



HyperIP enables VMotion over WAN

A DeepStorage.net Labs Validation Report

Much has been written about how x86 server virtualization reduces IT department capital expenditures on servers and the data center resources to house them. Organizations that have fully embraced the process discover that the savings in operational expenses from virtualization platform features such as VMotion, which allows administrators to move running virtual servers from host to host for maintenance or load balancing, are at least as significant.

For most users, this flexibility is limited to moving guest machines between hosts in the same data center as VMware's best practices require a dedicated high bandwidth, low latency, network for VMotion. NetEx's HyperIP provides WAN acceleration for TCP/IP protocols and applications by reducing the effect of network latency making more of the bandwidth useable.

NetEx has now packaged HyperIP as a VMware virtual appliance expanding the cost advantage of HyperIP over competing appliance based solutions. The virtual appliance implementation of HyperIP can accelerate up to 800Mbps of TCP/IP traffic.

Network Executive Software hired DeepStorage Labs to validate that a HyperIP virtual appliance can accelerate VMotion across typical WAN links.

The Bottom Line

Corporations with multiple geographically dispersed data centers want to easily migrate virtual machines between locations. While VMotion is the best solution for server migration, it is too slow over all but the shortest links. The problem isn't a lack of bandwidth but the inefficient way TCP connections handle latency. We tested NetEx's HyperIP WAN acceleration software and found:

- HyperIP accelerates VMotion across WAN links up to 10X
- HyperIP accelerated migration times over even low-cost networks are comparable to migrations within the data center
- HyperIP's virtual appliance architecture requires no changes to the underlying network simplifying implementation and reducing cost
- HyperIP is a valuable addition to the system architect's bag of tricks when data, or virtual systems, must be transferred between locations as in a DR scenario

What We Did

To see how sensitive VMotion is to WAN impairments like latency and data loss, we migrated a Windows Server 2003 virtual server between a pair of vSphere (ESX 4.0) servers connected through a WAN emulator. The emulator was programmed to provide bandwidth, (cont. next page) latency, packet loss and jitter levels consistent with the types of links medium and large enterprises would use between their primary and secondary data centers. The values for these line impairments were based on published service level agreements from carriers providing data transport services.

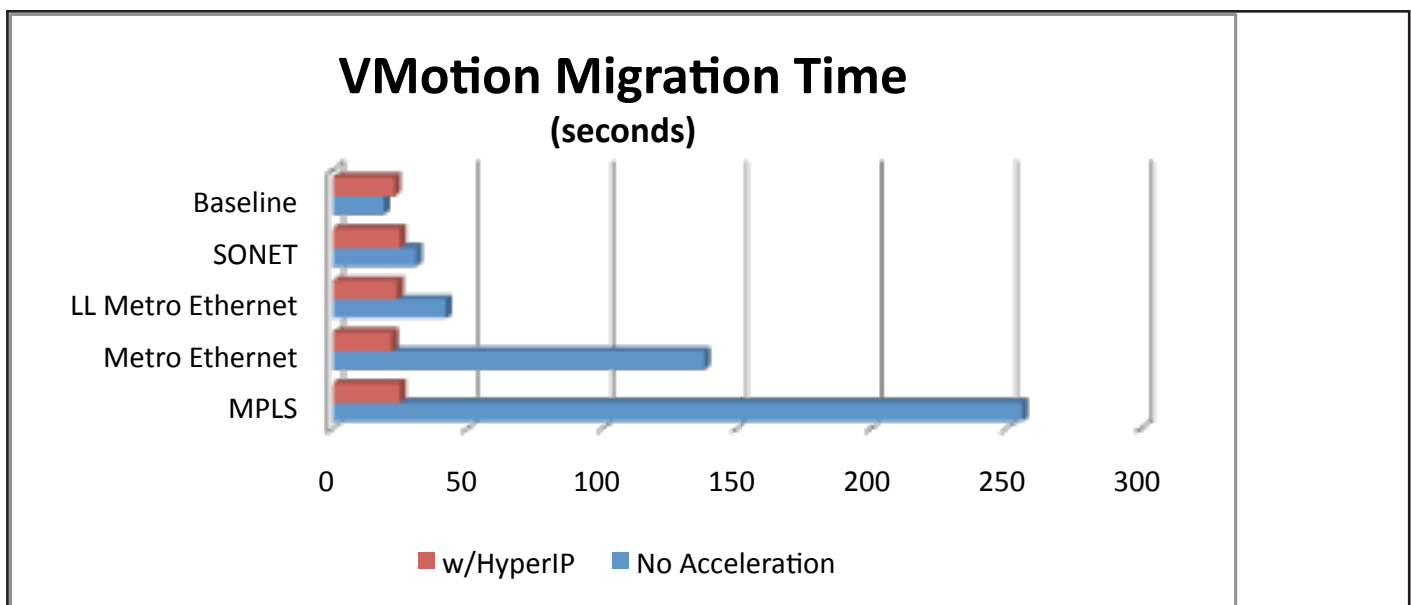
The migration of the test VM across each emulated WAN link was repeated four times to minimize the impact timing and other one time variations had on the results. We then repeated the process using HyperIP running on virtual machines in the same VMware hosts to accelerate the process, again repeating the migration four times for each emulated WAN link.

