

**Revised Edition.  
Includes updated  
SONET and IP pricing  
for Summer, 2003.**

The Economics  
of  
Large Scale Data Transfer

By: Andy Helland

[andyh@lightsand.com](mailto:andyh@lightsand.com)

*A White Paper describing the benefits of Fibre Channel  
and SONET as the basis for high-performance and cost  
effective data transport through the Wide Area Network.*



V2.0 June, 2003

## Introduction

The Internet and IP networks, in general, were designed to provide large-scale interconnection of computers and networks. This goal has been achieved with incredible success. The Internet provides rapid access for millions of client machines to hundreds of thousands of servers (mostly providing http, ftp, and email services). Yet, for all of its size and scalability, predictable quality of service (QoS) has continued to elude the Internet. As individual users continue to increase their requirements for wide bandwidth and low latency interconnections, alternative solutions must be developed that provide higher performance and greater cost efficiency in the movement of large amounts of data.

In this white paper, we will look at the economics of large-scale data transfer<sup>1</sup>. We will examine the real-world throughput rates that can be achieved using IP-based transfer over the WAN. We will normalize those throughput rates using the monthly recurring costs required for high data rate, routed IP service. As a comparison, we will describe a Fibre Channel over SONET based system for transferring large amounts of data across the WAN. We will analyze the effective throughput rates and normalize them against the monthly recurring costs.

In the end, we will demonstrate that Fibre Channel over SONET provides single user performance that dramatically exceeds that of IP based interconnect. Fibre Channel over SONET also provides a significant cost advantage when compared with IP-based data transfer in the movement of large amounts of data.

## The Price of Scalability in IP Networks

Given the tremendous success that IP routing has enjoyed in the development of the Internet, it is natural to assume that it is an ideal medium for interconnecting high performance data centers. Nothing could be further from the truth. IP routing protocols, including the random early discard (RED) algorithms of routers, were optimized to service the needs of millions of users who wished to move a small amount of data each. IP routing was not optimized to provide high throughput performance for a small number of users. This becomes clear when a detailed throughput analysis of TCP/IP routing is conducted.

---

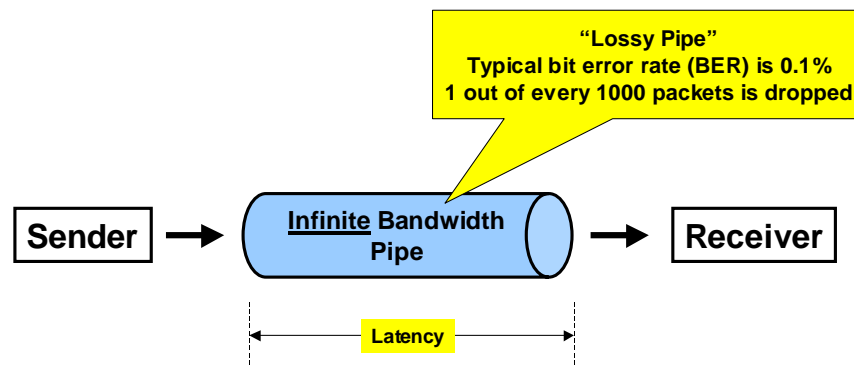
<sup>1</sup> In this paper, we address the need for continuous movement of large amounts of data. For a look at the economics of moving a single instance of a large database, the reader should consult "TeraScale SneakerNet: Using Inexpensive Disks for Backup, Archiving, and Data Exchange." Gray, Chong, Barclay, Szalay, and Vandenberg. Microsoft Corporation, May 2002.

One of the most often-cited research reports on the performance of TCP/IP is “The Macroscopic Behavior the TCP Congestion Avoidance Algorithm.”<sup>2</sup> In this seminal report, the behavior of the TCP in the face of latency (the time delay between the sender and the receiver) and packet loss rate is carefully analyzed. The result of this analysis is a relatively straightforward equation that explains the relationship between bandwidth, packet size (referred to as maximum segment size), round trip time, and packet loss rate.

$$BW \leq \frac{1.31 \times MSS}{RTT \sqrt{PkLoss}}$$

BW = bandwidth (b/s)  
MSS = Maximum Segment Size (bits)  
RTT = round trip time (seconds)  
PkLoss = packet loss rate

As we model the effects of real-world packet loss rate, we reach the inescapable conclusion that throughput (the total amount of data that is moved from point A to Point B in the WAN) is much more dependent on latency and packet loss rate than on the native bandwidth of the pipe. To demonstrate this point, let’s look at an interconnect pipe with *infinite* bandwidth that connects a sender and a receiver as shown in the figure below.



