



Tech Info Library

MacTCP 2.0.4 Release Notes and Read Me (8/94)

Article Created: 25 May 1994

Article Reviewed/Updated: 26 August 1994

October 21, 1993

MacTCP version 2.0.4 is an update release.

MacTCP version 2.0.3 was released internal to Apple.

MacTCP 2.0.4 contains the following improvements over MacTCP 2.0.2 and earlier versions:

Performance:

- Earlier versions of MacTCP had several problems associated with setting lower bounds for retransmission times on TCP connections. This was particularly apparent when using MacTCP over serial lines which tend to have both long round trip times and a tendency to corrupt packets. The lower bounding for retransmission timeouts was frequently set to 20 seconds or more. MacTCP 2.0.4 corrects these problems, and sets the retransmission lower bound to 1 second regardless of link type. This allows MacTCP to correctly calculate the retransmission time as $SRTT + 4(RTTV)$. Note that this may result in an increased number of retransmission on some connections, but the overall performance is much more acceptable.
- When set to large sizes, the disk cache included in System 7 and 7.1 does not work well with networking. Doing an FTP copy to a Mac with a 1 megabyte disk cache will result in cache flush delays of 30 seconds or longer, during which time the receive window will become full and no additional data can flow over the connection. We recommend that users of MacTCP set their disk cache size to 128k or less within the Memory control panel for optimal network performance. This should be fixed in a future release of the Mac OS.

NOTE: This problem has been fixed with System Software 7.5.

- Earlier versions of MacTCP incorrectly implemented TCP congestion avoidance by reopening the congestion window too quickly. MacTCP 2.0.4 corrects this problem.

Configuration:

- BootP requests are (again) sent to the Ethernet broadcast address as they were in MacTCP 1.x. MacTCP 2.0 and 2.0.2 sent BootP requests to an Ethernet multicast

address.

- If MacTCP is set up to run over Ethernet, and server configuration mode is chosen in the MacTCP control panel, MacTCP attempts to configure itself using BootP. If that fails, it will attempt configuration using RARP. Previous versions of MacTCP used RARP initially, and if that failed, attempted BootP.

Domain Name Resolver:

- If a domain name server cannot resolve a name but can provide references to more appropriate name servers, the MacTCP Domain Name Resolver (DNR) will recursively follow those references until it resolves the name, runs out of servers and references, or times out. Earlier versions of MacTCP did not recursively follow references.
- The DNR is now compliant with the domain name syntax portions of RFC 1123 with respect to decimal digits and hyphens. Previous versions of the DNR did not accept hyphens and assumed that names beginning with numbers were dot format internet addresses.
- In earlier versions of the DNR, the processing of MXInfo and HInfo replies leaked memory if more than one record was received in the answer section.
- Earlier versions of the DNR corrupted the higher order bits of MXInfo preferences greater than 255.
- When resolving an alias, the DNR will now follow up on any CNAMEs which are returned.
- The DNR manages cached resource record (RR) memory more efficiently than previous versions.

Program Interfaces:

- Earlier versions of MacTCP ignored the ULP timeout parameter in the TCPClose call.
- Previous versions of MacTCP dynamically allocated TCP and UDP ports in the range between 256 and 65535, inclusive. For consistency with most UNIX hosts, MacTCP 2.0.4 has been changed to first attempt to allocate ports in the range between 1025 and 5000. This does not impact programs which request specific ports.
- In earlier versions of MacTCP, it was possible for the UDP Create call to allocate a UDP port number which was already in use if a UDP Multiport stream existed.
- Previously, UDP did not call the application's completion routine if a driver error caused an IP Write call to fail.
- PBControl calls with csCode ipctlWrite now set the IOResult correctly.

- The UDP notify routine is called when an ICMP destination unreachable is received. In previous versions of MacTCP, the notify routine might not get called.

Odds & Ends:

- MacTCP 2.0.2 sometimes sent packets greater than 576 bytes long to hosts which were not on the same subnet. Earlier versions of MacTCP did not have this problem. It has been corrected in MacTCP 2.0.4.
- The TCP/IP protocol prohibits IP precedence level changes once a TCP connection is established. Some DEC Alphas running OSF/1 have been observed to do this when acting as FTP servers. MacTCP has always correctly sent a reset in response, but previous versions hung after doing so. Now MacTCP sends the reset but does not hang.

Article Change History:

26 Aug 1994 - Added note that the large disk cache sizes has been fixed
with System 7.5.

Support Information Services
Copyright 1994, Apple Computer, Inc

Keywords: knts

=====

This information is from the Apple Technical Information Library.

19960215 11:05:19.00

Tech Info Library Article Number: 15463