Data Expedition, Inc.®

A Brief Overview

of

Multipurpose Transaction Protocol®

technology

MTP™/IP

September 2008

877-292-2280

781-353-6403

http://www.DataExpedition.com/

U.S. Patents 7158479, 7313627, 7404003
Additional Patents Pending

Table of Contents

Summary

Data Expedition, Inc.® (DEI®) has implemented software technology that dramatically improves the performance and reliability of network data transportation. Applications built with DEI's Multipurpose Transaction Protocol® technology (MTP™/IP) see the following benefits relative to TCP/IP transport:

- Throughput on wide area networks is increased by two to seven times
- Transaction rates are increased by two to ten times
- Reliability is increased ten times over traditional TCP mechanisms.

MTP is fully compatible with existing network standards, requires no changes to the network or operating systems, and is transparent to the end user. Use of MTP provides network application developers and IT managers many benefits:

- Users spend less time waiting for the completion of large data transfers.
- Customers process a higher volume of transactions.
- Infrastructure and operations costs are greatly reduced.
- Applications make better use of existing network capacity, resulting in a deferred need to upgrade those connections.
- Applications place less load on their servers, lowering the cost of server infrastructure to support them.
- Applications can support more robust and varied features as a result of not having to accept the current limitations of the network.
- Critical functions are made more reliable in the face of network outages, third-party congestion, and denial of service attacks.

Network performance has become a critical issue for both network users and software developers. Even the most well designed data network systems are critically vulnerable to congestion problems. Traditional techniques for mitigating congestion and maximizing performance either require expensive end-user investments, or are limited to narrow application types.

This document explains why MTP is needed, what it does, and how developers and IT managers can make use of this technology.

Background

Industry Problem

Growing demands for high performance in data networks, including the Internet, private networks, and semiprivate Wide Area Networks, is a persistent and growing problem for both network users and the producers of network products.

> "The volume of new applications is saturating wide-area networks, causing numerous performance problems for mission-critical traffic. ... Enterprises must begin to focus on the optimization of their data networks."
>
> – META Group, June 2003
> "Bandwidth Compression and Optimization"

The demand for network performance is continuing to increase at an exponential rate. While this demand growth was once driven by increasing numbers of Internet users, several factors continue to accelerate demand even as the number of users is leveling off.

- The deployment of broadband access to home, office, and mobile devices continues to increase per-device bandwidth demand.
- The integration of wireless data services into a wide variety of devices and locations is further increasing the per-person usage of the Internet.
- Migration of traditionally local network applications to wide area networks, as driven by security and distributed workflow demands, is leading to new expectations for WAN performance.

Businesses are becoming critically dependent on all forms of data transport, from front-end public access, to back-end interoffice communication. Network applications, content, and expectation are growing more performance hungry.

> "It will be more economical for enterprises to improve WAN performance by using WAN optimization technologies, rather than adding bandwidth."
>
> – Gartner, December 2003
> "Enterprises Will Waste Money on Bandwidth in 2004"

The principal symptom of bandwidth demand exceeding supply has been chronic under performance. For the end user, under performance means slow transfer speeds, broken connections, and a loss of functionality. For businesses, slow or unreliable network service cuts into productivity, limits expansion, and increases the cost of operation. The perception that the wide area networks are too slow or unreliable presents a serious challenge to the developers of network software and solutions. There is a tremendous demand for even traditional applications to communicate and interact across wide area networks. Many new features and functionalities require reliable access to network bandwidth. But software developers have had no direct means to cope with network congestion.

Traditional Solutions

The most obvious solution for end-users is to purchase additional data capacity in the form of leased lines or private WANs. But such dedicated infrastructure can more than double bandwidth costs over that of an Internet based virtual private network, and still does not guarantee freedom from congestion.

> "Adding bandwidth is an expensive option that often does not solve the problem at hand."
> – Aberdeen Group, January 2004
> "WAN Optimization Techniques Improve Application Delivery"

As a result, several mechanisms have been developed to manage network performance. Such techniques include,

Mirroring:	Pre-distributing content to the network fringes,
Caching:	Preserving recent content near the network fringes,
Balancing:	Spreading the load across redundant paths and servers,
Prioritization:	Increase performance for some at the expense of other traffic, and
Compression:	Reduce the amount of data being sent.

Unfortunately, each of these solutions has severe limitations. Mirroring and caching are only effective for the one-way distribution of static content. Load balancing requires a client-server environment with central control and strong infrastructure investment. Prioritization is a zero-sum tradeoff and assumes a closed network. Compression has high overhead, can compromise data quality, and only works if there is significant redundancy in the data. Most existing mechanisms involve substantial setup and maintenance costs.

Underlying Problem

Traditional performance optimization techniques focus on demand mitigation: reducing the amount of data being transported without actually increasing the net carrying capacity of the underlying network. As a result, traditional solutions are only effective on those limited data types and use scenarios where the data volume can be satisfactorily reduced. Such techniques offer little or no benefit for the transport of original content that is already encoded or compressed.

