F5 BIG_IP Application Acceleration Manager in Data Replication Environments

F5 BIG_IP Application Acceleration Manager in Data Replication Environments

White Paper

Introduction

The Internet has been in common use for nearly two decades—long enough to reveal where bottlenecks are and highlight the throughput limitations of wide area networks (WANs). Data loads are constantly on the rise: Where transmissions once included only plain-text email and a few static web hits, people now send rich text email, dynamic web pages including multiple AJAX requests per page, replication data, and a host of other data types.

Replication data in particular is critical, especially to businesses. Businesses need their replication data to arrive at the appropriate destination, in a timely manner, and with no doubt of delivery. The slower and less reliable your replication is, the less useful it is to your original purpose.

Organizations' standard response to increasing demands on their Internet connection has been to purchase a bigger connection. Less used but still relatively common tactics include generic TCP optimization, rate shaping, compression, and optimizing application-specific protocols. These steps have reduced the burden on WAN connections, but none of them does the job completely. To assuage lingering fears of data loss, organizations still frequently end up leasing a larger connection, which increases overhead. Given increasing security concerns, data encryption of nearly every stream leaving the building is a must, and the process of encrypting data can adversely affect the performance of your applications and/or network connection.

These problems are certain to increase in the age of cloud computing. Whether you form an internal cloud across multiple data centers or connect to external cloud services, you will be sending more data over the WAN, further increasing the load on an already burdened Internet connection.

Data Replication and Redundancy: The Growing Need for Speed

With the frequent threat of virus attacks, hackers, and unforeseeable events that can lead to data loss, determining a data protection plan is serious business. If your organization trusts and relies on digital information, it's imperative to have the proper systems and technology in place to defend against data loss. Data downtime can have a devastating effect on your business, so if data loss does occur, you must be able to recover it as quickly as possible. Remote replication is a key component of data protection planning because it ensures there is an up-to-date copy of your data in a location geographically discrete from the original data.

The complexity and increasing regulation of data protection requirements have prompted the release of a wealth of new products and solutions. These newer solutions for redundancy and data protection generally have a common attribute: greater reliance on IP-based WAN services to implement the data protection scheme. Data protection solutions—particularly replication—require the ability to move data quickly and cost-effectively across the WAN, whether to another data center, the cloud, or a service provider. F5 BIG-IP Application Acceleration Manager (AAM) can make these solutions perform better, use less bandwidth, and use less of your systems' CPU cycles.

Seventy percent of surveyed customers valued the ability to move virtual machines between data centers when managing application traffic among multiple data centers using BIG-IP products.

Source: TechValidate TVID: A42-14D-7E1

Increasing Data Volumes

A primary factor driving the need for increased speed is the sheer volume of data that enterprises need to back up. For example, consider a small network. Assume the organization has assorted servers for files, email, and the web that hold 80 GB. These servers and their databases are networked and linked through disks and take up about 700 GB. Factor in three application servers that each hold about 400 GB. Then add 60 users' disk data with 20 GB of data each. This organization already has about three terabytes for off-site backup.

It would take about 50 hours to complete the data transfer on a gigabit Ethernet network backing up at 60 GB per hour. Considering that an internal network connection is generally faster than a WAN connection, and that the WAN connection is already in use by several applications, this difference in network performance characteristics can lead to a slow transfer of data between data centers.

As replication and other WAN traffic grows, congestion problems will become more urgent and trickier to manage. The problem is complex, and the solution will include many components, but it's clear that speed and efficient use of WAN links are a must. Any improvement in the WAN speed—whether actual growth of pipe or bandwidth reduction—will directly translate to more efficient data transfer, and therefore better data protection.

Shrinking Backup Windows

Data protection usually includes some requirement to migrate data off-site over the WAN. Some companies can migrate data during off hours. But as companies grow, this window of downtime shrinks. If usage of the WAN link is near capacity, running data protection applications during peak times can seriously degrade the performance of other applications—or slow your data protection transfers to a crawl.

Any bandwidth reduction on a WAN link can effectively serve as a simple multiplier of backup process throughput for off-site data protection or data merging. Sending less data means faster transfers. If your data moves across the WAN faster, a smaller backup time window is more viable. Increasing your bandwidth usage could also prevent you from having to pay for more bandwidth to solve data protection issues.

