Dragon Slayer Consulting

The Death of the WAN Optimization Hardware Appliance: R.I.P.

Marc Staimer, President & CDS of Dragon Slayer Consulting
marcstaimer@comcast.net 503-579-3763

Introduction

WAN optimization has been solving wide area networking performance issues for nearly a decade. Depending on the vendor and the technology, it can solve serious intractable throughput problems. So what could possibly be killing off WAN optimization? The answer is nothing. WAN optimization is not dying. However, the purpose-built or even white box hardware appliance that hosts WAN optimization software is another story. The demise of WAN optimization hardware appliances is a direct result of server virtualization and more specifically the virtual server appliance (VSA). Starting with VMware vSphere, and to a lesser extent Microsoft Hyper-V as well as Citrix XenServer, the virtual server appliance makes the WAN optimization hardware appliance quite frankly, as obsolete as the buggy whip. A brief examination of hardware appliances problems and how the VSA solves those problems clearly shows why the hardware appliance is dying.

WAN Optimization Hardware Appliance Problems

When running WAN optimization on a purpose-built or white box appliance, there are both operational problems and financial ones. The four primary operational problems include:

1. Complicated appliance availability;
2. Obsolescence;
3. Deployment;
4. Maintenance, hot fixes, upgrades.

❍ *WAN Optimization Hardware Complicated Appliance Availability Problem*

WAN optimization is usually implemented to support mission critical applications. Meaning that the WAN optimization appliance cannot go down otherwise that mission critical application (backup, replication, mirroring, file transfer, etc.) will likely fail. The net result is a minimum of two WAN optimization appliances required at each location. Unfortunately, it's not that simple.

The complexity of providing highly available WAN optimization appliances increases exponentially depending on the fault tolerance required. If there is "zero" downtime tolerance, the routing to the failover appliance must be nearly instantaneous. This is no easy task. Each appliance must communicate a heartbeat to the other and take over immediately when a failure occurs. It must also update the DNS router. Trivial, this is not.

If some time lag between an appliance failure and failover is acceptable, less heavy lifting has to be done by the appliances. There must still be a mechanism for the second appliance to take over for the failed one and there must be some mechanism to update the DNS router or for the failover appliance to take on the TCP/IP address of the failed unit.

In both cases, a second hardware appliance is required. The second hardware appliance also requires additional cables, switch ports, (possible additional switches,) rack space, floor space, power, and cooling. But it is the added manually intensive management of the additional appliance that makes it so maddeningly difficult. All implementations, operations, management, fixes, upgrades, etc. are duplicated as well as the additional tasks of setting up, operating, managing, and maintaining the HA.

❍ *WAN Optimization Hardware Appliance Obsolescence Problem*

The rationale behind purpose-built hardware is that it works faster, is more efficient, has higher throughput, etc., than standardized x86 server hardware. Historically, that has been true. Not any

longer. Advances in multi-core CPUs, integrated memory controllers, and silicon fabrication have essentially erased most of those advantages. More importantly, the specialized purpose-built hardware cannot match the ongoing billions in investment dollars by Intel and AMD in ongoing x86 advances. The fact that limited run specialized purpose-built hardware has no possibility of being as inexpensive as mass-produced x86 server hardware, makes it crystal clear that any advantages are fleeting at best.

The most common WAN optimization appliance option is the x86 white box. The argument most often advanced is that the white box has all of the x86 advantages and it costs less than brand name hardware. Although the facts are correct, the logic is flawed by omissions. White box x86 server hardware and software is often specifically configured to run the WAN optimization software. That specific configuration makes white box appliances severely limited as a potential server for anything else. Therefore it has no practical capability of ever being repurposed to run other software. Nor does the vendor allow any other hardware to be certified to run the WAN optimization software (takes too much time and money to certify and support ongoing), restricting the WAN optimization appliance to hardware supplied by the WAN optimization vendor. This is where the obsolescence factor comes into play.

For the WAN optimization vendor to capture any reasonable margins they must order substantial inventory or at least contractually commit to it. This puts off a new generation of hardware until the inventory is consumed (otherwise there is a significant financial write down). Even then new hardware generations must be tested and certified. The results are WAN optimization x86 white box hardware appliances that always significantly lag the latest generation of server hardware, regardless of vendor promises.

The third WAN optimization hardware appliance variation is the hybrid (combinations of x86 and purpose-built silicon). The hybrid conceptually takes the best of both worlds while theoretically limiting the downside. Makes a good story except it doesn't do either. The obsolescence issues of both come into play and unless there is a compelling performance advantage the hybrid will have both higher costs and even more of a lag behind standard x86 server hardware. New purpose-built hardware always takes longer to develop and test than x86. Custom silicon at some point requires new costly spins of fabrication because of errors. Testing and certification also takes significantly longer. Expect whole generations of x86 to be skipped. And the hybrid cannot be repurposed when it is replaced or upgraded.