At the core of this problem is TCP/IP, which is used to carry 99% of Internet traffic. The Transmission Control Protocol was created in 1974 to manage the reliable transportation of data between remote applications. Even relatively modern protocols, such as RTTP for audio/video streaming, have their designs deeply rooted in TCP concepts. TCP was designed based upon the characteristics of computer networking and applications which existed thirty years ago. The design assumed a network in which

- Congestion was very rare,
- The speed of the receiving computer was the limiting factor, and
- Very few users would be accessing the network at any one time.

These and many other assumptions fundamental to TCP's design are no longer true. Today's network is demand driven, with hundreds of thousands of simultaneous transactions, computers that can never receive data fast enough, and nearly continuous congestion. Even seemingly private networks are subject to variations in the underlying telecommunications infrastructure. Networks themselves now scale across a tremendous variety of technologies, resulting in orders of magnitude of variation in bandwidth, latency, and packet loss.

Numerous efforts have been made to "fine tune" TCP over the years. These efforts range from modifications to TCP's algorithms, to scaling adjustments which cause it to behave better under a narrow range of conditions. Several companies market variations of TCP derived protocols. But the fundamental limitations of the TCP data model severely limit the effectiveness of patchwork solutions. Tuned TCP stacks are neither scalable nor robust: they quickly break down when faced with real-world challenges. The Internet Engineering Task Force, which oversees TCP development, remains committed to making only incremental changes to TCP, but has stated that they will not make even this limited effort a priority until well after IPv6 is finalized.

DEI's Solution

Rather than trying to patch a fundamentally outdated design, DEI has chosen to address the data transport problem by creating a completely original and modern solution: the Multipurpose Transaction Protocol® technology, or MTP™/IP. MTP has been designed to handle the types of applications and network environments most prevalent in modern networking. Using patented and patents pending data models and flow control algorithms, MTP is able to avoid the debilitating effects of congestion while improving overall network performance and stability. As a result of this design, MTP transactions are able to make much better use of existing infrastructure. Tests of MTP software in real-world network environments have shown substantial performance improvements over TCP:

WAN Throughput:	Two to Seven times or more
Transaction rates:	Two to Ten times or more
Reliability:	Up to Ten times
Scalability:	Unlimited.

When MTP is used in a predominately TCP congested WAN environment, the speed of data transfer is typically improved by two to seven times. Transactions can be performed with no setup or tear-down costs, and very low server CPU utilization. This allows MTP to achieve two to ten times as many transactions per second in any network environment. Sophisticated congestion control and error recovery allows MTP to achieve a much higher degree of reliability than TCP. It is able to provide reliable service under packet loss conditions up to ten times worse than the point at which TCP simply fails to operate.

Unlike variations of TCP, MTP's results are sustainable, and scaleable. Unlike constant-rate UDP data-blasters, MTP adjusts to network conditions and does not make ambient congestion worse. And unlike hardware based bandwidth management solutions, MTP can be deployed with minimal cost and zero configuration.

MTP's performance advantage over TCP increases as congestion, concurrency, and path length grow. These advantages often mean that end-users can achieve the same or better performance using MTP over a VPN as they would using TCP over a private network, for half the cost.

Case Study: Motorola

At Motorola's *Advanced Product Technology Center* (APTC), located in Plantation, Florida, engineers manage a network of high-performance computing facilities sited around the world. These design centers run data intensive simulation and modeling software to assist Motorola and its partners in the development and testing of new technologies and products. With gigabytes in each typical data set, file transfer between end-users and the design centers became part of the product development critical path.

The problem faced by Motorola's engineers was that the 45 megabit per second leased network lines connecting the facilities were not living up to their performance potential. Users were spending hours each day waiting for data to move, time that could have been better spent on successive model runs.

That changed in December of 2004, when Motorola's *Commercial Government and Industrial Solutions Sector* (CGISS), parent of the APTC facility, purchased five site licenses for the *ExpeDat*™ file transfer software, from Data Expedition, Inc.® ("DEI®"). Based on DEI's Multipurpose Transaction Protocol® technology (MTP™/IP.), *ExpeDat* provided a drop-in replacement for Motorola's existing network of FTP servers and clients.

Engineers immediately observed a **six times reduction in file transfer times**. As *ExpeDat* has been deployed to Motorola facilities in the US and Asia, end-users are spending one half to one sixth as much time waiting for their modeling data. According to Motorola engineers, this is the **equivalent of increasing their total system investment by 40%**, at just a fraction of the cost.

This new-found ability to move data quickly and reliably is inspiring a quest for more ways to streamline data intensive tasks and improve the efficiency of Motorola's design centers. As DEI continues to improve its *ExpeDat*, *HyperGate*, and Software Development Kit products, we expect to bring even more savings to Motorola and all our customers.

How MTP/IP Works

Multipurpose Transaction Protocol® software (MTP™/IP.) is a transport protocol (OSI layer 4) that is capable of moving data much more quickly and efficiently than the traditional TCP transport protocol. MTP is able to achieve these results by taking a much more conservative and informed approach to flow-control, error recovery, and data modeling. The result is improved performance and reliability that is both robust and scalable.

The figure below illustrates TCP's oscillation behavior when attempting to move data across a network backbone.