The emulated WAN links used were:

Type of Link Emulated	Network Conditions Bandwidth, Round Trip Latency, Packet Loss, Jitter
Baseline	1Gbps, none, none, none
Low Latency Sonet link (NY-Jersey City)	1Gbps, 5MS, .001%, none
Low Latency Metro Ethernet	1Gbps, .5MS, .005%, none
Metro Ethernet	1Gbps, 12MS, .1%, none
MPLS VPN (NY-Chicago)	100Mbps, 25MS, .05%, 3MS

What We Found

As expected, we found that regardless of the bandwidth available, VMotion performance fell off dramatically as even small amounts of latency and/or packet loss were introduced to the emulated WAN link. HyperIP significantly boosted VMotion speed on all but our baseline link. In fact, HyperIP accelerated migration performance was comparable to the baseline – even on reduced bandwidth links with substantial impairments.



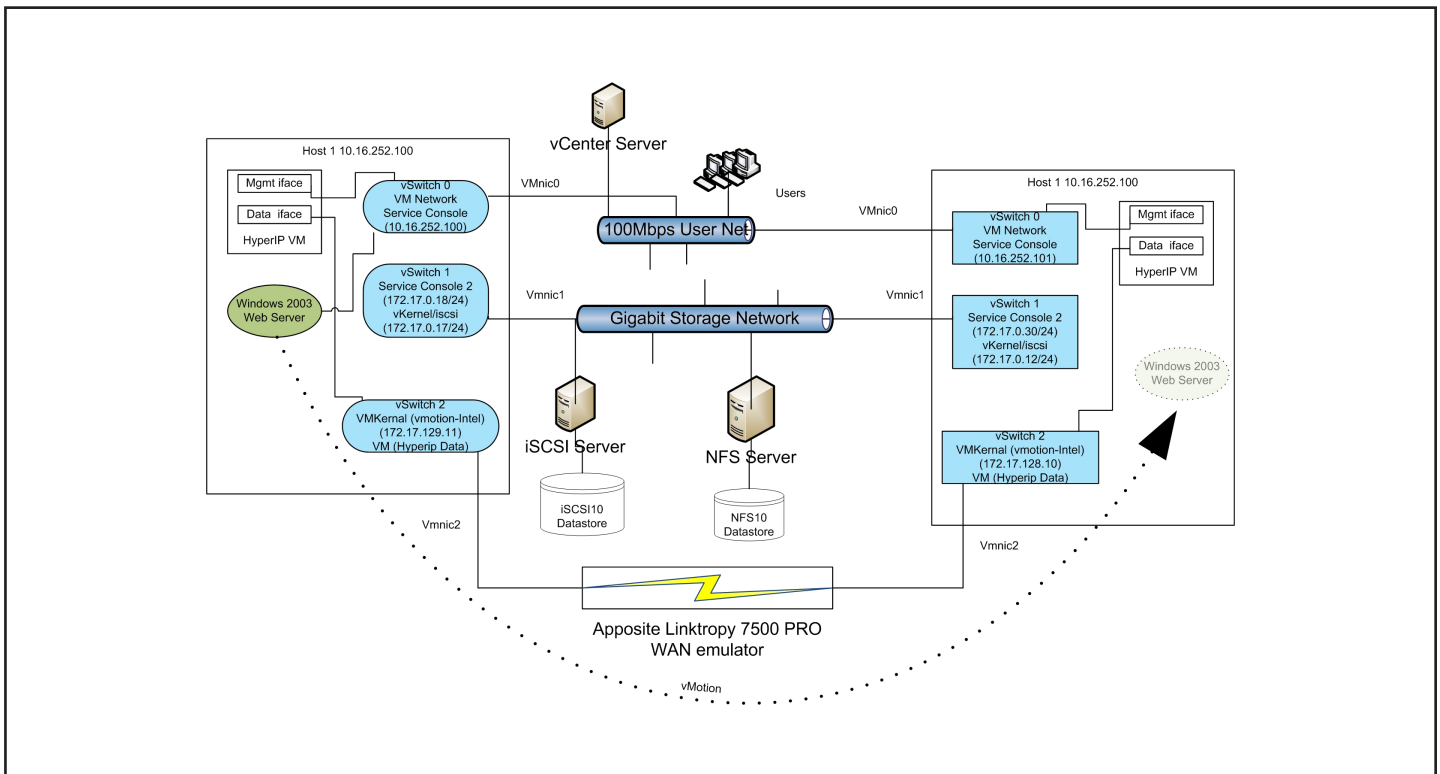
Other than on our baseline test, where processing the data through HyperIP added a few seconds to the migration time, HyperIP significantly reduced VMotion time across the board. Regardless of network conditions, HyperIP enabled VMotion migrations to complete in less than 30 seconds while keeping the virtual server running.