❍ *WAN Optimization Hardware Deployment Problems*

Deploying WAN optimization hardware appliances means the implementation of hardware, software, and onsite personnel. That's just the beginning. It is the manually intensive planning as well as the interdepartmental coordination and communication that is such a hugely stressful time consuming exercise. Implementation can only occur during scheduled periods of downtime that are few and far between. And when errors occur, which always do with highly coordinated multi-department efforts, they often have to be backed out. This is a major issue in most data centers. For remote offices, branch offices (ROBO) and especially exceedingly difficult to get too remote locations (such as oil rigs, ocean going vessels, countries with lengthy visa/custom processes, and facilities in desolate locations), it's a crisis of epic proportions. The vast majority of ROBOs and remote facilities lack dedicated IT personnel requiring headquarters staff to be dispatched for the duration of the planning and installation, which always goes longer than expected.

Of course scheduling the disruptive downtime, manually intensive tasks, interdepartmental coordination and communication is non-trivial requiring weeks of effort. In the end it is an unnecessary drain on limited IT resources.

- ***WAN Optimization Hardware Appliance Maintenance, Hot Fixes, Upgrade Problems***

Maintenance, hot fixes, and upgrades are likely to be the most onerous of the WAN optimization hardware appliance problems. All require disruptive scheduled downtime. If the applications cannot be taken down, then they must be pointed at a secondary or HA unit while the first is maintained, hot fixed, patched, fixed, or upgraded. Then pointed again at the primary while the process is repeated on the secondary unit. Even if there were no other reason to have HA before, this would be the convincing reason to do so.

Once again disruptive downtime and manually intensive tasks to keep WAN optimization hardware appliances running smoothly, must be scheduled, interdepartmentally coordinated, and communicated. This is about as much fun as a root canal.

The Virtual Server Appliance (VSA) Solution

The WAN optimization VSA solves each and every one of the problems and issues caused by hardware appliances. Because the WAN optimization software is running on a VM as a VSA it doesn't need a duplicate redundant VSA. It leverages the HA that comes with virtualization. If the WAN optimization needs to be moved while hardware is fixed, patched, or upgraded, it can be moved live, online automatically without ever disrupting the users.

When the WAN optimization needs to be implemented locally or at a remote facility, a VM is instantly provisioned (usually from a golden image), the WAN optimization is configured, and it's done. No muss, no fuss, no interdepartmental coordination or communication. The hardware is already in place. Everything is done online in real time from the central operations center. No one needs to helicopter onto an oil rig or desolate location. Everything is completed at the speed of light.

If the WAN optimization software requires a fix, patch, or upgrade, the VSA can be cloned off a golden image, upgraded then take the place of the original, without ever taking the users down for more than a few seconds. Even disaster recovery operations are tied automatically with the server virtualization plans.

Manually intensive operations are eliminated. Obsolescence issues are eliminated. Deployment issues are eliminated, especially remote deployments (all done comfortably from the central site.) Maintenance, hot fixes, operations, upgrades, and management issues are eliminated. And all of this comes at a much lower total cost of ownership.

Not So Fast…

But as most adults know, there is no such thing as a free lunch. So what's the catch? There are two. The first is that more budget will be spent on virtualized server hardware and software services. The good news is that's a double-edged sword whereas both the hardware and software expenditures are leveraged over many applications, and not just the WAN optimization VSA.

The second catch is that picking the right WAN optimization VSA is critical. WAN optimization VSAs are not all created equally. Nor do they all solve the same problems.

NetEx Virtual HyperIP VSA

The Virtual HyperIP VSA specializes in accelerating and moving bulk data for applications such as replication and backup, FTP, disaster recovery such as VMware's SRM (Site Recovery Manager), and the long distance movement of VMs to other data centers using VMware's VMotion or Microsoft's Live Migration. Virtual HyperIP provides the highest effective data throughput for these applications ranging from 2Mbps up to 800Mbps.

Most of the other WAN Optimization VSAs specialize in remote office for the centralization of servers and applications while providing local like performance. Because they cache files, they are able to reduce the amount of file-based traffic that can flow across a slow speed link.

However, when faced with a large database replication or backup, their dictionary caching algorithms typically are not able to find enough duplicated data in order to reduce the traffic across the WAN, thus all data is usually sent, but is not optimized at the transport layer. This creates bursts and gaps in the transmission, which can slow down the effective throughput of the data movement application, particularly across long distance WANs where latency and packet loss are inevitable. It is in these circumstances where Virtual HyperIP positively shines.