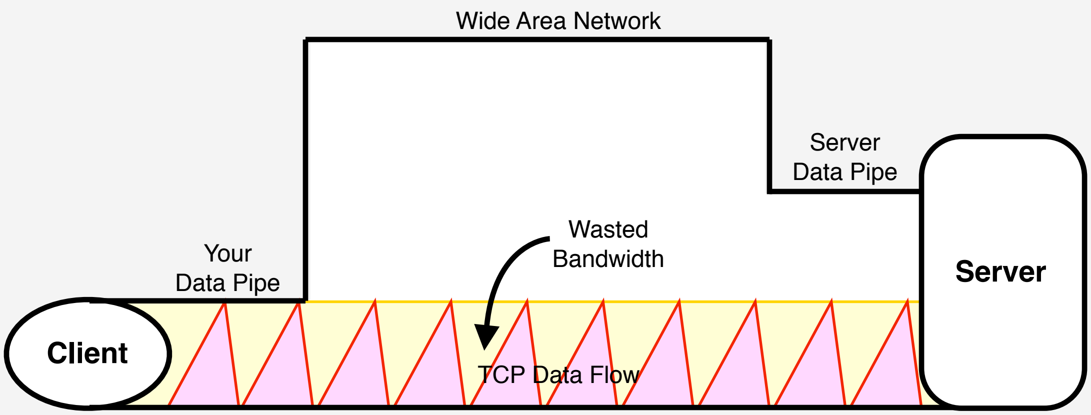

TCP can have trouble scaling, even on clear networks.
This leaves bandwidth underutilized.

TCP's primary mode of operation is to push increasing amounts of data into the network until congestion forces packet loss. At that point TCP pulls back and waits, then once again begins flooding the network. This creates the oscillations illustrated above in red. The result is that only a fraction of the available bandwidth is actually being used. This is represented by the white space above. Most modern TCP implementations include features to reduce this effect, but they are still subject to unstable flow and oscillation, particularly when TCP is operating across long, high latency, or lossy data paths.

The problem is greatly magnified when many TCP data streams interact along a network path. In a process known as "synchronicity", many different TCP streams begin to oscillate together. This causes the periods of congestion to be much worse, and the periods of underutilization to be much more pronounced. The figure below shows how one TCP stream is adversely affected by the presence of others, even though there is still "room" in the data pipe.

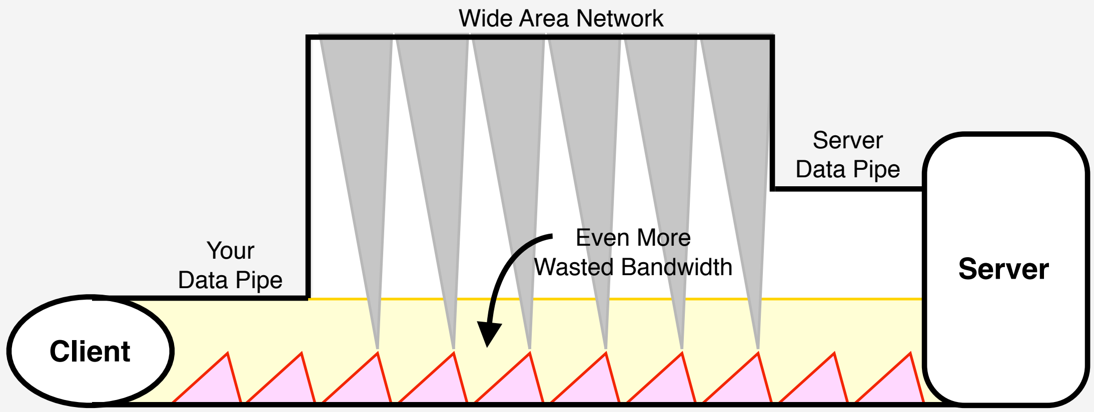

WAN traffic makes the problem worse.

Notice that the third-party traffic (shown in grey) is oscillating in synchronicity with the highlighted TCP data flow (shown in red). The coincidence of these congestion spikes (shown here as occurring along the backbone path) forces TCP to pull back much sooner than it otherwise would. In between the congestion spikes, the data path is underutilized, resulting in TCP data rates that are far below hardware capacity. Moreover, the constantly recurring congestion causes indiscriminate packet loss, which can further delay connections or disrupt connectivity.

MTP, on the other hand, is much more careful about how it adjusts its data flow. MTP observes network behavior to find the maximum sustainable data rate of the current path, and constantly adjusts as network conditions change from moment to moment. The following figure shows MTP operating in a TCP dominated environment.

