



Symmetric Optimization in the Cloud with BIG-IP AAM VE

As applications and storage move out to the cloud, performance of the corporate communications link between the data center and the cloud will become more important. Optimizing transfers to the cloud with symmetric WAN optimization technology improves WAN communications performance and contains costly bandwidth upgrades.

White Paper
by Don MacVittie



WHITE PAPER

Symmetric Optimization in the Cloud with BIG-IP AAM VE

Introduction

WAN optimization is designed to decrease the amount of traffic an organization transmits over its corporate WAN connection. This process, sometimes called data reduction, uses a combination of deduplication and compression to reduce the size of data transfers. The net result is a reduction in the amount of data physically sent over the wire. Once it's transmitted, the data is reconstituted on the receiving end.

Historically, WAN optimization has been point-to-point between two data centers or between a remote office and a data center. Because data reductions that occur during optimization had to be undone at the remote end, a device was required at the remote end to restore data to its original state. In a virtual infrastructure, this requirement is relaxed by virtual WAN optimization controllers like F5 BIG-IP Application Acceleration Manager (AAM) Virtual Edition (VE). With a virtual WAN optimization controller, the remote end of a connection can be in the cloud, with a physical BIG-IP AAM device in the data center and BIG-IP AAM VE devices wherever necessary in the cloud to meet the organization's needs. This lightens the burden on an organization's WAN infrastructure and helps to maintain their current level of bandwidth while sending more data over that bandwidth.

When Moving to the Cloud, Adaptability Is Key

The BIG-IP system can help organizations with cloud initiatives by enabling IT staff to choose which and how much traffic should be directed to cloud resources.

Ostensibly, the very same WAN optimization functionality that enables multiple data centers to communicate effectively by sending data back and forth efficiently and reliably could be used with cloud providers. But until now, an organization could run WAN optimization technology in the cloud only if their cloud provider happened to offer access to their hardware and if they used the same WAN optimization vendor that the organization had deployed in their data center—an unlikely scenario. The release of BIG-IP AAM VE clears this hurdle. It allows IT staff to extend WAN optimization functionality to cloud-based applications—even replication to the cloud can be handled via BIG-IP AAM VE. As long as the cloud provider offers the ability to "spin up" virtual machines, all of the WAN optimization and application delivery functionality in BIG-IP AAM developed for the corporate data center can be extended to the cloud.



WHITE PAPER

Symmetric Optimization in the Cloud with BIG-IP AAM VE

Multi-Data Center Capabilities Important for Cloud Providers

55% of IT organizations reported that the ability to redirect, split, or rate-shape application traffic between multiple data centers is valuable when choosing a cloud provider.

Source: TechValidate Survey of 109 F5 BIG-IP users TVID: 3D4-O64-27A

Extending the Data Center to the Cloud

There is more to this story than simple WAN optimization. In the course of deploying BIG-IP AAM, other Application Delivery Controller (ADC) functionality—from load balancing to SSL offloading—becomes available to the data center manager. The level of control that the BIG-IP system offers in the data center is extended to the cloud once BIG-IP AAM VE is installed at a given cloud provider.

Here are just a few of the things that BIG-IP AAM and the underlying BIG-IP Local Traffic Manager (LTM) functionality offer IT staff when both physical and virtual versions are available.

- **Intelligent routing** allows administrators to configure load balancing for "last choice" routing to the BIG-IP instance in the cloud, or for "first choice" routing to the cloud.
- **SSL offloading** can remove encryption burdens from virtual machines (VMs) and place them on the BIG-IP AAM VE instance, making all VMs more responsive.
- **Programmability** enables administrators to use advanced functionality like F5 iRules to intelligently route requests between the cloud and the data center. For example, they can stream video from a cloud provider, while servicing all other requests with servers in the corporate data center.
- **Data reduction** significantly enhances replication to the cloud because administrators can reduce the amount of bandwidth transferred in and out of the cloud with compression and deduplication. This increases performance and reduces the cost of cloud services that include a throughput element.



WHITE PAPER

Symmetric Optimization in the Cloud with BIG-IP AAM VE

When an organization is determining whether to place an application in the external cloud or in the data center, corporate infrastructure should enable IT's deployment choices, not limit them. Physical WAN optimization products limit choices by achieving performance gains only for applications that are between two physical locations that are under corporate control. WAN optimization requires symmetric deployment—a device on one end to compress and deduplicate the data going out over the WAN, and another to restore the data to its original state when it comes off the WAN at the destination. An organization's ability to deploy applications that require a lot of WAN transmissions to the cloud is severely limited by this type of infrastructure. BIG-IP AAM VE offers WAN optimization with the freedom to choose application deployment locations. Because it is deployable to the cloud and has all the same functionality as the physical edition of BIG-IP AAM, BIG-IP AAM VE provides cloud-hosted applications with the same level of control as the data center. And because BIG-IP AAM VE, like the physical edition, runs on top of BIG-IP LTM, administrators also enjoy the benefits of BIG-IP LTM.

