



Tech Info Library

System 7.5, Version 7.5.3: Open Transp Tech Information (6/96)

Revised: 2/17/97
Security: Everyone

System 7.5, Version 7.5.3: Open Transp Tech Information (6/96)

=====

Article Created: 21 June 1996
Article Reviewed/Updated: 28 June 1996

TOPIC -----

This article is the Open Transport Technical Information document which comes with System 7.5, Version 7.5.3.

DISCUSSION -----

Open Transport Technical Information

About This Document

This document contains technical information about Open Transport that may be useful to network administrators and managers. You do not need the information in this document in order to use Open Transport on your Macintosh computer.

Contents

Introduction

Files Added by the Open Transport Installer

What's New in Open Transport

- Static and Dynamic AppleTalk Address Allocation
- Use of Parameter RAM

Open Transport TCP/IP Features

- DHCP Server Support
- DHCP Address Lease Support
- Windows NT Advanced Server Support
- Local HOSTS File Support
- Human Interface Design
- MacTCP "Server" Addressing Support
- MacTCP "Dynamic" Addressing Support
- MacIP Support
- PPP Connectivity

Application Compatibility with Open Transport

- Backward Compatibility
- Performance
- Large Datagram Support

Introduction

Apple Open Transport is the modern networking and communications subsystem for the Mac OS. It is based on industry standards and brings a new level of networking connectivity, control, and compatibility to MacOS systems, while preserving and enhancing the hallmark of the Macintosh and Mac OS built-in support for easy-to-use networking.

Open Transport is designed to make it easier and more cost-effective to develop Macintosh-based applications for a wide variety of customers and markets. With Open Transport, Mac OS has built-in networking and communications based on cross-platform industry standards including the POSIX compliant X/Open Transport Interface (XTI), UNIX STREAMS and Data Link Provider Interface (DLPI).

Applications written to support Open Transport can directly support a wide range of networking environments (serial, dial-up network, LAN, and WAN), and multiple protocols (AppleTalk, TCP/IP, serial, and others) from a common code base. This capability is sometimes referred to as transport independence.

Open Transport 1.1 includes AppleTalk and TCP/IP protocols, and developer access to serial communications. Apple and third parties are working to add support to Open Transport for Point to Point Protocol (PPP), NetWare (NCP/IPX), Windows 95 (SMB/TCP/NetBIOS), DECnet, LAT, and X.25. Some of these additional capabilities may be incorporated or bundled with future releases of Apple Open

Files Added by the Open Transport Installer

The Open Transport Installer adds the following Extension files to Mac OS System Folder:

Open Transport Library contains the Open Transport code resource for 680x0-based Macintosh systems.

Open Tpt AppleTalk Library contains the code resource for AppleTalk communication protocol for 680x0-based Macintosh systems and AppleTalk compatibility support.

Open Tpt Internet Library contains the code resource for TCP/IP communication protocol for 680x0-based Macintosh systems and MacTCP compatibility support.

OpenTransportLib contains the code resource for PowerPC-based Macintosh systems.

OpenTptAppleTalkLib contains the code resource for AppleTalk communication protocol for PowerPC-based Macintosh systems.

OpenTptInternetLib contains the code resource for TCP/IP communication protocol for PowerPC-based Macintosh systems.

Ethernet (Built-In) contains the code resource allowing access to a built-in Ethernet port.

Serial (Built-In) contains the code resource allowing access to built-in serial ports.

Open Transport Guide Additions is an Apple Guide database that adds Open Transport instructions to Macintosh Guide

Open Transport also installs the new AppleTalk and TCP/IP configuration utilities (control panels) into the Control Panels folder in the System Folder.

What's New in Open Transport

The new Open Transport/AppleTalk and Open Transport/TCP both have been implemented as Open Transport STREAMS modules and as native code on Power Macintosh computers. They support the new XTI APIs, and their shared libraries can be dynamically loaded and unloaded as needed.

Both protocols also support dynamic reconfiguration (changed settings without requiring reboot), and feature new configuration applications offering Basic, Advanced, and Administrator tools. The new configuration applications -- AppleTalk and TCP/IP -- replace the older control panel implementations -- Network, MacTCP, and AdminTCP. To enhance ease of use and backward compatibility, the new applications are stored in the Control Panels folder in the System Folder.