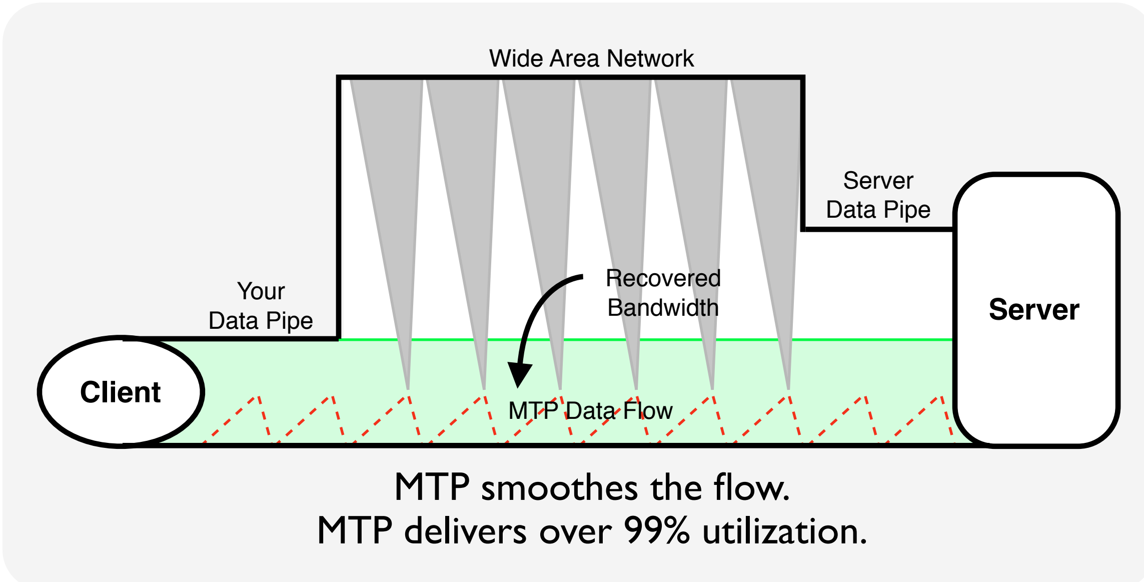

Note that unlike some UDP products which attempt to send data at a constant rate, MTP *does slow down* as the third-party traffic spikes. Failing to do so would only contribute to the congestion problem and increase packet loss. Instead, MTP rides down the congestion, then quickly recovers to take advantage of available bandwidth. Should the data path become truly over-utilized, MTP will quickly respond to keep from exacerbating the congestion as it constantly reevaluates network conditions.

This more conservative approach to flow control allows MTP to fully utilize the bandwidth that is available to it. It cannot send data any faster than the smallest pipe in the path, but unlike TCP, MTP can fill that pipe even as network conditions continue to change. MTP's throughput advantage comes from making efficient use of the resources that TCP wastes.

Application Benefits

MTP's unique design allows it to bring several key benefits to networked applications. In the previous chapter, we discussed how MTP's improved flow-control results in better throughput. Here we look at the benefits this provides plus the benefits that come from MTP's lower processing overhead and its ability to tolerate adverse network conditions.

Throughput

As discussed above, TCP tends to oscillate its data flow, resulting in an underutilization of network resources. MTP's more conservative algorithms produce a smoother data flow, which allows the previously wasted bandwidth to be recovered.

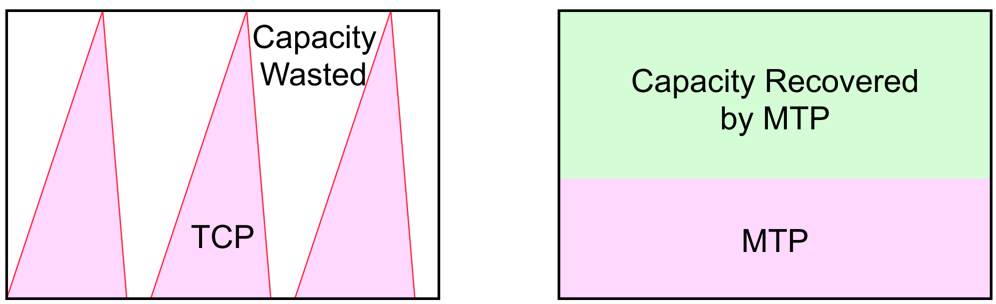

The increased throughput means that data dependent tasks will take less time. This reduction in waiting time improves productivity, increases service capacity, and brings large data exchanges closer to real-time. Another benefit is a reduction in bandwidth costs. With TCP, data pipes often must be oversized to compensate for the congestion which occurs at even low utilization. But with MTP, end users require less capacity to achieve the same performance.

Turnaround

TCP's design requires that it perform several time-consuming tasks in order to set up and tear down each connection. During the connection itself, TCP concentrates most of its processing overhead on the data server, where resources are the most scarce. This extra overhead increases the amount of time required for each data transfer and takes up valuable server side CPU cycles and memory. These factors become critical when processing many small data exchanges, such as for database queries, financial transactions, or network searches. MTP nearly eliminates setup and tear down time, and it greatly reduces the overhead burden placed on data servers.

