



Tech Info Library

Open Transport 1.1: Performance Q & A (3/96)

Revised: 3/8/96
Security: Everyone

Open Transport 1.1: Performance Q & A (3/96)

=====
Article Created: 4 March 1996
Article Reviewed/Updated: 8 March 1996

TOPIC -----

This article is the Reference Q & A (questions and answers) on performance for Open Transport 1.1.

DISCUSSION -----

Question: Is Open Transport native on PowerPC Mac OS? Does this make networking faster?

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

The compatibility services for existing AppleTalk and TCP/IP applications run as 680x0 code in emulation on Power Macintosh. This protects a customer's investment in networking applications, but also obscures - or in some cases, outweighs - the underlying performance increases from the native protocol implementations.

Question: Does Open Transport PowerPC native code include drivers for Macintosh onboard Ethernet?

Answer: PCI Power Macintosh systems ship with a PowerPC native DLPI driver for built-in Ethernet. Power Macintosh 6100, 7100, and 8100 models currently have 680x0 drivers.

Question: Will existing networking applications see performance improvements

with Open Transport on PowerPC Mac OS systems?

Answer: In general, current Mac OS networking applications are written for the 680x0 processor and use the classic networking programming interfaces. These are not likely to see performance boosts with Open Transport, as most of the performance potential comes from the move to native code for the PowerPC processor. Even for PowerPC native applications, use of the backward compatibility libraries offsets most of the performance gains in the low level protocol handling.

Users that select PowerPC 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 underlying processor speed of the system. Customers with demanding, network I/O intensive applications should give strong consideration to the higher performance PowerPC Mac OS systems.

However, even with existing applications using backward compatibility, TCP/IP users are 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.

Question: Will networking applications see performance improvements with Open Transport on 680x0 based applications?

Answer: Not in general, as most of the potential for increased performance with Open Transport comes from the move to PowerPC native code. However, users may find that Open Transport TCP exhibits superior performance to MacTCP, especially under adverse networking circumstances (slow links, lossy lines), due to its more robust implementation and more sophisticated error handling and recovery.

Question: When will new or updated applications that support the native Open Transport APIs become available?

Answer: Applications that are PowerPC native and Open Transport ready are available now. Users are urged to contact the specific third party vendor of interest for more details on their specific products.

Question: How is Open Transport performance being measured?

Answer: Apple has established plans for measuring the performance of Open Transport and related system components through four benchmark test suites:

- SpudPPC - this low-level benchmark tool focuses on measuring the raw throughput potential of Open Transport. It supports both AppleTalk and TCP/IP protocols, is PowerPC native code, and uses Open Transport programming interfaces. Because it measures point-to-point, memory-to-memory data transfers, it most directly measures the performance of Open Transport.

Because this test has the most direct access to Open Transport (the application

layer is "thin"), and because it is fully accelerated for PowerPC, this benchmark will generally indicate an upper bound on the performance potential of Open Transport.

- AppleShare file copy - this end-user benchmark focuses on measuring the throughput of the AppleShare client while drag-copying a file from the Mac OS desktop using the Finder. It is specific to the AppleTalk protocol suite, runs in 680x0 emulation, and depends upon backward compatibility services to access Open Transport networking. Because it measures user-perceived throughput of a complete application chain, this test only indirectly measures the performance of Open Transport.

Because this test depends upon emulated code, backward compatibility, and AppleTalk protocols, this benchmark will generally indicate a lower bound on the performance potential of Open Transport.

- Fetch 3.x - this end-user benchmark focuses on measuring the throughput of the popular ftp client, Fetch. It is specific to the TCP/IP protocol suite, runs as PowerPC native code, and uses Open Transport programming interfaces. Because it measures user-perceived throughput of a complete application chain, this test only indirectly measures the performance of Open Transport.

Because this test is PowerPC native, Open Transport ready, and is based on TCP/IP protocols, this benchmark will generally tend toward the upper bound on the performance potential of Open Transport.

- ZD Labs NetBench 4.0 - this suite of benchmarks is designed to test file server implementations. It is specific to the AppleTalk protocol suite, runs in 680x0 emulation, and depends upon the File System and backwards compatibility services to access Open Transport networking. Because it measures user-perceived throughput of a complete client-server environment, this test only indirectly measures the performance of Open Transport.

Because this test depends upon emulated code, backward compatibility, and AppleTalk protocols, this benchmark will generally indicate a lower bound on the performance potential of Open Transport. However, because it interacts with an AFP server - which may be PowerPC native and Open Transport ready - it can be useful in measuring the multi-client scalability of file server implementations built on Open Transport.

Only a combination of tests can provide good coverage, as user-perceived networking performance is heavily influenced by a combination of a number of Mac OS components including the file system, the Finder, driver code, and the applications used to move data across the network.

Question: How much faster will native Open Transport applications be?

Answer: Networking performance is influenced by many factors. As noted elsewhere in this document, customers will see the highest performance when using PowerPC native applications that fully support Open Transport APIs.

Performance potential 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, or ATM, both the performance of the network interface card (NIC) driver code and the number of allocated buffers are significant factors.

Open Transport has been clocked at over 9.3 Mbps using the SpudPPC tool. A pre-release version of a third party implementation of NFS was benchmarked at over 8.4 Mbps. Both figures approach theoretical maximum performance for 10 Mbps Ethernet.

Question: What about high-speed networking connections like fast Ethernet, ATM, and FDDI?

Answer: Benchmarking on fast Ethernet, FDDI, and ATM datalinks 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, using the SpudPPC tool. Other tests, such as the one conducted by Fore Systems on their 155 Mbps ATM cards, have shown even higher throughput. In one test, Fore was able to transmit UDP over their LAN-E stack on ATM (using 1500 byte datagrams) at over 100 Mbps.

AppleTalk performance is lower than TCP/IP performance due to the smaller DDP packet size and the ATP retry-acknowledgment algorithm. Current testing on fast Ethernet is turning in figures around 35-45 Mbps with a PowerPC native ATP test tool. This is a significant performance improvement, and further progress should be realized as application code is revised to take full advantage of Open Transport and PowerPC.

Question: Does Open Transport support the use of large datagram sizes? Does datagram size have an impact on network throughput?

Answer: 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:

08 Mar 1996 - Changed distribution status.

Copyright 1996, Apple Computer, Inc.

Tech Info Library Article Number:19436