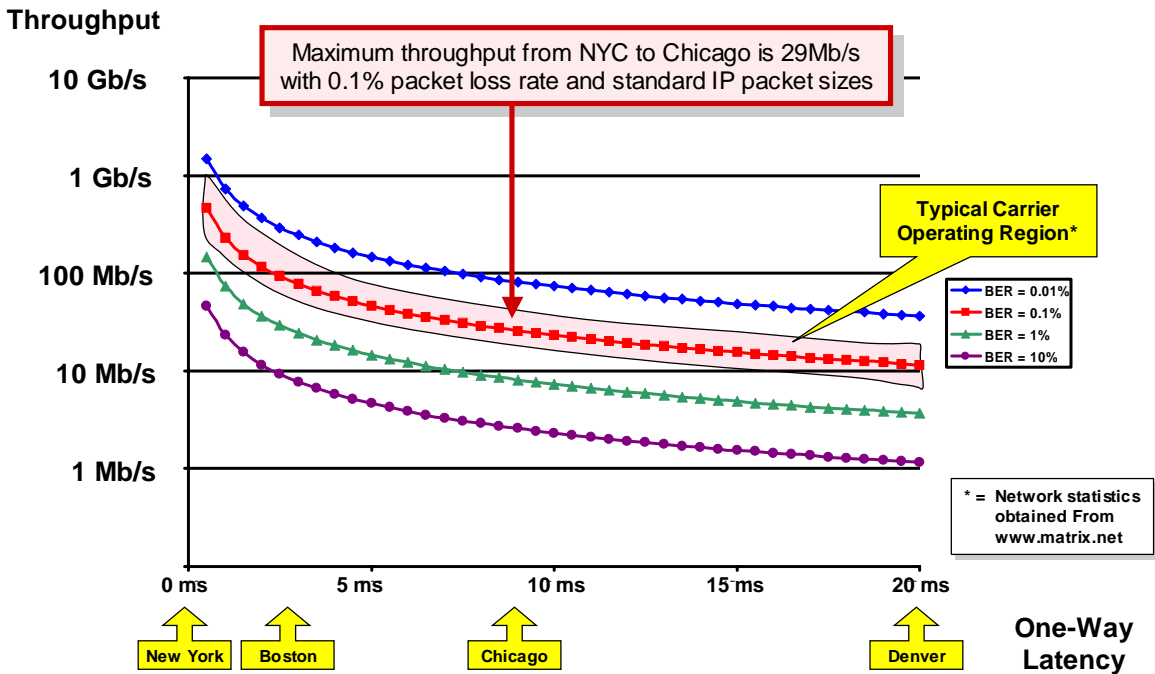
*To understand the impact of packet loss on TCP/IP throughput, we model a dedicated pipe with infinite bandwidth, non-zero latency and a typical packet loss rate.*

Even if we have a pipe that has infinite bandwidth, the fact that we must stop and retransmit data every time a packet is dropped quickly becomes the dominant factor that limits throughput. Furthermore, while we are retransmitting lost data, all of the

<sup>2</sup> “The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm.” Mathis, Semke, Mahdavi, and Ott. *Computer Communication Review*, a publication of ACM SIGCOMM, volume 27, number 3, 1997.

successfully transmitted data that came after the lost data must wait until the lost data has been retransmitted to prevent misordering the data.

Based on this analysis of TCP performance, we can plot a family of curves that describe the behavior of TCP as a function of latency and packet loss rate. This graph is shown below. Each curve shows the expected actual throughput when operating over an infinitely fast connection with a specific packet loss rate. The maximum length of an IP packet has been assumed to be 1500 bytes. This is the standard MTU (maximum transfer unit) for Ethernet. Note that as the error rates increase, the throughput decreases. Increasing the distance between the sender and the receiver also decreases the throughput.



*In a typical lossy IP network, throughput is limited more by latency and loss rate than it is by the bandwidth of the pipe.*

Remember that the bandwidth of the pipe in this model is infinite! The limiting factor in throughput is based on the need to retransmit dropped data and the amount of distance over which the data must be sent. To calibrate latency in terms of distance, we have assumed that the sender is in New York and the receiver is located in Boston, Chicago, or Denver. Packet loss rate is a critical parameter in this analysis. The average packet loss rate for the Internet at large<sup>3</sup> is between 1% and 2%. Service Level Agreements (SLAs) with major carriers typically guarantee only 0.5% packet loss. For the purpose of this

<sup>3</sup> Packet loss rates were obtained from [www.internetweather.com](http://www.internetweather.com), an independent source of IP network performance measurements statistics.

analysis, we will assume a packet loss rate of 0.1% (5 times *better* than premium SLA). With a packet loss rate of 0.1%, data that is sent between New York and Chicago can only achieve a maximum throughput of approximately 29 Mb/s (even if the actual pipe that interconnects the two sites has infinite bandwidth).