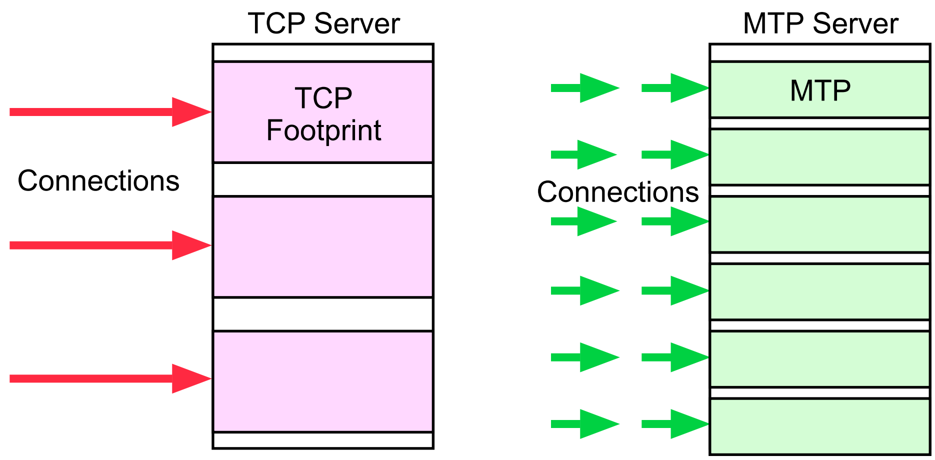

By reducing both the time and the resources needed to process transactions, MTP allows applications to achieve much higher turnaround rates (transactions per second) and reduces the need to spend money on expensive server hardware. This improved efficiency also allows software developers to consider new features and functionalities that are not practical with TCP.

Path Tolerance

TCP exhibits a well known aversion to long data paths, even without congestion issues. As path lengths increase, TCP's performance decreases. This is especially critical for very high bandwidth networks where even a few hops can begin to degrade TCP's performance. TCP is also very sensitive to packet loss. Drop rates as low as 5% can completely shutdown a TCP data flow. MTP is capable of scaling to paths of any length and it can continue to operate with drop rates as high as 50%. The following figure illustrates how TCP's performance drops off as the number of routers in the path increases, while MTP's performance remains constant.

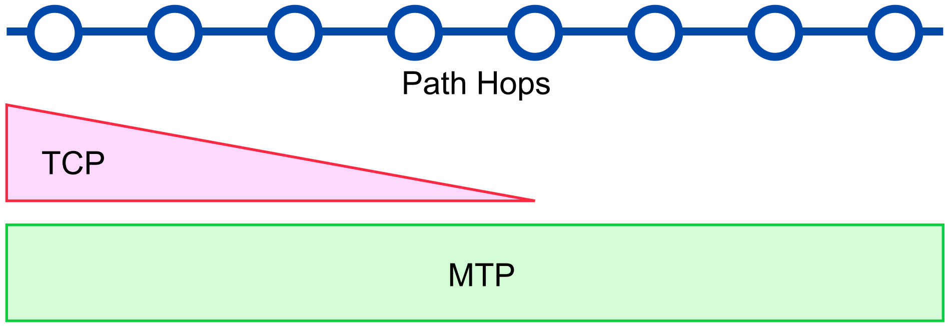

This path scaling difference has a profound effect on the costs for end-users who operate their own networks. With TCP, network complexity alone may limit its performance even if congestion is not an issue. To achieve high performance, end-users must often resort to privately piped, and very expensive, data distribution networks. MTP's ability to fully utilize existing Internet infrastructure allows it to achieve high performance without resorting to expensive private links. This can reduce connectivity costs by up to 50% without sacrificing performance.

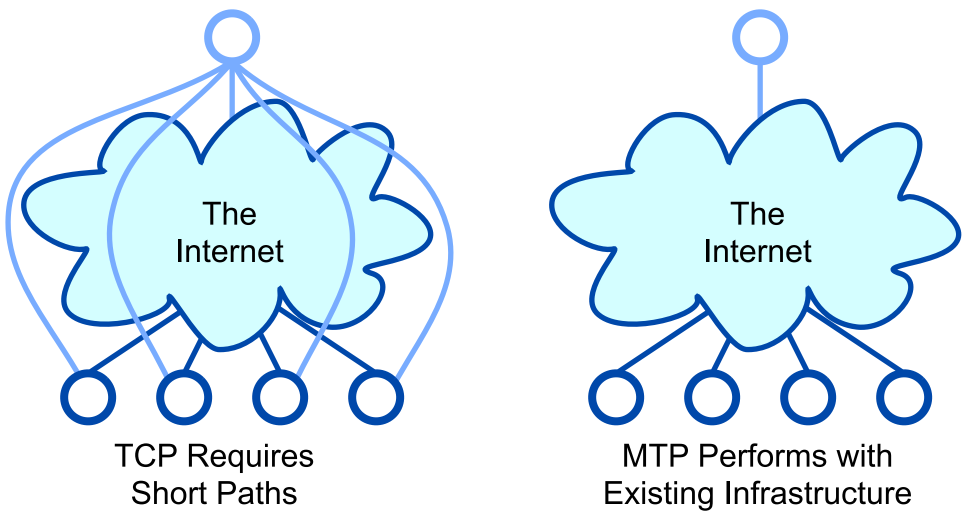

MTP/IP applications can offer faster throughput, higher volume, and reduced costs that no TCP application can match. With this new performance, developers and solution providers can offer new features and functionality the were not previously possible. Best of all is the fact that the end-user does not need to take any extra steps beyond what they would do for any network software. They don't even need to know MTP is there to gain the benefits of a modern transport protocol.

