HyperIP Accelerates Hyper-V Live Migration

A DeepStorage.net Labs Validation Report

Over the past few years organizations have been adopting server virtualization to reduce capital expenditures by consolidating multiple virtual servers onto a single physical server. Once they've begun the process, organizations soon discover that the ability to migrate running virtual servers from one host to another for maintenance or load balancing increases their operational efficiency as much as server consolidation reduces capex.

To date, this flexibility has been limited to moving virtual guest servers between hosts in the same data center as hypervisor vendors like VMware and Microsoft have only supported live migration over dedicated, high bandwidth, low latency connections.

NetEx's HyperIP is a virtual appliance that accelerates TCP/IP protocols and applications on WAN links by reducing the impact of latency and packet loss making more of the network's bandwidth usable. A HyperIP virtual appliance can accelerate up to 800Mbps of TCP traffic.

NetEx hired DeepStorage Labs to validate that HyperIP virtual appliances can accelerate the live migration of a virtual Windows server over typical WAN links between data centers.

The Bottom Line

Corporations with multiple geographically dispersed data centers want to easily migrate virtual machines between locations. While Hyper-V Live Migration is a good 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 Live Migration across WAN links up to 16X
- HyperIP accelerated migration times, even over low-cost networks, are comparable to migrations within the data center
- HyperIP's virtual appliance architecture requires minimal 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 are transferred between locations for disaster avoidance or other purposes

What We Did

To test how sensitive Hyper-V Live Migration is to WAN impairments like latency and packet loss, and how much HyperIP could ameliorate the impact thereof, we migrated a Windows Server 2003 web server between a pair of Hyper-V host systems across a WAN emulator. The WAN emulator was configured to provide bandwidth, latency and other conditions consistent with the links medium and large enterprises would use between data centers. The impairments used were based on published service level agreements from carriers providing these services.

The emulated 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

Note: These are the same conditions used in our previous report for NetEx using VMotion.

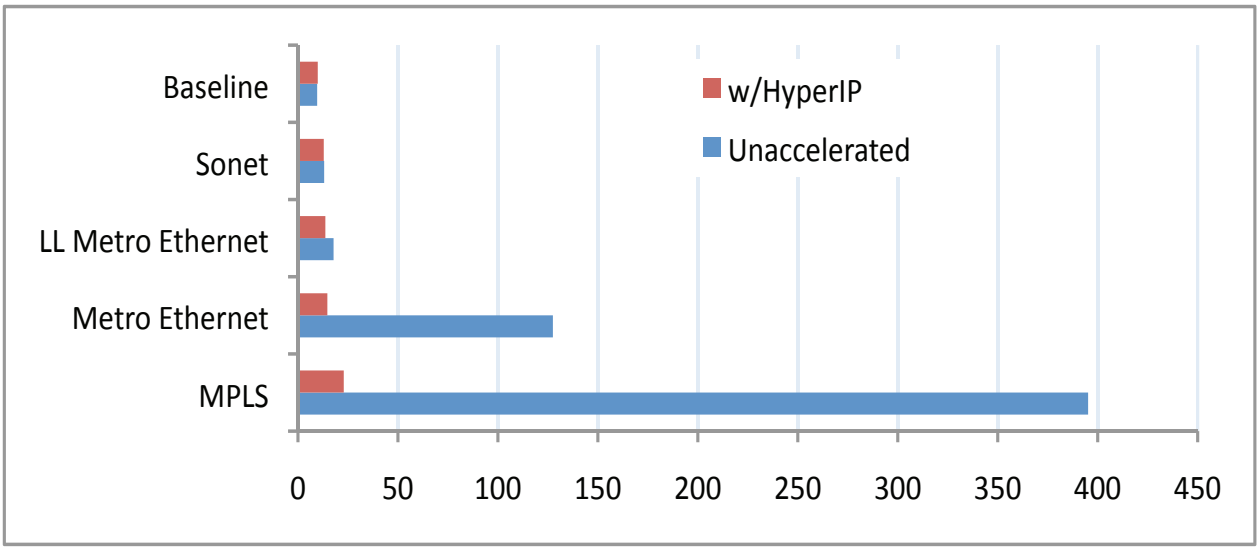

Live Migration Time
(in seconds)

What We Found

As expected, we found that regardless of the bandwidth available, Live Migration performance fell off dramatically as latency and/or packet loss were introduced to the emulated WAN link. HyperIP significantly boosted Live Migration speed on all but the best of our emulated WAN links. 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 reduced Live Migration time across the board. Regardless of network conditions, HyperIP enabled Live Migrations to complete in less than 30 seconds while suffering fewer and shorter outages than the un-accelerated case.

We were most impressed with the results on the 100Mbps emulated MPLS link where migrating our test server took 1/17th the time with HyperIP than it did without. While a higher bandwidth link would be preferred, this demonstrates that migration across slower links is possible and exceeds the vendor's claim that their product increases transfer rates by a factor of up to six.