This performance limitation of TCP/IP over lossy networks is nothing new to the supercomputing industry. For years, they have been trying to move large amounts of data across TCP/IP routed networks with little success. Los Alamos National Labs, for example, has resorted to sending data via tapes and Federal Express to their partner facility at Sandia National Labs. The real world throughput limitations of TCP/IP over lossy networks are so extreme that it is faster for them to send data via truck than to use their OC-3 IP wide area access<sup>4</sup>.

So, why do the carriers operate their networks with such a high loss rate and thwart the desire of large-scale single users? The answer is simple. They wish the SONET pipes that power their networks to be as full as possible. To make sure that happens, they oversell the bandwidth. Statistically, there is great variation in the amount of traffic that passes through an IP network. The carriers sell bandwidth based on their long-term averages—not on the short-term bursts of traffic that occur.

Major airlines operate the same way. They regularly oversell the number of seats on a given flight. Statistically, they know that some passengers will cancel at the last moment (without providing notice). In the event that all of the passengers that actually purchased tickets show up, they are forced to bump a passenger. Just like in the case of IP carriers, most customers never know that this took place. Only a few are inconvenienced.

This over-subscription business model works well if there are many little users whose requirements cannot be predicted in advance. However, if a single user attempts to move large amounts of data (passengers), this model begins to fail rapidly. Let's return to our airplane analogy. Suppose we have an entire football team that must be moved from one city to another. What if the airline decides to bump one person (say, the quarterback)? The rest of the team will arrive at the destination at the correct time. However, since the team cannot function until the quarterback is safely delivered, the entire team is forced to delay their game until he arrives. As the size of the team increases, the chances of being disrupted also increase. The same effect occurs with large-scale data transfer.

When single users attempt to transfer large amounts of data across the Internet, their chances of losing data increase as well. If the carriers drop one out of a thousand packets and the users are small, 999 customers will see great performance and only one customer will have difficulty. On the other hand, if a single user is moving large amounts of data, every time he sends a thousand packets, he will get hit. TCP windowing will force him to reduce his bandwidth and he will constantly experience frustration (retransmission and bandwidth reduction).

---

<sup>4</sup> [www.unix.mcs.anl.gov/discovery/wufeng.htm](http://www.unix.mcs.anl.gov/discovery/wufeng.htm)

Quite simply, routed IP networks are designed to let the carriers manage the aggregate behavior of many little users that need to route data to any destination. They are not well suited to the transfer of large amounts of data by single users.


## **Using Fibre Channel over SONET**

Fibre Channel over SONET (FC/SONET) represents an excellent solution that utilizes the available infrastructure and yet provides very high performance. Before we can talk about the benefits of FC/SONET as a data movement solution, let's review what SONET is and how it works.

SONET (Synchronous Optical NETWORK) is a network of fiber optic cables that carry data and voice throughout the United States. Data is passed through dedicated channels at rates from 155 Mb/s up to 10 Gb/s. Data paths are not changed rapidly like they are with routed IP networks. Rather, they are provisioned and remain very static. Each dedicated channel is given a guaranteed fixed bandwidth for the data. A good way to think about SONET would be to imagine a dedicated pipe with water flowing from one location to another. Every drop of water that enters the pipe on one side will emerge from the other side the pipe. The water is never lost along the way and the water always emerges in the same order that it entered the pipe. The same is true for SONET channels. All the data that enters the pipe will emerge from the far side of the pipe and it will exit in the same order that it entered the pipe.

SONET provides the universal infrastructure used by every carrier in the United States to transport data and voice. ATM and IP all ride on top of SONET. Intuitively, it makes sense to use the lowest layer of transport available that can accomplish a particular goal. This eliminates unnecessary overhead and streamlines the transport process. Because it is so universally deployed, SONET-based signals are readily passed from carrier to carrier to move across country.

Another benefit of SONET-based networks is that they are usually deployed in redundant rings. In the event that there is a failure (say, due to a backhoe digging up a cable), SONET add-drop multiplexers (ADMs) will automatically re-route the traffic within 50 milliseconds. This level of reliability is one of the key reasons that the telecom and datacom infrastructure of the world has worked so well for so many years.



*FC/SONET gateways directly couple Fibre Channel onto SONET—taking advantage of the large installed base of SONET networking equipment.*

SONET is like Fibre Channel in the sense that it inherently offers in-order, loss-less delivery of data. Because of this and because of its ready availability and ability to transit from carrier to carrier, SONET is the ideal mechanism to carry high performance FC data over distance.

A crucial aspect of a Fibre Channel based storage system is credit buffering. Credit buffering accomplishes two things—both related to the concepts of congestion and flow control. First, it prevents the switches in the fabric from becoming congested and dropping data. The receiving switch controls the pace of traffic as it moves from the sender to the receiver by issuing credits back to the sender. If the receiving switch is congested, it will not issue credits to the sender and the sender will slow down the pace of traffic. Since the receiving switch controls the pace, it can never be overwhelmed by traffic. This is a crucial difference between Fibre Channel based networks and IP based networks. At any instant, an IP router can be suddenly overwhelmed with traffic. If this occurs, it will be forced to drop data to protect itself.