Featured MTP/IP Applications

ExpeDat™

ExpeDat applications allow you to transfer data across many different platforms with several times greater speed and reliability than traditional FTP.

Fast	Reliable	Secure
Up to seven times better thoughput than FTP on Wide Area Networks. Instant response for directory listings and small downloads.	*Ten times* better congestion tolerance. *ExpeDat* stays online when others fail. Your data gets through even when network loss is as high as *50%*.	Advanced Encryption Standard (AES) security built-in. 128-bit key encryption available in US and select regions.

ExpeDat comes with both command line and graphical clients as well as an enterprise class file server. The server and command line clients are lightweight for easy installation, embedding, and scripting. Additional features:

Performance

- Up to seven times faster WAN speed
- Instant file browsing
- File sizes up to 1 Terabyte

Ease of Use

- Graphical and Command Line interfaces
- Zero-Configuration server deployment
- Support for all major platforms

Security

- AES Encryption
- User authentication
- Private, shadow, PAM, and NIS password support

Configuration

- Bandwidth regulation and management
- Scriptable command line client
- Server-side packaging and compression plug-ins

See the “Case Study” chapter for an example of *ExpeDat* deployment. Visit **http://www.DataExpedition.com/expedat/** for a demonstration or a free trial.

SyncDat™

SyncDat is a high-performance file replication and synchronization application. It rapidly compares and transfers files for backup, data migration, and mirroring.

Fast	Scalable	Flexible
Scans hundreds of thousands of files in minutes. WAN performance orders of magnitude beyond *rsync* and CIFS solutions.	Performs in any IP network environment, even congested WANs. No performance tuning required.	Synchronizes across all major platforms. Lightweight client and server. Easily scripted or embedded

Most synchronization software relies on LAN based file sharing protocols like CIFS or rsync. SyncDat is an MTP/IP application, capable of rapidly scanning and transferring files, even over congested wide area networks. Just name two directories, one local and one remote. SyncDat figures out what to do and does it fast. Additional features:

Performance

- Rapid directory scanning
- Up to seven times faster throughput
- Minimal server I/O load

Automation

- Detects changes and conflicts
- Move thousands of files at once
- Easily scripted or embedded

Security

- AES Encryption
- User authentication
- Private, shadow, PAM, and NIS password support

Configuration

- Sync master-slave or peer-to-peer
- Scriptable command line client
- Bandwidth management built-in

Visit **http://www.DataExpedition.com/syncdat/** for a free trial.

Building with MTP/IP

All MTP/IP Software Development Kits are written in plain C and distributed as static code libraries. As such, they should be compatible with any modern C, C#, or C++ compiler. The table below shows our current SDKs.

SDK	Applications	Description
MTP Core	High speed transaction processing, data servers.	Core MTP is a reliable, modular, request-response transaction protocol with support for static and dynamic content. Writing applications directly on top of the Core SDK affords maximum efficiency, flexibility, and control.
Document Exchange	File transfer, peer to peer, document exchange.	This high-level API provides applications with the ability to perform authenticated document transmission and retrieval. Little or no knowledge of network processes is required.
Stream	Dynamic data generation, TCP API adaption.	Emulating the sequential access of traditional FIFO API's, this package allows dynamic content generation to take advantage of MTP's performance. The stream API can also be used to provide quick integration of MTP into TCP legacy applications.
ExpeDat Client	Custom development of clients for the ExpeDat server.	Any application can fully integrate the functionality of an ExpeDat client. Allows direct access to the features of the ExpeDat server, including authentication, file management, and packaging.

MTP SDKs can be ported to any platform which supports UDP/IP and a modern C compiler. The most suitable targets are POSIX environments capable of running or building GCC 3.x or later. For the latest information about the platforms supported by MTP/IP, see:

http://www.DataExpedition.com/support/notes/tn0004.html

Frequently Asked Questions

0 General

0.1 What is MTP™/IP?

MTP/IP, or Multipurpose Transaction Protocol® technology, is transport software that moves data much more quickly and efficiently than traditional TCP/IP technology.

0.2 How is it used?

MTP/IP is built in to end-user software such as the ExpeDat™ file transfer application, the SyncDat™ directory replication application, and the HyperGate™ web accelerator. Because it is software built on top of existing network standards, it can be installed and running in seconds. These applications are also designed for easy embedding within other systems. Software Development Kits are also available, allowing application developers to greatly improve the speed, efficiency, and features of their own products.

0.3 Who can use MTP/IP?

MTP/IP applications can be used by anyone needing high performance, end-to-end data transport. Read the Motorola Case Study above for an example of how ExpeDat has improved the overall productivity of Motorola simulation clusters by 40%.

0.4 Who benefits from MTP/IP and how?

Application developers benefit by improving their product's performance and features and thus making it more attractive to their customers. End-users benefit from faster speeds, new functionality, and better reliability, and up to 50% savings on infrastructure costs. Bandwidth providers and consumers benefit from greater hardware utilization by fitting more traffic onto existing networks.

