



F5 White Paper

The ROI of Application Delivery Controllers in Traditional and Virtualized Environments

How modern offload technologies in Application Delivery Controllers can drastically reduce expenses in traditional and virtualized architectures, with a fast ROI.

By Lori MacVittie

Technical Marketing Manager, Application Services

By KJ (Ken) Salchow, Jr.

Manager, Technical Marketing



Contents

Introduction	3
<hr/>	
The Magic of Server Offload	3
<hr/>	
SSL Termination and Offload	5
<hr/>	
Compression Offload	7
<hr/>	
TCP Offload	9
<hr/>	
Caching In	10
<hr/>	
Virtualization and Consolidation	11
<hr/>	
Conclusion	14



Introduction

The concept of spending money to make money—often referred to as “investing” outside of the technology industry—is something just about every marketing campaign promises, but few deliver. The ROI calculations to prove how quickly an investment will reap return often come with a lot of conditions. For instance, it’s only valid on Tuesdays, under a full moon, and when applied to a specific version of software deployed on a (now) obsolete piece of hardware.

But solutions that provide a quick ROI along with significant technological benefits do exist. The trick is finding these solutions and proving that the ROI model is valid for almost every case.

It’s not magic. It’s simple math. In the following pages we won’t show you how to determine *if* there is a compelling ROI case for Application Delivery Controllers, but how to determine *how much* of a compelling case there really is.

The Magic of Server Offload

Let’s say you are in charge of a rather large data center for a rapidly growing web 2.0 company. And let’s say your “rather large” data center has approximately 1,000 servers. What if someone told you that you could reduce that server number by 40 percent without decreasing performance or availability? And what if that person told you that the solution capable of this magical feat would pay for itself in just 10 months? After you stopped laughing, you might want to hear more about the magic fairy dust that was going to reduce server count without impacting the application, so you could laugh some more.

Assume each server costs an average of US \$2,500, consumes 150 watts of power at an average cost of 10.6 cents per kWh¹, and costs the organization \$288² a year in administrative costs. As this paper will show, reducing the number of servers from 1,000 to 600, while servicing the same number of users at the same performance levels, results in a full return on a \$200,000 investment in about 10 months. The savings that achieve this ROI come from the reduction in power and management costs those 400 servers would have required. Future savings can be calculated by reducing the projected growth in server count and applying the same cost savings to those servers as well.



So, how can you realize these benefits? It's not magic or fairy dust; it's a technological concept called "server offload" that moves computationally intensive (CPU and memory) processing that would normally be handled by servers to an external platform. That external platform is commonly known as an Application Delivery Controller (ADC).

An ADC, in addition to performing commoditized functions like load balancing (which you probably already know about from scaling out your 1,000-server application) is also capable of offloading a variety of other functions such as SSL termination and compression. Both of these tasks are highly computationally intensive and CPU-bound, but are generally implemented at the server level instead of within the application code. This makes them ideal functions to offload to a device more efficient at handling such tasks. In addition, an ADC can add efficiency to the connections themselves—resulting in additional savings.

Whether you are looking to consolidate physical resources and create a virtualized data center, or you're sticking with a tried-and-true traditional architecture, the ability to forestall additional capital expenditures through the implementation of server offload techniques can only improve your financial efficiency—while maintaining or even improving availability, capacity, and performance. And with the anticipated growth of virtual machines per server, it is imperative to ensure that each application deployed within a virtual machine is as efficient as possible. The more concurrent users or transactions per second that can be processed, while limiting resources used, the more you can ensure that performance and capacity will not suffer as the number of virtual machines per physical server increases.

You may not run a data center of 1,000 servers; on the other hand, the reality is that typical enterprise servers actually cost a bit more than \$2,500, use more than 150 watts of power (since this is a typical idle draw), and have administrative costs much higher than \$288 a year. So, even with a moderate number of servers to manage, you will realize an excellent return on your investment in an ADC with server offload capability. A 2009 TechValidate survey indicates that a majority of customers see an ROI of 18 months or less on their investment in an F5 Application Delivery Controller.

Here's how they do it.

65% of organizations using F5 BIG-IP solutions reported a payback period of 18 months or less.