The second attribute of Fibre Channel credit buffering is that it allows the actual storage devices to communicate their levels of congestion to each other. Let's consider the example of a server that is suddenly faced with a large number of interrupts and is temporarily unable to service the needs of an in-process SCSI transfer from a disk. In this case, it will slow down the pace of credits that it issues to the Fibre Channel switch to which it is connected. By reducing the pace of credits that are issued, the server communicates backpressure to the Fibre Channel fabric. That backpressure will propagate through the fabric and ultimately cause the transmitting disk to slow its pace of traffic. The important theme here is that congestion is transmitted back through the network *without dropping data*. This is fundamentally different from the process used in TCP/IP routed networks that must drop data to communicate congestion.

These attributes of Fibre Channel are extremely important when Fibre Channel is extended outside the data center and onto the WAN. Once distance (latency) is added to the system, proper credit buffering becomes an even more crucial aspect of the system design. The WAN side of the extension gateway must use a credit-based system that has

been optimized for the distance and bandwidth of the link extension. The LightSand S-600 and S-2500 gateways were specifically designed to support loss-free transmission of large amounts of data over long distances. They deploy WAN-optimized buffering systems that ensure line rate data transfer for thousands of kilometers.


SONET has the most important features necessary to enable a high performance FC extension system. It has deterministic (low) latency and loss-less transmission of data. When coupled with a properly designed gateway, FC over SONET provides an optimal method of transferring large amounts of data over long distances.

For this analysis, we will look at the LightSand S-600 and S-2500 gateways, which encapsulate Fibre Channel over SONET. The S-600 uses one or two Fibre Channel inputs and a SONET OC-12c WAN port (622 Mb/s line rate). In practice, the actual data rate will be approximately 580 Mb/s because of the SONET encapsulation so we will use that number as the maximum single-user data rate for the system. In this case, the single user will saturate the SONET link that interconnects the sites.



*With its SONET OC-12c WAN port, the S-600 provides a single-user bandwidth of 580 Mb/s.*

The S-2500 has three Fibre Channel inputs (800 Mb/s each). These are multiplexed into a SONET OC-48c (2.488 Gb/s line rate). In practice, the S-2500 gateways will transmit approximately 2.32 Gb/s of usable data rate. The maximum single-user bandwidth in this case is 800 Mb/s (limited by the data rate of the Fibre Channel input port). A single-user is defined as any given flow of data from one site to another (i.e. from one disk to another disk or from a disk to a server).



*With its OC-48c WAN port, the S-2500 provides a maximum single-user bandwidth of 800 Mb/s and an aggregate bandwidth of 2320 Mb/s.*

Let's take a closer look at the real-world performance that can be expected when a standard IP WAN is compared with a FC/SONET system. We'll take a hypothetical example of moving data from New York to Boston, New York to Chicago, and New York to Denver. For each example, we will estimate the cost of service (both routed IP service and a dedicated SONET pipe). We will estimate the single user throughput and then calculate the cost efficiency of the link in dollars per Megabit per second (\$ per Mb/s). We will also look at the cost efficiency of data transfer when the links are fully utilized (with multiple concurrent transfers).

All of the analysis in this White Paper is based on the assumption that SCSI block size is not an issue in the transfer of data. While SCSI block size can seriously affect the performance of data transfer if not handled properly, it applies equally to any form of storage over IP-based and storage over FC/SONET-based WAN implementations. SCSI block size should be set as large as possible when transferring data over the WAN.

Before we get started, let's look at the how monthly recurring costs are estimated.

## **Cost Basis for Service**

In this white paper, we are addressing the recurring costs of service for interconnecting two sites. We have not addressed the capital costs of the equipment that would be used at the customer premises. In general, the cost of equipment such as IP routers and FC gateways is comparable. Regardless, the monthly cost of service is the dominant factor when compared with capital outlay.

As a basis for estimating the price of WAN connectivity, we consulted a major carrier and received quotes for both routed IP service and dedicated SONET point-to-point links. The quoted prices for routed IP service are listed below.

Routed IP Service	Monthly Recurring Cost	Source
OC-3	33,750	Carrier Quote
OC-12	66,000	Carrier Quote

*Monthly Recurring Costs for Routed IP Service.*

Note that routed IP service is not distance-sensitive in its pricing. These estimates are based on providing routed IP service to the customer premises. Since the customer can transmit his IP packets anywhere, there is no concept of distance in this price model. However, we will see below that distance does become a significant factor in calculating the *effective throughput* that a single user can achieve using IP transport.