Static and Dynamic AppleTalk Address Allocation

Open Transport/AppleTalk includes support for assigned (manually administered) protocol addresses. This allows AppleTalk nodes to be managed using protocol address as a unique and stable identifier. It can also reduce some of the network traffic associated with AppleTalk's dynamic address assignment features (AARP).

Dynamic addressing continues to be available for those customers who prefer the automated address allocation.

Network managers who prefer to have the network infrastructure automatically assign unique protocol addresses can continue to rely on AppleTalk Address Resolution Protocol (AARP). Network managers who find advantages in having fixed and well-known protocol addresses for each end-node can implement manual addressing.

When manual addressing is selected it is necessary to allocate and assign the initial protocol addresses, which will subsequently be "locked". Some administrators may prefer to do this allocation based on a central numbering plan, creating individual configuration templates (recommended or required settings) for each user. Others may prefer to allow the network to determine the initial address configuration (i.e., use dynamic addressing once), and then lock the uniquely assigned addresses after initialization.

It is important that all nodes on each individual AppleTalk subnet (a given

cable segment assigned a unique network number or network number range) be administered consistently -- either all with dynamic addressing or all with pre-assigned static addresses. This avoids a potential conflict between a new dynamic node acquiring an address assigned to an off-line, manually-addressed node. Administrators can enforce the addressing policy for a subnet by locking the addressing mode in the "dynamic" or in the "manual" state. As an administrative precaution, however, Open Transport/AppleTalk does continue to check for the presence of duplicate protocol addresses on the LAN when static addressing is configured.

Use of Parameter RAM

Under the classic AppleTalk networking architecture, AppleTalk's ON/OFF state, the selected network interface, the previous network (protocol) address, and the previous AppleTalk zone name were saved in persistent memory (parameter RAM) for reuse at boot time. To ensure backward compatibility, this information is still stored and retrieved on systems using Open Transport/AppleTalk. However, there are some changes made with Open Transport to accommodate the expanded capabilities of multiple, saved configuration files, and required network settings.

- At boot time, Open Transport reads the current AppleTalk configuration file to determine if AppleTalk should be set to ON or OFF. This value will override the value saved in parameter RAM.
- If the network interface specified in the current AppleTalk configuration file is locked (it is a required setting) and if the specified port is not available or cannot be initialized, AppleTalk compatibility will not automatically switch the port back to LocalTalk. Instead, AppleTalk will remain OFF. The user will receive a dialog notification in the event this occurs.

Open Transport TCP/IP Features

With the broad adoption of TCP/IP -- and the tremendous excitement and 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 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 version 2 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; and,
- multiple IP routers with fail-over, for increased robustness in mission critical applications.

DHCP Server Support

Apple's implementation conforms to the current versions of the applicable specification documents (RFCs). To date, Open Transport/TCP has been tested with the following DHCP server implementations:

- Competitive Automation,
- FTP Software (<http://www.ftp.com>),
- Hewlett Packard HP-UX (<http://www.hp.com>),
- Microsoft Windows NT Advanced Server,
- Silicon Graphics (<http://www.sgi.com>),
- Sun Solaris and SunOS (<http://www.sun.com>), and
- TGV (<http://www.tgv.com>).

DHCP Address Lease Support

Open Transport/TCP fully supports DHCP address leases. Open Transport/TCP will automatically attempt to renew any address lease that reaches its Renewal Interval, which defaults to half of the lease's lifetime. The Renewal Interval may be configured to a different value by making changes to the configuring DHCP server. Renewal will be attempted regardless of how many times the lease has already been renewed. Should an interface's IP address lease expire, the interface will be closed down.

Windows NT Advanced Server Support

With Open Transport 1.1, Mac OS clients are fully interoperable with the Windows NTAS DHCP server.

Macintosh clients running earlier versions of Open Transport (1.0.x) could experience some interoperability problems due to significant differences between the Microsoft implementation and that of a typical UNIX-based server.

Local HOSTS File Support

Open Transport/TCP supports a 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. This file is parsed when Open Transport/TCP is initialized. As in MacTCP, the supported HOSTS file features follow a subset of the Domain Name System Master File Format (RFC 1035).