High Availability Requirements

The replication process, when implemented properly, allows very high availability of critical applications and data. Data can be mirrored across disk networks and/or the WAN, essentially replicating the system at an off-site location. When this is done over the WAN, the connection being used to send this data becomes mission-critical to the IT department.

The requirements for redundancy and replication can be daunting for some industries. Maintaining high system availability, even in the case of site failure, means that you may need to rapidly move massive amounts of data so identical sites can be prepared to take over operations within minutes. When remote systems get out of synch due to downtime or connection overload, the amount of data that has to be transferred on re-connection is even larger.

Moving data to redundant systems to maintain high availability puts a premium on not only your link speed, but also the effectiveness of how you use all that bandwidth.

Data Protection Planning and Regulations

Companies must support their disaster recovery planning with the right processes and technology. An effective contingency plan should cover all the facets of your business operations, including personnel, customers, facilities, functions, and assets and records. It's critical to have technology that will simplify the management of data backup and recovery, and provide consistent, reliable data protection. To accomplish all this, companies need a straightforward, uncomplicated approach to resuming business operations in the event of a disaster.

Having an effective disaster recovery plan makes good business sense for any company; but it's also becoming a requirement. New government regulations and standards mean companies must also contain legal exposure. Companies are rethinking and retooling their data protection strategies to comply with these government requirements.

Haphazard solutions court disaster and now also risk governmental intervention. Data protection and backup suppliers have developed products that address data lifecycle management. Sophisticated tools from suppliers like Symantec, EMC, and Network Appliance are available to solve these complex storage issues within the data center—but they stop short of addressing the bandwidth on your WAN connection. While many of these products allow you to transfer data over the WAN, few of them help you improve the performance of such transfers.

Greater utilization of the links involved in the WAN portion of your protection plan will give you the flexibility you need to plan properly, and more efficient bandwidth usage helps you deliver a cost-effective solution. BIG-IP Local Traffic Manager (LTM) and BIG-IP Application Acceleration Manager (AAM) can help you deliver the effective WAN utilization that is a major component in redundancy and data protection plans.

When you transfer data across the Internet, security is always a concern. While there are many options for encrypting your data, the ability to send it through a secure, encrypted tunnel running over BIG-IP AAM without increasing the burden on your servers means more processing power and fewer server upgrades.

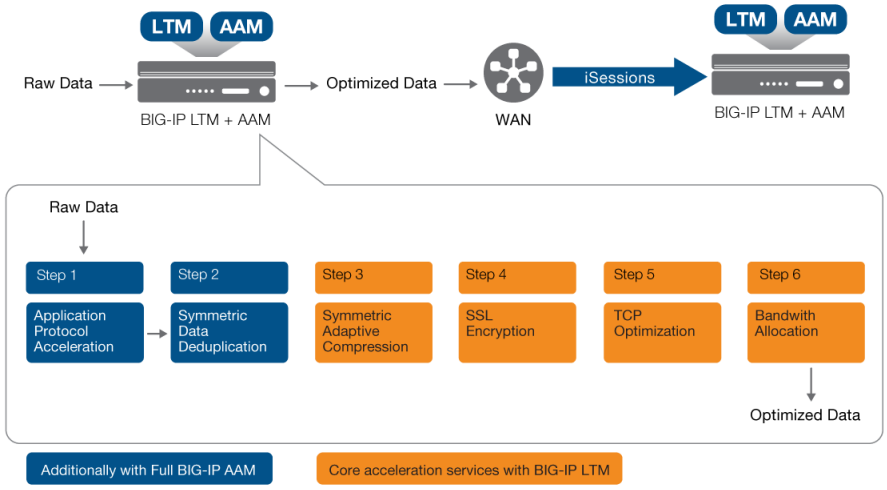

Figure 1: Transferring data over the WAN to gain high speed and high security.

Solution: BIG-IP Application Acceleration Manager

BIG-IP AAM achieves remarkable performance results for data transfers over WAN links by exploiting several opportunities:

- Most data is full of redundancies. By removing the redundancies, you can effectively increase the amount of data moved over the WAN by 10 times on average. Some types of data transfers can increase throughput up to 500

times.