Unlike the pricing of Internet access, the monthly recurring cost of SONET service is distance sensitive. Since the customer is contracting for a dedicated point-to-point link at a committed rate, the monthly recurring charge will be a function of the distance between the sites as well as the bandwidth. SONET price models are also complicated by differing rate structures that depend on the distance from the enterprise to the Carrier's POP (Point-of-Presence) as well as the distance between the POPs. SONET links that move across telephone LATAs (local access and transport areas) may also be subject to additional charges. We contacted a major carrier and received price quotes for SONET service between two hypothetical data centers. We assumed that the enterprise was located 5 miles from a carrier's POP in each of the representative cities. The pricing information that we used for this comparison (both routed IP and point-to-point SONET) assumed that the service was leased for a multi-year contract.<sup>5</sup>

The figures that are used for this analysis are just estimates (both routed IP service and point-to-point SONET). The actual pricing will be highly dependent on many variables such as the distance between sites, the availability of fiber in the local loop and the distance between the enterprise and the carrier's POP.

## **Moving Data from New York to Boston**

To examine the difference in performance between IP routing and FC over SONET, we first need to know the distances involved so that we can estimate the throughput of the network layer. Boston is 306 kilometers (190 miles) from New York. If we add 50% to allow for the non-straight path that fiber typically takes, we can use 459 kilometers as an estimate for the working distance. The speed of light in fiber is approximately 5 microseconds per kilometer. Thus, the one-way latency between the sites is approximately 2.3 milliseconds (459 Km x 5  $\mu$ s/Km).

---

<sup>5</sup> There are many research companies that conduct surveys of SONET pricing. Telegeography ([www.telegeography.com](http://www.telegeography.com)) is one such company. They publish comprehensive surveys and even offer some free samples of comparative SONET pricing. As of this writing, they have a free comparison of SONET OC-3 pricing between major metropolitan cities in North America and Europe. The current version that they have posted is dated April, 2002.

To interconnect two sites using the IP WAN of a major carrier, *each site* would require broadband Internet access. This will be factored into our cost model below. The actual costs for Internet Service will thus be twice the monthly recurring costs shown above. Once service is obtained, each site is free to transmit as much data as desired (up to the data rate of the routed IP service).

Using the equation above from [2], we can estimate the throughput that a single user can actually achieve between New York and Boston. For the purposes of this analysis, we assume that the round-trip-time (RTT) is 4.6 ms, the packet loss rate is 0.1% and the maximum data segment size is 1500 bytes. Note that the 1500 byte limit is based on the MTU (maximum transfer unit) limitations of Ethernet. Since all of this data originates at a server that is connected using Ethernet, this MTU is reflected throughout the system. For this combination of RTT, loss ratio, and MTU, the maximum single-user bandwidth is 108 Mb/s. This limitation applies regardless of the how much faster we make the access to the IP core. (i.e. The limitation stems from the packet loss rate and retransmission – not the maximum bandwidth of the routed IP network.) By comparison, FC/SONET uses the bandwidth of the channel much more effectively than TCP/IP. Data is never dropped to manage congestion, so precious bandwidth and time are never consumed with retransmitting lost data.

Once the SONET encapsulation is removed, the achievable data rate of the SONET link itself is approximately 580 Mb/s for the OC-12c SONET link and 2320 Mb/s for the OC-48c SONET link. As discussed above, we will use 580 Mb/s as the maximum single user bandwidth for the S-600 and 800 Mb/s for the S-2500. The 800 Mb/s limitation stems from the maximum speed of the FC input to the FC/SONET gateway. The table below shows how these different options compare.

	Line Rate	Effective Single-User Data Rate	Time to Move 1 TByte of Data	Monthly Recurring Cost of Service	\$ per Mb/s (Single-User Data Rate)
OC-12 Routed IP Service	622 Mb/s	108 Mb/s	23 Hours	\$ 132,000	1222
OC-12 FC/SONET	622 Mb/s	580 Mb/s	4.2 Hours	\$ 24,750	43
OC-48 FC/SONET	2488 Mb/s	800 Mb/s	3.1 Hours	\$ 69,000	86

*Cost Effectiveness of Data Transfer over the WAN (Single-User, NYC to Boston)*

Note the difference in the amount of time that it would take a single user to transmit a Terabyte of data using OC-12 Internet Access (23 hours) and an OC-12 FC/SONET link (4.2 hours). Note also that the cost of service is significantly less. The result is that the

cost per Mb/s of single-user bandwidth is more than an order of magnitude higher for IP service over OC-12 when compared with FC/SONET using the same OC-12 line rate! Clearly, the IP-based solution leads to very wasteful utilization of the bandwidth for the single user that simply wishes to move large amounts of data from one location to another.

