

# **Data Expedition, Inc.®**

---



## **HyperGate™**

**Performance, Reliability, and Savings**

for

**Web Based Document Transport**

January 2009

<http://www.DataExpedition.com/hypergate/>

© 2009 Data Expedition, Inc.  
All Rights Reserved

## Summary


HyperGate™ is a software package which greatly improves the speed and reliability of document transport for web clients and servers. Its primary purpose is to accelerate existing HTTP document downloads and uploads across networks that are challenged by high-latency, congestion, or packet-loss. HyperGate provides real throughput gains for all data types across all Wide Area Networks.

- Software based acceleration of all data types
- Transfers documents up to seven times faster
- Ten times greater reliability.

Because no hardware is required and no modifications need to be made to existing software, HyperGate can be installed in minutes and without disruption. For end-users, HyperGate can be as simple as clicking on a link. HyperGate works with existing web browsers, servers, and back-ends.

- No new hardware
- Compatible with all HTTP browsers and servers
- AES and SSL security

HyperGate is not a download manager or a compression engine: it actually moves data more quickly and reliably by making better use of network resources. The HyperGate package consists of two software components, *HGclient* and *HGserver*, installed near the web browser and server respectively. Documents can be designated for transfer by adding a HyperGate aware URL (web link), or via proxy configuration.




HyperGate Software Moves Web Documents Across the WAN

## Table of Contents

Summary.....	<i>ii</i>
Motivation.....	1
How HyperGate Works.....	3
How MTP Works.....	5
MTP FAQs.....	8
General.....	8
Technology.....	9
Deployment.....	11


## Motivation

The Internet promised fast and reliable global data connectivity with just a single point of entry. The classic metaphor was the “cloud”, into which data would flow and be routed invisibly around trouble spots to quickly reach its destination.



Internet Ideal: Seamless, Fast, Reliable


However, the business reality is that the Internet, as it stands today, is often too slow and too unreliable for large or critical document transport. Achieving acceptable performance often requires expensive work-arounds such as remote points of presence, third-party content delivery services, or acceleration appliances.



Internet Reality: Expensive Work-Arounds Needed

At the core of this performance problem is a data transport infrastructure that is approaching thirty years of age. While the computers, applications, and network hardware have all made stellar advances in speed and efficiency over the years, the software that governs the flow of data across IP networks remains little changed from its 1974 conception.

HyperGate solves this efficiency problem for Web based document transport by taking advantage of a new, modern transport technology called the Multipurpose Transaction Protocol® (MTP™). Built on top of existing Internet standards, MTP's more efficient algorithms are able to move data several times faster and many times more reliably across any standard IP network infrastructure.



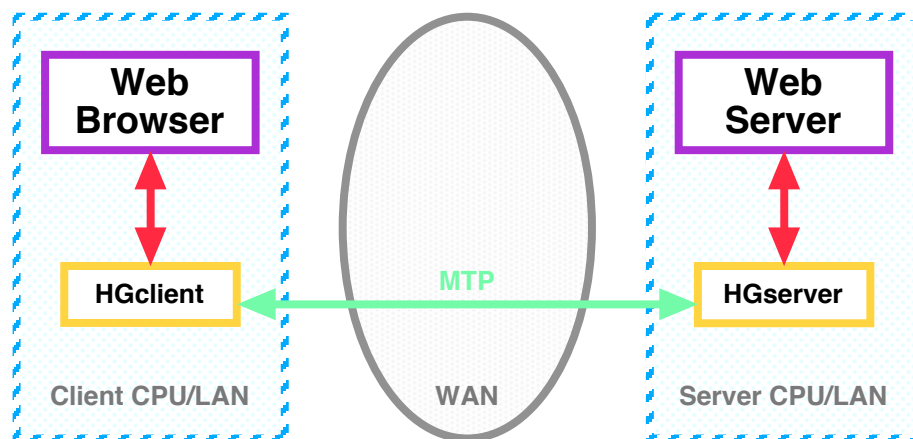
HyperGate: Helps the Internet Work the Way it Should

By moving your data using a modern, high performance, software infrastructure, HyperGate is able to restore the promise of the Internet. Reducing or eliminating the need for expensive performance work-arounds provides the dual benefits of increased business productivity and substantial cost savings.

