

SuperBASIC extensions

ini_cmd	SuperBASIC Procedure definition list
eth_init	SuperBASIC Procedure ETH_INIT to (Re)Initialize the CP2200 Ethernet controller Assuming that the Ethernet device driver was installed, an attempt is made to initialize the CP2200 again. Accepts an optional parameter for the duplex type. 0 (default) = Auto-negotiation, 1 = Full duplex, 2 = Half duplex
feth_init	SuperBASIC Function FETH_INIT of the above Procedure. Accepts an optional parameter for the duplex type. 0 (default) = Auto-negotiation, 1 = Full duplex, 2 = Half duplex Returns, 0 for OK 'Not Complete' for initialization timed out. 'Transmission error' for Auto-negotiation failed. 'Not Found' for uninitialized driver. 'Bad Parameter'
eth_mac\$	SuperBASIC Function ETH_MAC\$ to return the MAC address of the CP2200 as a dash seperated string. 'Not Found' for uninitialized driver. 'Bad Parameter'
arp_add	SuperBASIC Procedure ARP_ADD to add, or update an arp table entry. MACaddress in the format "aa-bb-cc-dd-ee-ff"
arp_remove	SuperBASIC Procedure ARP_REMOVE to removes one or all ARP table entries from the linked list. Without parameters, removes all entries
arp_list	SuperBASIC Procedure ART_LIST to lists all ARP table entries to #ch or default #1. In the form, IP Address MAC Address
eth_setip	SuperBASIC Procedure ETH_SETIP to set the Q68's IP address in the device definition block
eth_subnet	SuperBASIC Procedure ETH_SUBNET to set the Q68's subnet mask in the device definition block
eth_gateway	SuperBASIC Procedure ETH_GATEWAY to set the Q68's default gateway IP address in the device definition block
setip	Routine used by ETH_SETIP, ETH_SUBNET, and ETH_GATEWAY to do the setting of the IP addresses
eth_fgetip	SuperBASIC Function ETH_GETIP\$ to return the set IP address as a dot separated string
eth_fsubnet	SuperBASIC Function ETH_SUBNET\$ to return the set IP address as a dot separated string

eth_fgateway	SuperBASIC Function ETH_GATEWAY\$ to return the set IP address as a dot separated string
getip	Routine used by ETH_GETIP\$, ETH_SUBNET\$, and ETH_GATEWAY\$ to return the actual strings
eth_netname	SuperBASIC Procedure ETH_NETNAME to set the Q68's network name
eth_fnetname	SuperBASIC Function ETH_NETNAME\$ to return the Q68's network name as a string
eth_errno	SuperBASIC Function ETH_ERRNO to return the last driver specific error. This is not the same as a QDOS error, But an error code to indicate the last problem the driver encountered. ETH_ERRNO will clear the error code after it is read.
eth_ping	SuperBASIC Procedure ETH_PING. Sends 4 Pings to the supplied IP Address, Sending the results to the specified channel, or #1.
get1int	Fetch one Procedure/Function parameter integer and place it on the maths stack
get1str	Fetch one Procedure/Function parameter string and place it on the maths stack

Subroutine list

cp2200_init Initialize the CP2200 Ethernet controller.
 Entry
 D5 Required duplex mode. 0 = Auto-negotiation, 1 = Full duplex, 2 = Half duplex
 A3 Assumed start of driver definition
 Exit
 D0 error return. Possible errors -
 Not complete = Self Initialization timed out
 Transmission error = Auto-negotiation failed

cp2200_mac Initializing the MAC of the CP2200
 Entry
 A1 points at the base address of CP2200 registers
 A3 points at the definition of the device driver

cp2200_phlay Initialize the Physical Layer of the CP2200
 Entry
 A1 points at the base address of the CP2200
 Exit
 D0 error return. Possible errors -
 Transmission error, Self Initialization timed out

cp2200_WritePacket
 Write a packet to the Ethernet controller
 Interrupts are disabled during this routine to prevent it being recalled in the middle
 of writing a packet
 Entry
 D2.W number of bytes to send
 A1 base of buffer
 A2 base of CP2200 Ethernet controller
 A3 base of device driver definition block

 Exit
 D0 0 or QDOS error code
 buffer full transmit buffer is not empty, after waiting 1.5 sec
 transmission error if the last packet was not transmitted successfully
 D2 preserved
 A1 updated pointer to buffer

add_arp_rec Adds, or updates a record in the ARP table
 Entry
 D5.L top four bytes of the MAC address
 D6.W bottom two bytes of the MAC address
 D7.L IP address
 A3 base of device driver definition block
 Exit
 D0 0, or 'out of memory'

allocateport Allocate a free port from the managed table of ports. OPEN_IN and Binding requires a system selected port. This port will be selected from a pool of 256 ports between \$D200 to \$D300
The allocation is managed from a 32 byte port allocation map, where each bit identifies a port as being free (0) or in use (1)
There is a rotating port number pointer that is incremented each time a port is allocated. So if a port is used, then released, it will not be used again immediately.
Entry
D4 upper word is port supplied to the OPEN routine
A3 base of driver definition block
Exit
D0 error return. Possible errors -
Buffer full, no ports available
D4 lower word, selected port

check_IP_address
Checks the supplied IP address, If it's not on the local LAN, and a Default gateway has been set, then use the Default gateway IP address
Entry
D0 IP address of required computer
A3 base of driver definition block
Exit
D0 preserved, or the Default gateway IP address

arp_ip_request
Request a MAC address from a remote computer with the supplied IP address
Entry
D7 IP address of required computer
A1 base of buffer
A3 base of driver definition block
Exit
D0 0 or QDOS error code
buffer full transmit buffer is not empty, after waiting 1.5 seconds
transmission error if the last packet was not transmitted successfully
A1 updated pointer to buffer

check_mac_address
Scan the ARP table to see if we know the MAC address for the IP address in D0
Entry
D0 IP address of required computer
A3 base of driver definition block
Exit
D0 preserved
zero flag not set if successful, and
D5.L upper part of MAC address
D6.W lower word of MAC address

check_open_valid	<p>Checks to see if the required OPEN command can proceed</p> <p>Table for determining which OPEN commands are valid. The table takes the form of four bytes for disconnected from network and four bytes for connected to network. Each four bytes are for OPEN, OPEN_IN, OPEN_NEW, spare (OPEN_OVER)</p> <p>Return values are 0, invalid parameter, transmission error or format failed for an undefined open type .</p> <p>Entry</p> <p>D7 lower word, is open type</p> <p>A4 supplied IP address</p> <p>Exit</p> <p>D0 0, or an error code</p> <p>Table for determining which OPEN commands are valid. The table takes the form of four bytes for disconnected from network and four bytes for connected to network. Each four bytes are for OPEN, OPEN_IN, OPEN_NEW, spare (OPEN_OVER)</p> <p>Return values are 0, invalid parameter, transmission error or format failed for an undefined open type</p>
checksetport	<p>If supplied port number is in the range \$D200 to \$D300, and if already allocated. If so returns an 'in use' error. Otherwise flags port as in use</p> <p>Entry</p> <p>D3.W port to allocate</p> <p>A3 base of driver definition block</p> <p>Exit</p> <p>D0 0, or 'In Use' error</p>
cp_aneg	<p>Do auto-negotiation of the CP2200 +++++ unfinished ++++</p> <p>Entry</p> <p>A1 points at the base address of CP2200 registers</p> <p>Exit</p> <p>D0 error return. Possible errors -</p> <p>Transmission error, Self Initialization timed out</p>
ddlink	<p>Get the assumed start of Q68Net Driver Definition Block in A2, or set the zero flag for not found.</p>
deallocateport	<p>De-allocate a port from the managed table of ports, If in the a pool of 256 ports between \$D200 to \$D300</p> <p>Entry</p> <p>D0.W port to de-allocate</p> <p>A3 base of driver definition block</p> <p>Exit</p> <p>none</p>