We were most impressed with the results on the 100Mbps emulated MPLS link where migrating our test server took 1/10th the time with HyperIP than it did without. This exceeds the vendor's claim that their product increases transfer rates by a factor of up to six.

Test Results Data

	Baseline	SONET	LL Metro Ethernet	Metro Ethernet	MPLS
No Accel	00:00:19	00:00:31	00:00:42	02:16	00:04:16
w/HyperIP	00:00:23	00:00:25	00:00:24	00:00:22	00:00:24

The Test Configuration



The TCP Problem

The Transmission Control Protocol (TCP) provides assured delivery of data by having the receiving node acknowledge the receipt of data packets and having the sending station retransmit packets when it hasn't received an acknowledgement in the allowed time. While TCP provides reliable communications, it was designed in the 1970s and has difficulty effectively using today's networks, which are much faster, and less error-prone.

The biggest issue is that TCP can only send 64KB before it receives an acknowledgement from the destination station. Since 64KB takes only ½ a millisecond to transmit across a gigabit link, even 1ms of round trip latency could cut TCP performance across a metro Ethernet connection by a factor of four.

TCP's designers assumed that the network on which it would run would be both slow and congested. As a result, TCP sessions start slow and ramp up their data rate. TCP also assumes all packet losses are the result of congestion so it starts the slow ramp up again with each lost packet resulting in more idle time on the link.

This is reflected in our test results as adding 5ms of one-way latency and a moderate .005% packet loss extended vMotion times 56% even with a full 1Gbps link available.