In reality, a carrier (or the large enterprise) would not ignore all of the remaining bandwidth not consumed by the single-user. They would re-use the remaining bandwidth to service other users. Let's assume that they re-use it completely. If we take the bandwidth of the FC/SONET case (580 Mb/s) as the maximum achievable bandwidth for the SONET link, we can divide it by the single-user data rate using IP (108 Mb/s). Thus, we can increase the effective number of users by a factor of 5.4 (580/108). If we adjust the cost efficiency number the same way, we can determine the cost-per-Mb/s for a fully utilized link. Based on this calculation, we see that OC-12 Internet access can provide service for \$228 per Mb/s (132,000/580) when the link is fully utilized. Note that this is still significantly higher than the cost-per-Mb/s for FC/SONET. The case of FC/SONET for OC-48 access is similarly adjusted to reflect 3 users. OC-48 FC/SONET now becomes the most cost effective method of transferring data from one site to another (with three simultaneous users).

	Aggregate Data Rate (for Fully Utilized Link)	Number of Users to Fully Utilize Link	Monthly Recurring Cost	\$ per Mb/s (for Fully Utilized Link)
OC-12 Routed IP Service	580 Mb/s	6	\$132,000	228
OC-12 FC/SONET	580 Mb/s	1	\$24,750	43
OC-48 FC/SONET	2320 Mb/s	3	\$69,000	30

*Cost Effectiveness of Data Transfer over the WAN (Fully Utilized Link, NYC to Boston)*

From the above analysis, it is clear that FC/SONET is, by far, the most cost-effective method of moving data from site to site for the single user. Even when the link is fully utilized by other users, FC/SONET still maintains a cost effective edge.

What are the implications of this analysis for FCIP? Since FCIP operates as a single tunnel through the WAN, it is subject to the single-user analysis above. Even though there may be multiple flows of FC data through the FCIP tunnel, all of these flows will transit the same IP tunnel. FCIP will have significant difficulty in achieving the performance required by large-scale users when running through the public network. Furthermore, its cost-effectiveness will never approach that of FC/SONET.

## Moving Data from New York to Chicago

Let's look at the economic considerations for moving data from New York to Chicago. We can base our analysis on the above case for New York to Boston but substitute the additional latency for the longer distance. The distance is 719 miles (1158 kilometers). With our 50% fudge factor for cable length, we can use 1,737 kilometers as the distance. The one-way latency will be 8.7 milliseconds. With this increase in latency (and the same packet loss rate identified above), we see the single-user data rate drop substantially to 29 Mb/s. Notice that the monthly recurring cost for Internet access has remained the same and that the monthly cost for SONET has increased. This is to be expected since Internet access is blind to distance while SONET pipes are leased by the mile.

	Line Rate	Effective Single-User Data Rate	Time to Move 1 TByte of Data	Monthly Recurring Cost of Service	\$ per Mb/s (Single-User Data Rate)
OC-12 Routed IP Service	622 Mb/s	29 Mb/s	3.5 Days	\$ 132,000	4552
OC-12 FC/SONET	622 Mb/s	580 Mb/s	4.2 Hours	\$ 50,250	87
OC-48 FC/SONET	2488 Mb/s	800 Mb/s	3.1 Hours	\$ 176,250	220

### *Cost Effectiveness of Data Transfer over the WAN (Single-User, NYC to Chicago)*

Although the cost of service has remained the same for Internet access, the single-user data rate has plummeted to 29 Mb/s. The time required for a single-user to move a terabyte of data is now measured in days instead of hours. For FC/SONET, the performance has remained the same but the pricing has increased to reflect the increased distance in the SONET link. The disparity between the cost-effectiveness of the IP based solutions and the FC/SONET solutions has increased significantly for single-user data movement.

As with our example from New York to Boston, let's add back in the other users that would naturally fill the remaining bandwidth of the pipe.

	Aggregate Data Rate (for Fully Utilized Link)	Number of Users to Fully Utilize Link	Monthly Recurring Cost	\$ per Mb/s (for Fully Utilized Link)
OC-12 Routed IP Service	580 Mb/s	20	\$ 132,000	228
OC-12 FC/SONET	580 Mb/s	1	\$ 50,250	87
OC-48 FC/SONET	2320 Mb/s	3	\$176,250	76

*Cost Effectiveness of Data Transfer over the WAN (Fully Utilized Link, NYC to Chicago)*

Notice that the gap between the cost-effectiveness of IP networks and FC/SONET networks is closing although FC/SONET remains the more cost-effective method of transferring data between sites. This is to be expected since the monthly costs for Internet access are constant and the monthly costs for SONET circuits will increase with greater distance. More important, however, is the fact that we have had to increase the number of users to 20 in order to fully utilize the bandwidth of the link. For the enterprise that wishes to move large amounts of data through the WAN, this significantly complicates the issue. In order to move large amounts of data, we must break the data into 20 different flows and manage each one as a separate TCP connection. At the receiving end, we must reassemble the results from the 20 separate TCP connections in order to restore the data.

