:E6:60  ->  192.168.2.1
```

Here you can see that packets from four different machines were spotted by pkttool. Besides the expected ARP and IPv4 traffic there was IPv6 traffic, unidentified 802.3 style packets, and four "Extensible Authentication Protocol over LAN" packets.

Pkttool scans ARP traffic to try to build a mapping of MAC addresses to IP addresses. That table appears at the end. It is possible that you will have many more machines on your network than are identified in this table. The table only uses ARP traffic, which is the most reliable way to detect which MAC address has which IP address assigned. If you run pkttool for a long enough time you will have enough ARP traffic to build a fairly complete table.

Return codes

Pkttool sets the following return codes:

General return codes	
0	No errors
128	Usage error
129	Packet driver not found
130	Packet driver does not support the statistics command
131	Unknown error reading statistics from the packet driver
132	File I/O error
150	Miscellaneous error
Additional return codes for the Scan command	
<n>	Where n is the number of packet drivers found
Additional return codes for the Listen command	
0	At least one packet received
1	No packets received

Services

Introduction

Services is a utility that primarily provides an easy way to check if your mTCP version is up to date. Besides the version check it also can randomly serve you a quote, a fortune, or the time (Universal Coordinated Time). In the future more services (hopefully useful ones) are possible.

Services only sends your mTCP version date to the server; no other information is transmitted.

Using Services

The services command uses Netcat under the covers, so mTCP must already be set up and working before you can use it. If DNS is working and you can connect to the outside world then services should work.

Usage is like this:

```
services <command>
```

Commands are:

`help`	Show a help message
`version`	Perform a version check
`quote`	Display a randomly chosen quote
`fortune`	Display a Chinese fortune cookie style fortune
`time`	Tell you the current time (Universal Coordinated Time)

If you use a command that is not in the list services will still send it to the server, but it probably will just tell you that it is an unknown command. (Allowing that enables me to add new features on the server side that older clients can still try to use.)

Here is an example showing the version check:

```
C:\MTCP>services version

C:\MTCP>echo off

mTCP Services by M Brutman (mbbrutman@gmail.com) (C)opyright 2022
  Version: Jul  1 2022

Checking to see if your version is up to date.
Connecting to brutman.com on port 8088 ...

Server response:

Congrats! Your version of mTCP (2022-07-01) is still up-to-date.
```

The response you get will depend on the version number that you sent.

And here is an example that retrieves a quote:

```
C:\MTCP>services quote

C:\MTCP>echo off

mTCP Services by M Brutman (mbbrutman@gmail.com) (C)opyright 2022
  Version: Jul  1 2022

Getting quote ...

Server response:

"I have never let my schooling interfere with my education." -Mark Twain
```

Implementation notes

Services is a little unusual in that it is a DOS batch file that uses Netcat to do the communication with the server. It demonstrates how one can build up a useful utility without having to write a full program. You can use edlin or any other text editor to see how it works.

Netcat prints a startup message and possibly error messages to the screen. A new, not well documented environment variable called NC_SILENT can be set to suppress these messages, making it possible to use Netcat silently as a utility for other programs.

The batch file approach is simple, but it has two drawbacks:

- Running this BAT file from a floppy disk will be slow because of the disk I/O that is required to build up a small, temporary file (services.$$$). You can modify the BAT file to use a small RAM disk for the temporary file to speed things up. (Running from a hard drive should be fine as-is.)
- There is no good way to print an empty line using the echo statement that works on all versions of DOS. I've gotten around this by embedding a backspace character (^H) on echo statements that I use to print a blank line; this trick works on all versions of DOS from 2.0 forward, but it is not obvious. (edlin will show you the embedded backspace character as ^H.)

The serving side is a GoLang program that opens a listening socket, creates a new thread for each connection, and dispatches the request to one of four handlers (version, quote, fortune, and time). You could easily build a server side function with Perl, bash and netcat, or other simple scripting tools.

SNTP (Simple Network Time Protocol)

Introduction

SNTP is a small utility that lets you fetch the current time and date from an NTP server on the internet. This is useful if your machine does not have a battery backed clock/calendar or if you want to adjust for the drift that is inherent in most older clock/calendar solutions.

Using SNTP

Usage is like this:

```
sntp [options] <ntp.server>
```

Options are:

`-help`	Show a help message
`-port <n>`	Contact server on port <n> (default=123)
`-retries <n>`	Number of times to retry if no answer (default=1)
`-set`	Set the system time
`-timeout <n>`	Seconds to wait for a server response (default=3)

The NTP server parameter can be any NTP server on the network. I recommend using a well known pool of timeservers such as pool.ntp.org, time.google.com or time.windows.com.

The date and time that you get from the NTP server will be in Universal Coordinated Time and needs to be adjusted for your local timezone. The TZ environment variable tells the code what your local timezone is and how to make that adjustment. Before running SNTP the 'TZ' environment variable needs to be set. See Setting the TZ environment variable for instructions on how to set it.

SNTP uses the UDP protocol to contact the NTP server. UDP is not a reliable way to send packets on the network, so packets are often lost. The -retries option can be used to improve the chances that SNTP will be able to get a response from an NTP server.

By default SNTP only tells you the time that it received from the NTP server. If you want the time to be set on your system use the -set option.

Examples

```
set TZ=CST6CDT5,M3.2.0/02:00:00,M11.1.0/02:00:00   Set timezone to US Central Standard

sntp pool.ntp.org                       Get the time from pool.ntp.org
sntp -retries 3 pool.ntp.org            Same as above, but retry up to three times
sntp -set pool.ntp.org                  Set the system time to the NTP time
```

Return codes

SNTP will return the following return code using DOS ERRORLEVEL:

0	Completed with no errors
1	Unidentified error
2	Usage error
3	Configuration file error
4	No timezone set. (The TZ environment variable must be set before running SNTP.)
5	Network initialization failed
6	User aborted
7	DNS failure
8	Error setting the DOS time
9	Server timeout

These error codes allow you to determine if SNTP worked or failed in a batch file.

Miscellaneous notes

Accuracy will vary but on a modern network this SNTP client can set the time within a few dozen milliseconds of the time servers time.

There are two major sources of inaccuracy:

- A PC system clock is only accurate to 55ms. This is a limitation of the original IBM PC architecture from 1981 which uses an 8253 programmable interval timer.
- This program does not try to compensate for network latency. If you are on a reasonable good network using a large geographically diverse set of time servers that latency can be as low as 8 milliseconds.

Spdtest

Introduction

This program can be used to test the speed of your machine when sending and receiving data on TCP sockets. It has no practical use other than for benchmarking..

The basic idea is to send or receive data using the TCP library without doing any extra processing on the data. The data will not be copied, written to disk, or manipulated in any way. It will be fully processed by the TCP library though, just like normal data - no shortcuts like not computing the checksums are taken. This allows you to determine the maximum speed you should expect when sending or receiving data. Actual applications will be slower because they presumably do real work.

I use this program to test the speed of new versions of the mTCP code. I also use it to figure out the relative speed of different machines and Ethernet cards.

Using Spdtest

Spdtest uses the following syntax:

```
spdtest -target <ipaddr> <port> [options]
spdtest -listen <port> [options]
```

The first form of the command tells spdtest to create a socket connection to another machine. The second form of the command is used to make spdtest wait for an incoming connection from another machine.

Options are:

`-help`	Show basic help text
`-receive`	Do a receive only test
`-send`	Do a send only test
`-mb <n>`	Megabytes to send during a send test

Spdtest requires another machine to either send it data or to blast data at. To ensure a good benchmark the other machine should be far faster than the machine you are testing. For example, I use a 2Ghz Linux machine with wired Ethernet to test my slower PCs against.

The other machine should be running netcat. If spdtest is doing a send test then the other machine should be set to receive and throw away the data as quickly as possible. If spdtest is doing a receive test then the other machine should be sending a fixed amount of data, hopefully as fast as it can.

While any MTU size is supported, for testing throughput the MTU size should be set to 1500 bytes. This is the largest supported MTU and it will give the best performance. If for some reason the remote machine advertises a low MSS size (which is presumably derived from it’s MTU) then a warning will be issued.

Tips for getting accurate results:

- Both machines should be on the same network
- The network should be relatively quiet
- The machines should be plugged into a switch, not a hub
- The test should run at least 10 seconds, and possibly up to a minute.
- The longer you run the more accurate the results will be.
- Timing should be done with a stopwatch and on the faster machine. (Some Ethernet card and packet driver combinations hold off interrupts too long, distorting time on the machine running spdtest.)
- Check the ending messages for network errors. The test should run cleanly; network errors are a sign of other problems.

At the end of the benchmark statistics will be displayed including the amount of data sent or received, the elapsed time, the number of packets, etc. Statistics that track errors should be 0; you should not be having problems with bad checksums, dropped packets, etc. If you are seeing errors then your performance is being impacted; you should investigate the errors.

Here is the set of commands I use for speed testing:

DOS send testing:

On the DOS side:

```
spdtest -send -listen 2000 -mb 8
```

On the Linux side

```
time nc -d 192.168.2.241 2000 > delete.me
```

That instructs the DOS machine at 192.168.2.241 to listen on port 2000 for an inbound connection. When netcat on Linux makes the connection the DOS machine will send 8MB of data, which the Linux machine will write to a file called delete.me. (File I/O on Linux should be fast enough not to make a difference to the test results.) Timing will be performed on the Linux side, where it should be more accurate.

DOS receive testing:

On the DOS side:

```
spdtest -receive -listen 2000
```

On the Linux side

```
time nc -N 192.168.2.241 2000 < delete.me
```

That instructs the DOS machine at 192.168.2.241 to listen on port 2000 for an inbound connection. When netcat makes the connection the machine will send the contents of the delete.me file. (If we did the send testing first, then delete.me is conveniently at the same size as the send test.) Timing is performed again on the Linux side.

Telnet

Introduction

This is a small telnet client like other telnet clients that you have probably used before. Besides communicating with a telnet server using the telnet protocol it also emulates a standard 'ANSI' terminal and has uploading and downloading capability using Xmodem and Ymodem Batch. This particular telnet client features excellent screen updating performance and a small memory footprint.

Special hardware or software requirements

None

Note for genuine IBM CGA Users: see the section on "Removing CGA Snow" if you are bothered by glitches/noise caused by direct screen writes.

Using Telnet

Telnet uses the following syntax:

```
telnet [options] <telnet_server_addr> [port]
```

where <telnet_server_addr> is the name or numerical IP address of the telnet server you wish to connect to and [port] is an optional port to connect to use. By default the port is set to 23, which is the standard telnet server port.

Options are:

`-help`	Show basic help text
`-sessiontype <telnet|raw>`	Force telnet mode or raw mode

The default session type if telnet if you are using the default port of 23 or raw if you are using any other port. When the session type is telnet the telnet client will expect to do telnet options negotiation with the server. When the session type is raw no telnet options are negotiated. Use the -sessiontype flag to override the default settings as necessary.

For example:

```
telnet bbs.retroarchive.org
```

connects to a machine called 'bbs.retroarchive.org' using the standard telnet port (23). It expects to find a telnet server at that port and it will negotiate telnet options with that server.