fetchpacket	<p>Look to see if there is a packet waiting in the channels queue to be accessed. If so, check to see if it has yet to read, or has already been read. Linking in the next if needed. Otherwise return 'not complete'. Returns a 'transmission error' on a wrong MAC address.</p> <p>Must be in supervisor mode, and A6 pointing to the system variables.</p> <p>Entry</p> <p>A0 start of channel definition</p> <p>A3 base of driver definition block</p> <p>A6 base of system variables</p> <p>Exit</p> <p>D0 0, or an error code</p>
get_cdb	<p>Convert a channel ID in A0 to a pointer to the base of the channel definition block.</p> <p>Entry</p> <p>A0 Channel ID</p> <p>Exit</p> <p>A0 Points to base of channel definition block</p> <p>D0 Zero, or Channel not open error</p>
get_lang	<p>Get the system language, and return as English, French, or German in D0</p> <p>Exit</p> <p>D0 001 for English(US) 044 for English(UK) 049 for German 033 for French 039 for Italian</p>
init	<p>Create a device driver definition block, Initialize the CP2200.</p> <p>If successful link the block into the system</p>
int_serve	<p>Interrupt handler for reading data packets in supervisor mode on entry</p> <p>Entry</p> <p>D3 number of 50/60Hz interrupts</p> <p>A3 base of driver definition block</p> <p>A6 base of system variables</p> <p>A7 supervisor stack (64 bytes free)</p> <p>Exit</p> <p>everything preserved</p>
is_assign	<p>Try to identify the type of packet received, and assign it to a channel, or throw it away</p> <p>Entry</p> <p>A0 points to the start of the buffer</p> <p>A3 base of driver definition block</p>
lang_search	<p>Search the language table. Returns a pointer to the start of the required language line in the language table. If language code is not found, it defaults to English</p> <p>Entry</p> <p>D0 language code</p> <p>Exit</p> <p>A4 points at start of language table entry</p>

localWritePacket

Try to write the packet to the local host directing it to the correct open IP channel

Entry

D2.W number of bytes to send

A0 points at the channel definition block

A1 base of buffer

A3 base of driver definition block

Exit

D0 0 or QDOS error code

transmission error if there was a memory problem

D1.W number of bytes sent

D2 preserved

A1 updated pointer to buffer

nd_close Device driver channel close routine

nd_io Device driver I/O routines

nd_open Device driver channel open routines for UDP, TCP and SCK channels
(TCP not fully implemented as yet)

nd_getmac Check to see if the ARP table has been updated with the required MAC address.
If ARP table has not been updated, another ARP request is sent at the half time point,
and at the timeout.

Note this is not a subroutine, may not return to caller

Entry

D0 operation

A3 base of driver definition block

Exit

D0 preserved, or an error code

nd_gen_trans_csum

Generate a Transport layer checksum for the block of data pointed to by A1,
Length D2.W

Also creates the transport layer header in the channel definition block ready
for sending

Entry

A0 base of channel definition block

A1 points at the start of the data block

A3 base of driver definition block

D2.W number of bytes in data block

Exit

D0.W The required checksum

Zero flag set on an error

ndo_getbyte

ndo_getword Read a device name parameter, Converts an ASCII string into a number in D7

Entry

D5 number of digits to read

A4 pointer to start of ASCII number

Exit

D7 ndo_byte - byte value

 ndo_word - unsigned word

zero flag not set on error

nstr2long Check the IP address of a null terminated string, returning it as a long word in D7

Uses the str2long routine

Entry

D0 length of string

A2 Pointer to end of string

Exit

D0 0, or QDOS error code

 Bad Parameter

D7 IP address in network order

send_mac_request

Send a request on the network for a MAC address for the IP address in D0

Entry

D0 IP address of required computer

A3 base of driver definition block

Exit

D0 0 or QDOS error code

 buffer full transmit buffer is not empty, after waiting 1.5 seconds

 transmission error if the last packet was not transmitted successfully

str2long Converts an IP address string on the Maths stack and return it as a long word in D7

Entry

A1 Pointer to Maths stack

Exit

D0 0, or QDOS error code

 Bad Parameter

D7 IP address in network order

str2mac Check the MAC address QDOS string on the Maths stack and return it as a long word in D5 & a word in D6

The string should be in the format "aa-bb-cc-dd-ee-ff"

Entry

A1 Pointer to Maths stack

Exit

D0 0, or Bad Parameter

D5.L Top four bytes of MAC address

D6.W Bottom two bytes of MAC address

tcp_accept	<p>Deal with an IP_ACCEPT. Accept a connection for a socket specified by the channel ID supplied in D3</p> <p>Entry</p> <p>D3 channel ID of LISTENing channel</p> <p>D6 upper word is open type</p> <p>A0 start of device name - must be a SCK_</p> <p>A3 base of driver definition block</p> <p>A5 base of driver definition block</p> <p>A6 base of system variables</p> <p>Exit</p> <p>D0 0, or an error code</p> <p>A0 base of channel definition block</p>
tcp_close	<p>Do a TCP close connection sequence</p> <p>Entry</p> <p>A0 base of channel definition block</p> <p>A3 base of driver definition block</p> <p>A6 base of system variables</p> <p>Exit none</p>
tcp_connect	<p>Attempt to make a Three Way Handshake connection to a TCP server that is Listening for connection requests.</p> <p>Returns 'Transmission error' if a connection cannot be made</p> <p>This routine may be called from either the OPEN routine, or IP_CONNECT</p> <p>Entry</p> <p>A0 base of channel definition block</p> <p>A3 base of driver definition block</p> <p>A6 base of system variables</p> <p>Exit</p> <p>D0 0, or an error code ????</p>
uh_alloc	<p>Allocate an area in a user heap</p> <p>Entry</p> <p>D1 required space on the user heap</p> <p>A0 pointer to pointer to pointer to free space in user heap</p> <p>Exit</p> <p>D1 length allocated</p> <p>A0 base of user heap area allocated</p> <p>D0 0, or 'Out of memory'</p>
uh_setup	<p>Assign and set up a new User Heap</p> <p>Entry</p> <p>D1 required space for the user heap</p> <p>A1 pointer to pointer to pointer to free space in user heap</p> <p>A2 pointer to address to store base of allocated area</p> <p>Exit</p> <p>A0 undefined</p> <p>(A1) pointer to pointer to free space in user heap</p> <p>(A2) base of common heap allocated</p> <p>D0 0, 'Out of memory' or 'job does not exist'</p>

uh_rechp	Release an allocated area in a user heap Entry A0 base of space to free A1 pointer to pointer to free space in user heap Exit A0 undefined A1 undefined
valIPv4hdr	Validate an IPV4 network header by its checksum Entry A0 base of packets buffer A3 base of driver definition block Exit D0 error return
valICMPHdr	Validate an ICMP transport header by its checksum Entry A0 base of packets buffer A3 base of driver definition block Exit D0 error return
valTCPHdr	Validate a TCP transport header by its checksum Entry A0 base of packets buffer A3 base of driver definition block Exit D0 error return
valUDPHdr	Validate an UDP transport header by its checksum Entry A0 base of packets buffer A3 base of driver definition block Exit D0 error return

writepacket Adds the required headers to the payload in the transmit buffer, and writes it to the CP2200 for transmission