Source: Survey of 192 F5 BIG-IP users



TVID: 4F3-02B-15B



SSL Termination and Offload

SSL is the most ubiquitous security protocol used in conjunction with websites today. Data from Netcraft states that in January 2008 nearly 2.5 million sites on the Internet made use of SSL.³ SSL enables clients and servers to encrypt and decrypt the data they exchange, securing it from prying eyes and from manipulation while in transit over public networks.

SSL, like most mathematically complex algorithms, is CPU-intensive, requiring a lot of CPU resources to churn through the mathematical computations required to encrypt and decrypt large chunks of data. Additionally, because these complex computations are being executed on general-purpose CPUs, the process of encrypting and decrypting data can be a significant detriment to application and system performance.

One of the ways technology has addressed the problem of performance and resource consumption is through hardware acceleration. The use of specialized hardware—designed solely to perform the mathematical computations required of SSL operations—simultaneously reduces the resources required and increases the performance of such operations. Most of this kind of specialized hardware is found in offload devices such as load balancers and ADCs, like F5 BIG-IP® Local Traffic Manager™ (LTM). BIG-IP LTM offloads SSL processing by acting as a proxy for web and application servers. Because the offload device performs all of the SSL processing, the web and application servers can dedicate their resources to responding to application requests. (See Figure 1.)

Testing, empirical evidence, and conventional wisdom place the amount of CPU resources required for SSL processing (without hardware acceleration) at 30 percent of a typical server's resources. If you have one server running at 90 percent utilization, offloading the SSL processing to an ADC or to a load balancer will reduce that utilization to 60 percent. Similarly, if your application currently requires 10 servers to support 1,000 users, then offloading SSL to an intermediate device should reduce the number of servers required to seven, or increase the number of users you can support on those 10 servers to 1,300.

Gartner, Inc., Worldwide Server Forecast, 2002-2014, September 2009 estimates an annual average growth in server purchases of 4.48 percent for 2009-2012.⁴

Certificate management is another factor to consider when calculating the ROI for SSL acceleration. Using an ADC to offload SSL from servers means all associated certificates are managed in one place, on one device. This simplifies management, reducing operating expenses for an even greater ROI. And because SSL offload enables the organization to legally use one certificate per application regardless of how many physical or virtual servers are required to serve it, the organization saves even more on the cost of SSL certificates.




Figure 1: SSL offload with F5 BIG-IP LTM

- 1 BIG-IP LTM handles all SSL negotiations with the client. BIG-IP LTM receives the encrypted request and decrypts it, then chooses a server and forwards the request in plain text.
- 2 The server handles the request normally and returns it to BIG-IP LTM.
- 3 BIG-IP LTM encrypts the response and returns it to the client.

Deploying an ADC, and taking advantage of SSL offload, results in saving more than \$40,000 in power costs alone by simply turning off 30 percent of the servers. With an additional reduction in operating expenses in excess of \$85,000 from simplified server administration, our example data center achieves full ROI in just 13 months—for the entire cost of the ADC—based on the SSL offload capability alone. (See Table 1.) Additionally, the data center’s growth rate is effectively reduced, as it no longer requires four servers per month to support growing application demand. This reduces capital expenditures, as fewer hardware server purchases are required.

Size of Data Center	Cost of ADC ⁵	Payback ROI
Small (125 servers)	\$40,000	22 months
Medium (500 servers)	\$120,000	17 months
Large (1,000 servers)	\$200,000	14 months

Table 1. ROI for Application Delivery Controller using SSL offload

Based on savings realized due to reduction in administrative costs of \$288 per server per year, reduction in power usage from unnecessary servers using 150W at a cost of \$0.106/KwH, \$2,500 cost of acquisition per server, a 30% reduction in number of servers necessary, and a 4.48% year-over-year growth rate.

Again, it is important to note that this ROI is for the entire ADC solution, not just the acquisition of the SSL capabilities alone.

“Our business was able to save CapEx using F5 BIG-IP LTM to offload SSL certificates and SSL processing. This reduced the number of SSL certificates needed across multiple web servers, and reduced the overhead of those same web servers.”