## **Moving Data from New York to Denver**

Let's continue this process to an example with even greater latency between the sender and the receiver. The distance between New York and Denver is approximately 1629 miles or 2621 kilometers (one way). Applying our 50% fudge factor, we will use 3930 kilometers as the basis of our analysis. This translates into a one-way latency of 19.5 milliseconds. Recalculating the expected single-user throughput for the IP network, we can see that it has dropped to a mere 12.6 Mb/s. FC/SONET throughput has remained the same since a properly designed FC gateway (such as the LightSand S series) provides credit buffering and flow control specifically designed for long-haul communication. As before, the costs of routed IP service remain the same while SONET costs have increased in proportion to distance.

	Line Rate	Effective Single-User Data Rate	Time to Move 1 TByte of Data	Monthly Recurring Cost of Service	\$ per Mb/s (Single-User Data Rate)
OC-12 Routed IP Service	622 Mb/s	12.6 Mb/s	8.1 Days	\$ 132,000	10,500
OC-12 FC/SONET	622 Mb/s	580 Mb/s	4.2 Hours	\$ 107,000	184
OC-48 FC/SONET	2488 Mb/s	800 Mb/s	3.1 Hours	\$ 362,000	453

*Cost Effectiveness of Data Transfer over the WAN (Single-User, NYC to Denver)*

The trends that we observed from New York to Chicago continue as we look at moving data from New York to Denver. Notice that the “Time to Move a Terabyte of Data” has now increased to more than 8 days for the IP-based WAN. On the other hand, the FC/SONET interconnection remains constant at 4.2 hours and 3.1 hours for the OC-12 and OC-48 cases. The cost-effectiveness of IP-based transport for a single user has now plummeted even further and no one would consider dedicating an IP based link to moving a single user’s data. On the other hand, a single-user FC/SONET pipe could easily be established to move large amounts of data. It would remain very cost effective (even at large distances). It would also be a manageable process since the servers at each end would still only be managing a few connections.

Let’s look at the multi-user model that would fill the remaining bandwidth of the links and see what has happened to our cost efficiency and manageability. Since a single-user can now only expect 12.6 Mb/s, we must have 46 simultaneous connections to fully utilize the available bandwidth across the IP WAN. This model is perfectly normal for the Internet at large (many smaller users coexisting on the same links). However, it has dramatically increased the management overhead that would be required for a single user to realistically send data from one site to another. Imagine trying to push a fire hose of data through 46 drinking straws. It might be possible to break a data stream into 46 smaller threads and then reassemble it but it certainly is difficult.

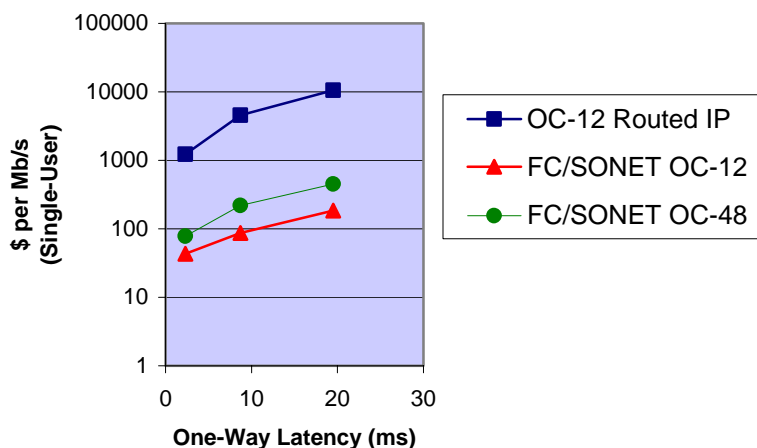
	Aggregate Data Rate (for Fully Utilized Link)	Number of Users to Fully Utilize Link	Monthly Recurring Cost	\$ per Mb/s (for Fully Utilized Link)
OC-12 Routed IP Service	580 Mb/s	46	\$ 132,000	228
OC-12 FC/SONET	580 Mb/s	1	\$ 107,000	184
OC-48 FC/SONET	2320 Mb/s	3	\$362,000	156

*Cost Effectiveness of Data Transfer over the WAN (Fully Utilized Link, NYC - Denver)*

Using the New York to Denver case, IP WAN transport cost efficiency (for a fully utilized link) has finally reached the same close order of magnitude as FC/SONET. However, it requires 46 separate threads that must be managed by the Sender and the Receiver. Despite the improvements in the cost-efficiency of the routed IP service, the FC/SONET data transfer model still provides the most cost-effective method of moving large amounts of data across country.