Entry

A0 base of channel definition block

A3 base of driver definition block

Exit

D0 0 or QDOS error code

There are also some other entry points into writepacket in addition to A0 & A3 above

translayer_udp, translayer_tcp, translayer_icmp

On entry D2.W is the number of bytes in the payload

 A1 is a pointer to the start of the payload

 (D1-D4 may be smashed)

netlayer_ip D2 is the number of bytes in the payload
 A1

phylayer D2
 A1

Supported System Trap calls

Trap #2

D0	Name	Notes
\$01	IO_OPEN	D3=0-2
\$01	IP_ACCEPT	D3=LISTENing channel ID
\$02	IO_CLOSE	

Trap #3

D0	Name	Notes
\$00	IO_PEND	
\$01	IO_FBYTE	
\$02	IO_FLINE	
\$03	IO_FSTRG	
\$05	IO_SBYTE	
\$07	IO_SSTRG	D2 is word sized, So should limit data size to 32K
\$48	FS_LOAD	
\$49	FS_SAVE	
\$51	IP_SEND	data size limited to 64K
\$52	IP_SENDTO	data size limited to 64K
\$53	IP_RECV	
\$54	IP_RECVFM	
\$58	IP_BIND	
\$59	IP_CONNECT	
\$5B	IP_GETHOSTNAME	
\$5C	IP_GETSOCKNAME	
\$5D	IP_GETPEERNAME	
\$64	IP_GETSERVBYNAME	
\$65	IP_GETSERVBYPORT	
\$6E	IP_GETPROTOBYNAME	
\$6F	IP_GETPROTOBYNUMBER	
\$72	IP_INET_ATON	
\$73	IP_INET_ADDR	
\$74	IP_INET_NETWORK	
\$75	IP_INET_NTOA	
\$76	IP_INET_MAKEADDR	
\$77	IP_INET_LNAOF	
\$78	IP_INET_NETOF	
\$7C	IP_ERRNO	

Device Driver Definition Block

\$00	ndd_eilk	link to next external interrupt
\$04	ndd_eiro	address of external interrupt routine
\$08	ndd_5ilk	link to next 50/60Hz interrupt
\$0c	ndd_5iro	address of 50/60Hz interrupt routine
\$10	ndd_silk	link to next scheduler interrupt
\$14	ndd_siro	address of scheduler interrupt routine
\$18	ndd_ddlk	link next device
\$1c	ndd_iolk	link to I/O routine
\$20	ndd_oplk	link to open routine
\$24	ndd_cllk	link to close routine
\$28	ndd_pmptr	Port map pointer, increments after each allocation
\$2A	ndd_last_err	Last IP error, cleared after reading
\$2C	ndd_chlist	Link to list of open IP channels
\$30	iod_cnam	Pointer to routine to make the channel name (QPAC2)
\$34	ndd_ipid	Network (IPV4) layer identification. Increments for every packet sent
\$36	ndd_q68e	Q68 Ethernet identity string
\$3a	ndd_base	base address of CP2200 direct registers
\$3e	ndd_etir	Q68 Ethernet interrupt register
\$42	ndd_mac	6 bytes of the MAC address of the CP2200
\$48	ndd_ip	IP address of this computer
\$4C	ndd_subnetmask	IP subnet mask
\$50	ndd_gateway	default gateway IP address
\$54	ndd_netname	computers network name. word + up to 26 characters
\$70	ndd_arp	start of ARP table of MAC to IP addresses
\$74	ndd_queue_base	base of packet queue user heap
\$78	ndd_queue_p2p	pointer to the pointer to the user heap free space (packet queue)
\$7C	ndd_arp_base	base of ARP table user heap
\$80	ndd_arp_p2p	pointer to the pointer to the user heap free space (ARP table)
\$84	ndd_txpackets	number of packets sent by the CP2200
\$88	ndd_txbytes	number of bytes sent by the CP2200. Includes header bytes
\$8C	ndd_rxpackets	number of packets received by the CP2200
\$90	ndd_rxbytes	number of bytes received by the CP2200. Includes header bytes
\$94		12 spare bytes
\$A0	ndd_portmap	32 byte port allocation map
\$D0	ndd_buffer	1514 byte buffer for interrupt routine packet handling, or other buffering
	ndd_endi	ndd_buffer+1514 End of definition block
	ndd.leni	ndd_endi-ndd_eilk Length of definition block

ARP table linkage block

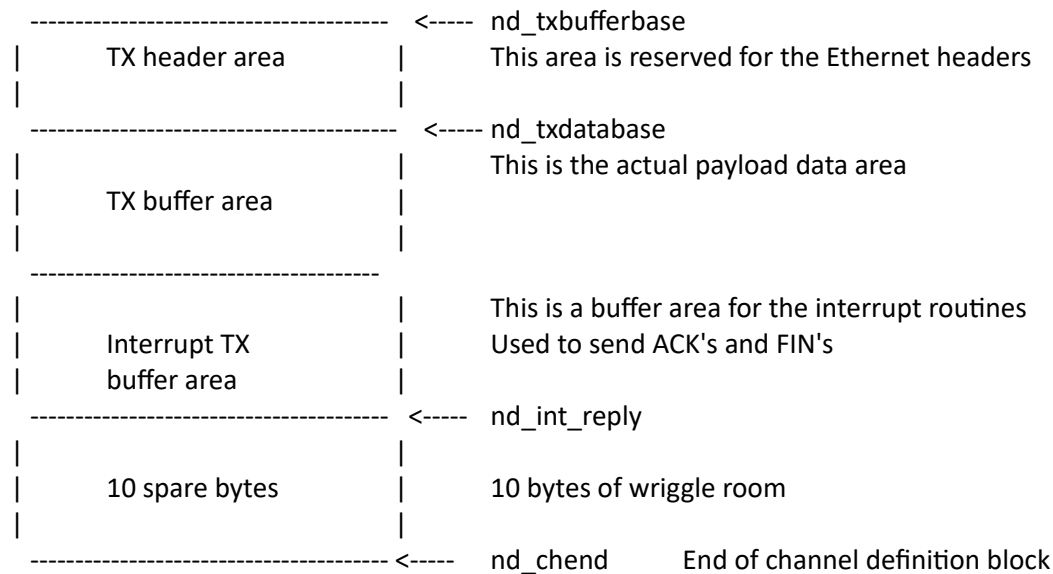
\$08	arp_next	pointer to next link
\$0C	arp_ip	IP address
\$10	arp_mac	6 byte MAC address
\$16	arp_free	2 spare bytes
\$18	arp_end	end of entry

