



Tech Info Library

Open Transport 1.1: TCP/IP Features Q & A (7/96)

Article Created: 2 January 1996

Article Reviewed/Updated: 2 July 1996

TOPIC -----

This article is the Reference Q & A (questions and answers) on TCP/IP features for Open Transport 1.1.

DISCUSSION -----

Question: What are some of the features of Open Transport/TCP?

Answer: With the broad cross-platform adoption of TCP/IP - and the tremendous visibility of the Internet - Apple has made a significant investment in bringing a workstation-class implementation of TCP/IP protocols to Mac OS. As with Apple's earlier MacTCP, Open Transport/TCP is a full 32-bit stack.

Open Transport/TCP adds support for:

- dynamic path MTU discovery, for more efficient network use in heterogeneous network topologies;
- Dynamic Host Configuration Protocol (DHCP), for centralized IP address configuration management. DHCP is an Internet Engineering Task Force (IETF) standards-track protocol;
- IP multicast, for participation as an Mbone client;
- simultaneous TCP connections limited only by installed memory and processor power, for increased functionality as a Internet or other TCP/IP network server;
- a new, more robust and standards-compliant domain name resolver (a caching stub DNR);
- support for developer access to raw IP services, as well as TCP and UDP;
- Ethernet_SNAP and IEEE 802.3 framing, for better interoperability with a wider range of TCP/IP hosts;
- implicit and explicit domain name search paths, for increased control of domain name resolution,

- multiple IP routers with fail-over, for increased robustness in mission critical applications.

Question: How does the new support for Dynamic Path MTU discovery work?

Answer: Open Transport/TCP sets the "don't fragment" bit in the IP datagram header on transmission unless the packet size is larger than the MTU for the network. Intermediate routers are required by current RFCs to send back an "ICMP can't fragment" error when presented with a "don't fragment" packet that cannot be forwarded without fragmentation with that MTU size. In that event, Open Transport/TCP moves to the next smaller MTU size for that path and re-sends the packet, again with the "don't fragment" bit set. This process automatically results in using the largest supported MTU size for off-subnet traffic.

Question: How does the new Open Transport/TCP domain name resolver (DNR) work?

Answer: The new DNR implements name-to-address (A), address-to-name (PTR), system CPU and OS (HINFO), and mail exchange (MX) queries. It does not implement negative caching, depending on a local full service resolver to provide this facility. The DNR will always request recursion, but will follow references if recursion is not available.

The DNR caches name-to-address and cname-to-name mappings. It does not cache host info (OS and CPU type information) nor does it cache Mail Exchange/Preference information. It does not save name server references after a query is resolved; further queries begin anew at the configured name servers.

Fully qualified domains names or FQDNs (those ending with a "."), and provisional FQDNs (those containing at least one "." internally but not ending with ".") are submitted for resolution without manipulation. Otherwise the name is assumed to be a partially qualified domain (PQDN).

The first - but optional - step in PQDN resolution is the use of an Implicit Search Path. To be used, Implicit Search must first be configured using the Advanced or Administrator view by entering values in the "Implicit Search Path: Starting Domain Name" and "Ending Domain Name" fields. When so configured the DNR will attempt to change the PQDN to a FQDN for resolution by concatenating the PQDN with domain names in the ancestor hierarchy delimited by the Starting and Ending Domain Name values (that is, searching for a PQDN of joe could result in a search for joe.hardware.support.apple.com, joe.support.apple.com and joe.apple.com). Implicit searching stops when the FQDN is resolved, or when the Ending Domain Name value has been tried and fails (that is, joe.com would not be tried, assuming an Ending Domain Name value of apple.com).

If the PQDN has not yet been resolved (including the case where an Implicit Search Path was not configured), explicit Additional Search Domains are searched. For each Search Domain configured, name server(s) are contacted in the order specified in the Name Servers field. If the name is resolved in the first search domain from which an answer is returned other Search Domains will not be checked. Note that at least one Search Domain (roughly equivalent to MacTCP's

Default Domain) must be explicitly configured in order to resolve any PQDNs.