- The behavior of TCP slows down many WAN applications. If you reduce the TCP overhead and manage TCP behavior, you can exploit 100 percent of your bandwidth and speed up your applications.
- Most data transferred over the WAN is highly compressible. Excepting already-compressed data such as zipped files or image formats that include compression, almost everything else can be compressed down on the way out of the building and decompressed at the destination.

For data replication applications, most of the performance increase is a result of compression and TCP optimization. But BIG-IP AAM takes it a step further: When TCP latency increases (normally from distance between sites), BIG-IP AAM TCP management enables you to move data over long distances at full link speed.

Symmetric Data Deduplication: The Optimal Bandwidth Expansion Solution

Some people call it bandwidth expansion; others call it data reduction. But the goal is the same: to squeeze more data through costly wide-area pipes. The basic technique is also the same: remove redundant bit patterns as a data stream enters the WAN and restore them as the stream exits the WAN. Done well, bandwidth expansion lets you support more off-site backups, handle decreasing backup windows, and even provide superior performance for redundant applications.

F5's solution for bandwidth expansion is a combination of technologies called symmetric adaptive compression and symmetric data deduplication, which operate at the session layer of the OSI model. Symmetric data deduplication examines application data streams before they merge, so it finds and removes more redundancies than Layer 3 methods. Symmetric data deduplication operates on all application data streams, and therefore significantly optimizes more types of traffic than protocol-specific Layer 7 techniques.

How Symmetric Adaptive Compression and Symmetric Data Duplication Work

Symmetric adaptive compression automatically identifies all application data streams that can benefit greatly from compression. For other streams that can't benefit, symmetric adaptive compression simply passes the streams along at wire speed. Symmetric adaptive compression ensures that important data is backed up or restored as gracefully and efficiently over the WAN as on the LAN.

Symmetric data deduplication catches intra- and cross-stream duplications before the streams are merged, yielding a higher benefit than merged-stream forms of deduplication.

Symmetric data deduplication technology also dramatically reduces the amount of bandwidth consumed across a WAN link for repeat data transfers. For example, if the same 10 MB Word document is transferred across a WAN link by 100 different users, it would consume 1000 MB of bandwidth without symmetric data deduplication. With symmetric data deduplication, it would consume less than 100 MB—in this case, more than a 90 percent reduction in WAN traffic volume. Because symmetric data deduplication has nothing to do with file systems and operates solely on the data passing through the BIG-IP AAM appliance, the name, location, and other file attributes have no effect on the level of savings introduced. The reduction is based on duplicate bit patterns regardless of file attributes, not on esoteric file data.

Symmetric data deduplication benefits:

- There are no complex policies to set. Simply enable the functionality on the profiles you want to apply it to.
- It's transparent. Client- and server-side applications are unaware that symmetric data deduplication is being applied and therefore require no modification.
- It's protocol-agnostic. Symmetric data deduplication works with any TCP-based protocol including FTP, HTTP, CIFS, and MAPI.
- It's file type–agnostic. Symmetric data deduplication works for any file type, whether proprietary or not, including Word, Excel, PowerPoint, PDF, CAD, database, and more.
- It provides data reduction even for already-compressed data such as ZIP archives and images.
- It avoids the problem of stale data. Unlike traditional file caching algorithms, with symmetric data deduplication, all file requests are actually fulfilled by the server that has the file, but data blocks comprising a file are served from the local BIG-IP AAM appliance. As a result, users can even change a file's name or type and still experience the performance benefits of symmetric data deduplication.
- Symmetric data deduplication provides dramatically higher bandwidth savings than traditional bandwidth expansion techniques.

The benefits of BIG-IP AAM vary according to traffic mix and link speed, but with symmetric data deduplication you will be able to send 5 to 10 times more data through your WAN. Given the right circumstances, some F5 customers have experienced capacity increases in excess of 50 times.

Adaptive TCP Optimizations

The automatic, self-tuning technique of BIG-IP AAM dramatically increases WAN throughput and does not require any changes to your servers. These optimizations are completely transparent to participating applications, the network, and computers.