Channel Definition Block

\$18		2 spare bytes
\$1A	nd_ARPtmr	ARP request timeout
\$1E	nd_MACbad	when set, The destination MAC address is bad, \$7F after 20 sec
\$1F		1 spare byte
\$20	nd_destmac	destination MAC address
\$26	nd_desip	destination IP address
\$2A	nd_destport	destination IP port
\$2C	nd_myip	my IP address - used for BINDing a channel
\$30	nd_myport	my IP port
\$32	nd_devicetype	device type -1=SCK, 0=UDP, 1=TCP
\$33	nd_protocol	device protocol - eg 17 for UDP
\$34	nd_acces	access mode (D3 on open call)
\$35	nd_sock_state	socket status
\$36	nd_flagsOffset	IPV4 flags and offsets
\$38	nd_sequence	TCP sequence number
\$3C	nd_acksequence	TCP acknowledge sequence
\$40	nd_offsetResFlags	UDP flags and offsets : TCP flags and things - needs sorting
\$42	nd_windows	TCP windows - needs sorting
\$44	nd_urgent	TCP urgent - needs sorting
\$46	nd_tcp_opt_len	length of option part of TCP header
\$48	tcb_SND.UNA	oldest unacknowledged sequence number
\$4C	tcb_SND.NXT	next sequence number to be sent
\$50	tcb_SND.WIN	send window size
\$52	tcb_RCV.NXT	next sequence number to be received
\$56	tcb_RCV.WND	receive window size
\$58	nd_SEG.ACK	next sequence number expected by the receiving host
\$5C	nd_SEG.SEQ	first sequence number of a segment
\$60	nd_SEG.LEN	the number of octets of data in the segment
\$62	nd_SEG.LAST	last sequence number of a segment SEG.SEQ+SEG.LEN-1
\$66	tcb_mss	host maximum segment size
\$68	tcb_ws	window scale
\$69	tcb_TCP_STACK	TCP SACK permitted true/false
\$6A	nd_TCP_packcount	the number of packets to send before waiting for an ACK
\$6B	nd_listenQ	length of IP_LISTEN backlog queue (LISTEN channel only)
		8 spare bytes
\$74	nd_ddbase	assumed start of device definition block
\$78	nd_nextch	link to next open IP channel
\$7C	nd_packqueue	link to linked list of received queued data packets
\$80	nd_txptr	transmit buffer pointer running pointer
\$82	nd_txendptr	transmit buffer end pointer
\$84	nd_rxptr	receive buffer pointer running pointer
\$86	nd_rxendptr	receive buffer end pointer
\$88	nd_rxdatabase	address of start of current rx packet's payload
\$8C	nd_txbufferbase	transmit buffer base (transmit header area)
	nd_txdatabase	nd_txbufferbase + 78 bytes Start of transmit buffer data area
	nd_int_reply	pointer to end of transmit buffer for interrupt routine
	nd_chend	nd_txdatabase + 1600 bytes End of channel definition block

The end of the channel definition block, from `nd_txbufferbase` onwards is used as the transmit buffer for the channel. The data packet is composed in this area.

There are two pointers, `nd_txptr` and `nd_txendptr`, used to track the current data position and the end of the available buffer space.



Dummy channel definition block

A dummy channel definition block (in the user heap) is used by LISTEN to handle the 3 way handshake of a connection request. This dummy channel definition block is the same as a normal channel definition block, only shorter. It has a small transmit buffer area, as it only has to send a SYN,ACK.

The dummy channel definition block is linked into a list of connection requests maintained by the LISTENing channel. And also the linked list of open IP channels.

When `IP_ACCEPT`, accepts the established connection, then the dummy channel definition block is unlinked from the two lists, and copied into the real channel definition block, and the dummy one is then deleted.

Some channel definition block entries are re-tasked for the dummy channel definition block

<code>\$00</code>	<code>dmy_base</code>		8 byte user heap header (don't touch)
<code>\$08</code>	<code>dmy_next</code>		link to next dummy channel definition block
<code>dmy_owner</code>	<code>nd_ARPtmr</code>	<code>long</code>	address of owner listing channel

PING

Some channel definition block entries are re-tasked for ICMP, Ping

Normal	Re-assignment		

nd_ARPtmr	ping_timeout	long	time to wait for ping reply
nd_sequence	ping_ident	word	ping identifier
nd_sequence+2	ping_sequence	word	ping sequence number
nd_acksequence	ping_startTime	long	start time for loop travel time
nd_urgent	ping_type	byte	ICMP transport layer type
nd_window	ping_ttl	word	Time To Live
nd_txdatabase+\$40	ping_myIPtext	20 bytes	my IP address as a string
nd_txdatabase+\$54	ping_gatewayIPtext	20 bytes	default gateway IP address as a string
nd_txdatabase+\$68	ping_targetIPtext	20 bytes	target IP address as a string
nd_txdatabase+\$7C	ping_TTLtext	8 bytes	time to live as a string
nd_txdatabase+\$84	ping_triptimes	4 words	4 trip times in ms
nd_txdatabase+\$8C	ping_received	word	number of received reply's

Receive data buffering

The CP2200 Ethernet controller can only buffer up to 4K bytes of received data, or up to 8 data packets. Whichever come first.

There is an interrupt routine that constantly monitors the Ethernet controller for data packets being received.

The basic operation of the routine is that, If the reception of a data packet is detected, Then a buffer is allocated in memory, and the data packet is copied into it.

The content of the packet is then examined, and a scan of the opened IP channels is made to see if the packet is intended for one the open channels. If a match is found, then the data packets buffer is linked onto the end of a queue of data packets intended for that channel.

If no match can be found, or the routine does not know what to do with the received packet, Then the buffer is deleted, throwing the data packet away.

Receive buffer format

\$00	rxp_base	heap allocation header, don't use (8 bytes)
\$08	rxp_next	link to the next receive data buffer
\$0C	rxp_datastart	offset from start of buffer to start of payload
\$0E	rxp_datalen	length of payload data
\$10	rxp_lifetime	number of read attempts left. If 'rxp_ok2read' is not true, then this is the number of times the channels I/O (timeout) will try to read this packet before it gives up and deletes the incomplete fragmented packet
\$12	rxp_ok2read	true if the packet is ready to be read. False if the packet is an incomplete fragmented packet
\$13	rxp_sockstate	status of packet for server connection
\$14	rxp_start	start of the data packet

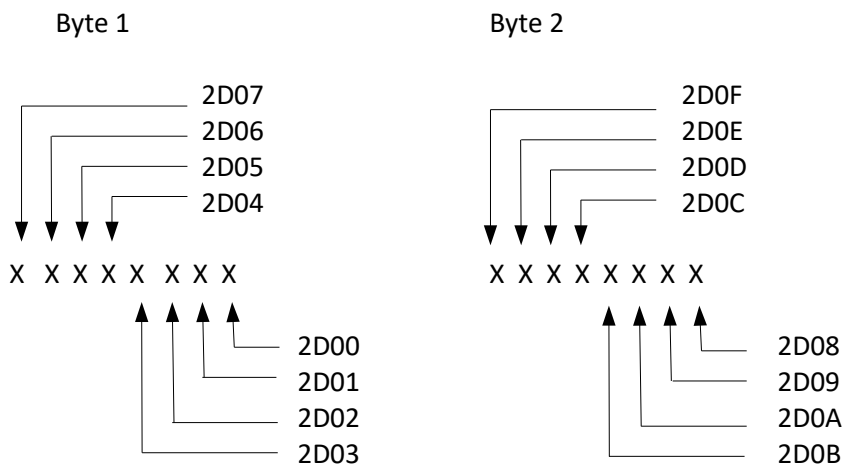
Managed port area

Sometimes the driver has to choose a port to receive data on. Rather than try to manage all 65536 ports, and to try to keep the load on the system resources down, the driver will only pick ports from a managed area. It uses a pool of 256 ports, from \$D200 to \$D300. There is a 32 byte port map in the driver definition block 'nnd_portmap', where each bit represents one of the 256 ports.

This does not mean that you cannot use ports outside of this area. It just means that the driver will not accidentally try to reuse a port in this area, that is already in use.

Each time a port is allocated from the managed port area, a pointer 'nnd_pmptr' is incremented to prevent a port being used twice in a row.

The following diagram shows the relationship between the port map and the addresses
A bit set to '1' indicates that the port has been allocated.



Background packet reading routine (INT_SERV)

The interrupt driver background reading routines are responsible for reading data packets from the CP2200 Ethernet controller. Analyse them, act on, or allocate the data packet to a channel.