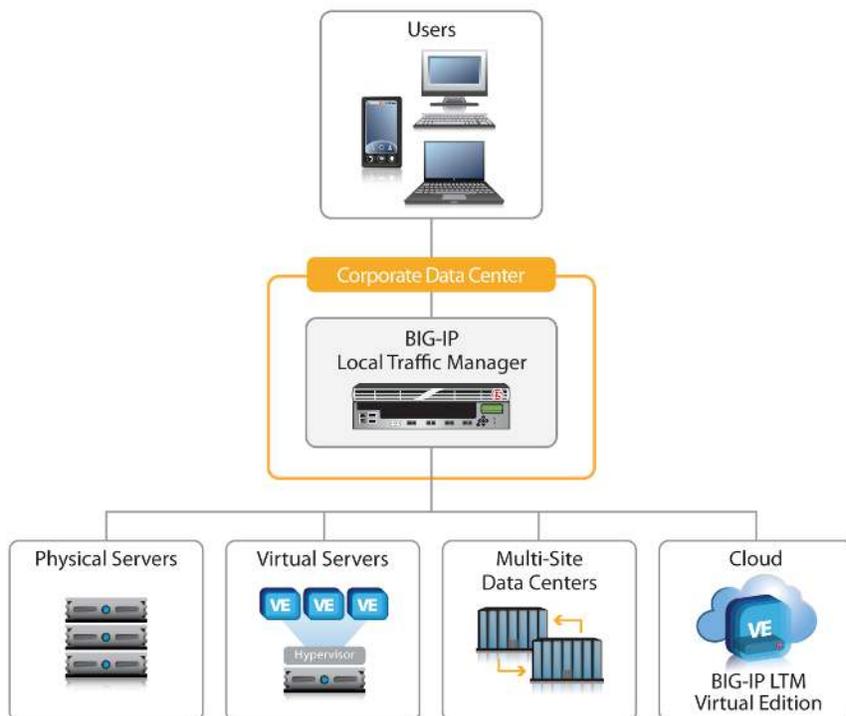


Figure 1: Physical, virtual, remote, and cloud applications—the BIG-IP system supports them all.



WHITE PAPER

Symmetric Optimization in the Cloud with BIG-IP AAM VE

Extending WAN optimization and the entire Application Delivery Network (ADN) to the cloud, which makes how and where requests are serviced invisible to the user, is enabling technology that will drive corporate application delivery needs now and in the future. With decreased overall bandwidth in and out of the cloud and the ability to redirect users to wherever makes the most sense, all within a common infrastructure platform and using SSL offloading, IT can offer organizations more flexibility than ever before.

The Future, Evolving Now

The options for application deployment have grown over the last few years—physical, virtual, internal cloud, external cloud, hybrid cloud, and hosted application deployment are all viable options. Increasingly, organizations are choosing a deployment model for a given application before beginning implementation.

Whether an organization is deploying an application to traditional hardware, a virtual machine array, an internal cloud, or an external cloud will depend on the application, the target user group, availability, and security concerns. One issue that limits application architects' choices is how much data the application in question must pass back to the data center or between data centers. Today's architecture is hampered by a lack of WAN optimization and limited Application Delivery Controllers in external cloud environments. In this model, only a finite subset of applications that is nearly self-contained or requires little data transfer between the application server and the core data center can be deployed to the cloud. Enabling WAN optimization and all other ADC functionality in the cloud with BIG-IP AAM VE allows IT organizations to deploy applications to the cloud that would normally have to be deployed locally. This makes it easier for administrators to put the applications best suited for cloud usage into the cloud without the risk of overburdening the corporate WAN connection or the bandwidth being consumed by cloud-based applications.

Complementary Data Reduction Technologies, Multiplicative Results

Data reduction over any given WAN link is a function of how thoroughly the systems used to reduce transferred bits are implemented. While simple ZIP-style compression can yield impressive results, it cannot eliminate duplications across multiple streams in the same connection. It also cannot reduce the overhead created by protocols like TCP and CIFS.