```
telnet example.com 5500
```

connects to a machine called example.com on port 5500. This is not the standard telnet port so no telnet options negotiation will be done. If that machine really were running a telnet server at that port you could force telnet

mode using:

```
telnet -sessiontype telnet example.com 5500
```

To make the screen performance tolerable on older machines telnet will take incoming data and render the current screen on a virtual buffer before trying to repaint the real screen. This approach uses a little bit more memory but it dramatically improves the performance of the screen handling, especially when scrolling large amounts of data. You will notice that the screen will pause and stop updating while the machine is getting flooded with incoming data - this is normal and it is much more pleasant than waiting for a slow display adapter to scroll. A nice side-effect of the virtual buffer approach is that you get backscroll capability for free - if something does scroll past the screen you can hit Page Up and Page Down to browse around.

The following special keys are recognized while telnet is running:

PgUp or PgDn	Go up and down through the backscroll buffer
Alt-H	Show a combined help and status screen
Alt-B	Toggle sending DEL chars when the Backspace key is used on and off
Alt-E	Toggle local echoing on and off
Alt-N	Toggle between sending [Enter] as CR/NUL, CR/LF, CR or LF
Alt-W	Toggle automatic wrapping around the right margin on and off
Alt-R	Refresh the screen from our local virtual buffer. (Should not be needed unless I have a screen drawing problem)
Alt-D	Download a file using Xmodem or Ymodem Batch
Alt-U	Upload a file using Xmodem or Ymodem Batch
Alt-F	Drop to DOS to run some commands (make it quick!)
Alt-X	Exit the program (Ctrl-Break does this too) Telnet will wait for the server to close the connection so that you don't lose any last bytes that were being sent. If this is taking too long hit Alt-X again and mTCP Telnet will exit immediately.

When you use one of the toggle options a single beep means it was turned off, while a beep followed quickly by a higher pitched beep means it was turned on.

The following special keys are sent to the server to handle:

- Cursor keys (Up, Down, Left Right)
- Home
- Insert
- Back-Tab (shift Tab)

Telnet protocol features and limitations

This isn't a fully compliant telnet client implementation. It has the following limitations:

- SGA (Server Go Ahead) must be enabled on both sides of the connection. Most servers expect this behavior, as a TCP/IP socket is a full duplex connection.
- This client will not do local line editing.
- This client does not support the telnet Data Mark command. This is not a big issue, but it means that you don't have an easy way to tell the server to squelch a flood of output if it happens. At Ethernet speeds this should not be a big problem.

It does properly support the following telnet options:

- Binary transfer (used during file transfers)
- Echoing - it will request or let the server offer to do echoing. You can also turn on local echoing.
- Terminal type - it will send the terminal type if asked for it
- Window size - it will send the window size in rows and columns if asked

Other options such as passing the environment variables are not supported, but that is not a big deal considering that you are connecting from a DOS machine. (If you have a burning desire for a missing option let me know.)

Uploading and Downloading files

In ye old days one would use a terminal emulation program with a modem to gain access to a BBS or multi-user time sharing system. If you wanted to download a file you used something like Xmodem, Kermit, or Zmodem. Downloading was slow and agonizing, which made the reward for getting a good file that much more precious.

Now people use HTTP and FTP to transfer files over broadband. It just works. It's not interesting. Most people don't even know that FTP or HTTP exist - they are just "getting a file."

In order to give you a taste of what it was like to transfer files back in the old days I have added Xmodem and Ymodem to telnet. Many telnet BBSes support uploading and downloading through telnet connections and Unix systems have supported it using the sx/sb and rx/rb commands for years. It is not as fast as FTP but you can do it right from telnet without quitting to run another program.

Protocols

Here are your options for uploading and downloading:

Xmodem: This is the classic protocol designed by Ward Christensen in 1977. It sends your file in chunks of 128 bytes; if your file size is not a multiple of 128 bytes it will be padded to the next 128 byte multiple. Xmodem does not preserve file sizes or timestamps. You may only transfer one file at a time. The data is protected with a one byte checksum which is probably completely inadequate for unprotected data links but is not necessary over telnet as TCP/IP does its own error detection and correction.

Xmodem CRC: This is a slight improvement on Xmodem. Instead of protecting your data with a 1 byte checksum it uses a 2 byte CRC. It can be used with Xmodem packets that are 128 bytes in length (the standard) or 1 KB in length (an extension).

Xmodem 1K: This is major performance improvement on standard Xmodem. By sending 1KB at a time instead of 128 bytes throughput is dramatically improved. The speed of Xmodem is limited by how quickly the receiving side can tell the sending side "go ahead, I got that last packet" or "that last packet looked

bogus, please resend." Xmodem 1K reduces that overhead by a factor of eight, resulting in much less "dead air" on the wire. Xmodem 1K is almost always used with the CRC option.

Ymodem Batch: Ymodem is an improvement over Xmodem in several ways. It sends the filename, file size and timestamp to the remote side in a special header packet so that you don't have to re-enter the name or worry about the file size and timestamp being altered. Ymodem is a batch protocol so it can be set up to send multiple files at a time. Ymodem uses the CRC option; normally 1KB packets are used but 128 byte packets might be seen near the end of a file transfer.

To download a file (or files) give the server the command to start sending the file(s) and then press Alt-D. You will need to select the protocol to use - be sure this matches what the other side is going to use. If you have a mismatch the file transfer will probably fail. If you choose one of the Xmodem variants you will have to enter the filename to save the file to and the saved file size will be rounded up to the next 128 byte size. The file timestamp will also be set to the current time, as Xmodem does not preserve the file timestamp. If you download using Ymodem Batch just select the protocol and the transfer will start. The filename is automatically received, the file size will be exact, and the original file timestamp from the server size will be preserved.

If you try to download a file that already exists on your system you will be prompted to see if you really want to write over the existing file. If the file can not be over written for some reason then the transfer is aborted. This can happen if the filename happens to be the same as a directory on your local machine.

To upload a file give the server the command to start receiving the file and then press Alt-U. Select a protocol to use, and enter the name of the file to send.

When uploading and downloading use the ESC key to abort the file transfer. Xmodem and Ymodem transfers are normally canceled by sending a special character (Ctrl-X) at least two times. mTCP Telnet will do this for you, but if you need to cancel a file transfer manually just hit Ctrl-X a few times and be patient. (You might need to do this if you started the file transfer on the other side and changed your mind before starting it on the mTCP Telnet side.)

All file operations happen in the current directory. See the section on "Shelling to DOS" for details on how to change directories and manage files.

Limitations on file transfers

When receiving using Ymodem you can tell the other side to send more than one file at a time. But when sending you can only set up mTCP Telnet to send just one file at a time.

Filenames must be in 8.3 format. The sending and receiving of paths is not supported.

There is no Ymodem G. Ymodem G support on other systems is often implemented wrong, or does not exist at all. It is just not worth the hassle.

Notes for Linux/Unix

If you have lrzsz package installed then you have the following commands available:

- rx - receive on the server side using Xmodem
- rb - receive on the server side using Ymodem Batch

- sx - send from the server side using Xmodem
- sb - send from the server side using Ymodem Batch

Sometimes these commands write to STDERR, which will show up in the terminal data stream and confuse Xmodem and Ymodem. If your version of the lrzsz package does that then you can pipe the output of STDERR to a file or to /dev/nul. For example:

```
sx dosfile.exe 2> stderr.txt
```

That will send the file dosfile.exe from Linux to the mTCP Telnet client, piping the output on stderr to a file called stderr.txt. (This assumes you are using bash for your shell. If you use a different shell figure out how to pipe stderr using that shell.)

The send commands (sx and sb) might not use 1KB packet sizes or CRC by default. Not using CRC is not an issue over telnet, but not using 1KB packets can slow you down quite a bit. Use the "-k" option to set up for 1KB packets. (And using CRC never hurts either.)

Another possible step you might need to take is to put your terminal session in “raw” mode. That will prevent the terminal driver from interpreting characters and injecting or removing extra characters from the data stream while your transfer is in progress. To enable raw mode use "stty raw", and to return things to normal use "stty sane".

This sample sequence illustrates the safest way to transfer a file using Unix:

```
stty raw; sx myfile 2> stderr.txt; stty sane
```

That sequence puts the TTY in raw mode and starts a file transfer using Xmodem while piping the extraneous output to a file called stderr.txt. After the transfer the terminal will be put back into sane mode.

```
stty raw; sb -k myfile 2> stderr.txt; stty sane
```

This sequence is similar, except that Ymodem Batch with 1KB blocks is used. This should be much faster than the standard block transfer size, which is 128 bytes.

Notes for Synchronet

Synchronet BBS is great, but it has a bad habit. When using Ymodem batch it will append extra bytes to your file to insert an advertisement for Synchronet. This kind of defeats the purpose of preserving the original file size. While I appreciate the effort that goes into maintaining Synchronet, this is bad behavior.

Shelling to DOS

Sometimes you may need to suspend telnet and execute DOS commands while still connected to a server. You might need to do this manage files that you have just transferred. The Alt-F (File) key in telnet allows you to do this.

When you use this function telnet will be suspended and you will be dropped at a DOS prompt. You can execute commands here; I would limit it to just a few seconds and only using simple commands. For example, if you need to delete a partially downloaded file, change directories, or check a directory for files this function

is ideal. Do not do anything that might alter the screen mode; the screen contents will be restored but telnet is not expecting the mode to change.

While you are at the DOS shell telnet is not processing any incoming Ethernet packets. If you receive something from the other system the connection will time out and you will be disconnected! Keeping that in mind, do not use this function if you are expecting data from the other side. When you are sitting at a Unix command prompt it is probably safe; doing it while a directory listing is coming down is definitely not safe.

Use the "exit" command at the DOS prompt to return to telnet.

ANSI Emulation Notes

Full blown ANSI terminal emulation is a complicated mess. All of the different variations of ANSI emulation over the years have not helped anything either. To put it politely, I've done the best I can and you might still see screen rendering problems.

That being said, I've tried to do a reasonable good job of interpreting ANSI escape sequences and rendering them properly. The ANSI emulation does not depend on ANSI.SYS or any other console device driver - it is all internal to the program. I used a few sources to determine the required set of escape sequences, including the Linux TERMCAP entries for a few ANSI-like terminals, including the generic ANSI terminal definition and a few related PC flavors. The emulation was good enough for me to write this original documentation using VI, browse in text mode with Lynx, run the 'screen' program to get a few virtual terminals, use the 'info' command to browse the terminfo write-up, and run some of the more complicated 'system-config-*' commands provided with Fedora. However, there may be bugs.

Here are a few notes to improve your experience:

Unicode support

Telnet supports Unicode using the UTF-8 encoding which is the preferred encoding on modern Unix-like servers. It is not enabled by default.

If enabled Unicode characters will be mapped to characters in your local code page. Characters that can not be displayed locally will be represented by the “tofu” character (a white block that looks like “■”).

Characters generated on your keyboard that are not 7 bit ASCII will be mapped to Unicode characters before being sent to the server. These include accented Latin alphabet characters, math symbols, characters generated by holding the Alt key and entering a three digit character code on the numeric keypad, etc.

Direct entry of Unicode code points can be done by pressing Alt-Minus, then entering a four digit hexadecimal code for the Unicode code point. For example, entering **`Alt-Minus 2 6 3 A`** will result in a “happy face” emoji (“☺”). Entering **`Alt-Minus 2 2 2 9`** will result in the math intersection symbol (“∩”). If the Unicode code point can not be displayed it will show as the “tofu” character but sent to the server correctly.

See Unicode support for the details on how to enable Unicode support.

Skipping Unicode support

If you choose not to use the Unicode support then incoming UTF-8 codes will look like corrupted text on your display. You can try to avoid this by telling the remote system not to use UTF-8 encoding. On Linux systems do this by setting the LANG environment variable to a non-UTF-8 encoded variant. For example, on my recent Linux boxes the LANG environment variable looks like this when I first login:

```
echo $LANG
en_US.UTF-8
```

Set it to en_US instead to disable UTF-8 encoding on the server side:

```
export LANG=en_US
```

BBSes and Unicode

If you are using Telnet to connect to BBS systems then you may need to enable or disable Unicode support on a per-system basis:

- If the BBS does not have UTF-8 support do not enable Unicode in Telnet when connecting to that BBS. The BBS is probably going to send high-bit characters which will get garbled by the UTF-8 parsing code.
- If the BBS requires UTF-8, then enable UTF-8 in Telnet or you will see a lot of garbage-looking characters.
- If the BBS gives you a choice, generally it is best to not use UTF-8 and to disable it in Telnet. This is especially true of the BBS will let you use Code Page 437. Using UTF-8 won't add anything and just makes things slower.

Code pages

The standard code page built into the MDA or CGA cards is code page 437. No other code pages are possible if you are using these two video cards. Set Lynx to use code page 437 to avoid getting weird characters.

EGA and VGA users are able to load other code pages using DOS internationalization support.

Terminal Types

By default if the telnet server asks type type of terminal is connecting this code will report back as 'ANSI'. This happens during telnet option negotiation. If you want to experiment with other terminal types supported by your system then you can change the string that gets reported. See the section entitled 'Advanced Set up' for instructions on how to do this.

If you just want to make a temporary change most Unix systems have a TERM variable in the environment that you can alter.

Depending on your system there may be several suitable terminal types to explore. Older Linux systems have over 20 variants of ANSI terminals to choose from. Newer Linux systems might just have 'ANSI' and 'PCANSI' to choose from. Look in /usr/share/terminfo/ for possible termcap definitions to play with. (The infocmp command and the terminfo man pages are particularly useful.) If you want to make a custom termcap for this

program check the source code to see what I implemented.

ANSI terminal emulation supports color but it is not as full-featured as an xterm. For example, function keys, the DEL key, and other keys are not mapped and will appear as dead keys.

Monochrome Adapter Users

The monochrome adapter does not display color, but it does have the ability to underline characters properly. If you are on a true monochrome display adapter (MDA) then you should set your terminal type to 'pcansi-mono', or something similar that tells the server that you have a terminal with true underlining capability. The standard ANSI setting will work, but will not enable underlining.

Removing CGA Snow

CGA “Snow” is static or noise that is caused by writing directly to the video buffer memory on CGA cards that do not have dual ported memory. Only the original IBM CGA adapter and close clones are affected by this, and it can be very distracting.

If your CGA card is displaying CGA Snow you can use an environment variable to tell IRCjr to change the way it writes to the screen, thus eliminating the CGA Snow. It will cause screen updates to be slower but it removes the noise.

To enable this feature set the MTCP_NO_SNOW environment variable to any value:

```
SET MTCP_NO_SNOW=1
```

Color rendering

ANSI color supports 8 colors for the foreground and background. A “bold” attribute was often implemented as a brighter shade of the base color, allowing for 16 foreground colors. The supported colors are black, red, green, yellow, blue, magenta, cyan and white.

CGA text mode supports all of those colors except for yellow, which is rendered as brown. Yellow is available as brown with the bold attribute turned on. (“Bright brown.”) This limitation appears on EGA and VGA screens as well because IRCjr doesn't try to do any programming to correct the limitation.

In practice the color substitution is mostly harmless. Just be aware that brown is really yellow, and yellow is the bold version of brown.

Strict DOS ANSI.SYS Mode

Normal ANSI terminals do not home the cursor after clearing the screen, but DOS ANSI.SYS does. If you are playing a BBS door-game or something similar and notice that the screen is being drawn incorrectly after the screen is cleared try setting this environment variable:

```
SETENV TELNET_STRICT_DOS=1
```

That will tell Telnet to home the cursor after clearing the screen, copying the behavior of ANSI.SYS.

Special note: Telnet BBSes and MUDs

Not everything out there claiming to be a telnet server is actually running a compliant telnet server that does telnet option negotiation. A lot of telnet BBSes and multi-user dungeons fall into this category. If you are having trouble trying to connect to something that is not a standard telnet server, try using the "-sessiontype raw" option. You will still get ANSI emulation but all of the telnet protocol negotiation code will be disabled. This is done automatically when you connect to anything other than port 23.

Assuming that works, you might need to turn on local echoing using Alt-E.

[Enter] key handling

The telnet standard spells out that a "newline" character is comprised of a CR (carriage return) followed by a LF (line feed). But it is not clear what it means when a user presses "Enter" at a terminal; that is traditionally just a CR that gets sent to the server. This makes the standard ambiguous.

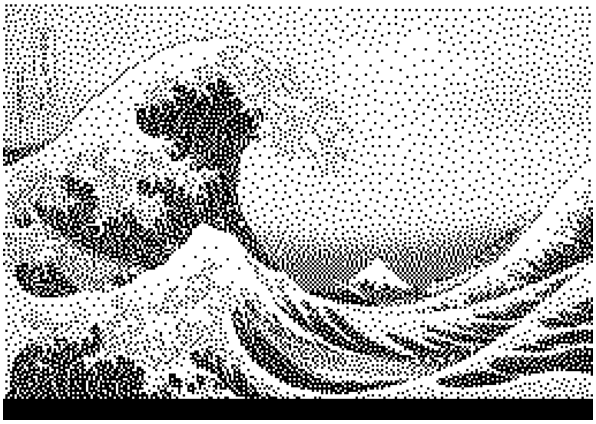
After two years of struggling with the ambiguity I think I have found something that works for most telnet servers. By default pressing the "Enter" key will send a CR/NUL pair. If you need something different for a non-standard telnet server you can use Alt-N to toggle between the following options: CR/NUL, CR, LF, CR/LF. The help screen will show you the current setting.

Regardless of this setting you can always send a CR by pressing Ctrl-M and send a LF by pressing Ctrl-J.

RLE Graphics

This version of Telnet supports CompuServe RLE graphics directly in the program; if you have a graphics capable display and the other system sends an RLE file Telnet will switch to graphics mode and start rendering the RLE file automatically. This is nearly a completely useless feature in 2023 but it does give you a taste of what graphics were like in the CompuServe era.

RLE graphics files can generated from JPEGs and other modern formats using Linux tools such as pbmtocis, allowing you to display graphics right from a Linux machine. The quality of the conversion is limited by the graphics resolution of the RLE format, but it is still pretty fun to play with.



"The Great Wave off Kanagawa" rendered by mTCP Telnet

There is a server at brutman.com, port 8086 that you can connect to and see RLE graphics in action:

```
telnet brutman.com 8086
```

See http://www.brutman.com/RLE/RLE_Graphics.html for more information on CompuServe RLE graphics.

Sixel Graphics

This version of Telnet also has support for Sixel graphics. If you have a graphics capable display and the other system sends a Sixel file Telnet will switch to graphics mode and render the file.

Only black and white Sixel graphics are supported. To prepare a picture for display, use a tool like "convert" on Linux to resize the image. Then use the img2sixel tool with the -e option to convert the image to Sixel. The maximum horizontal resolution is 640 pixels. There is no maximum vertical resolution, as the display will scroll the image as it needs to.

The 640x200 graphics mode has a strange aspect ratio; the dots are nearly 2x taller than they are wide. Be sure to adjust for this when resizing your image.



Buddy the cat, Sixel image rendered by mTCP Telnet

The server at brutman.com, port 8086 also has some sample Sixel graphics files that you can look at:

```
telnet brutman.com 8086
```

Advanced set up

There are some options that you can specify in the MTCPCFG file to override default behavior, such as the initial state of toggles. Below is the list:

`TELNET_VIRTBUFFER_PAGES <n>`

Telnet makes screen access fast by writing to a virtual screen buffer in memory, and then updating the real screen as needed. <n> is the number of virtual screen pages to store in memory. Valid values are from 1 to 8.

`TELNET_CONNECT_TIMEOUT <n>`

<n> is the number of seconds to wait for a connection to the Telnet server. The default is 10 seconds.

`TELNET_AUTOWRAP <n>`

Set <n> to 0 to specify no autowrap or 1 to turn autowrap on. The default is to turn autowrap on.

`TELNET_SENDBSASDEL <n>`

Set <n> to 1 to send a DEL character (ASCII 127) when the Backspace key is hit instead of sending a Backspace character (ASCII 8). The default is to send DEL.

`TELNET_TERMTYPE <termtype>`

Report <termtype> as your terminal type during Telnet option negotiation. The default termtype is 'ANSI'. By convention terminal types reported during telnet option negotiation use upper case letters and the telnet client will automatically convert lower case letters to upper case letters.

`TELNET_SEND_NEWLINE <chars>`

Choose which characters will be sent when you press your Enter key. The valid values for TELNET_SEND_NEWLINE are "CR/LF", "CR", "LF", "CR/NUL", or "AUTO". The default is AUTO, which uses CR/NUL unless telnet binary mode is enabled. If binary mode is active then just a CR will be sent. (You don't have any direct control over binary mode so don't worry about it.) See the section on [Enter] key handling above for details.)

Examples:

```
TELNET_SEND_NEWLINE AUTO          Default: use CR/NUL or just CR automatically
TELNET_SEND_NEWLINE CR/LF         Send both CR and LF at all times
TELNET_SEND_NEWLINE CR            Send just a CR
TELNET_SEND_NEWLINE LF            Send just an LF
TELNET_SEND_NEWLINE CR/NUL        Send CR/NUL
```

Advanced Topics

Unicode support

Introduction

mTCP programs generally only use only the 7 bit US ASCII character set and some line drawing or graphics characters. This is sufficient for programs like PING, DHCP, SNTP, etc. While DOS allows you to set a code page and keyboard layout, the mTCP programs are not aware of those settings and will blindly pass along whatever characters are provided to them including the high bit characters in the current code page. This can cause problems with IRCjr or Telnet which connect to other systems that are probably not using the same code page or encoding settings. (Most modern servers expect Unicode via the UTF-8 encoding.)

This version of mTCP adds some basic Unicode support to IRCjr and Telnet, hopefully making things more pleasant for people who need better extended language support. The support is based on a mapping of Unicode code points to local code page characters supplied in a text file.

When a Unicode code point encoded using UTF-8 is received it is decoded and mapped to a character in the current code page. If no mapping was found the "tofu" character (a white block that looks like "■") is used. While somewhat limited, this allows you to use all of the extended characters in the current code page and it is more pleasant than the "mojibake" that you get from trying to interpret text with the wrong encoding.

Similarly, if you enter a non-US ASCII character on your keyboard the character will be mapped to a Unicode code point and transmitted using UTF-8. These characters might be accented letters on a non-US keyboard or a character generated using the Alt key and a three digit sequence on the numeric keypad.

Set up

Unicode support is disabled by default. Without Unicode the programs are limited to the 256 characters or symbols available in the current code page. For early IBM compatibles using MDA or CGA display adapters these are the symbols in code page 437, which includes the standard US ASCII characters, some Latin-based characters, line drawings, graphics characters and mathematical symbols. Later machines using EGA or VGA can load alternative code pages, swapping out the upper 128 symbols with a different set of characters and symbols. Whether you use code page 437 or something else there are still only 256 possible characters and symbols that can be displayed when using a text mode program.

The Unicode support in these programs allows you to map incoming Unicode code points to one of the 256 characters or symbols that are available in your current code page. The standard US ASCII letters, numbers, and symbols will always map correctly and are handled by the programs automatically. The upper set of 128 characters and symbols in the code page will always have a mapping to a Unicode code point but the reverse is not true as Unicode has over 140,000 code points. Some Unicode code points can be mapped to things that are close to the local character set, while most Unicode code points will have no suitable mapping.

IRCjr and Telnet do not provide a default mapping from Unicode code points to a code page character as there are many code points, many code pages and the mapping process can be subjective. The programs let you choose your own mapping, allowing you to map Unicode code points to the characters and symbols available to you in your selected code page. This gives you the flexibility to decide on a strict mapping of Unicode code points to local characters and symbols or tolerate some "close" but imperfect mappings.

A sample mapping for code page 437 is provided. If you are using a different code page you will need to create

a text file that defines the mapping from Unicode to your selected code page. The standard 7 bit ASCII characters and symbols do not need to be explicitly mapped as those are standard and a perfect mapping already exists.

Here are the rules for the file:

- A line starting with a '#' is a comment.
- Blank lines are allowed.
- A mapping from a Unicode code point to the local code page 8 bit value is defined by putting the hexadecimal value for the Unicode code point followed by the local 8 bit value on a single line.

Here are some lines from the sample mapping file:

```
# Unicode to CP437 mapping table
# 2023-01-21 M. Brutman

# Values 0x00 to 0x1f and 0x80 to 0xff from
# https://github.com/torvalds/linux/blob/master/drivers/tty/vt/cp437.uni.
# Duplicates are listed on separate lines,
#
0x263a 0x01
0x263b 0x02
0x2665 0x03
0x2666 0x04
0x25c6 0x04
```

In this sample we see five Unicode code points being mapped to four characters in code page 437. Note that in this sample two different Unicode code points (♦ and ◆) are being mapped to the same local character (0x04, a black diamond character.) In Unicode those are slightly different black diamonds and should be treated differently, however for display purposes code page character 0x04 is close enough.

You do not need to provide mappings for Unicode values 0x0000 to 0x007F, as those are the same as the 7-bit US ASCII character set. A minimal mapping file will provide mappings from Unicode to the 128 high bit characters in a code page. A more complete mapping file will include additional mappings that are close enough for display purposes. While up to 511 different mappings from Unicode code points to code page characters in a file are allowed, for performance reasons limit it to 400 or less mappings.

To tell mTCP to enable Unicode support and use the mapping file set the MTCP_UNICODE_MAP environment variable to the location of your mapping file before running IRCjr or Telnet. For example:

```
SET MTCP_UNICODE_MAP=C:\MTCP\CP437.TXT
```

If the environment variable is not set or the file is invalid Unicode will not be enabled.

Limitations

Only UTF-8 encoding is supported. Other forms of Unicode encoding are not available. UTF-8 is universally supported and preferred for IRC and Telnet, so this is not a significant limitation.

Internally the Unicode support is limited to 16 bit code points. This means that only Unicode Plane 0, the Basic Multilingual Plane is supported. This should not be a problem as the Basic Multilingual Plane covers anything

displayable by a DOS text program.

Only a simple mapping from a code point to a character is supported. The more advanced features of Unicode, such as spacing modifiers, combining diacritical marks, etc. are not supported. A system that sends these characters will cause the screen to be rendered incorrectly.

When a Unicode code point does not have a mapping defined a white square block will be displayed. This is commonly known as the “tofu” character and is displayed as a white block that looks like “■.”

A mapping file may provide up to 511 mappings, but for the best performance do not use more than 400 mappings.

Debugging problems

Introduction

Welcome to the wonderful world of DOS networking! If you are reading this you are either bored or you have a problem. Hopefully this will help you get past your problem.

Keep in mind there is a wide variety of hardware out there including emulators on modern machines which are technically not even hardware! This guide will give you some pointers on how to get past problems but it is going to take some work.

This guide is broken up into two parts - [1] debugging hardware and packet driver problems and [2] debugging mTCP problems.

Besides this guide check out the instructions for the PKTTOOL.EXE. PKTTOOL can talk to a loaded packet driver and get statistics from it, allowing you to see if the packet driver is seeing any packets "on the wire." There is also a packet sniffer tool that will let you look at the traffic on your network. (Hint: if you never see any traffic, then there is something wrong!)

Common symptoms and causes

Below are some common symptoms and possible causes:

Red Flags:

- Is the packet driver reporting the MAC address correctly? It should match a sticker on the card. If it does not or if the MAC address is 00:00:00:00:00:00 or FF:FF:FF:FF:FF:FF then something is wrong with the packet driver or the card.
- Checksum errors in the traces: see the section below entitled “Checksum errors”

No packets are being received:

- The hardware IRQ for the Ethernet adapter is incorrectly set. When this happens you will be able to send packets but not receive them.
- Something is wrong with your cabling. If you have a card with more than one connection type (e.g AUI and twisted pair) on it check the jumpers or packet driver settings.
- Try a different switch or hub. Examine the status lights on the hub to see if it detects the cable being plugged in.
- Ensure the card is compatible with your connection type. For example, there is a variant of the WD 8003 series that has an RJ45 connector but it is Lattisnet compatible, not Ethernet compatible. Without knowing this it would look like a malfunctioning card.
- Is the network address that you have chosen correct? Is it possibly being used by a different machine?

Connections are slow or freeze:

- Is the network address that you have chosen correct? Is it possibly being used by a different machine?
- Are you running DOS 5 or later on a slow machine on a large hard drive partition? If so then DOS might be pausing to compute the free space on the partition. The larger the partition the longer the wait will be.
- Reduce your MTU to 576. Ethernet MTU is 1500 bytes but IP gateways are only required to support 576 byte packets. It is possible that fragments are being created at a gateway; reducing your MTU will

ensure that you don't send packets that are too large for that gateway. (You may still have problems though because the fragment reassembly code is fairly slow compared to the fast-path normal code.)
- Flow control is difficult. There will naturally be a pause every time a packet is lost in transit. Depending on the TCP/IP stack of the other machine, a single dropped packet can easily ruin a connection by basically making every packet go through the re-transmit path. mTCP doesn't implement the selective acknowledgment (SACK) TCP option but it does temporarily shrink the TCP receive window when more than one packet in a row is lost, enabling the connection to recover. (This feature is new in the 2019 release.)

FTP problems
- FTP won't send or receive files! Passive mode is the default and it should work in nearly all cases. However, I do know of a buggy FTP server implementation (ProFTPD) where early versions could not send files reliably in Passive mode. Try Port mode if Passive doesn't work.
- FTP is corrupting my data! Did you specify BIN (binary) mode before starting your data transfer? The FTP specification leaves the choice of the default data transfer mode up to the server, not the client. If you transfer a binary file and forget to set the transfer type to BIN or IMAGE before starting, all of the carriage returns and line feeds are going to be altered. This is not an FTP bug; it's working as designed. Make sure you specify BIN before starting binary file transfers.

Hardware and packet driver problems

Note: Some of this might not apply to PCI cards. Parallel port devices like the Xircom PE3 series require the parallel port to be correctly configured. Serial line users (SLIP or PPP) need to ensure that their COM ports are set up correctly. Do what makes sense for your specific setup.

mTCP will not work if your Ethernet card is not set up correctly. Here is a partial list of things that can trip you up:

I/O ports

Like any PC hardware, your Ethernet card needs to be set up to use I/O ports. The I/O ports can not be shared with other cards, either on purpose or by accident. If there is a conflict with another card you will get flaky operation or maybe even random crashes.

If your packet driver loads and reports the correct MAC address for your Ethernet card you probably have the I/O port set correctly. (You might still have conflicts with other cards though.)

A MAC address is a unique identifier that is 48 bits long usually written in hexadecimal notation with a colon separating each byte. For example, 02:31:BE:D1:2F:08. The MAC address can usually be found printed on a sticker on the card. A MAC address that is all 00s (00:00:00:00:00:00) or all FFs (FF:FF:FF:FF:FF:FF) is probably not correct and indicates that the packet driver is not talking to the card.

If the MAC address reported by the packet driver does not match the MAC address printed on the card then it is possible that it has been changed on the card, but it is more likely that you have the wrong I/O port address.

Hardware Interrupts (IRQ)

Your Ethernet card probably needs a hardware interrupt assigned to it so that it can tell the packet driver when a new packet has arrived on the wire. The IRQ has to be set on the card so that it is not shared with any other

hardware in your system, including devices that might be on your motherboard.

PC and XT class systems generally have very few interrupts available. IRQ 0, 1, and 6 are usually always used or unavailable. IRQ 2, 3, 4, 5, and 7 may be available depending on what else you have in your system.

AT class systems have more choices for interrupts. Unlike an XT class system, IRQ 2 is not available on an AT class system.

If the IRQ is not set correctly then your card may be able to send packets but each time it receives a packet it will not be able to tell the packet driver about it. This will look like your card is never receiving packets.

Shared Memory

All Ethernet cards use some memory to store incoming and outgoing packets. Simple Ethernet cards keep that memory on the card in a private area that the rest of the computer can not see. Better cards expose this memory to the computer which improves performance but requires you to find a memory range that does not conflict with other cards or motherboard settings.

The NE1000 card is an example of a card that does not use shared memory. (The NE1000 just needs I/O ports and an IRQ.) The Western Digital 8003 series is an example of card that does required shared memory. It might need to show 8KB or 32KB (depending on the card) to the system, requiring that much space free in your memory map.

Cards that use shared memory will generally need space in the area reserved for expansion cards which ranges from C000:0000 to F000:0000. Be sure to pick a range that does not conflict with other devices.

Cabling selection

Besides the basic resource allocation settings above there may be other settings that you have to make on the card. One good example is setting the cable type - some cards make you set jumpers to choose between the AUI port, Thinnet connector, or RJ45 connector. If this is set wrong the card will appear to work normally but will not be able to send or receive data on the wire.

Bad cabling or network connection

Even with a properly configured card you can still encounter problems.

Believe it or not, cables go bad. If you suspect a cabling problem find another cable that you know works and use that instead. Cables are cheap so there is no reason to wrestle with a potentially bad one.

Most Ethernet cards, hubs and switches have indicator lights that tell you if a good connection has been made. Examine your indicator lights and ensure that your hardware thinks there is a connection. Use a known good port on your switch or hub if you even get a hint of a problem with the port you are using.

Some cable modems are configured to only hand out a limited number of addresses and will refuse to hand out more. You may need to reset the cable modem to clear its list of clients. Switches and hubs are generally reliable but sometimes cycling the power (turning them off then on) clears up problems.

Older Ethernet cards pose special challenges. Before modern Ethernet standardized on CAT 5 cabling there was ThickNet and ThinNet cabling. ThickNet (the original Ethernet) was often supported by using the AUI port on

a card with an external transceiver, while ThinNet was often directly connected to the card. The card would have to be configured to use one port or the other port and choosing the wrong port will make it look like you have bad cabling.

Even on cards that look like they support twisted pair wiring (CAT5) the signaling might not be correct, and thus not work on modern switches and hubs. For example older WD8003 series cards have both an AUI port for an external transceiver and an RJ45 jack which accepts a CAT5 cable, but it actually uses Lattisnet which is not compatible with Ethernet signaling. It is close enough to make a hub or switch light up the port, but not close enough to enable data transfer.

Packet driver

First, ensure you have the right packet driver for your card! Packet drivers are very specific to models of Ethernet cards. I always try to look for the packet driver from the manufacturer that was shipped with the card. If I can't find that I might use a packet driver from a card in the same chipset family.

A good source for packet drivers for older cards is the Crynwr packet driver collection at http://crynwr.com/. They defined the packet driver specification and did a lot of the early work to make packet drivers available for ISA cards. Source code for their work is available, which is an added bonus.

If your Ethernet card is set up correctly you should know the I/O port and other settings to pass to the packet driver on the command line. Packet drivers can detect some configuration problems or mismatches but not all of them, so be sure you have the settings that match your card.

The PKTTOOL.EXE program might provide some insight as to how the packet driver is behaving. See PktTool for instruction on how to use it.

Virtual machine problems

Virtual machines are usually a great way to experience other operating systems, but sometimes they add complexity that is difficult to debug.

In particular most modern Ethernet hardware includes support for computing TCP and IP checksums, allowing the host operating system to offload those functions to the Ethernet hardware. This (and other acceleration features) can fool mTCP because the host OS might present packets to mTCP with the checksums missing.

If you suspect this is happening to you try to turn off features of your Ethernet card that alter the packet. There will be a small performance cost but it may fix the problem.

Collecting mTCP traces

I would rather have you believe that mTCP is perfect, but we know that nothing ever is. But we can strive ...

In general, the mTCP code is too complex for a casual user to try to debug. If there really is a problem in the code we are going to need to take a trace so that I can see what is going wrong.

Even if mTCP is acting normally you might still be having problems related to your network or setup. To make debugging and fixing problems easier I have included some tracing features in the library.

To enable tracing you need to set two DOS environment variables before running an mTCP program. The LOGFILE variable sets the name of a text file where trace output will be written to. The DEBUGGING variable sets the desired level of tracing to be used. For example:

```
set DEBUGGING=127
set LOGFILE=logfile.txt
```

Those two lines specific log level 127 (to be explained next) and that logfile.txt should hold the trace.

Tracing can be slow so mTCP lets you choose how verbose you want the tracing to be. The number used on the DEBUGGING environment variable controls this. Here are the useful values for DEBUGGING that you can use:

1	Warnings only
3	Warnings and application messages
127	Full tracing except for packet contents
255	Full tracing with some packet contents

The bigger the number the slower the machine will get. Traces can also get quite large. On slower machines tracing will cause a noticeable overhead because of the extra disk I/O.

The number is actually based on a bitmask, so if you want to be selective you can turn on specific types of trace messages. The bit positions are:

0x0001	Warnings
0x0002	General - used by applications
0x0004	ARP - used by ARP
0x0008	IP - used by the IP/ICMP layer
0x0010	UDP - used by UDP
0x0020	TCP - used by TCP
0x0040	DNS - used by DNS
0x0080	Packet dumping

Now that you know the bitmask you may find it easier to set the DEBUGGING variable using a hexadecimal value. For example:

```
set DEBUGGING=0x7F
```

turns on all tracing except for packet dumping, and is equivalent to the "127" level shown in the previous example.

With the debug trace turned on it is possible that you will have a crash and not see all of the debug messages. This is because the debug messages are buffered and depending on the timing the buffer may not be fully flushed at the time of the crash. If you want to be absolutely sure that all trace messages are flushed and visible there is a bit to flip that will aggressively flush the buffers after each trace point is written:

```
0x8000 Turn on aggressive buffer flushing
```

So for example, if you want full packet tracing with normal trace file buffer behavior use:

```
set DEBUGGING=0xFF
```

And if you are paranoid and you want to make sure that you see every trace message, flip the extra bit and use:

```
set DEBUGGING=0x80FF
```

If you have a trace that you would like me to read send it along with a brief explanation of what you were doing and what your network setup looks like. I have been able to help many people get up and running this way and I have fixed a few bugs based on what I discovered in user supplied traces.

Checksum errors

If you have tracing turned on and you get a checksum warning this is very bad. A checksum error indicates that your machine or the last machine to touch the packet has a problem.

Under normal circumstances the IP, TCP or UDP checksums are computed and checked at each “hop” in the network. Routers may have to alter the packet, which forces the checksum to be recomputed. (This happens when Network Address Translation (NAT) is used or fragments are generated.) Ethernet protects the Ethernet frame on the current physical network so you will never see a corrupted frame being delivered at the packet driver. However, the Ethernet CRC can't protect against the higher level software making mistakes or hardware failures on the machines. If somehow a bad checksum is detected by mTCP it means that it was set badly by the last machine or corrupted on the current machine.

While it is true that one can recover from a bad checksum by just re-transmitting the affected packet, the concern is that the source of the corruption is not known. It could be bad memory which is affecting the packets, the stack, or even the code that is running. This makes it very hard to debug what is going on.

If you see checksum errors run the diagnostics on your machine, including memory testers. (Let them run for a few hours if you can.) Also try replacing the network card. If the problem does not look to be on your machine try connecting it a different way to bypass whatever gateway or router is directly upstream, as it may be that device that has the problem.

PPP and SLIP Set up

Introduction

While mTCP is designed for Ethernet cards it is possible to use it over a serial port using PPP (Point-to-Point Protocol) or SLIP (Serial Line Internet Protocol). The trick is to use a PPP or SLIP packet driver that emulates a class 1 packet driver (Ethernet) while it actually uses the serial port.

When using PPP or SLIP you will assign addresses by hand using a static configuration. The DHCP client can not be used.

Prerequisites

- It is assumed that you are familiar with mTCP configuration and that you have some knowledge of Linux system administration. It is possible to use other operating systems by mapping the networking concepts presented here to those other operating systems.

- A known good connection between the machines. This is usually done with a null modem cable. You should verify that the link works using terminal emulation software before you start configuring PPP or SLIP; it makes debugging much easier if you start with a known good setup.

- A packet driver that implements PPP or SLIP while emulating Ethernet:

  PPP: http://www.ibiblio.org/pub/micro/pc-stuff/freedos/files/net/dosppp/dosppp06.zip
  SLIP: https://github.com/mbbrutman/EtherSLIP

  Other packet drivers for PPP and SLIP are available; I list these two as working examples.

Picking addresses for the connection

My router assigns addresses to the machines inside of my home and receives one public address to use from my Internet provider. As with most IPv4 Internet connections my router uses NAT to keep track of the active connections to the outside world and route packets correctly. Normally I use Ethernet connections on my DOS machines so they look like any other PC in the house and can connect to the router through the Ethernet network. My usual network looks something like this:

```
Public Addr <-> My Router   <-> Internal Addresses
a.b.c.d         192.168.2.1     192.168.2.i
```

A serial port can be used with PPP or SLIP to network a machine, but it is slower and it needs another machine to serve as a gateway to the rest of the network. When you add a DOS machine using PPP or SLIP and a Linux machine as a gateway for it the network looks like this:

```
Public Addr <-> My Router   <-> My Computers
a.b.c.d         192.168.2.1     192.168.2.i
                                192.168.2.j (Linux) 192.168.2.x <-> 192.168.2.y (DOS)
```

Note that the Linux machine has to have two addresses to serve as a gateway for the DOS machine. One address is the normal IP address on the Ethernet network assigned by the router and the other is used exclusively to connect to the DOS machine.

Assuming that your Linux gateway machine already has an IP address assigned (192.168.2.j), we need to come up with two more addresses: 192.168.2.x will be a secondary address for the PPP or SLIP connection on the gateway machine and 192.168.2.y will be an address for the PPP or SLIP connection on the DOS machine.

Pick two addresses that do not conflict with other addresses on your network. As long as 192.168.2.x or 192.168.2.y are not in use they will work for establishing the connection. The netmask for the connection will be 255.255.255.255, which signifies a point-to-point connection.

Next, ensure the serial port that you are going to use on the Linux machine is not in use. It is possible that a getty process is running and providing a login prompt on that serial port; it will need to be disabled before using the serial port for PPP or SLIP. The serial port must be dedicated to the PPP or SLIP connection.

Option 1: Establish the PPP connection to the Linux machine

On the Linux side you will be using the PPP daemon. The important parameters to specify are:

- The serial port that will be used.
- The speed of the port.
- "local" to set the flow control.
- "lock" to ensure that other programs do not interfere with the port.
- "passive" or "silent" to tell PPPD to wait for a connection instead of timing out. (Passive and silent are slightly different, but both should work.)
- "proxyarp" to make your machine look as though it is visible on the network.
- A reasonable MTU size. 576 should work fine.
- “noauth” as this is a private connection and you should not need to worry about authentication.
- IP addresses for the Linux side and the remote side.

For example, on my Linux system I use the following command line:

```
sudo pppd ttyUSB0 9600 local lock passive proxyarp defaultroute noauth
mtu 576 192.168.2.91:192.168.2.92
```

(All of that should be on the same line.)

Depending on your hardware you may also need to use the nocrtscts option to disable hardware flow.

In this setup I'll be running at 9600 bps over a USB device providing a serial port. The Linux machine will get a new network interface called ppp0 with address 192.168.2.91 and it will be talking to the DOS PC which will have an address of 192.168.2.92.

If successful the Linux machine will return silently to the command line but the PPP daemon will be running in the background waiting for a new connection from the DOS machine. You can verify this using the ps command.

On the DOS machine we need to run EPPPD to complete the connection. Here we need to specify the serial

port and the speed to use. I use this command on my DOS machine:

```
EPPPD com1 9600 local
```

EPPPD is the PPP driver with Ethernet emulation. This is necessary for mTCP because mTCP is written with Ethernet in mind and can not use a serial link directly.

My DOS machine will be using the first serial port at 9600 bps using flow control suitable for directly connected machines.

After a few seconds of delay EPPPD responds with:

```
Installed packet driver handler at vector 0x60
```

Now you have a live link between the two machines. If you go back to the Linux machine and run "ifconfig" you will be able to see the new network interface:

```
ppp0: flags=4305<UP,POINTOPOINT,RUNNING,NOARP,MULTICAST>  mtu 576
        inet 192.168.2.91  netmask 255.255.255.255  destination 192.168.2.92
        ppp  txqueuelen 3  (Point-to-Point Protocol)
        RX packets 3  bytes 48 (48.0 B)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 5  bytes 71 (71.0 B)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
```

EPPPD will also create a batch file called IP-UP.BAT which sets environment variables suitable for WATTCP programs. mTCP won't use those environment variables but some of their values are useful to us to configure mTCP.

Option 2: Establish the SLIP connection to the Linux machine

Note: These instructions were written with a Linux machine serving as a router in mind. Devices such as "WiFi" modems may also use SLIP, and are basically another computer embedded in the device. If you have such a device follow its specific set up instructions instead, as use these instructions as background information.

On the Linux side you will be using the SLIP daemon (slattach). The important parameters to specify are:

- the protocol to use (slip or compressed slip)
- the serial port to use
- the speed to use

Here is an example:

```
slattach -p slip -L -s 9600 /dev/ttyUSB0
```

That command says to use plain SLIP (not compressed) at 9600 bps over the serial port labeled ttyUSB0.

By default slattach will try to use modem control signals to manage the serial port. If you are using a null modem connection you will probably want to disable that beavior, which is what the -L option does.

slattach creates a network interface but nothing else. Next we have to tell Linux what IP address to use for that

interface. The ifconfig command is used for that:

```
ifconfig sl0 192.168.2.91 pointopoint 192.168.2.92 mtu 576 up
```

Here I am using 192.168.2.91 for the IP address on the Linux side and the DOS side is going to use 192.168.2.92. (The spelling of "pointopoint" is deliberately missing a "t".) The MTU size for the connection is 576, which is a safe value.

If you are using the more modern “ip” command instead of ifconfig these are the equivalent commands:

```
sudo ip addr add 192.168.2.91 peer 192.168.2.92 dev sl0
sudo ip link set sl0 mtu 576
sudo ip link set sl0 up
```

Now you have a live link between the two machines. If you go back to the Linux machine and run "ifconfig" you will be able to see the new network interface:

```
sl0: flags=4305<UP,POINTOPOINT,RUNNING,NOARP,MULTICAST>  mtu 576
        inet 192.168.2.91  netmask 255.255.255.255  destination 192.168.2.92
        slip  txqueuelen 10  (Serial Line IP)
        RX packets 21  bytes 1216 (1.1 KiB)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 8  bytes 480 (480.0 B)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
```

On the DOS machine we need to run our SLIP driver to complete the connection. Using EtherSLIP as an example:

```
ethersl 0x60 4 0x3f8 9600
```

That tells EtherSLIP (ethersl) to install using software interrupt 0x60 on COM1 (IRQ4 and port 0x3f8) at 9600 bps.

(Note: Do not use the EtherSLIP version 11.7 packaged with the older Crynwr packet driver collection; it has a serious bug that can cause crashes. Version 11.8 posted at https://github.com/mbbrutman/EtherSLIP has that bug fixed.)

EtherSLIP doesn't let you just name a COM port to use - it wants you to specify the hardware interrupt and I/O port address of the COM port you are going to use directly. Here are the common port assignments:

```
COM1: IRQ 4, Port 0x3F8
COM2: IRQ 3, Port 0x2F8
```

(Your hardware may be different.)

If EtherSLIP is successful it will print out some status information, including the "MAC" address of the simulated Ethernet card.

Configuring mTCP

mTCP configuration is done as usual but with a few changes:

1. You need to set the MTCPSLIP environment variable to something. The actual value does not matter; it just needs to exist. At the DOS command line this command will do the trick:

   SET MTCPSLIP=true

   Note that this same environment variable is used for both PPP and SLIP connections.

2. You should ensure that your MTU setting in the mTCP configuration file matches what you told the SLIP daemon on Linux. I recommend using a value of 576 unless your software requires otherwise.

3. You need to set up the IP addresses statically because you will not be using DHCP in this setup. In particular you need to set the IPADDR, NETMASK, and GATEWAY variables in the mTCP configuration file:

   IPADDR: This is second IP address you used on the ifconfig command in Linux. In this example it is 192.168.2.92.

   NETMASK: Use 255.255.255.255 to indicate a point-to-point link where ARP packets will not be used.

   GATEWAY: If you plan to send packets past your Linux machine to other machines the Linux machine will serve as your gateway. Put its SL0 address here. In this example it is 192.168.2.91.

If these are set correctly you should be able to run PING on the DOS machine and get a response back from the Linux machine. At 9600 bps on a 386 class machine the time for a standard 32 byte PING packet should be around 150ms.

[The math is pretty straight forward - at 9600 bps you get about one byte per millisecond. The standard payload for ping is 32 bytes and you need an IP header, an ICMP header, and four more bytes for the ping identification number and sequence number. The total bytes sent is that overhead (28 bytes) plus the ping payload bytes (32 default) for a total of 60 bytes. PPP adds at least another 8 bytes of framing overhead to that, so the total being sent per packet is closer to 68 bytes. 68 bytes at 9600 bps takes 70.8ms to send and you need to wait for the response which is the same size. So the minimum possible time with no additional overhead is 141ms.]

If you ping from Linux to the DOS machine while it is sitting at the DOS prompt you will not get a response; remember mTCP only responds when an mTCP application is running. If you want to try to ping from Linux run something simple on the mTCP machine. Netcat in "listen" mode will do the trick:

```
nc -listen 2000
```

That will enable you to test PING from Linux to DOS.

Now you have basic connectivity between the two machines. Telnet, FTP, IRCjr and other applications will work if the target machine is your Linux machine. Next we need to enable some advanced routing to go beyond

your Linux machine to other machines on your network and the outside world.

Routing beyond the Linux machine

So far you have a "point-to-point" connection between the two machines and the DOS machine can not see past the Linux machine. In this section we are going to do a few steps to give your DOS machine connectivity to the rest of the world.

SLIP connections: Enable Proxy ARP

(This step is only required for SLIP connections. The PPPD program handles this for you when using PPP.)

The first step is to turn on "proxy ARP" which is a trick that tells the Linux machine to answer ARP requests from other machines on behalf of your DOS machine, which is not directly connected to the network. The "arp -s" command is used to set this up:

```
sudo arp -Ds 192.168.2.92 wlan0 pub
```

Here we are telling the Linux machine that if it should see an ARP request for IP address 192.168.2.92 (the IP address of the DOS machine) that it should answer with its own MAC address. The “-s” adds the ARP table entry while the -D option tells it to use the hardware address of the wlan0 network interface on the Linux machine. Change the IP address and the network interface name as needed for your setup.

PPP and SLIP: Enable IP forwarding

The next step is to tell the Linux machine to forward packets between the two network interfaces. (Even though the new connection has an address that looks like it belongs on your Ethernet network, it is still treated as a separate network.)