IT Manager
Medium Enterprise Computer
Software Company



TVID: 420-53F-8C5



Compression Offload

Compression is commonly enabled on web servers as a means to lower costs by reducing bandwidth utilization. It is also used to improve application performance. Compression, like SSL operations, is mathematically intensive and is typically CPU-bound. When used to compress dynamic content for which local server-based caching is not available, compression can consume 4 to 30 times the CPU resources that a server serving the same content without compression would utilize. This is true for both Microsoft IIS and Apache web servers. (See Table 2.)

It is important to note that the decrease in bandwidth is significant and is therefore providing benefits despite the increase in CPU utilization. Compression typically affords applications a 3:1 reduction in size and improves application response time dramatically, especially over high latency or bandwidth-constrained connections.


	File Size	Bandwidth Decrease	CPU Utilization Increase
IIS 7.0	10KB	55%	4x
	50KB	67%	20x
	100KB	64%	30x
Apache 2.2	10KB	55%	4x
	50KB	65%	10x
	100KB	63%	30x

Table 2: Effect on CPU utilization of compression for dynamic web application content⁷

But the hit on the CPU is significant enough to impact the overall capacity and performance of the application (and any other applications deployed on the same server). The magic of compression may well be negated by the need to deploy additional servers to compensate for the reduction in processing power available.

On the other hand, the magic of offload can eliminate that need by taking on the task of providing compression for the web/application server. ADCs and load balancers can apply compression to content and generally take advantage of hardware-assisted compression, which has a higher compression ratio (4:1 instead of 3:1) and provides more bandwidth savings than enabling compression on the web/application server (See Figure 2.)

Offloading the task of compressing content—particularly dynamic content—achieves a slightly better compression ratio while freeing the web/application server CPU resources that would have been used to perform the task. Based on available



ADCs that can intelligently apply compression to content only when it would benefit performance and resource consumption provide further efficiencies on the ADC itself. No cycles or memory are wasted on content and connections that would not benefit from compression.

Figure 2: Compression offload using F5 BIG-IP Application Delivery Controller

- 1 BIG-IP LTM receives a web request and checks the client’s bandwidth, then chooses a server and forwards the request.
- 2 The server handles the request normally and returns it to BIG-IP LTM.
- 3 BIG-IP LTM takes into consideration the available bandwidth and the type of content and determines whether it will be a performance plus or negative to apply compression. It then acts on the decision and returns the response to the client.

data regarding the impact of compression on CPU utilization and current average sizes of web pages, we can assume that an average of 20 percent of a server’s resources are consumed by the process of applying compression. When offloaded to an external device such as an ADC, those resources can be refocused on the server’s primary task of serving applications.

As with SSL offload, this results in either a reduction in servers needed to support capacity needs, or in an immediate increase in capacity. (See Table 3.) If 1,000 users are being supported by 10 servers, offloading compression to an ADC should result

50% of IT organizations report that they reduced annual OpEx budgets by 10% to 20% or more by deploying F5 BIG-IP solutions.

Source: Survey of 200 BIG-IP users



TVID: 5DD-99E-9B6



in those same 1,000 users being supported by only eight servers, or in total user capacity being increased to 1,200 using 10 servers.

Size of Data Center	Cost of ADC ⁵	Payback ROI
Small (125 servers)	\$40,000	34 months
Medium (500 servers)	\$120,000	25 months
Large (1,000 servers)	\$200,000	21 months

Table 3. ROI for Application Delivery Controller using compression offload

Based on savings realized due to reduction in administrative costs of \$288 per server/per year, reduction of power from unnecessary servers using 150W at a cost of \$0.106/Kwh, \$2,500 cost of acquisition per server, a 20% reduction in number of servers necessary, and a 4.48% year-over-year growth rate.

TCP Offload

The use of the term “TCP offload” to describe TCP multiplexing is somewhat of an anomaly. Other offload technologies actually offload complete functionality from servers, while TCP offload uses optimization of resources to offload TCP overhead and dramatically increase capacity of servers.

TCP offload, more often referred to as TCP multiplexing, is an optimization technique common to ADCs that exploits the nature of persistent connections to achieve higher utilization of TCP connections by sharing them on the back end, across users. Because the full-proxy architecture of ADCs requires two separate networking stacks, user TCP connections are made to the intermediary (the ADC), while server connections are maintained between the intermediary and the servers. This allows the intermediary to maintain a much higher number of user connections than may actually be supported by the server infrastructure, effectively increasing the capacity of the servers.