The background packet reading has to operate autonomously with no direct feedback to the user of any problems. The only feedback is via the **ETH_ERRNO** S*BASIC function. Whenever the background packet reading routines, don't know what to do with a data packet, it just quietly throws it away.

The background packet reading is handled by both the 50/60Hz interrupt, and a hardware interrupt.

When an interrupt occurs, A test is made to see if a data packet is available in the CP2200 Ethernet controller. If there is no data packet available, then the interrupt ends (is_leave).

If a data packet is available (is_dopacket). A buffer is allocated in the 256K user heap, and the data packet is copied from the CP2200 Ethernet controller to the user heap.

The type of the data packet is now tested in (is_assign). If it is an ARP request it is dealt with in (is_doarp). If it's an IP packet, it is dealt with in (is_doip). Otherwise the data packet is just quietly thrown away (is_delpacket).

ARP requests (IS_DOARP)

The ARP packet is examined to see if it's a reply to a request we made, A request for a MAC address, or a general announcement that a computer has joined the network.

The appropriate action is preformed. Either store the supplied MAC address in the ARP table user heap, or send an ARP packet with the Q68's MAC address to the requester.

The ARP packet is then deleted.

IP packets (IS_DOIP)

The packets protocol is checked to see if it is either, UDP (is_doudp), ICMP (is_doicmp), or TCP (is_dotcp).

UDP packets (IS_DOUDP)

The packet is checked to see if it fragmented. If this is the first fragment of a group, then a new data packet is created in the user heap that is large enough to hold all the fragments of the group. And as further fragments arrive, they are inserted into this new packet. So you end up with one complete packet for the whole group of fragments.

A search is made of all the open IP channels looking for match of protocols and ports. If a channel is found, then the packet is added to the end of a linked list of packets waiting to be read.

TCP packets (IS_DOTCP)

A search is made of all the open IP channels looking for match of protocols, IP addresses and ports. If a channel is found, the TCP flags are analysed to decide what kind of TCP packet it is (is_tcp_decode).

Depending on these flags, and the status of the connection. Different actions take place. The actions may result in the packet being added to the end of the linked list of packets waiting to be read by a channel. Or data packets being created and sent back to the sender.

ICMP packets (IS_DOICMP)

If the packet is a Ping request (is_ping_req), Then the request is patched into a reply and sent back to the sender.


If the packet is a Ping reply (is_ping_reply) from the **ETH_PING** command, Then the requesting SMSQ/E channel is found. The round trip time is calculated in 25nS intervals, and saved in case there are any delays before getting back to S*BASIC. The packet is then linked to the SMSQ/E channel to be dealt with by the **ETH_PING** S*BASIC command.

Any other ICMP packets are discarded.

TCP Support

There is a pseudo TCP implementation in this driver. With very little TCP error handling. So all data packets must arrive complete and in the right order.

TCP name meanings



TCP Protocol Operation (Parts taken from Wikipedia)

TCP protocol operations may be divided into three phases. Connections must be properly established in a multi-step handshake process (*connection establishment*) before entering the *data transfer* phase. After data transmission is completed, the *connection termination* closes established virtual circuits and releases all allocated resources.

A TCP connection is managed by an operating system through a resource that represents the local end-point for communications, the *Internet socket*. During the lifetime of a TCP connection, the local end-point undergoes a series of state changes:

LISTEN

(server) represents waiting for a connection request from any remote TCP and port.

SYN-SENT

(client) represents waiting for a matching connection request after having sent a connection request.

SYN-RECEIVED

(server) represents waiting for a confirming connection request acknowledgement after having both received and sent a connection request.

ESTABLISHED

(both server and client) represents an open connection, data received can be delivered to the user. The normal state for the data transfer phase of the connection.

FIN-WAIT-1

(both server and client) represents waiting for a connection termination request from the remote TCP, or an acknowledgement of the connection termination request previously sent.

FIN-WAIT-2

(both server and client) represents waiting for a connection termination request from the remote TCP.

CLOSE-WAIT

(both server and client) represents waiting for a connection termination request from the local user.

CLOSING

(both server and client) represents waiting for a connection termination request acknowledgement from the remote TCP.

LAST-ACK

(both server and client) represents waiting for an acknowledgement of the connection termination request previously sent to the remote TCP (which includes an acknowledgement of its connection termination request).

TIME-WAIT

(either server or client) represents waiting for enough time to pass to be sure the remote TCP received the acknowledgement of its connection termination request. [According to RFC 793 a connection can stay in TIME-WAIT for a maximum of four minutes known as two maximum segment lifetime (MSL).]

CLOSED

(both server and client) represents no connection state at all.

Keys used by the driver for socket status

0	sts_none	
1	sts_listen	LISTEN
2	sts_syn_sent	SYN-SENT
3	sts_syn_recv	SYN-RECEIVED
4	sts_estab	ESTABLISHED(c & s) connection is established
5	sts_fin_wait1	FIN-WAIT-1
6	sts_fin_wait2	FIN-WAIT-2
7	sts_close_wait	CLOSE-WAIT
8	sts_closing	CLOSING
9	sts_last_ack	LAST-ACK
10	sts_time_wait	TIME-WAIT
11	sts_closed	CLOSED

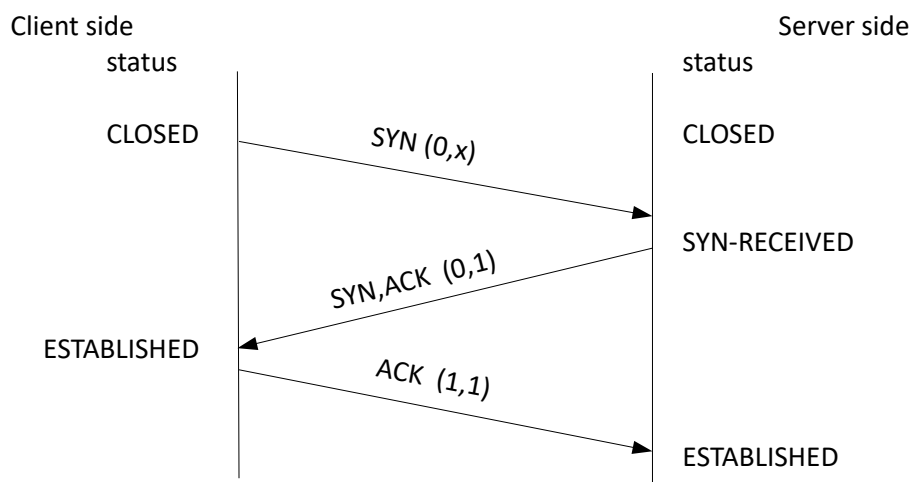
Connection establishment

To establish a connection, TCP uses a three-way handshake. Before a client attempts to connect with a server, the server must first bind to and listen at a port to open it up for connections: this is called a passive open. Once the passive open is established, a client may initiate an active open. To establish a connection, the three-way (or 3-step) handshake occurs:

1. **SYN:** The active open is performed by the client sending a SYN to the server. The client sets the segment's sequence number to a random value A.
2. **SYN-ACK:** In response, the server replies with a SYN-ACK. The acknowledgement number is set to one more than the received sequence number i.e. A+1, and the sequence number that the server chooses for the packet is another random number, B.
3. **ACK:** Finally, the client sends an ACK back to the server. The sequence number is set to the received acknowledgement value i.e. A+1, and the acknowledgement number is set to one more than the received sequence number i.e. B+1.

At this point, both the client and server have received an acknowledgement of the connection. The steps 1, 2 establish the connection parameter (sequence number) for one direction and it is acknowledged. The steps 2, 3 establish the connection parameter (sequence number) for the other direction and it is acknowledged. With these, a full-duplex communication is established.