WHITE PAPER

Symmetric Optimization in the Cloud with BIG-IP AAM VE

BIG-IP AAM uses three features to reduce traffic on the corporate WAN: two that are specialized functions of WAN optimization and one that BIG-IP AAM implements via calls to the underlying BIG-IP LTM because the functionality is generally applicable to much more than just WAN optimization. It is noteworthy that these three features perform their tasks in a single pass through the data, with each handing off the results to the next. The three features are:

- **TCP optimizations.** BIG-IP AAM, via BIG-IP LTM, reduces the overhead introduced by TCP when establishing connections and recovering from errors. These optimizations prevent TCP from flooding the WAN connection with unnecessary acknowledgements and resend requests.
- **Symmetric adaptive compression.** BIG-IP AAM utilizes compression in much the same manner as a disk-based compression program does—by compressing data before it leaves the original location, and restoring it once it arrives at the destination data center. Because this compression is context-sensitive, it is performed on one stream at a time, yielding significant data reduction for each stream, but none across the streams.
- **Symmetric deduplication.** To reduce the amount of duplicate data transferred across all of the streams, BIG-IP AAM applies deduplication before compression to determine whether a string of bits has been sent before; if it has, BIG-IP AAM reduces it. Because this occurs across streams, if the same resource is requested from two different application instances (for example, two user connections or two separate virtual servers), only the first is sent; subsequent sets of the same data are replaced by a token in the stream where the duplication occurred. This greatly reduces the amount of data that compression routines must process and send, but the benefits of symmetric deduplication are dependent on the installation's configuration and data flow. In short, the more data flowing through a connection, the greater the benefits of deduplication. And if data is repeated over multiple streams, deduplication can achieve astounding results even without compression.

Together, compression and deduplication ensure the data passing through an organization's WAN connection receives the maximum amount of data reduction both cross-stream and intra-stream, while TCP compression reduces the effect of latency and transmission failures. Multiplicative data reduction and latency mitigation in a single package help IT manage a corporate, multiple-site architecture and keep the sprawling overhead costs of bandwidth and cloud throughput in check.

These three features, together with the load balancing, security, and dynamic redirection benefits of BIG-IP LTM, give IT management access to a platform that helps achieve corporate architecture and deployment goals. The following is a sample application deployment with and without BIG-IP AAM VE.

Application A is on the schedule for implementation. To determine the best solution to a given business problem, analysts have turned up these key points:



WHITE PAPER

Symmetric Optimization in the Cloud with BIG-IP AAM VE

1. The application will host critical customer data, and therefore must be highly secured.
2. The application will need access to databases hosted in the data center to query user and inventory information for order placing and tracking.
3. The application will have highly variable connection requirements, with near zero connections per second most of the time, but peaks of up to thousands of connections per second.

In an environment without BIG-IP AAM VE, the second point would likely limit the application to local deployment. The amount of data an application could push over the WAN connection might overburden the connection. The alternative possible resolution to this problem is to increase the bandwidth of the WAN connection. This may or may not be viable depending on the application, and it almost always introduces conflict around who should pay for it.

In a BIG-IP AAM VE environment, an architect can consider point two, and with a little research, can classify Application A as viable for cloud deployment. After deduplication and compression have been applied to result sets of similar data, the WAN connection should not be significantly affected. BIG-IP AAM VE also assists with the first point by providing secure tunnels between the physical BIG-IP AAM in the data center and BIG-IP AAM VE in the cloud; this way, corporate data is both tunneled and secured over the Internet.

The third point indicates that IT staff should direct Application A toward the cloud where the amount of processing power can expand and contract to meet the needs of the application very fluidly. Alternatively, an organization could achieve the same result in a highly proficient virtual environment using BIG-IP LTM. However, because a high volume of connections would be flooding the corporate network at peak times, network administrators would have to account for the largest potential daily usage. So the cloud is the better option for Application A. BIG-IP AAM VE can reduce the volume of data being transferred over the wire and secure the data while it is on the wire, leaving the analysts to focus on whether the application itself can be secured in the cloud.

Conclusion

As application deployment options increase, IT will need to have the architecture in place to extend current data center functionality out to the cloud. BIG-IP AAM VE gives symmetric WAN optimization a termination point in the cloud to go with the physical BIG-IP AAM device in the data center, offering all of the ADC functionality of BIG-IP LTM and state-of-the-art WAN optimization in a single box with a single pass on the data being transferred.



WHITE PAPER

Symmetric Optimization in the Cloud with BIG-IP AAM VE

Optimizing the corporate WAN connection to the cloud constrains the cost of doing business in the cloud, reduces throughput, and delays the point at which the Internet connection must be upgraded by reducing traffic flow over the WAN without requiring application changes. BIG-IP AAM enables faster applications, more versatile architecture, and more timely backups with dynamic infrastructure controlled from the organization's core data center. It is even possible to set up a redundant deployment where the BIG-IP system in the data center directs traffic to local instances or cloud instances depending on workload or other parameters set by network administrators and application architects.

F5 BIG-IP AAM VE enables an IT manager to broaden deployment possibilities, recoup wasted bandwidth, and jump-start the organizational cloud initiative while keeping the throughput costs of cloud usage down.

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

Americas
info@f5.com

Asia-Pacific
apacinfo@f5.com

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

Japan
f5j-info@f5.com