The next chapter illustrates how HyperGate can be quickly and seamlessly deployed alongside your existing web based document transport infrastructure.

## How HyperGate Works

HyperGate works with your existing web clients and servers to transfer data more quickly and reliably than with HTTP alone. It does this by acting as an intermediary between the web browser and web server, transporting their HTTP data using the more efficient Multipurpose Transaction Protocol®.



HyperGate Software Works Between Web Browsers and Servers to Transport their Data More Efficiently

As shown above, there are three stages of communication in HyperGate. Your web browser begins each HyperGate transaction by connecting to the HyperGate client. The *HGclient* forwards this transaction across the Wide Area Network to the HyperGate server. The *HGserver* then passes the connection along to your web server as a standard HTTP transaction.

Both the Web Browser and Web Server communicate using the same HTTP protocol they have always used. There is no need for plug-ins, scripts, libraries, or other configuration changes. All of your existing software, including CGIs and other back-end infrastructure, remains unchanged. The only difference is in the middle, where the data is carried across the network using the more efficient Multipurpose Transaction Protocol. This WAN portion can also be protected with AES encryption, providing high security to your data. This works even if your web server does not support SSL.

The key to the seamless integration of HyperGate into your web infrastructure is the HyperGate URL. All that is required to transport a document or form using HyperGate is to add a link as shown below.

A typical web link takes the following form:

`http://server/document-path`

To access this same document with HyperGate, using the same web browser and web server, the URL could be changed to

`http://127.0.0.1:8088/server;document-path`

HyperGate can also be used to accelerate entire servers by acting as an HTTP proxy. For managed environments, this can be achieved simply by updating or distributing a standard Proxy Auto Configuration file.

Because these configuration changes take place behind the scenes, the end-user sees no difference in interface or work-process. There is no need to retrain users or modify existing software. HyperGate URLs work with any software that uses HTTP to Get or Post document data. They can be bookmarked, saved, or transcribed just like regular URLs. When used in proxy mode, even the URLs remain the same.

The example above assumes that the HyperGate software is running on the client and server CPUs. But more complex network environments can also be accommodated. For example, the HyperGate client can be run on a network gateway, allowing private network clients to take advantage of HyperGate without exposing them to the Internet.

HyperGate is available for numerous computing platforms including Windows, Linux, Mac OS X, FreeBSD, NetBSD, and Solaris.

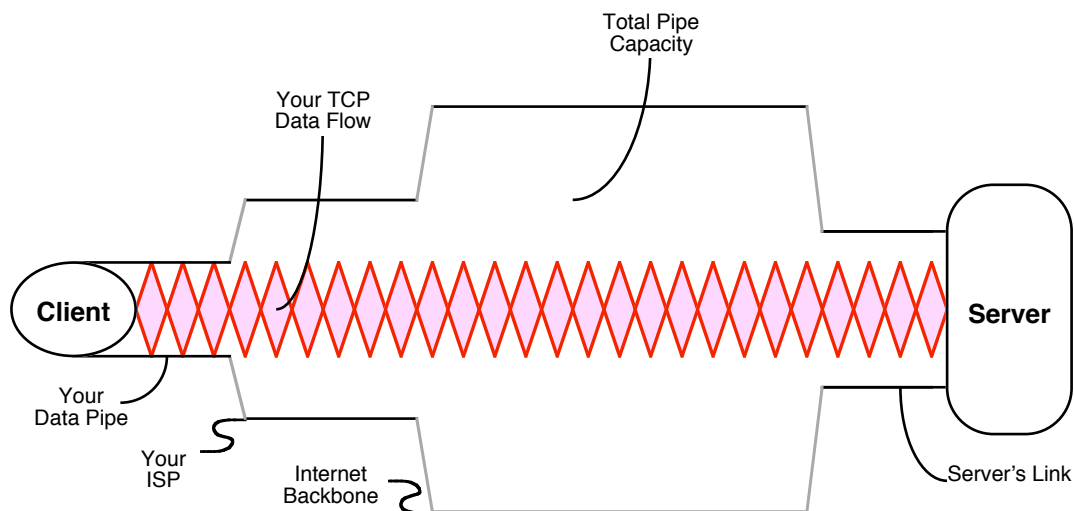
For complete details about how HyperGate is deployed, configured, and used, please visit <http://www.DataExpedition.com/HyperGate/> where you can browse the complete installation manual and download a free trial.