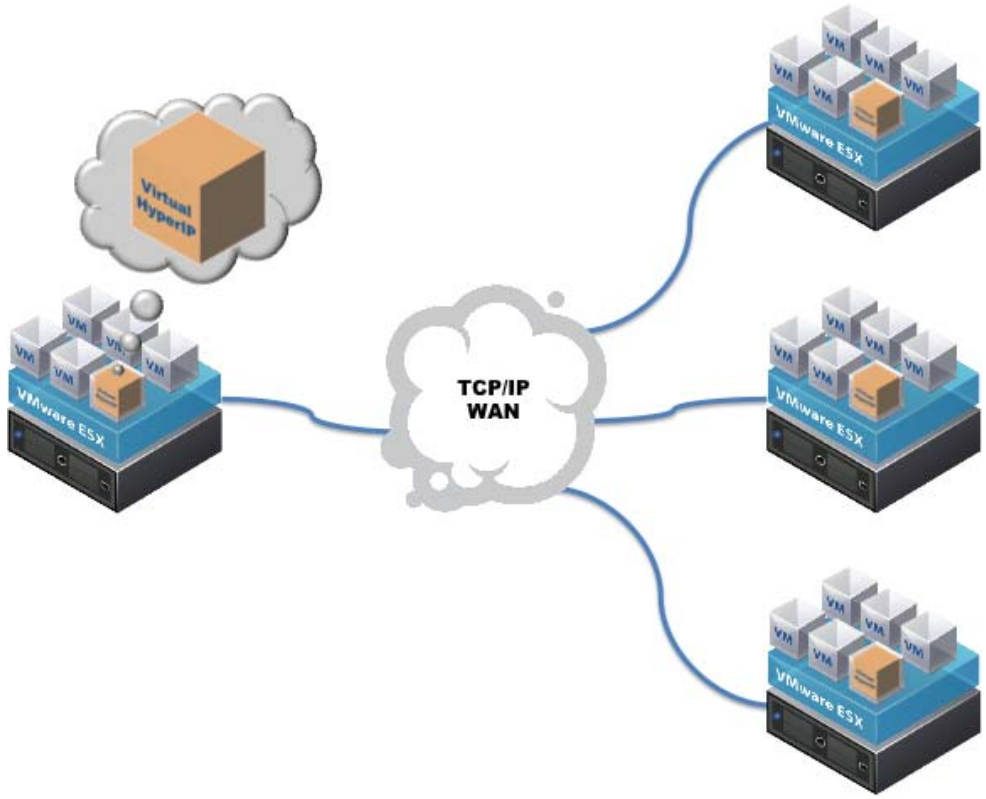

Figure 1: NetEx Virtual HyperIP Appliance

Virtual HyperIP is a data transport optimizer that significantly reduces bandwidth for these mission critical requirements, freeing up bandwidth for chatty applications, such as those using CIFS, NFS, SSH, and HTTP. It is architected to alleviate and mitigate the effects of data packet loss, latency, and TCP/IP data packet loss management to the point where it becomes unnoticeable. It's able to do this because of Virtual HyperIP's unique algorithms.

When there is packet loss, Virtual HyperIP retransmits and reorders only the lost packets instead of all of the packets that follow the lost ones, as standard TCP/IP requires. The impact on effective data throughput on lossy networks is quite impressive. Virtual HyperIP eliminates the horrible compounding throughput effects of TCP/IP latency by terminating TCP/IP at each end of the transmission demonstrably increasing effective data throughput up to 46,000 circuit miles (distance from the earth to a geosynchronous orbiting satellite and back to earth.)

In addition, Virtual HyperIP is easy to deploy, works in many-to-many or many-to-one network configurations, provides block level compression from 2:1 to 10:1 (up to OC3) depending on the data's compressibility and session level deduplication, and is a good citizen on a TCP/IP network by rate-limiting from 1 Mbps to OC12 speeds including (by time of day). Virtual HyperIP is not application or protocol specific (such as only NFS, or CIFS, or SSH, etc.) it is completely transparent, thus it works with ALL TCP/IP applications. And since the licensing of Virtual

HyperIP is reasonably priced and tied to the total bandwidth optimized, its TCO is impressively low.

Conclusion

The advent of server virtualization and virtual server appliances has made WAN optimization hardware appliances obsolete. VSAs eliminate all of the shortcomings of WAN optimization hardware appliances regardless of whether they're purpose-built, x86 white box, or hybrids. It is the operational gains of seriously reduced manually intensive tasks that require scheduled downtime that is perhaps the biggest VSA positive in addition to the much lower total cost of ownership. And VSAs also eliminate the hiding of software performance deficiencies that can be disguised via hardware tweaking. It levels the WAN optimization playing field allowing decisions to be made on which solution best solves the organization's requirements at the lowest cost. Picking the correct WAN optimization VSA is crucial to success.

NetEx's Virtual HyperIP VSA provides a complete answer to increasing the WAN effective bandwidth throughput for mission critical applications moving lots of data. VMotion, SRM, Live Migration, replication, backup, etc. are all examples of where Virtual HyperIP is implemented, proven, and getting the results required.

For more information contact:

info@netex.com or www.netex.com