In any architecture, each client connection requires a matching connection on the server. This usually means a range of two to six connections per user are consumed on the server, whether virtual or traditional. Using TCP multiplexing, the user still opens two to six connections to the “server” but the intermediary brokers those connections and instead opens only one to the server and then reuses that connection across the user session. The ADC will also reuse that same connection

“We’ve been able to leverage the F5 (solution) in ways that we didn’t expect when we purchased it.”

IT Architect,
Global 500 Professional
Services Company



TVID: 956-3C2-AD6



for additional users, opening new connections to the server only when necessary to maintain application availability and performance.

This magical “offload” has dramatic results: a reduction of 66 to 90 percent⁶ in server-side connections with an improvement in performance as measured by Time To First Byte (TTFB).

What that means in practical terms is that you can serve the same number of concurrent users with one-third the physical hardware—or one-third the same number of virtual instances of the application. We’ll keep our estimates for ROI purposes on the low end, using a 66 percent reduction for calculations. (See Table 4.)

Size of Data Center	Cost of ADC ⁵	Payback ROI
Small (125 servers)	\$40,000	10 months
Medium (500 servers)	\$120,000	7 months
Large (1,000 servers)	\$200,000	6 months

Table 4. ROI for Application Delivery Controller using TCP multiplexing

Based on savings realized due to reduction in administrative costs of \$288 per server/per year, reduction of power from unnecessary servers using 150W at a cost of \$0.106/Kwh, \$2500 cost of acquisition per server, a 66% reduction in number of servers necessary, and a 4.48% year-over-year growth.

Cashing In

The individual value to your organization from any one of these offload technologies can be significant, but putting them all together amplifies their value. If we take the 30 percent reduction from SSL offload, apply the 20 percent compression improvement, and apply a 66 percent reduction through TCP optimization, we get a roughly 81 percent overall resource reduction. In this pristine example, our 1,000-server data center would see an ROI of its entire initial ADC deployment in five



months—just from the offload technologies alone. (See Table 5.) This doesn't even begin to consider the increased uptime and operational savings. While compelling, it's reminiscent of that magic fairy dust. However, if we assume even half of that overall reduction, still with very moderate power, server, and management costs, the ROI is still quite compelling.

Size of Data Center	Cost of ADC ⁵	Payback ROI
Small (125 servers)	\$40,000	16 months
Medium (500 servers)	\$120,000	12 months
Large (1,000 servers)	\$200,000	10 months

Table 5. ROI for Application Delivery Controller using combined offload technologies

Based on savings realized due to reduction in administrative costs of \$288 per server/ per year, reduction of power from unnecessary servers using 150W at a cost of \$0.106/ Kwh, \$2,500 cost of acquisition per server, and a 40% reduction in number of servers necessary.

Virtualization and Consolidation

Many organizations look to server virtualization and data center consolidation to achieve the same kinds of OpEx and CapEx savings we have already demonstrated can be achieved using ADC offloading technologies. While the savings from virtualization can be substantial, additional investment is required, not to mention the challenges of maintaining operations as the organization moves from a one-to-one physical server environment to a virtual one. The value of an ADC to a virtualized infrastructure can be quite significant, even without using offload technology.

However, it is important to note that while traditional ADC features (load balancing and traffic management) enable a seamless transition between physical and virtual machines, ADC offload technologies are even more relevant in the virtualized environment, where the dynamic nature of the environment makes managing SSL certificates and TCP connection setups and teardowns more challenging. By centrally managing SSL certificates and application-wide compression profiles, you can alleviate the need to do this at the virtual machine (VM) level every time you spin up or spin down another server, further reducing the management costs of those VMs. Using an ADC, each VM has more processing power available for the application, effectively increasing its capacity and alleviating the need to deploy more VMs that can constrain the architecture and increase management costs. This can drastically

“When virtualizing, set up a planning scenario and acquire the latest advanced tools. Such tools include capacity, placement and performance to automate and ‘metricize’ the planning process. Use total-cost-of-ownership analytics to ascertain optimum virtual machine placement and to determine the benefits.”