## Comparing the Options

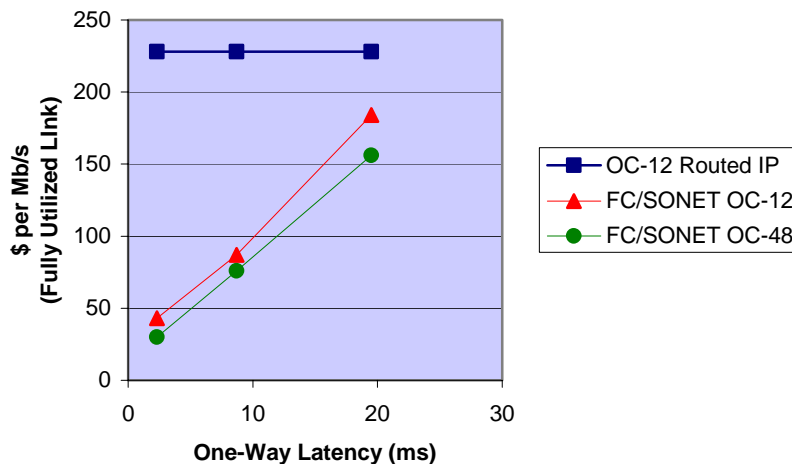
Let's summarize the options that we have examined for moving data from New York to Boston, Chicago, and Denver. The chart below contains the single user cost-efficiency estimates for OC-12 Internet access as well as the single-user cost efficiency estimates for FC/SONET (OC-12 and OC-48 dedicated channels). Note that the vertical scale is logarithmic.



*Cost Effectiveness of Single-User Data Transfer as a Function of Inter-site Latency*

As we discussed above, FC/SONET offers a significant advantage for the single-user that wishes to transfer large amounts of data from one site to another. Although it becomes more expensive (on a cost per Mb/s basis) as the distance increases, the cost-per-Mb/s of FC/SONET increases much less rapidly than Routed IP transport over the WAN. The FC/SONET OC-12 link is more cost-effective than the FC/SONET OC-48 link because the OC-12 link is fully saturated and a single user is consuming all the bits-per-second being purchased.

If we look at the cost efficiency of fully utilized links, we see some other interesting trends as well.



*Cost Effectiveness of Data Transfer as a Function of Inter-site Latency*  
*Cost Effectiveness of Multi-User Data Transfer as a Function of Inter-site Latency*

Here, we see the cost efficiency of routed IP transport remaining constant (on a cost per Mb/s basis). This makes sense since the pricing for Internet access is independent of distance. The price paid per Mb/s for FC/SONET, of course, increases since the cost of a dedicated SONET channel goes up with the increasing distance between the cities. Even with the increase in price-per-Mb/s, FC/SONET remains more cost effective than routed IP transport. Furthermore, in order to obtain the high utilization factor for IP routed networks, it is necessary for the enterprise to maintain numerous simultaneous connections (20 for a fully utilized link to Chicago, and 46 for a fully utilized link to Denver).

Recent WAN-based demonstrations of IP storage<sup>6</sup> seem to indicate that IP can provide wide area connectivity with high data rates at transcontinental distances. A closer look at the results indicates that this might not be so. The test results<sup>7</sup> include a snapshot of the IP carrier's performance statistics during the time of the test. At the instant of the testing, the carrier's network was operating with 0% packet loss. Since this packet loss rate is

<sup>6</sup> See [www.nishansystems.com/promontory](http://www.nishansystems.com/promontory)

<sup>7</sup> See [www.nishansystems.com/promontory/Promontory\\_Performance\\_Note\\_01.pdf](http://www.nishansystems.com/promontory/Promontory_Performance_Note_01.pdf)

artificially low, it produces an unreasonably high estimate for the real throughput of the system. The same chart also shows that the 7-day average packet loss rate was 0.1%, the 30-day average packet loss rate was 0.07% and the 90-day average packet loss rate was 0.58%. Had these tests been conducted when the network was operating at its long-term average packet loss rate, the results would have shown a significant reduction in available throughput. Note that these numbers are averages. By definition, the packet loss rate will be even worse than these numbers during a large portion of the time. This reduces the system throughput even further.

## Conclusions

We have examined the impacts of real-world packet loss rates on the large-scale transfer of data across the IP WAN. As we have seen, routed IP networks provide amazing flexibility and scalability but they do so at the cost of single-user performance. Single users who wish to move large amounts of data are continuously thwarted in their attempts to do so. Although IP routing offers the illusion of inexpensive data transfer over long distance, we have actually seen that the ability of a single user to move data significantly decreases as the distance and the amount of data are increased.

On the other hand, Fibre Channel over SONET provides a high performance and cost effective method of transferring large-scale data across the WAN. The secret to Fibre Channel's success really isn't such a secret—don't drop data. If you haven't dropped the data in the first place, there is no need to waste the precious bandwidth and time of the system by re-transmitting it.

In the end, as the requirements for large-scale data transfer increase, economics and practicality will dictate the solutions chosen by the end-user. Fibre Channel over SONET provides single user performance that dramatically exceeds anything available over the IP backbone. Furthermore, because it does not waste bandwidth and time on retransmission, it provides a significantly more cost-effective method of transferring large-scale data across the WAN.