If an authoritative answer that the "name-does-not-exist" is returned, the DNR immediately begins the search in the next configured Search Domain. The search continues through the configured Search Domains.

The DNR has an overall time-out of 2 minutes, after which it will abandon the search.

Question: Does Open Transport/TCP support a local HOSTS file?

Answer: Open Transport/TCP supports one or more HOSTS file, stored in the System Preferences folder, that may be used to supplement and/or customize the domain name resolver's initial cache of information. The selected file is opened and parsed when Open Transport/TCP is initialized. As with MacTCP, the supported HOSTS file features follow a subset of the Domain Name System Master File Format (RFC 1035).

Supported features include blank lines, comments (indicated by a semicolon), and data entry. Comments may begin at any location in a line; they may follow data entry on the same line. A comment extends from the semicolon to the end of the line. Data entry must follow the format:

```
<domain-name> <rr> [<comment>]
```

where <domain-name> is an absolute or Fully Qualified Domain Name, and where

```
<rr> = [<ttd>] [<class>] <type> <rdata> OR [<class>] [<ttd>] <type> <rdata>
```

The only <class> currently supported is IN (Internet Domain); <ttd>, time to live, indicates the record's configured lifetime in seconds; and <type> can be A (host address), CNAME (canonical name of an alias), or NS (name server). If <ttd> is not present the entry is assumed to have an infinite lifetime; this may also be indicated by specifying a value of minus-one (-1). \$INCLUDE and \$ORIGIN are not supported.

Open Transport/TCP is more stringent regarding the content and format of the HOSTS file than was MacTCP, which permitted violation of the FQDN requirement for <domain-name>. For instance, this format:

```
charlie A 128.1.1.1
```

which was acceptable to the MacTCP DNR, is no longer permitted because of the use of domain search lists in Open Transport/TCP (charlie could potentially exist in any or all of the configured domains). To accomplish the same effect, use this format instead:

```
charlie CNAME myhost.mydomain.edu
```

```
myhost.mydomain.edu A 128.1.1.1
```

This associates the local alias charlie with the fully qualified domain name

myhost.mydomain.edu, and resolves it to the address 128.1.1.1. Use of local aliases is limited to CNAME entries; NS and A entries must use fully qualified domain names.

If a HOSTS file is used, every effort should be made to keep it as small as possible and to only include entries that will be accessed frequently. This reduces the total memory footprint required to cache the DNS information and minimizes the need to maintain and update the HOSTS files as system information changes over time.

In order to activate a HOSTS file, the Advanced or Administrator mode must be used to select the desired file. The text file must already exist; it could have been created with any text editor or word processor. The HOSTS file is tied to the selected configuration. An administrator might, for example, specify different HOSTS files for use when connecting via Ethernet to the campus LAN and when dialing-in from a remote location.

Question: What are some of the changes to the human interface for Open Transport/TCP?

Answer: The Open Transport/TCP configuration application represents a complete overhaul of the human interface from the MacTCP software it replaces. In addition to generic new features noted elsewhere (multiple saved configurations, recommended and required settings, on-line documentation, and so on), key new features include:

- direct entry of IP addresses and subnet mask in standard "dot notation";
- explicit selection of desired configuration method, now including Manual, RARP, BootP, MacIP, and DHCP;
- support for attachment to networks using Classless InterDomain Routing (CIDR);
- support for multiple entries in the router, name server, and explicit domain search lists; and
- improved support for large AppleTalk networks when using MacIP server/gateways.

Question: Does Open Transport/TCP support MacTCP "Server" addressing?

Answer: MacTCP Server mode addressing is a combination of the Bootstrap Protocol (BootP) and Reverse Address Resolution Protocol (RARP) configuration methods. When Server mode is selected, MacTCP will use BootP to attempt to acquire an IP address. If BootP fails to provide a valid address it then tries RARP. Whichever protocol is successful is stored as a preference, and is used first on next system startup. While this "fall-back" approach adds a degree of robustness from the users point of view, it also adds a degree of unpredictability from a network administrators point of view.