```
echo 1 | sudo tee /proc/sys/net/ipv4/ip_forward
```

Raspberry Pi WiFi users: Enable promiscuous mode on wlan0

If you are using the wireless interface of a Raspberry PI like I am you also need to enable “promiscuous mode” on the wireless interface:

```
sudo /sbin/ip link set wlan0 promisc on
```

This step is not needed if you are using the wired interface of the Raspberry Pi.

After these steps are done your DOS machine will behave as though it is just another Ethernet connected machine on the network. ARP packet handling and packet forwarding on the Linux machine hide the fact that it is on a separate point-to-point network.

Advanced options

What has been presented above is the bare minimum to get a DOS machine running mTCP on a local home network. Your home network might not look like this, requiring more work. Some other options include:

- Running IP masquerading (NAT) on the Linux machine.
- Running the PPP or SLIP network on a different subnet and having Linux routing to that subnet. This requires advanced routing on other machines in your network if you are using non-routable networks like 192.168.x.x.

These are advanced topics and are not addressed here.

DOS Idle/Power control

Introduction

Part of the mTCP design requires the application programmer to periodically check for new packets that have arrived so that they may be processed. That check is supposed to be done when the application is not actively working on something else; it is a core part the idle loop in an event driven program.

On old classic hardware running DOS such as the IBM PC or PCjr this polling behavior does no harm. But it does keep the CPU busy and that can show up when using mTCP applications in a virtual environment such as VirtualBox or VMWare. A side effect is higher power usage which is an important consideration for laptop users.

How mTCP minimizes CPU usage

Starting with the 2012-04-29 version mTCP applications will try to minimize unnecessary CPU usage. When in the idle loop each application will check for packets and work to do and if nothing is found they will try to "sleep" for a short period. The mechanisms to do this are:

- Calling the DOS "idle" interrupt 0x28. Traditionally this allows TSR programs to come active while DOS is waiting for a keystroke but it is also a general purpose way for DOS to signal that the machine is generally idle.
- Calling int 0x2F function 0x1680. This is the "VM yield time slice" signal.

The first method (int 0x28) is controversial; many people will claim that it is not needed. It has existed since DOS 2.x and it is documented as being used by DOS to signal to TSRs that there is idle time. After doing a lot of research and testing there is no reason why an application should not use it too, especially if the application does not use DOS for keyboard input (thus depriving DOS the ability to call int 0x28 itself.)

If power saving is enabled in mTCP (which is the default) the calls to int 0x28 will always be made during idle loops.

The second method (int 0x2f 0x1680) is well documented and should work in other environments such as DOS boxes in Windows, OS/2, and other virtual environments. I've tested it in Windows XP when using SwsVpkt.

When mTCP first starts up it will use int 0x2f function 0x1680 to see if the function is supported. If it is supported mTCP will use it during idle loops.

Notes on specific power saving programs

POWER

The standard power management utility for MS-DOS and PC-DOS is POWER.EXE which is installed as a device driver and controlled from the command line.

While POWER seems to work to reduce CPU consumption at the DOS command line, it does not help with the

mTCP idle loop. So it does not harm anything, but it doesn't help either.

POWER works correctly on real hardware and in VMWare. But the current version of VirtualBox seems to have a bug that causes the FTP client to hang up and freeze when POWER is loaded. If you see weird behavior like this under VirtualBox get rid of POWER and use a different utility.

FDAPM

FDAPM is an open source replacement for POWER that ships with FreeDOS and can be downloaded separately.

The FDAPM power control utility detects the calls to int 0x28 to see when DOS is idle. This has a very noticeable effect when running mTCP applications with FDAPM loaded; instead of making a processor fully busy the processor will idle at a low CPU utilization. Here is the command line I use to enable FDAPM:

```
fdapm apmdos
```

If you run that FDAPM will detect calls to int 0x28 and try to adjust your hardware to save power. For more information on FDAPM go to: http://help.fdos.org/en/hhstndrd/base/fdapm.htm

DOSIDLE

DOSIDLE is an older power management utility that can replace POWER. I've not experimented a lot with it, but it works so well that it forced mTCP to save power even when you have told mTCP not to try to save power!

DOSIDLE can be found at http://maribu.home.xs4all.nl/zeurkous/download/mirror/dosidle.html.

Disabling power control

In the unlikely event that you need to disable the power control calls in the mTCP idle loop you can set an environment variable. The name of the variable is "MTCPSLEEP" and an example of how to set it to disable power control is shown:

```
set MTCPSLEEP=0
```

If you do this mTCP will not attempt to call int 0x28 or int 0x2f function 0x1680 while it is in an idle loop.

Connectivity problems with specific machines

This section may help you if you can ping or connect to some machines but not others.

At a minimum you should always be able to ping your network gateway. If ping doesn't work for you then look at the previous debugging topics.

Firewalls

You may find that you are not able to connect to particular machines. The cause is almost always the firewall configuration on the target machine. Modern operating systems have very aggressive firewall rules that block incoming traffic, even if you are on the same physical network.

To see if a firewall is the source of the problem try to disable the firewall temporarily. If everything magically starts to work then you know you need to add some firewall rules to allow your mTCP machine to connect. Disabling the firewall is also possible but that might leave you exposed to other threats. (And never disable the firewall on a machine that connects to public networks.)

Ethernet card acceleration options

Modern Ethernet cards can offload the IP and TCP checksum calculations from the main processor, improving performance. While this is generally a good thing it can cause problems when mTCP is hosted on a computer in a virtual machine or emulator; depending on the implementation the packets from the host machine may be delivered to mTCP without checksums being filled in which would look like bad packets to mTCP. (mTCP always fills in checksum fields, so outgoing packets will be fine but incoming packets will fail.) Connectivity to other machines will be fine but connectivity to the host machine will seem broken.

This happened to me on Windows 7 using a DOSBox build with NE2000 support. The work-around was to turn off the IP and TCP checksum offload functions on the Ethernet card driver, which reduced performance slightly but enabled mTCP to be able to see correct packets.

Application specific problems

Applications often implement additional rules that control what Ethernet interfaces, ports or protocols they will use. If you can ping a machine but not connect to a specific application then check the firewall rules for the machine and any application specific configuration.

Network Address Translation (NAT)

NAT originally was designed to work around limited IPv4 address space. The nature of the work-around also improves security by masking machines that are behind the router implementing NAT - outgoing connections are allowed but incoming connections are blocked. This works great for client machines which usually make only outgoing connections but not so well on applications where incoming connections need to be supported. This kind of problem can usually be worked around by implementing routing rules on the NAT device.

On a home network NAT is generally not a problem because all of the machines on the home network are behind the firewall and should have unrestricted access to each other. Virtual machines often allow you to set

up NAT for the virtual machine, which can break things. When creating virtual machines use a bridged connection unless you absolutely need NAT for some reason.

NAT can also cause problems with protocols like FTP because it has to inspect packets and possible rewrite them. For example, FTP using port mode (not “passive” mode) sends IP addresses and port numbers to the server, which must be rewritten because the external server can not see your private network. Devices that implement NAT should be able to do this but they are expecting to use well-known ports. Using FTP with non-standard ports is effectively broken as a result. This kind of problem can not be fixed with additional routing rules.

NetDrive Server

NetDrive Server

Introduction

The mTCP NetDrive server is a program written in Go that provides remote disk services for the NetDrive DOS device driver. The server has the following features:

- A text user interface console window so you can monitor the server and interact with it.
- The ability to handle many concurrently connected clients.
- Supports journaling of disk images allowing clients the ability to "undo" recent writes.
- Allows the sharing of a single disk image (read-only mode). A single disk image may also be shared in a simulated "read-write" mode where the writes from a particular client are only visible to that client, and only while that client is connected.
- The ability to create new floppy disk and hard disk images.

The server is cross-platform and runs on Windows 10, Windows 11, and Linux (64 bit x86 or Arm). The server requires no special permissions to run.

Running the server

The server has three major commands:

- create: create disk images (floppy or hard drive)
- image: perform maintenance on disk images
- serve: serve disk images to DOS clients

The server side program is named netdrive.exe (Windows) or netdrive (Linux). NetDrive takes a command, optional flags, and possibly positional parameters depending on the command being used.

The command line format is:

```
netdrive [<common_flags>] command [<command_flags>] <command_args>
```

Flags (common or command specific) are optional. There are two common flags that apply to all commands:

-log_file <filename>	Append log messages to the specified file (default: no logging)
-log_level <level>	Set the logging level to DEBUG, INFO, WARN or ERROR (default: INFO)

Run "netdrive help" to see the list of commands and common flags. Each command also has a help feature.

The create command is discussed in Creating disk images. The image command is discussed in Advanced disk images: Session scoped images and Journaled images.

Flags for serving images

The serve command starts the code that serves read and write requests from DOS clients. To start the server code use the following command:

```
netdrive <common_flags> serve <server_flags>
```

The optional server flags are:

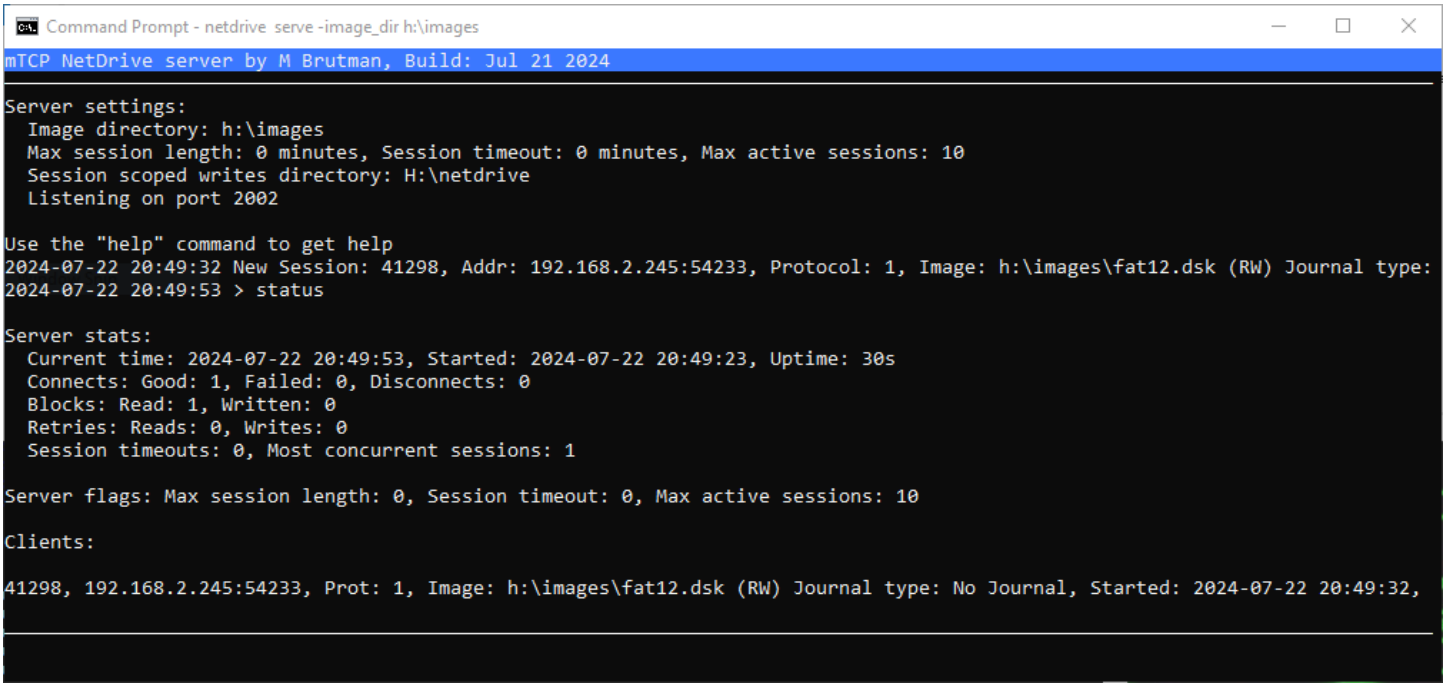
-port <n>	Serve using UDP port n (default is port 2002)
-bind_addr <ip_addr>	Only listen for connections on ip_addr (default: all addresses)
-image_dir <dir>	Use <dir> when looking for disk images (default: current dir)
-headless	Run in “headless” (non-interactive) mode
-timeout <n>	Set the session timeout in minutes (default: unlimited)
-max_time <n>	Set the max session duration in minutes (default: unlimited)
-max_active_sessions	Set the maximum number of concurrent sessions (default: 20)
-session_scoped_writes_dir <d>	Set the directory to use for Session scoped disk image journal files

Example command line:

```
netdrive -log_file serving.log serve -port 8086 -image_dir my_images
```

will start the server using UDP port 8086, log messages to serving.log, and look for images in the my_images subdirectory.

Run the program at the Windows command prompt (cmd.exe) or a terminal window in Linux. When the server starts it will draw a screen that consists of a messages area and a command input line at the bottom.



Messages are displayed as machines connect, disconnect, or other notable things happen. Output from commands also appears in the messages area. The messages area holds up to 10,000 lines in a backscroll buffer.

If the window is not wide enough some messages may look truncated; just make the window wider to see them.

There are four commands available:

help	Show help text
kill <session>	Force session number <session> to end
set	Set options (set run by itself displays help text)
status	Show the status of the server
quit	End the server

The following keys are also recognized:

Esc	Clear the command input area
Ctrl-L	Refresh the screen
Ctrl-C	Quit (Use it twice to confirm it.)
PgUp and PgDown	Scroll the messages are up or down

In headless mode all you will see is a message telling you that headless mode was specified. Ctrl-C is used to terminate the server in headless mode.

The server must be running before you try to connect a DOS machine to a virtual hard drive.

The -timeout flag can be used to protect the server against clients that crash before they can terminate their session cleanly. The server checks for inactive sessions every minute. If a client is inactive for longer than the session timeout period the server assumes that the client has crashed or gone away and the session is terminated. If that should happen and the client is still viable then the next operation that the client sends will result in an error because the session was terminated. If this should happen the client can use the DOS NETDRIVE.EXE command to disconnect and reconnect again.

The -max_time flag is different in that it sets the maximum amount of time that a client may be connected. This is useful when running a public server as it helps to prevent a user from monopolizing all of the available connections to the server.

The -max_active_sessions flag sets the limit for the number of clients that can be connected at the same time. This is also useful on a public server, as it helps the server protect itself from excessive loads. While the default limit is 20 clients, serving client requests is not that expensive so it is possible to serve hundreds of clients using a single server.

The -session_scopes_write_dir flag is used to set the directory where the server will create the per-session private journal files for images using that feature. See “mTCP NetDrive disk images and journals” for information on Session scoped disk images.

Serving standard disk images

If a disk image is read-write then only one DOS client may open and use the disk image at a time. This is because the file systems used by DOS (FAT) and DOS assume they have total control of the emulated disk and there is no way to coordinate writes from other possible clients. (Allowing multiple clients write access to a disk image at the same time will probably corrupt the file allocation table.)

When opening an image the server will check to see if the image is read-only. If it is then the DOS client will be able to read blocks from the image but writes will be blocked by the server's operating system. This is useful for protecting imaged floppy disks or for doing a forensic examination on a hard drive image.

Another advantage of read-only disk image files is that several DOS computers can mount the same read-only disk image file at the same time. Normally a disk image file can only be used by one computer at a time because writes to the disk image file can't be properly coordinated. Read-only disk image files do not have that problem and thus the server will allow several DOS machines to mount the same read-only disk image file at the same time.

Serving Journaled disk images

Note: See "Advanced disk images: Session scoped images and Journaled images" for more information on Journaled disk images and how they are different from standard image files.

Journaled disk images are effectively the same as read-write disk image files, so the same limitation of only one connected client at a time applies to them too. It does not matter if the base disk image is marked read-only or read-write, as writes will go to the journal file and not the base disk image. The journal file must be read-write.

Serving Session scoped disk images

Note: See "Advanced disk images: Session scoped images and Journaled images" for more information on Session scoped disk images and how they are different from standard image files.

Session scoped disk images are a special case of journaled images where the journal file is private to a specific client session and temporary, lasting only as long as the client session. As a result, Session scoped disk images are effectively read-only even though the client has the illusion of being able to do writes during their session. (The writes are lost at the end of the session.) This means that multiple clients can concurrently connect to Session scoped disk images, just like with read-only disk images.

The temporary session journal files will be created in the image directory by default, or in the directory specified by the -session_scoped_writes_dir flag if that flag was used. (These files are allowed to be placed separately from image files because of their transient nature.)

Session scoped disk images are useful for running software on images and allowing that software to make writes if it needs to, but without having to worry about backing up and restoring the original image after each connection. A good example of this would be if you were sharing a games archive you would want people to be able to play the games and modify high score files; on a read-only image they would get an error if a write was attempted but with session_scoped_writes the high score file can be created or updated while the shared image remains untouched.

Disk images

A disk image is a file that looks like a DOS disk; instead of blocks on a mass storage device there are blocks in the file. The disk image can be any format that is valid for DOS as long as 512 byte blocks are used. (DOS can use different block sizes but 512 byte blocks are the standard for disk storage.) Under normal operation reads and writes go directly into the disk image file, simulating what would happen on a real mass storage device. Like with a real mass storage device, writes persist across connections.

mTCP NetDrive uses standard disk images, which are plain files where each 512 byte block of the file can be treated as a sector or block of a mass storage device. There is no metadata in the file and no compression is used. These files are the same format used by Linux 'dd' and can be mounted by Linux systems using the loopback mount option. These files can be converted to and from other formats, such as those used by VirtualBox, QEMU or VMWare.

Disk image formats (FAT variations)

The mTCP NetDrive server basically just does read and write operations into disk image files on behalf of a connected DOS client. While the server really doesn't care about the contents of the file, DOS absolutely does care and DOS requires that the image file contents look like a valid FAT (File Allocation Table) based filesystem.

There are three versions of FAT supported by DOS:

- FAT12:
  - Required for DOS 2.x.
  - Supports hard drives up to 32 MB. (65535 sectors max)
  - Maximum number of files: Just under 4K
  - Max format: Just under 4K clusters at 16 sectors per cluster
  - Used by floppy disks and small hard drives
- FAT16 (also known as original FAT16):
  - Introduced with DOS 3.x.
  - Supports hard drives up to 32 MB, but more efficiently than FAT12. (Also 65535 sectors max)
  - Maximum number of files: Just under 64K
  - Max format: 64K clusters at 1 sector per cluster
- FAT16B (also known as final FAT16):
  - Introduced with Compaq MS-DOS 3.31, widely available in DOS 4.x and up
  - Supports hard drives up to 2GB
  - Max format: 64K clusters at 64 sectors per cluster

While FAT does not specify the sector or block size, it should be 512 bytes, which ensures maximum compatibility. (As of this writing mTCP NetDrive only supports 512 byte blocks.)

The number of sectors used for the File Allocation Table (FAT), the number of Root directory entries, the number of FAT copies, and other parameters can be varied. In general one should choose enough FAT sectors to be able to use the full size of the drive, 1 copy of the FAT, and a reasonable number of Root directory entries:

- FAT12: Each entry in the table is 12 bits. A full table requires 12 512 byte sectors.
- FAT16 and FAT16B: Each entry in the table requires 16 bits. A full table requires 256 512 byte sectors.
- Root directory entries: 32 bytes each, 16 per 512 byte sector.

The size of the table determines the maximum number of files across the entire hard disk. You can use less than a full table if you know that you do not need the maximum number of files.

DOS looks at the number of clusters to determine if the FAT is using 12 bit entries or 16 bit entries. If you use less than 4085 clusters DOS uses 12 bit FAT entries. If you use 4085 or more clusters then DOS uses 16 bit FAT entries. The decision between FAT16 and FAT16B is made by looking at the total number of sectors in the BPB; if it is 65535 or less then FAT16 is in use. If the total number of sectors is 0 then a 32 bit sector count

from a field in the extended BPB is used and FAT16B is being used.

If you create a floppy image with NetDrive you will get one of the standard floppy images. If you create a hard disk image with NetDrive you can choose the FAT type, the number of FAT copies and the number of root directory entries. The size of the FAT will always be the maximum allowed and the cluster size will be the minimum necessary for the given disk size.

Creating disk images

You can use any tool that creates a standard disk image to create a disk image usable with mTCP NetDrive. (In Linux I've used dd and mkfs.vfat.)

The mTCP NetDrive server also has a command line function that creates new floppy or hard disk images. Floppy disk images are restricted to 8 well known formats, while hard drive images have more options.

To create a floppy image use the following command:

```
netdrive create floppy [create flags] <size_in_KB> <image_name>
```

where <size_in_KB> is a diskette capacity size in KB and <image_name> is the name of the image file you want to create. The following sizes are supported:

160 or 320	DOS 1.x 160 and 320K images (40 tracks, 8 sectors per track, 1 or two sides)
180 or 360	DOS 2.x and up 180 and 360K images (40 tracks, 9 sectors per track, 1 or two sides)
1200	DOS 3.0 and up 1200K images (80 tracks, 15 sectors per track, two sides)
720	DOS 3.2 and up 720K images (80 tracks, 8 sectors per track, 2 sides)
1440	DOS 3.3 and up 1440K images (80 tracks, 18 sectors per track, 2 sides)
2880	DOS 5.0 and up 2880K images (80 tracks 36 sectors per track, 2 sides)

The following optional flag can be used when creating floppy images:

-image_dir <dir>	Use <dir> when looking for disk images (default: current dir)

Example:

```
netdrive create floppy -image_dir floppy_images 1440 empty.dsk
```

will create a standard, 1.44 MB floppy image in the floppy_images directory.

To create a hard drive image use the netdrive create hd command:

```
netdrive create hd [create flags] <size_in_MB> <fat_type> <image_name>
```

Size may range from 1 to 2047 MB and fat type is either FAT12, FAT16, or FAT16B. The cluster size will be automatically selected based on the image size.

(Hint: to ensure the most efficient use of space choose hard drive sizes that are just under a power of 2. For example, 255 MB, 511 MB, 1023 MB, etc. The cluster size has to double near the power of 2 boundaries, so being just under a power of 2 makes the best use of the space and cluster size.)

The following optional flags can be used when creating hard drive images:

-image_dir <dir>	Use <dir> when looking for disk images (default: current dir)
-num_fats <n>	Create <n> copie of the File Allocation Table. (n must be 1 or 2)
-root_dir_entries <n>	Create a root directory with <n> entries

The default number of root directory entries is 512. By default two FAT copies are made although only one is actually ever used. (Blame Microsoft.)

Example:

```
netdrive create hd -num_fats 1 511 FAT16B bigdisk.dsk
```

will create a 511 MB FAT16B drive image with just one FAT copy.

Manipulating disk images with Linux

Linux makes it fairly easy to create and work with disk images:

- The “dd” command can be used to create an empty file for a new disk image. Use /dev/zero as the source.
- The mkfs.fat command can write a filesystem on the new disk image.
- The mount command can mount images in read-only mode for inspection or read-write mode for manipulation. (Use the -o loop,check=normal -t msdos options. You might also need to use the uid option to be able to make changes.)
- mtools can be used to manipulate image files.

Remember that two systems can’t write to the same image at the same time. This includes a remotely connected DOS system and a Linux system mounting the image using the loopback option.

Caution: Linux tools are not capable of manipulating mTCP NetDrive Journaled disk images because they are not aware of the journal file that mTCP NetDrive uses. You can use Linux tools to create new image files but if you turn an image file into a Journaled disk image you should not use Linux tools to manipulate it.

Manipulating disk images with Windows

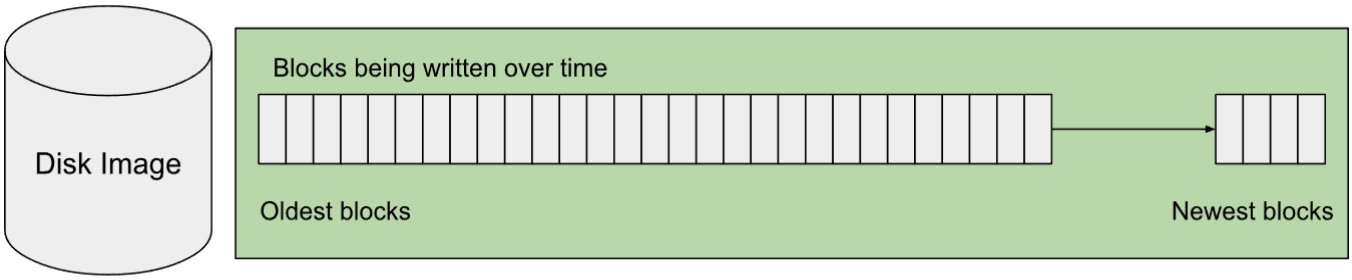
(I’m not as familiar with manipulating disk images with Windows, so if you have a tip please send it to me.)

- WinImage seems to be able to open and manipulate floppy and hard drive disk images. (I tried it a little bit during a trial period; it was awkward but it might be usable.)
- qemu-img can convert hard drive image formats such as VMDK, VDI, VHD, and raw (what NetDrive uses.)
- Virtual machines such as VirtualBox can mount floppy images in their emulated floppy drives.

Caution: Windows tools are not capable of manipulating mTCP NetDrive Journaled disk images because they are not aware of the journal file that mTCP NetDrive uses. You can use Windows tools to create new image files but if you turn an image file into a Journaled disk image you should not use Windows tools to manipulate it.

Advanced disk images: Session scoped images and Journaled images

In computing, a journal is a data structure or a file where pending transactions are recorded. When a journal file is used with a disk image, newly written data goes to the journal file and the disk image is left unmodified. When a block is read the journal file is checked first to get the newest version of the block. If the block is not in the journal file (meaning the block has not been written to) then the block is fetched from the disk image.



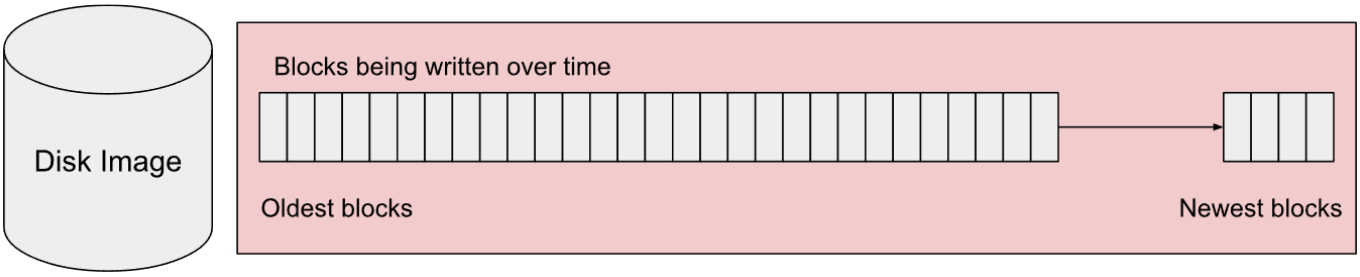
The journal file (green) takes all writes instead of the disk image.

Using a journal file for writes allows one to quickly undo those writes by just simply deleting the journal file. You can also "undo" a subset of writes (the newest writes) by truncating the journal file, effectively wiping out writes past the truncation point.

This property of journal files is used to implement two additional disk image types in mTCP NetDrive, each of which has an advanced feature.

Session scoped disk images

When a DOS client connects to the server it starts a session, which is used to track the state of the client and the disk image while the client is connected. When a Session scoped disk image is used a private journal file for the session is added to a disk image, allowing the client to have the illusion of disk writes that are only visible during that session. When the session ends the journal file is deleted, making it look as though those writes never happened.



Deleting the per-session journal (denoted in red) throws away any writes the connected client made.

Since all writes go to the private journal file the disk image itself is never altered, effectively making the disk image read-only. As the disk image is effectively read-only, multiple clients can connect to and use the same disk image when it is a Session scoped disk image.

This can be used to add a "demo mode" disk to images, allowing people to interact with them (including writes) without worrying about the possibility of permanent changes to the disk image. A good example of this is a public server with a DOS shareware game collection on it; a Session scoped disk image lets you share one copy

of the disk while giving everybody the illusion of being able to write to it. (At least temporarily during a session.)

To mark a disk image as a session scoped disk image create a new file in the same directory as the image with the same name, but add ".session_scoped" to it. For example, if your disk image is called "dos_utilities.dsk" then create "dos_utilities.dsk.session_scoped" in the same directory to turn on session scoped writes for that disk image. The file can be empty - it just has to exist.

To disable the session scoping feature on a disk image just delete that file.

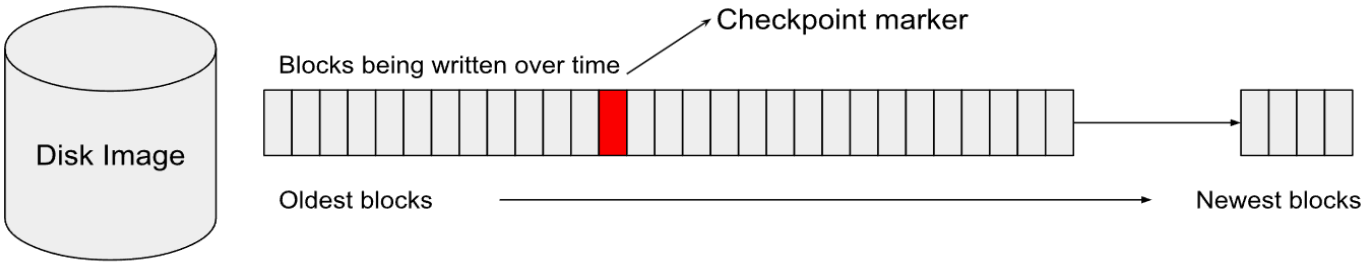
When a client connects to a Session scoped disk image new writes go to a journal that is private to the session. The filename for the journal has the session number as the name and '.log' as the extension, and it is only created if the client attempts to write data. At the end of the session the journal file will be deleted.

By default the journal files for Session scoped disk images will be in the same directory as the disk images. Use the -session_scoped_writes_dir flag when starting the server to put the Session scoped disk image journal files in a different directory.

No maintenance operations are needed for Session scoped disk images.

Journaled disk images

A Journaled disk image is a disk image that has a single, persistent journal file that survives across sessions. This allows a DOS client to undo writes that are in the journal file, effectively "going back in time" on the disk image, including writes that were made during earlier sessions. As the journal file for the image persists across sessions there can only be one journal file, so this resembles a normal read-write image but with the special undo feature. Disk images with this support are known as Journaled disk images.



Deleting new blocks after the checkpoint marker block effectively "goes back in time" to when the checkpoint marker was written.

(Note: While Session scoped disk images use a journal file, the implementation is simpler than the full journalling implementation used by Journaled disk images. As a result GoBack/Undo is not available when using a Session scoped disk image.)

Enabling journaling on a disk image

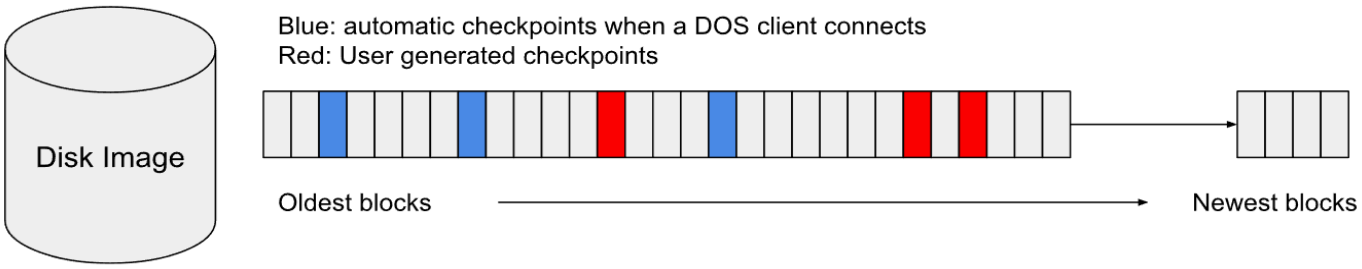
To enable journaling on a disk image create a new file in the same directory as the image with the same name but add ".journal" to it. For example, if your disk image is called "dos_utilities.dsk" then create "dos_utilities.dsk.journal" in the same directory to turn on journaling for that disk image. The file must start as an empty file with 0 length.

If the journal file exists, writes to the image will be appended to the journal file. Do not delete this file - if you do you will lose those writes.

In addition to the ".journal" file a new file with the same name as the disk image and ".journal.map" appended to the end will also be created in the directory. This file (the "map" file) is used to speed up the process of re-connecting to an image. If this file gets deleted by accident it will be recreated and no data will be lost.

Showing checkpoint markers in a journal

Checkpoint markers are metadata written to a journal that are used to mark points in time that are available for GoBack function. Each time a client connects to a journaled disk image a generic "session start" checkpoint marker is written to the journal. The user can also add their own checkpoint markers at any time.



To see the checkpoint markers in a journal use the netdrive image list_cp command.

While the DOS client also has a list_cp command, it can only show a limited number of checkpoint markers because it has to transfer them over the network and buffer space is limited. The version of list_cp in the netdrive server does not have that limitation.

Undo writes (GoBack)

(This function is the same as the goto_cp function in the DOS client.)

A journaled drive supports to GoBack feature, which allows you to "go back" in time with respect to writes, making the drive look like an earlier version of itself. This is effectively an "undo" feature where you select how far back you want to go.

Use the "netdrive image goto_cp" command to delete all of the writes past a specified checkpoint marker. (Use the "netdrive list_cp" command first to see the available checkpoint markers.) This operation can not be un-done - if you use it those writes are lost and can not be recovered. The journal file will be truncated at the point you specify, throwing those writes away.

Compacting journal files

All writes to a journaled disk image get appended to the journal file, causing the journal file to grow in size. Only the newest write to a specific block is valid, as the previous writes represent stale data that was valid at an earlier point in time. While the GoBack feature makes use of this behavior it can result in a significant amount of dead space inside of the journal file. This is especially true when the same set of blocks is repeatedly written.

The compaction process scans a journal file and copies the active journal records to a new journal file. Checkpoint markers are preserved and there should be no difference between a journal file before and after

compaction except for the size, which will be smaller.

To run compaction on a journal file use the “netdrive image compact_journal” command.

Committing journaled writes

As discussed above, all writes to a journaled disk image get appended to the journal file, causing the journal file to grow in size. A large journal file should not be a problem for the system but eventually you will want to make the journalled writes permanent in the disk image file.

To do this use the “netdrive image commit_writes” command. The command will let you commit all writes in a journal file, or only writes up to a specified checkpoint marker. After the command is run the journal will either be empty (all writes committed) or will have only writes made after the specified checkpoint marker.

Disabling journaling on a disk image

The journal file on a disk image contains your writes to the disk image, so do not just delete it - if you do you will lose those writes.

To disable journaling while preserving your changes you need to commit those changes to the disk image file:

- Make sure the disk image is not in use by a connected client.
- Use the “netdrive image commit_writes” command to move the writes in the journal file into the disk image. This makes the writes permanent in the disk image.
- After the commit_writes operation the .journal file should be empty (0 length). It is now safe to delete the .journal and .journal.map files.

To disable journaling while throwing away any changes in the journal file just delete the .journal and .journal.map files.

After the .journal and .journal.map files are done the image will be treated like a standard disk image and any writes will be made inside the disk image file.

Other operations on disk images

Copying and renaming disk images

Disk images are just normal files so you can copy them or rename them as needed.

If a disk is journaled be sure to copy or rename the .journal file too, otherwise you will lose the writes to the image. You should also copy the .journal.map file but if you forget it will be recreated.

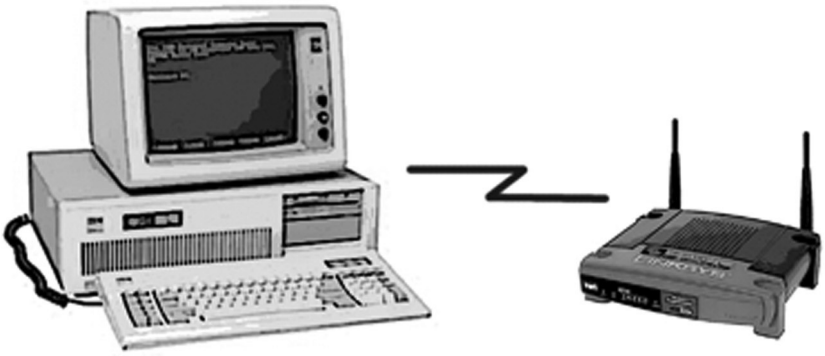
Comparing disk images

The “netdrive image compare” command can be used to compare two images for equivalence. Journaled and non-journaled images are supported, and the journals can be of different sizes.

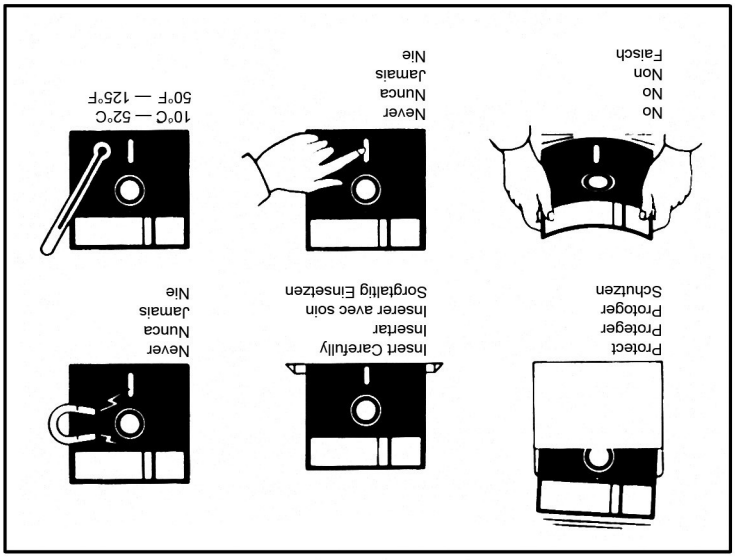
This command is basically only useful as a safety check after committing writes to a disk image. If you make a backup copy of a disk image and (and its journal) before committing writes to the disk image you can compare the backup to the new version of the disk image and they should be equivalent.

Bonus Materials!

Make your own floppy disk jacket!



TCP



For extended media life –
here's how to take care of your flexible disk

Old style customer comment card!

(Print these next two pages on a double sided printer, then scan and email it to me.)

The DOS Networking
Software Library

Product Comment Form

mTCP 8675309

Your comments assist us in improving our products. Brutman Heavy Industries may use and distribute any of the information you supply in anyway it believes appropriate without incurring any obligation whatever. You may, of course, continue to use the information you supply.

Comments:

If you wish a reply, provide your name and address in this space.

Name _____

Address _____

City _____ State _____

Zip Code _____



NO POSTAGE
CAN HELP GET
THIS MAILED
IN THE
UNITED STATES



BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO. √-1 KLATU BARADA NIKTO

POSTAGE WILL BE PAID BY ADDRESSEE

BRUTMAN HEAVY INDUSTRIES
SALES & SERVICE
P.O. BOX 0xDEADBEEF
THE BIT BUCKET, 10101

Fold here

Tape Please do not staple Tape