3 way handshake



The numbers in brackets indicate offsets from a randomly generated numbers on each side, of SEQ and ACK numbers, where x is undefined. In practice x will be zero.

Connecting to a server

Connecting to a server involves the SMSQ/E Open channel routine calling the 'TCPCONNECT' routine.

This routine will try to make a 'Three Way Handshake' connection to a TCP server that is 'Listening' for connection requests.

At this point, as far as SMSQ/E is concerned, The channel has not yet been opened, so no normal I/O requests can be processed. The area that will be the channel definition block is loaded with data to send a SYN message, and then sends it.

The routine then waits for the interrupt driven background packet reading routines to receive the SYN,ACK message. This is handled by the 'IS_TCP_DECODE' routine, which sets the socket status byte in the channel definition block (nd_sock_state) to 'ESTABLISHED'

The routine then sends an ACK message, completing the connection.

Server accepting a connection


The server has a 'Listening' channel that waits for incoming connection requests. When a 'SYN' is received by the background packet reading routines, for a listening channel. Then a dummy channel definition block is created to handle the 3 way handshake (is_sendSYNACK). This dummy channel definition block is then added to a list of queued requests, and the linked list of open IP channels.

When an IP_ACCEPT system trap is called (tcp_accept), the request queue of the supplied listening channel is scanned for the oldest queued request of dummy channel definition blocks. This dummy block is then unlinked, and a new channel definition block (that will be the real one) is created. Data is copied from the dummy block to the real one, and the dummy block is then removed.

Closing a connection

Closing a connection involves a '4 way handshake', or a '3 way handshake' process. It's a bit more complicated than making a 'connection', and involves the socket (channel) going through a number of states. Depending on which side initiates the close. And one side may leave a channel open as far as SMSQ/E is concerned.

Below is shown the sequence of messages to terminate a TCP connection. As taken from the TCP Wikipedia page.




The '3 way handshake' involves the middle two messages being combined into one message. The 'Close channel' routine will send anything left in its buffer, then call the 'TCP_CLOSE' routine.

The 'TCP_CLOSE' routine co-operates with the 'IS_TCP_DECODE' routine, as in making a connection. To move the channels socket status through the various stages, depending on which side initiates the close.

Note that in real packet transfers, extra flags may be set in the messages that are sent. For example, the first message sent may be, FIN, or FIN,ACK, or PSH,FIN,ACK.

The 'IS_TCP_DECODE' routine tries to account for all these combinations, and also the state that the socket is currently in.

There is also the possibility that both ends of the TCP connection will try to initiate the close at the same time. In this situation, the process is slightly different



The numbers in brackets indicate sample SEQ and ACK numbers

Ethernet Header Formats

<u>Physical (Ethernet) layer header format</u>		14 bytes
\$00	6 bytes	destination MAC address
\$06	6 bytes	source MAC address
\$0C	word	length/type
		<\$0800 Length of the packet
	\$0800	Ethernet IPV4 datagram
	\$0806	ARP Frame
	\$0835	RARP
	\$8100	IEEE802.1 Q tag 10/100 VLAN Frame
	\$86DD	IPV6
	\$8808	10/100 Control Frame

<u>Network (IPV4) layer header format</u>		20 bytes	Used by ICMP, IGMP, TCP, UDP, ENCAP, OSPF, SCTP
\$00	byte	Version/HL, Upper nibble=Version, Lower nibble=IHL	
\$01	byte	Type of service, Bits 7-2=DSCP, Bits 1-0=ECN	
\$02	word	Length	
\$04	word	Identification	
\$06	word	Flags and offset, Bits 15-13=Flags, Bits 12-0=Fragment offset	
\$08	byte	TTL Time to live	
\$09	byte	Protocol	
\$0A	word	Checksum	
\$0C	long	Source IP address	
\$10	long	Destination IP address	

Description of header format

Version

The first header field in an IP packet is the four-bit version field. For IPv4, this is always equal to 4.

Internet Header Length (IHL)

The IPv4 header is variable in size due to the optional 14th field (options). The IHL field contains the size of the IPv4 header, it has 4 bits that specify the number of 32-bit words in the header. The minimum value for this field is 5,[26] which indicates a length of $5 \times 32 \text{ bits} = 160 \text{ bits} = 20 \text{ bytes}$. As a 4-bit field, the maximum value is 15, this means that the maximum size of the IPv4 header is $15 \times 32 \text{ bits}$, or $480 \text{ bits} = 60 \text{ bytes}$.

Differentiated Services Code Point (DSCP)

Originally defined as the type of service (ToS), this field specifies differentiated services (DiffServ) per RFC 2474 (updated by RFC 3168 and RFC 3260). New technologies are emerging that require real-time data streaming and therefore make use of the DSCP field. An example is Voice over IP (VoIP), which is used for interactive voice services.

Explicit Congestion Notification (ECN)

This field is defined in RFC 3168 and allows end-to-end notification of network congestion without dropping packets. ECN is an optional feature that is only used when both endpoints support it and are willing to use it. It is effective only when supported by the underlying network.

Total Length

This 16-bit field defines the entire packet size in bytes, including header and data. The minimum size is 20 bytes (header without data) and the maximum is 65,535 bytes. All hosts are required to be able to reassemble datagrams of size up to 576 bytes, but most modern hosts handle much larger packets. Sometimes links impose further restrictions on the packet size, in which case datagrams must be fragmented. Fragmentation in IPv4 is handled in either the host or in routers.

Identification

This field is an identification field and is primarily used for uniquely identifying the group of fragments of a single IP datagram. Some experimental work has suggested using the ID field for other purposes, such as for adding packet-tracing information to help trace datagrams with spoofed source addresses,[27] but RFC 6864 now prohibits any such use.

If IP packet is fragmented during the transmission, all the fragments contain same identification number. to identify original IP packet they belong to.

Flags

A three-bit field follows and is used to control or identify fragments. They are (in order, from most significant to least significant):

bit 0: Reserved; must be zero.[note 1]

bit 1: Don't Fragment (DF)

bit 2: More Fragments (MF)

If the DF flag is set, and fragmentation is required to route the packet, then the packet is dropped. This can be used when sending packets to a host that does not have resources to handle fragmentation. It can also be used for path MTU discovery, either automatically by the host IP software, or manually using diagnostic tools such as ping or traceroute. For unfragmented packets, the MF flag is cleared. For fragmented packets, all fragments except the last have the MF flag set. The last fragment has a non-zero Fragment Offset field, differentiating it from an unfragmented packet.

Fragment Offset

The fragment offset field is measured in units of eight-byte blocks. It is 13 bits long and specifies the offset of a particular fragment relative to the beginning of the original unfragmented IP datagram. The first fragment has an offset of zero. This allows a maximum offset of $(2^{13} - 1) \times 8 = 65,528$ bytes, which would exceed the maximum IP packet length of 65,535 bytes with the header length included ($65,528 + 20 = 65,548$ bytes).

The fragment offsets are calculated from the start of the transport layer

Time To Live (TTL)

An eight-bit time to live field helps prevent datagrams from persisting (e.g. going in circles) on an internet. This field limits a datagram's lifetime. It is specified in seconds, but time intervals less than 1 second are rounded up to 1. In practice, the field has become a hop count—when the datagram arrives at a router, the router decrements the TTL field by one. When the TTL field hits zero, the router discards the packet and typically sends an ICMP Time Exceeded message to the sender. The program traceroute uses these ICMP Time Exceeded messages to print the routers used by packets to go from the source to the destination.