The Test Configuration

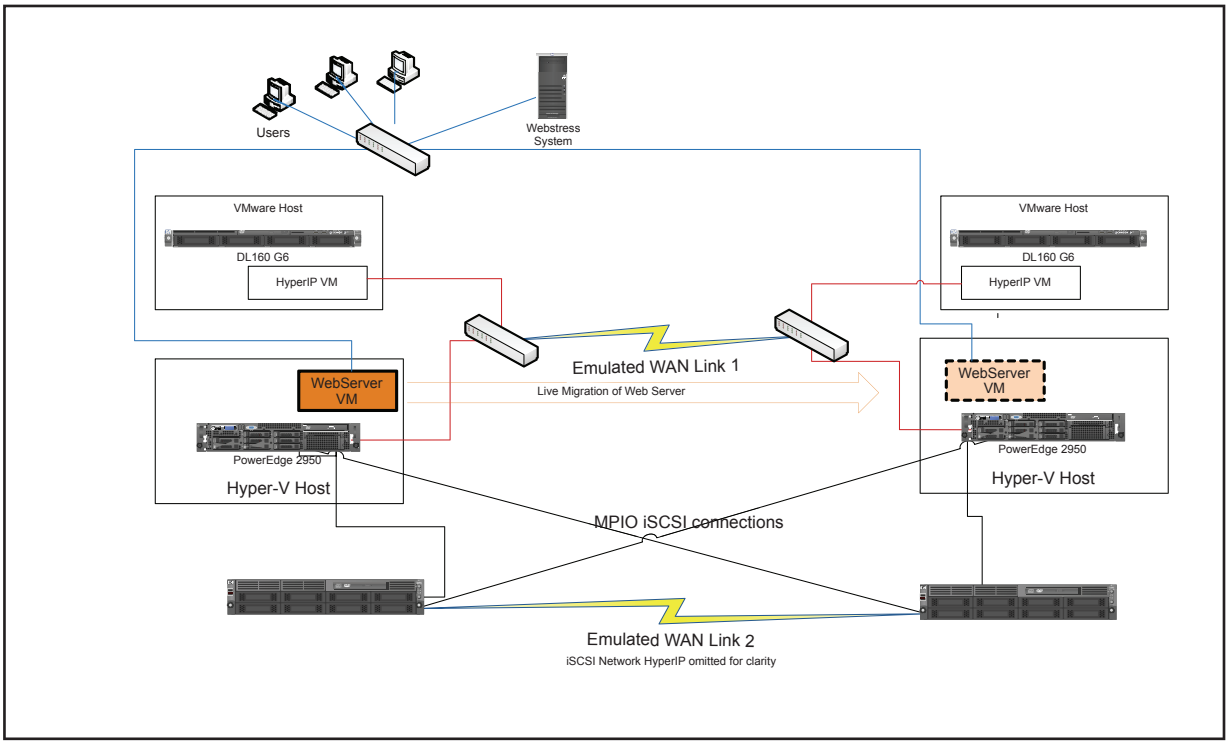

Test Result Data
(in seconds)

	Baseline	Sonet	Low Latency Metro Ethernet	Metro Ethernet	MPLS
Unaccelerated	9.6	13.1	17.8	127.5	395.2
w/HyperIP	9.9	12.9	13.7	14.7	22.9

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 typically only send 64KB before it receives an acknowledgement from the destination station. Since 64KB takes only ½ 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 Live Migration 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 software virtual appliances using UDP, which is not latency sensitive. NetEx adds a proprietary data delivery protocol to UDP which adjusts the amount of data "in flight" and the block size to accommodate changing WAN conditions.

Rather than assuming that all packet loss is due to link congestion, HyperIP manages the link adjusting the rate at which it sends data to match varying line conditions.

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.

How HyperIP Works

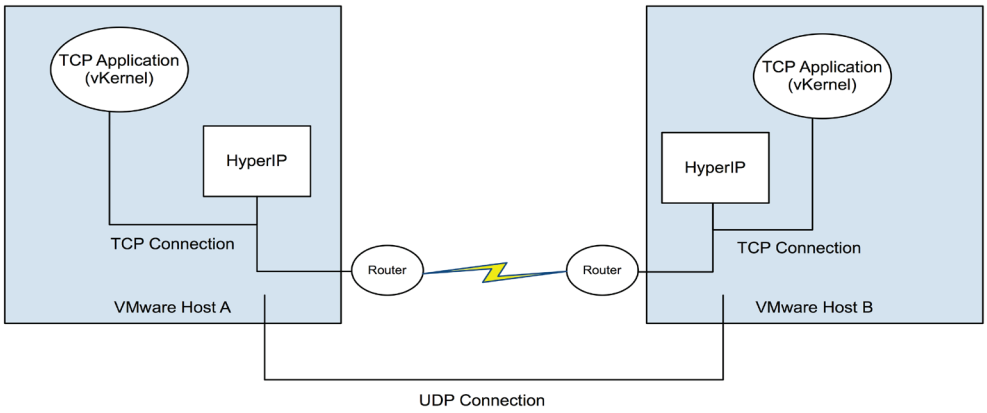

About DeepStorage.net

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 the latest networking gear, including FCoE, and multi-processor 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. 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 Lyndhurst, NJ. DeepStorage.net determined all test parameters and no NetEx personnel were present during testing.

A pair of Dell PowerEdge 2950 servers were configured as a Windows Server 2008 R2 cluster with Cluster Shared Volumes. Each was configured with separate gigabit Ethernet controllers for the user network, iSCSI storage and Live Migration. Storage was provisioned on a Dell PowerEdge 1850 server and an HP DL380 G5 using the StarWind iSCSI target for Windows creating a high availability cluster.

An Extreme Summit 7i gigabit Ethernet switch was used with multiple vlans. Apposite Linktropy 7500 PRO was used to emulate WAN links for both the Hyper-V Live Migration link and the iSCSI network link between the two simulated data centers. A system running Paessler's Web Server Stress Tool was continuously accsessing the web server as it was migrated across the emulated link.