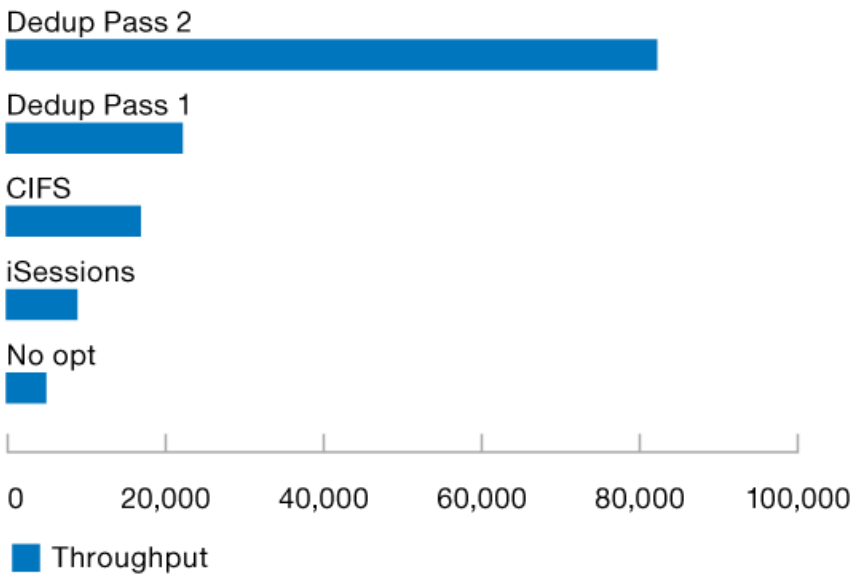

Figure 2: Effective throughput with BIG-IP AAM optimizations.

BIG-IP AAM selects the TCP window size that achieves the highest possible throughput and minimizes retransmission in case of packet loss on the WAN link. To avoid degrading the performance of other applications sharing the same link, throughput is automatically adjusted based upon parameters you set using Bandwidth Controller, according to application behavior and WAN resource availability.

BIG-IP AAM TCP acceleration uses well-known techniques for improving TCP performance, such as local TCP acknowledgments, Selective ACK, and TCP Fast Start. When coupled with BIG-IP AAM Bandwidth Controller capabilities, these techniques enable you to perform powerful adaptive TCP acceleration. For example, you can tell BIG-IP AAM to prioritize replication data and give it a significant portion of your WAN bandwidth. As a result, all data, whether compressed or uncompressed, moves over the WAN faster than ever before, and replication has the edge that allows it to use the bandwidth it needs to complete.

TCP Optimizations and Symmetric Data Deduplication: Multiplicative Performance Increases

With TCP optimizations, the latency effect is almost completely eliminated. That is to say that extra distance adds little or no time to WAN transfers because the latency is overcome with increased throughput, so WAN protocol performance becomes LAN-like. Symmetric adaptive compression and symmetric data deduplication reduce the sheer volume of data that must be physically moved from point A to point B.

When combined, these BIG-IP AAM technologies provide a multiplicative effect where deduplicated and compressed data moves over a highly tuned WAN transport mechanism. The net result is that BIG-IP AAM solutions provide unrivaled performance gains for data-intensive applications such as backup, replication, file transfer, and wide area file services. Add in the native encryption capabilities of BIG-IP AAM, and you have the ability to send highly optimized, secure data over the WAN at increased performance.

Conclusion

In today's complex world of increasing data center connectivity, increasing regulation, and increasing data transfer requirements, you need to get the most from your data center WAN connections. Secure is no longer good enough—now you need secure and optimized. The more you can push through your existing connection, the less you will spend each month for increasing bandwidth, and the more responsive and reliable your systems will be during peak interconnect load times.

The premier solution for speeding, securing, and optimizing your WAN connections, F5 BIG-IP AAM offers a manifold increase in throughput, extension of existing connections even over high-latency links, and data encryption in-flight over the Internet.

F5 Networks, Inc.
401 Elliott Avenue West, Seattle, WA 98119
888-882-4447 www.f5.com

Americas
info@f5.com

Asia-Pacific
apacinfo@f5.com

Europe/Middle-East/Africa
emeainfo@f5.com

Japan
f5j-info@f5.com

©2015 F5 Networks, Inc. All rights reserved. F5, F5 Networks, and the F5 logo are trademarks of F5 Networks, Inc. in the U.S. and in certain other countries. Other F5 trademarks are identified at f5.com. Any other products, services, or company names referenced herein may be trademarks of their respective owners with no endorsement or affiliation, express or implied, claimed by F5. WP-3103 0113