Protocol

This field defines the protocol used in the data portion of the IP datagram. The Internet Assigned Numbers Authority maintains a list of IP protocol numbers as directed by RFC 790. Tells the Network layer at the destination host, to which Protocol this packet belongs to, i.e. the next level Protocol. For example protocol number of ICMP is 1, TCP is 6 and UDP is 17.

Some of the common payload protocols are:

Protocol Number	Protocol Name	Abbreviation
1	Internet Control Message Protocol	ICMP
2	Internet Group Management Protocol	IGMP
6	Transmission Control Protocol	TCP
17	User Datagram Protocol	UDP
41	IPv6 encapsulation	ENCAP
89	Open Shortest Path First	OSPF
132	Stream Control Transmission Protocol	SCTP

Header Checksum

The 16-bit IPv4 header checksum field is used for error-checking of the header. When a packet arrives at a router, the router calculates the checksum of the header and compares it to the checksum field. If the values do not match, the router discards the packet. Errors in the data field must be handled by the encapsulated protocol. Both UDP and TCP have checksum fields.

When a packet arrives at a router, the router decreases the TTL field. Consequently, the router must calculate a new checksum.

Source address

This field is the IPv4 address of the sender of the packet. Note that this address may be changed in transit by a network address translation device.

Destination address

This field is the IPv4 address of the receiver of the packet. As with the source address, this may be changed in transit by a network address translation device.

Options

The options field is not often used. Note that the value in the IHL field must include enough extra 32-bit words to hold all the options (plus any padding needed to ensure that the header contains an integer number of 32-bit words). The list of options may be terminated with an EOL (End of Options List, 0x00) option; this is only necessary if the end of the options would not otherwise coincide with the end of the header. The possible options that can be put in the header are as follows:

Field	Size (bits)	Description
Copied	1	Set to 1 if the options need to be copied into all fragments of a fragmented packet.
Option Class	2	A general options category. 0 is for "control" options, and 2 is for "debugging and measurement". 1 and 3 are reserved.
Option Number	5	Specifies an option.
Option Length	8	Indicates the size of the entire option (including this field). This field may not exist for simple options.
Option Data	Variable	Option-specific data. This field may not exist for simple options.

Note: If the header length is greater than 5 (i.e., it is from 6 to 15) it means that the options field is present and must be considered.

Note: Copied, Option Class, and Option Number are sometimes referred to as a single eight-bit field, the Option Type.

Packets containing some options may be considered as dangerous by some routers and be blocked

Network (ARP) layer format			28 bytes
\$00	word	HDR	Hardware type, \$0001 ethernet
\$02	word	PRO	Protocol, \$0800=ethernet internet protocol
\$04	byte	HLN	MAC address length, usually 6
\$05	byte	PLN	IP address length, usually 4
\$06	word	OP	Operation, 1=request, 2=reply
\$08	6 bytes	SHA	Sender MAC address
\$0E	long	SPA	Sender IP address
\$12	6 bytes	THA	Target MAC address
\$18	long	TPA	Target IP address

Description of header format

HDR Hardware type

This field specifies the type of hardware used for the local network transmitting the ARP message; thus, it also identifies the type of addressing used. Some of the most common values for this field

1	Ethernet (10Mb)
6	IEEE 802 Networks
7	ARCNET
15	Frame Relay
16	Asynchronous Transfer Mode (ATM)
17	HDLC
18	Fibre Channel
19	Asynchronous Transfer Mode (ATM)
20	Serial Line

PRO Protocol Type

This field is the complement of the *Hardware Type* field, specifying the type of layer three addresses used in the message. For IPv4 addresses, this value is 2048 (0800 hex), which corresponds to the EtherType code for the Internet Protocol.

HLN Hardware Address Length

Specifies how long hardware addresses are in this message. For Ethernet or other networks using IEEE 802 MAC addresses, the value is 6.

PLN Protocol Address Length

Again, the complement of the preceding field; specifies how long protocol (layer three) addresses are in this message. For IP(v4) addresses this value is of course 4.

OP Opcode

This field specifies the nature of the ARP message being sent. The first two values (1 and 2) are used for regular ARP. Numerous other values are also defined to support other protocols that use the ARP frame format, such as RARP, some of which are more widely used than others

- | | |
|---|---------------|
| 1 | ARP Request |
| 2 | ARP Reply |
| 3 | RARP Request |
| 4 | RARP Reply |
| 5 | DRARP Request |
| 6 | DRARP Reply |
| 7 | DRARP Error |
| 8 | InARP Request |
| 9 | InARP Reply |

SHA Sender Hardware Address

The hardware (layer two) address of the device sending this message (which is the IP datagram source device on a request, and the IP datagram destination on a reply, as discussed in the topic on ARP operation).

SPA Sender Protocol Address

The IP address of the device sending this message.

THA Target Hardware Address

The hardware (layer two) address of the device this message is being sent to. This is the IP datagram destination device on a request, and the IP datagram source on a reply)

TPA Target Protocol Address

The IP address of the device this message is being sent to.

Transport (UDP) layer format

Pseudo header for checksum calculation. Not to be included in actual header 12 bytes

\$00	long	Source IP address
\$04	long	Destination IP address
\$08	byte	zero
\$09	byte	\$11 (17) UDP
\$0A	word	UDP length from actual header

Actual header 8 bytes

\$00	word	Source port
\$02	word	Destination port
\$04	word	UDP length, payload length + 8 bytes of header
\$06	word	UDP checksum

Transport (TCP) layer format

Pseudo header for checksum calculation. Not to be included in actual header 12 bytes

\$00	long	Source IP address
\$04	long	Destination IP address
\$08	byte	zero
\$09	byte	\$06 TCP
\$0A	word	TCP length, Actual header length + options + payload

Actual header 20 bytes

\$00	word	Source port
\$02	word	Destination port
\$04	long	Sequence number
\$08	long	ACK Sequence number
\$0C	word	offset/res/flags
\$0E	word	Window
\$10	word	Checksum
\$12	word	Urgent pointer
\$14		Options bytes

Description of header format

Source Port

The source port number.

Destination Port

The destination port number.

Sequence Number

The sequence number of the first data octet in this segment (except when SYN is present).

If SYN is present the sequence number is the initial sequence number (ISN) and the first data octet is ISN+1.

Acknowledgment Number

If the ACK control bit is set this field contains the value of the next sequence number the sender of the segment is expecting to receive. Once a connection is established this is always sent.

Data Offset: 4 bits

The number of 32 bit words in the TCP Header. This indicates where the data begins. The TCP header (even one including options) is an integral number of 32 bits long.

Reserved: 3 bits

Reserved for future use. Must be zero.

Control Bits: 9 bits (from left to right):

NS: ECN - nonce - concealment protection
CWR: Congestion Window Reduced
ECE: ECN - Echo has a dual role
URG: Urgent Pointer field significant
ACK: Acknowledgment field significant
PSH: Push Function
RST: Reset the connection
SYN: Synchronize sequence numbers
FIN: No more data from sender

```

O O O O R R R C C C C C C C C C C
                        |  |----- FIN
                        |----- PSH
```

Window

The number of data octets beginning with the one indicated in the acknowledgment field which the sender of this segment is willing to accept.

Checksum

The checksum field is the 16 bit one's complement of the one's complement sum of all 16 bit words in the header and text. If a segment contains an odd number of header and text octets to be checksummed, the last octet is padded on the right with zeros to form a 16 bit word for checksum purposes. The pad is not transmitted as part of the segment. While computing the checksum, the checksum field itself is replaced with zeros.

