



Application Note: HyperIP WAN Optimization Improves Replication for Dell EqualLogic SANs

Introduction

The HyperIP® WAN optimization virtual appliance that operates on a VMware ESXi or Microsoft Hyper-V server to boost the performance of replication applications from vendors such as Dell EqualLogic. HyperIP mitigates TCP performance issues that are common when moving storage data over WAN connections because of bandwidth restrictions, latency due to distance and/or router hop counts, packet loss, and network errors.

TCP link bandwidth issues and limitations

Several characteristics of TCP/IP cause applications to perform poorly over lossy bandwidth and long distances:

Window Size

Window size is the amount of data allowed to be outstanding (in-the-air) at any given point-in-time. The available window size on a given bandwidth pipe is the rate of the bandwidth times the round-trip delay or latency. Using a cross-country OC-3 link (approximately 60 ms based on a total 6000-mile roundtrip) creates an available data window of $155\text{Mbps} \times 60\text{ms} = 1,163\text{Kbytes}$. A DS3 satellite connection (540 ms roundtrip) creates an available data window of $45\text{Mbps} \times 540\text{ms} = 3,038\text{Kbytes}$.

When this is contrasted with standard and even enhanced versions of TCP, there is a very large gap between the available window and the window utilized. Most standard TCP implementations are limited to 64Kbytes windows. There are a few enhanced TCP versions capable of using up to 512Kbytes or larger windows. Either case means an incredibly large amount of "dead air" and very inefficient bandwidth utilization.

Acknowledgement Scheme

TCP causes the entire stream from any lost portion to be retransmitted in its entirety. In high bit-error-rate (BER) scenarios this will cause large amounts of bandwidth to be wasted in resending data that has already been successfully received, all with the long latency time of the path. Retransmissions may be subjected to the performance penalty issues of "Slow Start".

Slow Start

TCP data transfers start slowly to avoid congestion due to possible large numbers of sessions competing for the bandwidth, and ramp-up to their maximum transfer rate, resulting in poor performance for short sessions.



Session Free-for-all

Each TCP session is throttled and contends for network resources independently, which can cause over-subscription of resources relative to each individual session.

The net result of these TCP limitations is very poor bandwidth utilization. The typical bandwidth utilization for large data transfers over long-haul networks is usually less than 30% and more often less than 10%.

Implications for EqualLogic Replication on TCP/IP networks

TCP/IP is included in most operating systems and has become the preferred transport for moving data over the WAN. Like many other replication applications EqualLogic Replication uses TCP/IP for moving data over the WAN.

Latency and packet loss will affect the application's throughput and overall performance.

It should not be a surprise to see the application performance degrade when TCP is used for many of the reasons mentioned above. It should be noted, it is not a fault of the Dell EqualLogic Replication application but more a symptom of using the TCP transport technology.

The cost effective high performance solution for iSCSI SAN Recovery: Dell EqualLogic Replication and HyperIP

HyperIP was designed to move large amounts of data over big bandwidth and long distance, and to be highly efficient regardless of the BER congestion, or jitter. HyperIP is a standard TCP/IP network node requiring no modifications to LAN / WAN infrastructures requiring no proprietary hardware. It provides transparent "acceleration" across WANs.

HyperIP provides the following benefits:

Window Size

The HyperIP transport protocol keeps the available network bandwidth pipe full. The results are over 90% efficient link utilization. It eliminates the discrepancy between maximum available bandwidth and the results provided by native TCP/IP.

Acknowledgement scheme

HyperIP transport protocol retransmits only the NAK'd segments and not all the data that has already been successfully sent.

Slow Start

Configuration parameters allow HyperIP to start transmissions at a close approximation of the available session bandwidth.



Dynamic Adjustments

When feedback from the receiver in the acknowledgement protocol is received, HyperIP quickly "zeroes-in" on the appropriate send rate for current conditions.

Session Management Pipeline

HyperIP allows traffic from multiple TCP sessions to be aggregated over a smaller set of connections between the HyperIP devices, enabling a more efficient use of the bandwidth and less protocol overhead acknowledging many small messages for individual connections.

Adaptive Block-level Compression of Data

HyperIP applies, as required, block-level compression of the optimized data allowing more data to be shipped over the WAN link, without having to upgrade the circuit. This software-based algorithm is much more efficient than packet-level compression and allows the Veeam Backup & Replication application to offload compression for more efficiency in the replication process.

Real-world installation

Decision Lens Inc. a provider of desktop and web-based collaborative decision-making software, implemented HyperIP WAN optimization virtual appliances to improve the performance of its IP SAN system's volume replication function to duplicate production system data to its cross-country disaster recovery site.

Decision Lens works with organizations to bring structure and quality to key decisions. Its solutions manage the strategic alignment of goals and priorities with investments in people, projects, products and suppliers for a more accurate, repeatable and transplant process for decision making.

The company was using Dell EqualLogic SAN replication in an attempt to inexpensively duplicate its all Linux-based virtual machines in a VMware infrastructure cluster in Ashburn, Virginia to a DR site in Phoenix, Arizona. Latency issues created long replication times resulting in the DR site always being several hours out of sync with production.

The company discovered that the SAN replication on its own was ill-suited to the 80+ms latency of the cross-country trip across the public internet. It would not make efficient use of the available bandwidth and would only move data at about 9Mb/s.

Decision Lens now replicates two 1,500GB Dell EqualLogic SAN system volumes, as well as 500GB of data between the Linux backup systems. Replication performance between the sites with HyperIP's virtual appliances was greatly improved, so much that Decision Lens is rarely ever more than an hour out of sync at the DR site. The company pays for bandwidth based on a 95th percentile average, so being able to control the



bandwidth usage between the SAN systems, via the HyperIP rate limiting function, allowed them to recoup their annual investment in four months in bandwidth cost savings alone.

There was no need for additional hardware - each end is served by a VM that lowers the cost of management - plus HyperIP leverages the built-in reliability of its own VMware cluster. The company plans to expand its use of HyperIP as it scales the amount of data it synchronizes and as the number of worldwide locations grows.

Conclusion

Decision Lens discovered that HyperIP WAN Optimization virtual appliance and Dell EqualLogic can significantly enhance their replication performance by minimizing the effects of network disruptions for improved data replication across shared WAN connections.

Deploying HyperIP to optimize replication data between a company's storage system and a geographically dispersed DR site helps ensure that their business-critical information is protected and moved on time, every time. Decision Lens was able to improve their replication performance and bandwidth usage by implementing HyperIP, making data transfers extremely fast and secure.

NetEx has provided high-end networking tools for over 25 years to some of the world's largest and most sophisticated organizations. Today, NetEx focuses on virtual appliance-based WAN optimization software that is affordable and practical for solving WAN throughput issues for users, solution providers and IT service companies. The company is based in Minneapolis, MN. For more information about NetEx, visit www.netex.com or call +1-763-694-4300.