The next chapter discusses some of the technical details underlying the Multipurpose Transaction Protocol technology that HyperGate uses to move document data.

## How MTP Works


The Multipurpose Transaction Protocol® (MTP™) 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'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. This oscillation is most prevalent when TCP is operating across long data paths (many routers) with a broadband connection.

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.



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. Typically these oscillations occur about once per second, resulting in TCP data rates that are far below the actual available network capacity. Moreover, the constantly recurring congestion causes indiscriminate packet loss, which can further delay connections or disrupt connectivity. This is typical behavior on any backbone path.

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 the underutilized lull between TCP packet storms. 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 as long as the network has capacity to spare. MTP's throughput advantage comes from making use of resources that TCP wastes.

The next chapter addresses some of the frequently asked questions about MTP.

## 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 HyperGate™ web accelerator. Because it is software built on top of existing network standards, it can be installed and running in minutes. Software Development Kits are also available, allowing application developers to greatly improve the throughput of their own products.

#### 0.3 Who can use MTP/IP?

MTP applications can be used by anyone needing high performance, end-to-end data transport. MTP toolkits can be used by software developers to create highly efficient network applications that pass performance benefits on to their customers.

#### 0.4 Who benefits from MTP and how?

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

#### 0.5 Does it really work?

Yes. MTP is a robust technology that has been tested in real-world environments and applications. Free trials of all MTP/IP applications are available for immediate download at [DataExpedition.com](http://DataExpedition.com).

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

Other technologies claiming to improve performance have severe 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 a high degree of data repetition and does nothing for unique or dynamic data flows;
- *Prioritization*, which sacrifices the performance of some data for the benefit of other data;
- *TCP tuning*, which sacrifices adaptability to improve performance in a particular stable environment; and
- UDP streamers, which send data packets at a fixed rate and, at best, react to flooding after the fact.

MTP is none of these: it is a unique, patented transport protocol designed to adapt to any network path for maximum efficiency with minimum disruption.

#### 0.7 How can I learn more?

Read the rest of the FAQ, visit our website <http://www.DataExpedition.com/>, call us at 877-292-2280, or send a message to [info@DataExpedition.com](mailto:info@DataExpedition.com).

## 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 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 compared to TCP?

- *Two to seven* times faster throughput (megabits per second) than TCP across high speed Wide Area Networks is typical, with faster speeds possible,
- *Two to six* times as many transactions per second as TCP on all networks,
- *Forty times* higher latency tolerance;
- *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 data reduction: 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 ambient 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 allocated.

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 as well. Many factors can adversely affect TCP performance, making TCP difficult to predict. The best way to know how MTP will perform in your environment is to try it.

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. Contact DEI for assistance in assessing the health of your network.

1.9 I installed 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 10 gigabits per second, at latencies as low as 1 millisecond and as high as 20,000 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 such protocols simply dump data into the network as fast as possible and hope that some of it will get through. Whatever doesn't get through, they try to send again later. Flooding the network like this causes massive congestion and packet-loss. Ultimately, such protocols spend almost as much time correcting for lost data as they save by flooding. Plus they severely degrade third-party traffic in the process. While MTP also happens to be built on top of UDP/IP, MTP performs very careful flow control and error recovery. MTP does not flood the network, actually 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.

## 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.

### 2.2 Does the end user need to install anything (kernel drivers, DLLs, applications, etc.)?

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.

### 2.4 Does MTP require special hardware?

No. MTP works with any hardware that meets IPv4 or higher standards.

### 2.5 Does MTP work with firewalls, NAT, DHCP, PPPoE, etc.?

Yes. The network perceives MTP as standard UDP traffic and MTP requires no unusual handling. Firewalls may need to be configured to allow MTP traffic to pass through, or MTP can use any UDP port that is already available.

### 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. In the case of HyperGate, the client application can be deployed on network gateways and act itself as a proxy.

### 2.7 Does MTP require that a specific UDP port number be used?

No. You are free to choose whatever port numbers are most convenient for your application or network environment.

### 2.8 Does MTP need to be adopted as a new 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. In the case of HyperGate, the client and server software must be installed near the web browsers and servers.

### 2.10 Can I try it?

Yes! Free trials of all MTP/IP applications are available at [DataExpedition.com](http://DataExpedition.com).