0.5 Does it really work?

Yes. MTP/IP is a robust technology that has been tested in real-world environments and applications. You can find out for yourself with any of our free trial downloads at *DataExpedition.com*.

0.6 I've heard others make such claims before. Why should I believe you?

Several companies have marketed products to improve network performance, each with significant limitations:

- *Compression*, which has a lot of overhead, can compromise data quality, and only works if the data isn't already compressed;
- Caching, which requires on a high degree of data repetition and can result in stale or out-of-synch data feeds;

- *Prioritization*, which sacrifices the performance of some data for the benefit of other data;
- *TCP derivations*, which require complex configuration and carefully controlled environments; and
- *UDP data blasters*, which flood the network and then spend almost as much time recovering from the mess as they saved by trashing the network.

MTP is none of these: it is a unique, patents pending, completely original transport protocol designed for modern networks and applications.

0.7 How can I learn more?

Visit our website `http://www.DataExpedition.com/`, call us at 877-292-2280, or internationally dial US 781-353-6403

1 Technology

1.0 Is MTP/IP compatible with existing networks?

Yes! MTP/IP is built on top of the existing UDP/IP standard, which is supported by all hardware and operating systems that are compliant with Internet standards. MTP adds sophisticated flow-control, error-recovery, and session management layers on top of UDP/IP. This allows it to provide new performance and efficiency in software that works with existing systems.

1.1 Where does MTP/IP fit in with other acronyms like HTTP, PPP, TCP, etc.?

The Internet is often organized into four protocol layers: application, transport, network, and link. This is called the Protocol Stack (similar to the 7 layer OSI stack).

- The *application* layer refers to protocols that manage data content and are not concerned with the details of how the data gets moved. Examples are web (HTTP), email (SMTP), file transfer (FTP), and news groups (NNTP).
- The actual data movement is usually left up to the *transport* layer, most commonly the Transmission Control Protocol (TCP). The transport layer handles error correction and flow control, but leaves the routing of data across the network to the underlying network layer.
- The Internet is pretty much defined by the use of the Internet Protocol (IP) at the *network* layer to route data across the many links, or "hops", which may lie between machines.
- The protocols used to physically communicate data across those links (ethernet, FDDI, PPP, v.90 etc.) make up the *link* layer. Most of these are related to physical hardware and many can be layered on top of each-other.

MTP is a transport protocol, so it sits between applications and the network, providing fast and reliable delivery of data.

1.2 How much faster is MTP?

- *Two to seven* times faster throughput (kilobytes per second) than TCP across high speed Wide Area Networks,
- *Two to six* times as many transactions per second as TCP on all networks,
- *Ten times* longer path tolerance (hop count), and
- *Ten times* more congestion (packet loss) tolerance than TCP.

1.3 Why is MTP faster than TCP?

TCP's design is over thirty-years old and it makes a lot of assumptions about the network that are now just plain wrong. As a result, TCP wastes a lot of time and bandwidth causing congestion that it must then correct. MTP is much more careful and is able to fully utilize the resources that TCP wastes.

1.4 Is MTP a compression scheme?

No. Applications which use MTP may choose to compress the data they send, but MTP itself does not rely on reducing the size of the data: it actually sends the data faster.

1.5 Is MTP a custom TCP implementation?

No. MTP is completely unrelated to TCP.

1.6 Don't all your gains come at the expense of third-party traffic?

No. While any new network traffic will impact existing traffic, kilobyte for kilobyte MTP has less of an effect on TCP data flow than TCP itself. Much of MTP's performance gain comes from its better utilization of otherwise wasted resources. MTP even has bandwidth management features built-in, giving you the option to precisely control how resources are allocate

1.7 Is MTP always faster than TCP?

For transaction processing, yes. For throughput, not always: MTP can't move data faster than your network hardware. If you have a slow connection, a very short network path, and there is no congestion, then TCP might perform well. Many factors can adversely affect TCP performance, making TCP difficult predict.

1.8 Is MTP ever slower than TCP?

No, not in a properly functioning IP network. However, misconfigured routers, firewalls, or other misbehaving network components can affect all network performance in unexpected ways.

1.9 I downloaded a .DLL that says it's faster than TCP. Why is MTP better?

What you downloaded is a "tuned" version of TCP that has been adjusted to perform better under specific conditions of link speed, latency, and loss. But

TCP does not scale well: if those exact conditions are violated, it will not only lose the performance advantages, but it may perform worse than a standard TCP. The more TCP is tuned to work well in one set of circumstances, the worse it will fail when those circumstances change. MTP is very scalable.

1.10 How scalable is MTP?

It has been tested and performs well at link speeds from 14.4 kilobits per second to 2.5 gigabits per second, at latencies as low as 1 millisecond to as high 2,000 milliseconds, and at packet loss rates from 0% to 50%. (If you've got a network outside this range, we'd love to test on it!)