The checksum also covers the 96 bit pseudo header conceptually

Urgent Pointer

This field communicates the current value of the urgent pointer as a positive offset from the sequence number in this segment. The urgent pointer points to the sequence number of the octet following the urgent data. This field is only be interpreted in segments with the URG control bit set.

Options: variable

Options may occupy space at the end of the TCP header and are a multiple of 8 bits in length. All options are included in the checksum. An option may begin on any octet boundary. There are two cases for the format of an option:

Case 1: A single octet of option-kind.

Case 2: An octet of option-kind, an octet of option-length, and the actual option-data octets.

The option-length counts the two octets of option-kind and option-length as well as the option-data octets.

Note that the list of options may be shorter than the data offset field might imply. The content of the header beyond the End-of-Option option must be header padding (i.e., zero).

A TCP must implement all options.

Currently defined options include (kind indicated in octal):

Kind	Length	Meaning
----	-----	-----
0	-	End of option list.
1	-	No-Operation.
2	4	Maximum Segment Size.
3	3	Window scale.
4	2	TCP SACK permitted.

Specific Option Definitions

End of Option List

Kind=0

This option code indicates the end of the option list. This might not coincide with the end of the TCP header according to the Data Offset field. This is used at the end of all options, not the end of each option, and need only be used if the end of the options would not otherwise coincide with the end of the TCP header.

No-Operation

Kind=1

This option code may be used between options, for example, to align the beginning of a subsequent option on a word boundary. There is no guarantee that senders will use this option, so receivers must be prepared to process options even if they do not begin on a word boundary.

Maximum Segment Size

```
+-----+-----+-----+-----+
|00000010|00000100|   max seg size   |
+-----+-----+-----+-----+
Kind=2  Length=4
```

Maximum Segment Size Option Data: 16 bits

If this option is present, then it communicates the maximum receive segment size at the TCP which sends this segment. This field must only be sent in the initial connection request (i.e., in segments with the SYN control bit set). If this option is not used, any segment size is allowed.

Window Scale

```
+-----+-----+-----+-----+
|00000001|00000011|00000011|   scale   |
+-----+-----+-----+-----+
Kind=3  Length=3
```

The scale factor is the number of bits to left shift the 16 bit window size (ignored in SYN message)

TCPASCK permitted

```
+-----+-----+-----+-----+
|00000001|00000001|00000100|00000010|
+-----+-----+-----+-----+
Kind=4  Length=2
```

Padding: variable

The TCP header padding is used to ensure that the TCP header ends and data begins on a 32 bit boundary. The padding is composed of zeros.

<u>Transport (Ping) layer format</u>		ICMP	8 bytes
\$00	byte	Type, 8=IPV4 request, 0=IPV4 reply	Type of ICMP message
\$01	byte	Code, 0	
\$02	word	Header checksum - including payload	
\$04	word	Identifier	
\$06	word	Sequence number	
		32 byte payload	

<u>Transport (IGMP) layer format</u>		8 bytes
\$00	byte	Type, Membership Query, membership Report, Leave group
\$01	byte	Max response time, only in Membership Query messages
\$02	word	Checksum
\$04	long	Group address, Behaviour of this field varies by the type of message sent:
		Membership Query: (set to)
		General Query: All zeroes
		Group Specific Query: multicast group address
		Membership Report: multicast group address
		Leave Group: multicast group address

Fragmented Packet layout

First packet :-

Standard Physical layer.

Standard Network layer, other than Flags/offset = \$2000

Transport layer has a length (and checksum?) for the entire unfragmented payload.

Second packet:-

Standard Physical layer.

Standard Network layer, other than Flags = %001 or %000 on the last fragment.

Offset = offset from start of the payload, divided by 8.

There is no transport layer

ARP Handling

The handling of acquiring and supplying MAC addresses works as follows.

When the background packet receiving routine receives an ARP packet. If the packet is a Gratuitous request, or a reply to a request we made. Then the details are entered into the ARP table.

If it is a request for our MAC address, then a reply is sent to the sender.

When the systems IP address is set, or changed. Then a Gratuitous packet is broadcast over the network.

When the required MAC address, for an IP address, is not available from the ARP table. The method for obtaining it goes along these lines.

When a channel is opened with OPEN_IN, If there is not a MAC address entry in the ARP table, then a ARP request is sent. A 40 second timer, and a flag, are initialized in the channel definition block to indicate that the MAC address is invalid. The channel open, then finishes normally.

What happens next depends on how quickly the ARP request is replied to, and how much time passes between opening the channel, and trying to do any I/O.

Assuming a worst case scenario, where there is never an ARP reply received, and channel I/O starts straight after opening the channel. The process will be as follows.

When any I/O is done to the channel, there will be a wait for 20 seconds, followed by another ARP request, In case the first one way lost. Then a further wait of 20 seconds, followed by a final ARP request. The timer will be reset to another 40 seconds, ETH_ERRNO will be set to error 34, and finally a system 'transmission error' will be returned.

If an ARP reply is received within the 40 seconds, then the above procedure will terminate, and I/O will continue normally.

IP Error messages

Receive errors

iperr.rxomem	1	Insufficient memory to read packet into
iperr.rxcsum	2	Read packet validation failed, checksum mismatch
iperr.badmac	3	Received packet from unexpected MAC address
iperr.ctrl	4	Unable to process TCP control bits
iperr.segerr	5	Unexpected TCP segment numbers
iperr.seqerr	6	Unexpected TCP sequence number

Send errors

iperr.txerr	8	Last packet was not transmitted successfully
iperr.txnogo	10	Timeout waiting for transmit buffer to empty
iperr.txprot	11	Protocol not found when creating a packet
iperr.txarpf	12	Failure sending ARP request
iperr.txpingf	13	Failure sending Ping reply
iperr.txnoak	14	Attempt to send too many packets with received ACK's (TCP)
iperr.txSYAC	15	Failure sending a SYN,ACK
iperr.maxt	16	Too many attempts made to send a packet

Open errors

iperr.inparm	20	Invalid parameters for requested OPEN type
iperr.noport	21	No managed ports available
iperr.nomac	22	Unable to acquire MAC address for specified IP address
iperr.portna	23	Requested port already in use
iperr.noserv	24	No response from requested TCP server for a SYN request

I/O errors

iperr.txcsum	30	Error creating checksum for transmission
iperr.toobig	31	Attempt to send more than 64K bytes
iperr.scklen	32	Bad sockaddr length
iperr.fragto	33	Fragmented packet incomplete before life ran out
iperr.macto	34	I/O timeout while waiting for a MAC address
iperr.outrng	35	IP_LISTEN queue out of range
iperr.nottcp	36	Not a TCP channel
iperr.isomem	37	Out of memory creating a dummy channel definition block
iperr.lqfull	38	The queue for a listening channel is full

Close errors

iperr.closto	40	Timed out waiting for a TCP close acknowledgement
iperr.wait1	41	While in FIN-WAIT-1, Timed out waiting for an ACK
iperr.wait2	42	While in FIN-WAIT-2, Timed out waiting for a FIN (last ACK)
iperr.wait3	43	While in LAST-ACK, Timed out waiting for final ACK
iperr.uxfin	44	Unexpected FIN received
iperr.wait4	45	While in CLOSING, Timed out waiting for final ACK
iperr.nofin	46	Failed sending a FIN

IP Trap errors

iperr.dbrec	50	Error reading a database entry
-------------	----	--------------------------------

CP2200 Initialization errors

iperr.inito	100	Timed out while waiting for controller to reset
iperr.aufail	110	Timed out while waiting for auto-neg in Physical layer