Based on customer feedback, Open Transport/TCP has been designed to allow a network administrator to explicitly specify the single method they prefer to use. Thus while both RARP and BootP are individually supported, Server mode does not appear as a choice in the Open Transport/TCP configuration utility.

Question: Does Open Transport/TCP support MacTCP "Dynamic" addressing?

Answer: No. MacTCP "Dynamic" mode addressing was based on an Apple-proprietary extension to TCP/IP protocols, which applied the address negotiation and assignment rules used by the AppleTalk protocols to TCP/IP networks. This made it very easy to set-up a Mac OS only TCP/IP network, but could create additional work for a network administrator in more typical heterogeneous TCP/IP networks.

The Internet community (the IETF) has since developed a multivendor standard for the dynamic assignment of IP addresses, known as Dynamic Host Configuration Protocol (DHCP). Apple has adopted the industry standard DHCP and dropped support for the earlier Apple "Dynamic" mode addressing with Open Transport/TCP.

Question: What is MacIP?

Answer: MacIP, sometimes also referred to as KIP (Kinetics Internet Protocol), is a protocol specification developed as a method for carrying TCP/IP traffic on AppleTalk only networks - originally these would have been LocalTalk networks. MacIP is today frequently used in conjunction with AppleTalk Remote Access Protocol (ARAP) to provide mobile users access to TCP/IP network services. MacIP specifies encapsulation of TCP/IP datagrams in AppleTalk packets for transmission over such connections.

Use of MacIP requires a gateway. AppleTalk encapsulated IP packets are sent to the gateway using AppleTalk protocols (DDP). The gateway strips off the encapsulation and places the IP packet on the TCP/IP LAN. When packets are destined for a MacIP end-node, the gateway provides the needed encapsulation services.

MacIP gateway support is most frequently offered as an integrated service within a multiprotocol router. The gateway (router) attaches to both an AppleTalk and a TCP/IP network, acting as a middleman between the MacIP end-node and the appropriate TCP/IP based hosts on the LAN or WAN.

Open Transport supports MacIP end-nodes. It is selected using the TCP/IP configuration utility by choosing "AppleTalk (MacIP)" in the "Connect via:" pop-up menu. The user (or network administrator) must also specify which zone contains the desired MacIP gateway. Once selected, TCP/IP will be encapsulated in AppleTalk and will be sent to the gateway via the NIC selected using the AppleTalk configuration utility.

Question: How is MacIP support improved with Open Transport/TCP?

Answer: Open Transport/TCP offers new features in the human interface for

selecting the MacIP server, including:

- AppleTalk zones are now displayed in a scrolling list in a movable window. This display is easier to view compared to MacTCP's pop-up menu, especially when there are a large number of zones in the network.
- The Zone list window now supports selection using the mouse, the arrow keys, and/or "type-select", allowing the user to more quickly select a specific zone from the list.
- There is an option to display only those AppleTalk zones containing MacIP servers. When selected, this creates a background search task which when completed filters the zone list display to show only those zones containing active MacIP servers.
- There is a short cut "Current Zone" option which causes the Mac to check the current AppleTalk zone for a MacIP server without requiring the user to select a specific zone. This can be a time-saver for the user and a potential bandwidth-saver on the network, especially when there are mobile users that connect in different locations to a enterprise-wide network for MacIP services.

This article is one of many available through the Apple Fax center. For a complete list of available fax documents, search the Tech Info Library for Apple Fax Document Index or call the Apple Fax line at 1-800-505-0171 and select document number 20000 (Apple Fax - Document Index - Product Support Literature).

The Apple Fax center is available free of charge 24 hours a day, 7 days a week.

Article Change History:

- 02 Jul 1996 - Added Fax Doc word
- 20 May 1996 - Corrected reference to Ethernet SNAP.
- 08 Mar 1996 - Changed distribution status.
- 04 Mar 1996 - Updated to latest information.

Copyright 1996, Apple Computer, Inc.

Keywords: kfax

=====
This information is from the Apple Technical Information Library.

19960924 07:31:19.00

Tech Info Library Article Number: 19133