1.11 XYZ Corp is marketing a UDP based protocol. Why is MTP faster?

Most UDP protocols dump data into the network as fast as the CPU can send it. Flooding the network like this causes massive congestion and packet-loss. These protocols spend almost as much time correcting for lost data as they save by flooding, plus they trash the network in the process. While MTP is also built on top of UDP/IP, it performs very careful flow control and error recovery. MTP does not flood the network, causes less congestion than TCP, and because of its greater efficiency, is able to reliably move data faster than flooding mechanisms.

1.12 Is MTP a reliable protocol?

Yes, by default MTP guarantees delivery of data. MTP also has unreliable modes and APIs for applications which value performance over reliability.

1.13 Is MTP meant to replace TCP?

For the vast majority of applications and network conditions, MTP provides superior performance to TCP. There are only a few legacy applications in which TCP has any advantage (for example, telnet). However, MTP and TCP can peacefully coexist in the same network and even in the same application. An application can easily support both, automatically using MTP when communicating with a system that supports MTP, and otherwise falling back on TCP.

2 Deployment

2.1 What is needed to use MTP?

Simply install an MTP application. Any two computers on an IP network can communicate using MTP/IP software, such as our ExpeDat FTP replacement, or our HyperGate web accelerator.

2.2 Does the end user need to install anything special?

No. The MTP SDK's consist of statically linked code libraries which are compiled directly into the application. MTP applications require no additional DLLs, drivers, or hardware.

2.3 Does MTP require changes or upgrades to the operating system?
No. Any OS that supports the TCP/IP stack will support MTP/IP.

2.4 Does MTP require special hardware?
No. MTP works with any hardware that meets IPv4 standards.

2.5 Does MTP work with firewalls, NAT, DHCP, PPPoE, etc.?
Yes, provided that these devices are functioning correctly. The network perceives MTP as standard UDP traffic and MTP requires no unusual handling. As with all network traffic, firewalls will need to be configured to allow MTP traffic to pass through. MTP can use any UDP port number.

2.6 Is MTP compatible with proxies?
Yes. Proxy servers only deal with TCP traffic for specific application types and will ignore UDP based traffic, such as MTP. Some MTP applications, such as HyperGate, can be configured to run on a proxy machine, providing their own proxy functionality.

2.7 Does MTP require that a specific UDP port number be used?
No. You are free to choose whatever port number is most convenient for your application or end-user environment.

2.8 Does MTP need to be adopted as a standard before it can be used?
No. Because it is built on top of existing Internet standards, MTP is already supported by IP networks.

2.9 Does MTP need to be present at both ends?
Yes, since MTP is a communications protocol, it must be present at both ends of a path in order to accelerate that path. As noted above, MTP can co-exist with TCP to provide legacy support. For example, the HyperGate web gateway works with and along-side existing TCP based web browsers and servers.

2.10 Can I try it?
Yes! You can download trial versions of our MTP applications at *DataExpedition.com*. For a trial of our Software Development Kits, please call or email.

3 Security

3.1 Is MTP more secure than TCP?
Yes. MTP is at least as secure as TCP in all respects and is less vulnerable to several types of attacks.

3.2 Does MTP support encryption?
Yes. DEI's ExpeDat, SyncDat, and HyperGate applications all support AES encryption and HyperGate uses SSL. The SDKs support the use of any encryption mechanism.

3.3 Does MTP work with encrypted VPNs and IPsec?
Yes. MTP will operate over any VPN which supports UDP/IP. Note that SSL and other tunneling VPNs do not fully support UDP/IP. IPsec VPNs are recommended for maximum performance and security.

3.4 Is MTP vulnerable to any of the Denial of Service attacks used against TCP?
No. Due to its completely different architecture, attacks designed against TCP will not affect MTP.

3.5 Is it possible for a Denial of Service attack to be designed against MTP?
Yes. It is not possible to completely eliminate the threat of DoS attacks when operating on a public network. However MTP's low-overhead design and more robust packet validation does minimize the potential impact of such attacks. MTP also provides application developers with better tools for identifying and squelching suspicious activity.

3.6 Does MTP have any vulnerabilities similar to the recent TCP session termination issue?
No. Unlike TCP, MTP datagrams carry information specifically designed to verify that each arriving datagram belongs to a valid transaction.

References

Aberdeen Group, “WAN Optimization Techniques Improve Application Delivery”, January 2004.

APM Advisors, “Traffic Management”, August 2004.

Ashton, Metzler & Associates, “The Three Components of Optimizing WAN Bandwidth”, March 2003.

Ashton, Metzler & Associates, “Why Performance Management Matters”, Oct 2003.

Ashton, Metzler & Associates, “The Changing Branch Office Network”, May 2004.

Business Communications Review, “Taming Your WAN”, October 2003.

Enterprise Management Associates, “Increasing IT Maturity by Automating Application Monitoring and Optimization”, 2004.

Gartner, “Network Impacts of Web-Enabled Enterprise Applications”, Dec 2003.

IDC, “Worldwide WAN Optimization Management 2004-2008 Forecast and Analysis”, May 2004.

META Group, “Bandwidth Compression and Optimization”, June 2003.

META Group, “Accelerating Web Application Performance”, June 2004.

VNetSoftware, “Network Testing of MTP”, February 2003.