Should a HOSTS file be 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, Open Transport/TCP must be configured using either the Advanced or Administrator mode to select the appropriate file. The

text file must already exist, and can be created with any text editor or word processor. Also note that the HOSTS file selection is tied to the selected configuration. An administrator might, for example, specify different HOSTS files for use when a user connects via Ethernet on the campus LAN and that same user when dialing-in from a remote location.

Human Interface Design

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, etc.), 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.

MacTCP "Server" Addressing Support

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 uses BootP to attempt to acquire an IP address. If BootP fails to provide a valid address, MacTCP 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 user's point of view, it also adds a degree of unpredictability from a network administrator's point of view.

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

MacTCP "Dynamic" Addressing

Open Transport 1.1 does not support MacTCP "Dynamic" addressing. MacTCP "Dynamic" mode addressing was based on an Apple-proprietary extension to TCP/IP protocols. It applied the address negotiation and assignment rules used by the AppleTalk protocols to TCP/IP networks, making it very easy to set-up a Macintosh-only standalone TCP/IP network. Use of this Dynamic Addressing method in other scenarios, however, could create additional work for a network administrator.

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 "dynamic" mode addressing with Open Transport/TCP.

MacIP Support

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 LocalTalk networks. MacIP is today frequently used 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 MacIP gateway using AppleTalk protocols (DDP). The gateway strips off the AppleTalk encapsulation and places the IP packet on the TCP/IP LAN. When packets are destined for the MacIP end-node, that 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 includes end-node support for MacIP. A end-node is configured to use MacIP using the TCP/IP configuration utility by selecting "AppleTalk (MacIP)" in the "Connect via:" pop-up menu. The user (or network administrator) must also specify where on the network (in which zone) to look for the MacIP gateway. Once selected, TCP/IP will be encapsulated in AppleTalk and will be sent out the "Connect via:" interface selected using the AppleTalk configuration utility.

Open Transport/TCP includes a new human interface for selecting the MacIP server which offers the following new features:

- AppleTalk zones are now displayed in a true scrolling list in a movable window. This display is easier to view compared to MacTCP's pop-up menu 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.
- Network Interface Options

PPP Connectivity

PPP (Point to Point Protocol) connectivity is not bundled with Open Transport at this time. AppleTalk and TCP/IP connectivity over Open Transport/PPP is planned

to first be offered as a feature of upcoming Apple Remote Access products. ARA products would include necessary Open Transport components.

Apple currently plans to fully merge ARA client and personal server capabilities with basic Open Transport capabilities to offer an integrated communications package for LANs, WANs, and remote networking. These integrated capabilities are also expected to be delivered as a part of a future update to Mac OS. This timetable has not been finalized; details will be announced at a later date.

Application Compatibility with Open Transport

Apple has defined three levels of interoperability with Open Transport. The first -- known as Open Transport Compatible -- is used to describe a network application originally developed for "classic" AppleTalk or MacTCP, that now takes advantage of Open Transport Compatibility Services. These applications automatically gain the benefits associated with the new Open Transport configuration utilities. However, they will not realize a significant performance increase on Power Macintosh systems, nor can they take advantage of Open Transport's transport-independence capabilities.

Open Transport Ready applications are those that have been modified to adopt the new Open Transport APIs (XTI). They are PowerPC native, in addition to running on 680x0-based Macintosh systems. Open Transport Ready applications not only benefit from the new configuration utilities, but have the opportunity for a significant performance boost when running on Power Macintosh.

The third (highest) category of interoperability is referred to as Open Transport Enhanced. In addition to adopting the new Open Transport APIs and being Power PC native, these applications have been modified to exploit the transport-independent capabilities of Open Transport, i.e., they can be dynamically configured to support AppleTalk, TCP/IP, or serial communications.

Backward Compatibility

Applications that rely on undocumented APIs or examine private data structures in current the AppleTalk or MacTCP may not be fully compatible with Open Transport. Examples include the MacSNMP AppleTalk and TCP/IP Agents (however, MacSNMP and the Macintosh System Agent are compatible), the Apple Internet Router 3.x and some utilities like MacTCP Watcher and MacTCP Spy. Updated versions of these software products will be required for full compatibility.

Performance

Open Transport is written to take advantage of the PowerPC processor -- it is native code. This provides the necessary foundation for increased networking performance in Mac OS. To realize performance gains, however, it is equally important that networking applications also be accelerated for Power Macintosh, and that applications adopt the new Open Transport XTI-based programming interfaces.

