HyperIP Frequently Asked Questions

Q: Why isn't TCP/IP adequate for storage applications over the WAN?

A: TCP employs very inefficient error recovery and session management techniques that provides less than 20% effective data throughput for most storage and data movement applications. This effective data throughput will degrade rapidly when there is packet loss, latency, network jitter, router buffer overflows, or bit errors. This results in major production window disruptions.

Q: How much will my effective data throughput improve with HyperIP?

A: Production implementations of HyperIP have shown effective data throughput increases ranging from 3 to 10 times than that of native TCP/IP. The performance increases realized with HyperIP is especially dramatic on circuits where packet loss, poor line quality conditions, or high latency exist. Native TCP/IP is incredibly inefficient under these conditions.

Q: Why is HyperIP's performance so radically better than that of native TCP/IP?

A: TCP is not an efficient protocol for high performance data transport, especially with any kind of packet loss, network jitter, bit errors, congestion, or distance. It gets worse as these types of network disruptions increase.

HyperIP leverages a production hardened time-proven protocol that minimizes the effects of packet loss, jitter, bit errors, and the distance latency. It utilizes dynamic and optimized window sizing, selective acknowledgement, retransmissions of lost packets only (versus TCP resending the entire data stream after the lost packet,) and adaptive compression that compresses concatenated blocks of data (versus individual TCP packets.) When combined with HyperIP dynamic real-time adjustments, protocol overhead is significantly reduced while data integrity is guaranteed end-to-end.

Q: How simple are HyperIP appliances to configure?

A: Very simply. You define HyperIP as a standard IP Gateway/Proxy

No complex and costly network installation

Uses your "existing" network IP infrastructure

No changes to storage or data movement applications

No special network equipment

No special network software

No special network management

HyperIP is completely transparent to all TCP based applications and typically implemented in pairs, with a unit at each end point of the network. Point-to-multipoint is also configurable with multiple appliances.

Q: How does HyperIP's production hardened "*shield*" help my application performance?

A: The performance of TCP and many UDP applications degrade significantly when circuits experience problems such as packet loss, bit errors, excess latency, jitter or changes in available bandwidth. Because the upper layer protocol of HyperIP has been designed and optimized for high-speed data transport, it includes features that minimize retransmission and dynamically adjust to variations in circuit conditions. Although the performance of HyperIP will degrade if circuit conditions become excessively bad, HyperIP's underlying protocol has proven to be very robust and far less sensitive to degrading circuit conditions than TCP and most UDP protocols. Even if variations in circuit conditions are rare, they can be very disruptive when they occur. HyperIP has been tested to 6% packet loss and up to 46,000 miles of latency with outstanding results.

Q: Once HyperIP is installed, what do I have to do to my network to add additional storage or TCP applications?

A: Nothing. Application aggregation improves ROI and lowers TCO of additional / add-on Storage Solutions. Users can invest in new storage applications with NO additional network investments required. All applications can concurrently leverage a common highly efficient IP network connection.

Q: How is HyperIP defined to an existing environment?

A: HyperIP is defined as a gateway or proxy at the server level. It is assumed that all data passing between servers would want to be optimized. If there are particular applications on a server that the user does not want to be optimized, they can be filtered out by IP address, well known port, etc.

Q: Are any application changes required to take advantage of HyperIP?

A: No. HyperIP is totally transparent to the applications, other than the fact they are achieving superior end-to-end performance throughput. Unlike implementations of TCP/IP Extensions, that require extensive system tuning and application changes to be made, application performance can immediately be improved with HyperIP, regardless of platform, operating system or network vendor.

Q: What makes HyperIP compression better than TCP router compression?

A: Compression is standard with all HyperIP appliances and significantly increases effective data throughput up to OC3. Compression ratios depend on the type of data that is being compressed. Production results have demonstrated a range of 2:1 up to 15:1. Compression throughput decreases at speeds above OC3.

HyperIP compression compresses data blocks (versus TCP packets) and efficiently aggregates data into blocks, then applies the compression. This puts more data into the HyperIP accelerated pipe.

The bottom line is that it takes less bandwidth to get the job done with HyperIP than without it. HyperIP compression reduces TCO and increases ROI.

Q: Compression is often ineffective at circuit speeds above 6 – 10 mbps. What circuit speeds can HyperIP compression scale to?

A: HyperIP has the unique ability to compress data in "blocks" of 4,000 to 64,000 bytes instead of packets of 1,500 bytes. HyperIP's efficiency in data aggregation and block level compression results in compression that is very effective even at speeds up to OC3 (155 Mbps).

Q: Are there any "preprocessing" data requirements with HyperIP?

A: No, there are no preprocessing steps required by HyperIP. Data directed to HyperIP is enhanced on the fly and passed on to the WAN router or switch for transmission over the network.

Q: Why can't I simply add bandwidth to achieve the same results as HyperIP?

A: Native TCP isn't a very efficient protocol for packet loss, high bit error rates, high latency, or long. Packet loss caused by high bit error rates, network congestion, and latency caused by physical distance of the circuit, will place a maximum limit on TCP file transfer rates. For example, at distances greater than 200 miles, TCP users typically see file transfer rates of approximately 4 to 6Mbps whether on a DS3 (45Mbps) or OC3 (155Mbps). Adding bandwidth doesn't overcome these physical and protocol limitations.

