

CA ARCserve Backup Application Note

Introduction

HyperIP® for VMware is a software-based virtual WAN Optimizer that operates on a VMware ESX or ESXi server to boost the performance of backup and replication applications from vendors such as CA. HyperIP mitigates TCP performance issues that are common when moving storage data over wide area network connections because of bandwidth restrictions, latency due to distance and/or router hop counts, packet loss, and network errors.

CA ARCserve Backup is a comprehensive, distributed storage management solution for distributed and multiplatform environments. The application can back up and restore data from all the machines on your network, (including machines running Windows, UNIX, NetWare, and Linux) using optional client agents. CA ARCserve Backup also provides media and device management utilities.

TCP link bandwidth issues and limitations

Several characteristics of TCP/IP cause it 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. Each retransmission is additionally 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 issues 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 ARCserve Backup on TCP/IP networks

New regulations, business continuity, and disaster recovery has led to a surge of replication applications over the WAN. Man-made (9/11, blackouts, human error) and natural (hurricanes, earthquakes, tornados, firestorms) disasters have driven demand. TCP/IP has become the preferred WAN protocol of choice. The market perceives bandwidth is essentially available or “free”, because the TCP/IP WAN bandwidth already exists for interactive traffic. Conventional wisdom is that replication occurs at night or on weekends. This is when the majority of users are not utilizing the network. This allows already existing TCP/IP bandwidth to be leveraged by the replication applications without negatively impacting current applications. Dedicated, separate replication WANs are not required. This also eliminates separate WAN management. Additional bandwidth implemented for the replication applications could be shared by the interactive TCP/IP desktop-based applications. The facts show that TCP/IP bandwidth is neither free nor is there typically enough to accomplish the replication task in the window of time allotted. TCP/IP bandwidth utilization and network issues are rarely taken into account in calculating bandwidth requirements. The most likely result is a bandwidth shortfall. This means either the replication task cannot complete within the window of time allotted, or there is a perception that the user must buy more bandwidth, neither of these scenarios is acceptable.

The cost effective solution: ARCserve Backup and HyperIP®

HyperIP® was designed specifically for large amounts of data over big bandwidth and long distance, 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 and 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 backup and recovery application to offload compression for more efficiency in the offsite backup process.

HyperIP's results with ARCserve Backup over IP WANs

Performance results with ARCserve backup have been outstanding. ARCserve Backup, over HyperIP, achieves bandwidth utilization consistently exceeding 90% from distances of hundreds of miles (with high bit error rates on dirty lines) to as far as transcontinental distances.

❑ *What does this mean to ARCserve users?*

ARCserve on native Ethernet TCP/IP fabrics with HyperIP® is now one of the highest possible WAN throughput options for offsite backup. ARCserve Backup is no longer limited by TCP throughput issues, shortening backup window times. Packet loss and Out-of-order packets are virtually eliminated by HyperIP allowing for full utilization of the WAN link. Latency is mitigated to almost nothing making the WAN perform as well as a LAN. Compression on HyperIP allows the users to move more data over the WAN, without adding more bandwidth.

Summary and Conclusion

NetEx and CA have a common goal in ensuring total protection and recovery of data using leading edge, but proven technology, in a cost effective solution delivery.

The CA ARCserve Family of Products, from CA Technologies, gives you control over your changing business by delivering total protection, recovery and availability for systems, applications and data. The CA ARCserve Family of Products supports many customer initiatives including: Managing data growth caused by day to day operations; Application deployment, upgrade or expansion; Server virtualization; Remote office expansion; and Disaster Recovery planning.

With more than 800 people worldwide committed to delivering innovative products, our goal is to provide a complete strategy to manage your backup and recovery responsibilities in both virtual and physical environments. For more information about ARCserve products, visit <http://arcserve.com>.

NetEx Software, Formed in 1999 as a spin-off of Storage Technology Corporation (StorageTek®), privately-held NetEx provides the world's fastest WAN optimization software in the industry, along with guaranteed data delivery. As a VMware Technology Alliance Partner, NetEx's VMware Ready HyperIP WAN optimization software is leading the way in demonstrating impressive performance results for supercharging VMware applications worldwide.

NetEx works with some of the world's largest and most sophisticated organizations, including some of the most prestigious providers of financial, transportation, government, and telecommunications services. For more information about NetEx, NetEx/IP or HyperIP, visit www.netex.com.