The built-in "backward compatibility" support for existing AppleTalk and TCP/IP applications must continue to run as 680x0 code in emulation on Power Macintosh. This protects a customer's investment in networking applications, but also

obscures underlying performance increases from the native protocol implementations.

In general, current Mac OS networking applications are written for the 680x0 processor, and use the "classic" (68K-based) networking programming interfaces. These applications are not likely to see performance boosts with Open Transport, as most of the performance potential is based on the move to native code for the PowerPC processor. Even for Power Macintosh native applications, a continued use of the Open Transport backward compatibility libraries offsets most of the performance gains in the low level protocol handling.

Users who select Power Macintosh native applications that are Open Transport-ready will realize the greatest performance gains. Performance of specific network applications may also be significantly influenced by the processor speed of the system. Customers with demanding, network I/O intensive applications should give strong consideration to the higher performance PowerPC-based Macintosh systems.

However, even with 68K emulated applications using backward compatibility, TCP/IP users are more likely to see some performance improvements with Open Transport. This is because of the differences in the way compatibility is provided for MacTCP vs. AppleTalk, and differences in the two protocol architectures.

Performance will be greater with protocols that use larger datagram sizes -- such as TCP/IP -- than with AppleTalk (which has a fixed and limited datagram size). On high-speed datalinks such as fast Ethernet, FDDI, and ATM, the performance of the network interface card (NIC) driver code is also a significant factor. Comparative shopping for NICs -- based on price, service, reliability, and performance -- will be in order.

Open Transport running on the built-in Ethernet of the Power Macintosh 9500 has been clocked at 9.3 Mbps throughput using low-level TCP/IP benchmarks. A pre-release version of a third party Open Transport-native implementation of 'NFS' protocols was benchmarked at 8.4 Mbps. Both figures approach theoretical maximum performance for 10 Mbps Ethernet. AppleTalk performance is somewhat lower, with low-level benchmarking coming in at a bit over 7.5 Mbps throughput.

The Open Transport engineering team is continuing to work with NIC developers to realize high-performance DLPI drivers for high-speed datalinks. This is a cooperative effort, with work being done on both driver code and on Open Transport. We expect that high-speed datalink NIC drivers based on Open Transport 1.1 can be fully competitive with other PCI networking products. Of course, performance tuning will be an ongoing priority, as Apple intends to always offer a platform capable of industry leading network performance.

Benchmarking on these types of datalinks, with pre-release Open Transport 1.1, has been underway for some time. Some sample results include:

- 48 Mbps with a Rockwell fast Ethernet card and driver (1500 byte block size)
- 72 Mbps with a Rockwell FDDI card and driver (4K block size)
- 93 Mbps with an Interphase ATM-155 card and driver (8K block size)

These tests were performed using Open Transport/TCP 1.1 beta software, running on a Power Macintosh 9500/132, and reflect memory to memory, point to point transfers (the "SpudPPC" test tool) on a dedicated test bed. In addition to these numbers, we've seen even higher throughput from Fore Systems on their ATM card, using UDP.

AppleTalk performance is expected to be lower than TCP/IP performance due to various technical issues, including DDP packet size and the ATP retry-acknowledgment algorithm. Current testing on fast Ethernet is turning in figures around 35-40 Mbps with a PowerPC native ATP test tool.

Large Datagram Support

Maximum allowable datagram size is dependent on both the selected datalink and the selected protocol stack. Open Transport supports the use of large datagram sizes as appropriate to the protocol and datalink in use.

Because Open Transport/AppleTalk is based on the Phase 2 architecture, datagram size for AppleTalk is limited to a maximum of 617 bytes. Open Transport/TCP supports larger datagrams; up to 1,500 bytes on Ethernet and fast Ethernet, and up to 16K on token ring. Even larger block sizes can be used on FDDI and ATM.

Block size does play a role in maximum throughput; the larger the block size used, the greater the potential end-to-end throughput. Users demanding the highest network throughput may find FDDI to be a more attractive alternative than fast Ethernet because it can support larger block sizes at the same bit rate.

Article Change History:

28 Jun 1996 - Change status at product release.

Copyright 1996, Apple Computer, Inc.

Tech Info Library Article Number:20088