Q: Does HyperIP support standards-based IP infrastructures?

A: Yes, HyperIP runs over standard IP networks.

Q: I get 72% "utilization" on my lines today, so why do I need HyperIP?

A: It is far more accurate to measure data transfer efficiency in terms of "effective data throughput" rather than "line utilization". "Effective data throughput" represents actual data received at the end point. "Line utilization" statistics contain all the inefficient data retransmits that packet loss, high bit error rates, congestion, and the latency of long distances generate. It is not an accurate measurement of "pure data" transmitted.

Q: How can HyperIP save my organization money?

A: Because HyperIP increases effective data throughput 3 to 10 times, organizations are increasing throughput without adding bandwidth. Or as many HyperIP customers realize: "You mean I can get 3 to 10 times the performance of my existing DS3 network without adding additional, expensive circuits?" How much can you save?

Q: How does HyperIP handle congestion between two sites? For example, what will happen if the link/network between two sites is not dedicated or there is no bandwidth guaranteed?

A: HyperIP will modify its send rate, backing down to the receive rate of the remote HyperIP. Congestion will cause the receive rate to go down, since congestion will cause packets to either get dropped or be delayed. The sending side will try to recover from the congestion slowdown by cautiously, but continuously, attempting to increase the send rate. This Dynamic Rate Adjustment has the effect of maximizing overall WAN throughput by accelerating data on the HyperIP segment of the circuit.

Q: How does HyperIP complement my storage environment?

A: HyperIP accelerates data transfer for storage networking applications. It is particularly suited for data replication situations where the "other end" is located in another state or across the country. HyperIP has been proven with EMC's SRDF, Celerra and Centera, NetApp Snapshot applications, Symantec Volume Replicator (VVR), and NSI's Double Take, plus many others. HyperIP also complements and enhances standard file sharing protocols like File Transfer Protocol (FTP).

Q: What storage applications are tested and supported with HyperIP?

A: The HyperIP supported application matrix is found on the NetEx website.

Q. If Data Replication over long distances is available today, why do I need HyperIP?

A: Remember, HyperIP doesn't make the data go any faster; it increases performance by increasing *effective data throughput*. Data replication still requires volumes of data to be transferred. HyperIP moves these volumes faster by using all available bandwidth. HyperIP's ability to shield applications from the effects of variations in circuit conditions also makes the performance of Data Replication more consistent and robust.

Q: What compression algorithm does HyperIP utilize?

A: HyperIP uses an adaptive compression algorithm that has the ability to continually test the data to see if it lends itself to compression. If the data is already compressed, HyperIP will not re-compress it. The ratio of data compression depends on the compressibility of the data transmitted. An average compression ratio is 2.4X, but HyperIP has demonstrated the ability to provide compression ratios up to 17X for some real production data streams.

Q: How does HyperIP maintain data integrity?

A: Standard UDP or TCP protocol checksums are used when receiving or transmitting data between the user application node and the local HyperIP unit. Additionally, the transport protocol between the HyperIP units uses a proven block-numbering acknowledgement and retransmission mechanism along with UDP checksums to ensure data is not lost or corrupted over the wide area network.

Q: What management support is available for HyperIP (web GUI, telnet, etc.)?

A: A web browser GUI is available for configuring the unit. SNMP support is available for gathering NIC interface and TCP/IP stack statistics. A HyperIP Private Enterprise MIB is also available.

Q: Does HyperIP support any type of IP Class of Service (COS) or QOS?

A: Many enterprises express the difficulty of implementing and monitoring COS. HyperIP's Dynamic Rate Adjustment provides for efficient sharing of HyperIP transfers and other applications utilizing the same WAN circuits. This results in a type of COS/QOS that does not require changes or adjustments to network configuration parameters in other network elements or applications software. HyperIP also has a Rate Limiting feature that can limit the maximum bandwidth that HyperIP will send data over a selected circuit. Rate limiting insures that high

volume transfers do not block other application traffic over circuits. Rate limiting can be configured to operate by time of day.

Q: How is HyperIP sold?

A: HyperIP is licensed for use via via two pricing components. The first component is the Initial License Fee (ILF). The ILF sets up the product license and is invoiced upon shipment. The second pricing component is the Annual Right to Use (RTU) fee. The RTU fee allows usage of the product. The first Annual RTU is invoiced upon shipment, and is then invoiced annually thereafter. Refer to NetEx's HyperIP License Agreement.

Q: How is HyperIP Supported?

A: For over two decades, NetEx has met the needs of the world's most demanding businesses. Our customers depend on trouble-free high performance data transport. And they demand the most responsive support and a single point of contact for all performance, service and enhancement requirements. These have always been the hallmarks of NetEx professional and customer services.

NetEx provides:

Experienced NetEx technical resources are on call from 8 AM to 5 PM U.S. Central time, M-F, to provide toll-free assistance with product questions and issues.
After regular business hours, on-call NetEx technicians assist in emergency situations.

NetEx offers professional services whenever needed on any and all product-related issues. NetEx professional services range from installation assistance to system and network configuration.

Network Executive Software, Inc. (NetEx)
6420 Sycamore Lane North Suite 300
Maple Grove, MN 55369
+1 763-694-4300 or 1-888-604-5573

www.netex.com
mailto:hyperip@netex.com