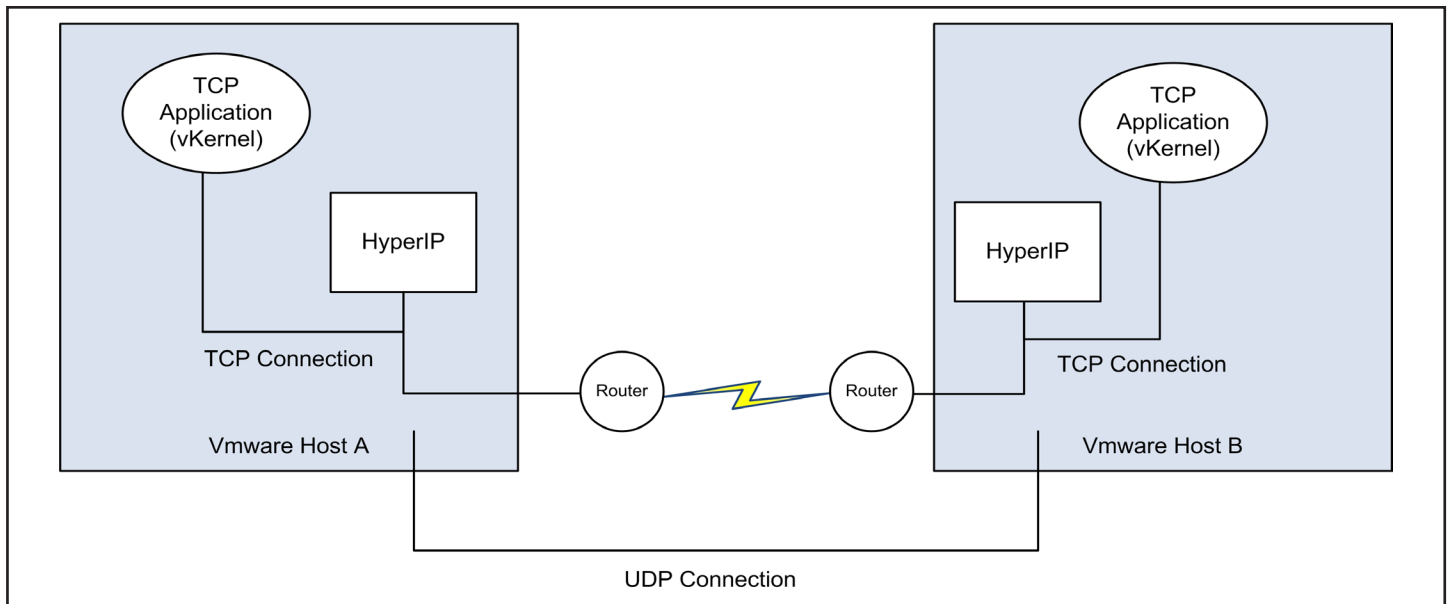
How HyperIP Works

HyperIP boosts the effective bandwidth of WAN links using two basic techniques. First it terminates the TCP connection from the applications at each end of the connection at the HyperIP virtual appliance. Since these TCP connections run across short, and essentially error-free LAN links, TCP can fill the available bandwidth.

Second, it sends the data across the WAN link between a pair of HyperIP appliances using UDP, which is not latency sensitive. NetEx adds a proprietary error correction protocol to UDP which adjusts the amount of data "in flight" and block size to accommodate changing WAN conditions.

Rather than assuming that all packet loss is due to link congestion, HyperIP manages the amount of data it sends avoiding congestion by adjusting the data sent rate.

In addition, HyperIP compresses data before segmenting it into packets to send across the net so each packet can be sent at the optimal size. It also monitors the effectiveness of the compression algorithm disabling compression when the compressed data is more than 80% the size of the source data.



About DeepStorage.net

DeepStorage is dedicated to providing independent information on storage and other inside the data center technologies to help organizations solve real world problems. Our 25 years of experience testing, writing about and most importantly implementing systems in the real world informs everything we do, leading to the development of our hands-on philosophy.

Our lab is at the heart of the hands-on philosophy providing a facility where we evaluate hardware and software products away from the influence of the vendors. We perform validation, usability, power use and performance testing in individual and competitive tests for publication and the private use of our clients. Our reviews have appeared in major publications including *Network World*, *Network Computing* and *InformationWeek*.

The lab is equipped with Fibre Channel and iSCSI storage, gigabit and 10 gigabit networking gear and multiprocessor servers from IBM, Dell, and HP allowing us to test applications under Windows and Linux, including virtual server environments. We believe a mix of real world and synthetic testing reveals more than either alone so, we also use dedicated test equipment from Spirent, Apposite, Shunra and others to provide repeatable testing.

The Test Environment

All testing was performed at DeepStorage Labs in Purchase, NY. While NetEx personnel provided configuration assistance via telephone, DeepStorage.net determined all test parameters and no NetEx personnel were present during testing.

A pair of Dell PowerEdge 2950 servers with dual quad-core Xeon 5510 processors and 16GB of memory were used as the vSphere ESX servers. Each was configured with separate gigabit Ethernet controllers for the VM network, storage and VMotion. Storage was provisioned on Dell PowerEdge 1850 servers with StarWind for iSCSI and Red Hat Enterprise Linux for NFS.

An Extreme Summit 7i gigabit Ethernet switch was used with multiple vlans. WAN emulation was provided by an Apposite Linktropy 7500 PRO.