Gartner, Inc.
Ten Helpful Hints for Reducing Server Infrastructure Costs
November 2008



accelerate and amplify the ROI of your virtualization and consolidation efforts.

Let’s assume that the 1,000-server data center from our example implements a 100 percent virtualization effort, converting all existing physical servers into VMs at a fairly respectable rate of 15 VMs to one physical machine. This requires some significant investment in new hardware and licensing, but, over time, generates significant savings. (See Table 6.)

Year	Consolidate				Savings	
	Hardware	Software	Mgmt.	Power	Total	Cumulative
1	\$ (378,000.00)	\$ (196,000.00)		\$ 126,138.36	\$ (447,861.64)	\$ (447,861.64)
2	\$ (16,200.00)	\$ (8,400.00)		\$ 131,852.91	\$ 107,252.91	\$ (340,608.73)
3	\$ (21,600.00)	\$ (11,200.00)		\$ 137,567.47	\$ 104,767.47	\$ (235,841.26)
4	\$ (16,200.00)	\$ (8,400.00)		\$ 143,839.54	\$ 119,239.54	\$ (116,601.72)
5	\$ (16,200.00)	\$ (8,400.00)		\$ 150,390.37	\$ 125,790.37	\$ 9,188.65

Table 6. Virtualization savings

Based on a 15:1 physical machine reduction with new hardware costs of \$5,400 per platform, \$2,800 in virtualization software per platform, using 300W per new platform at a cost of \$0.106/KwH, no reduction in the actual number of “servers” or VMs resulting in no management savings, and a 4.48% year-over-year VM growth rate.

While these calculations are certainly not intended to provide guidance in the costs and savings of real-world virtualization and consolidation efforts, and do not reflect the myriad variables associated with such a drastic architectural change (like the reduced cost of cooling due to reduced BTU output or the savings in not having to build new facilities), they are sufficient to show the significant impact that adding server offload technology to your virtualization efforts can provide.

Based on these calculations, it is well into the fifth year of operation before the effort breaks even and the organization starts to see a positive ROI. After that point, the organization begins to see a substantial savings year-over-year.

However, using the same scenario, but adding an ADC with server offload technology to reduce the number of VMs by 40 percent, can accelerate the organization’s ROI in two ways. (See Table 7.)



Year	Consolidate/Offload				Savings	
	Hardware	Software	Mgmt.	Power	Total	Cumulative
1	\$ (426,800.00)	\$ (117,600.00)	\$ 120,384.00	\$ 133,943.60	\$ (290,072.40)	\$ (290,072.40)
2	\$ (10,800.00)	\$ (5,600.00)	\$ 125,568.00	\$ 139,936.92	\$ 249,104.92	\$ (40,967.48)
3	\$ (10,800.00)	\$ (5,600.00)	\$ 131,328.00	\$ 146,208.99	\$ 261,136.99	\$ 220,169.51
4	\$ (10,800.00)	\$ (5,600.00)	\$ 137,088.00	\$ 152,759.82	\$ 273,447.82	\$ 493,617.33
5	\$ (10,800.00)	\$ (5,600.00)	\$ 143,424.00	\$ 159,589.41	\$ 286,613.41	\$ 780,230.75

Table 7. Savings on virtualization using Application Delivery Controller with offload technologies

Based on a 15:1 physical machine reduction with new hardware costing \$5,400 per platform, \$2,800 in virtualization software per platform, using 300W per new platform at a cost of \$0.106/KwH, a 4.48% year-over-year VM growth rate, an initial \$200,000 ADC investment, and 40% reduction in VMs due to server offload technology.

First, notice that the break-even point is a full two years sooner (cumulative savings become positive in year 3). Despite the additional upfront costs of the ADC, the initial costs are about 50 percent lower as a result of the reduction in hardware and VM licenses needed, lower VM management costs, and additional power savings.

Second, because that investment is significantly reduced on a year-over-year basis, and there is an associated ongoing reduction in management costs (because the number of new VMs grows more slowly), the positive ROI is more than doubled on a yearly basis.

Again, the real costs and savings for any particular virtualization effort depend on a great number of variables that are beyond the scope of this discussion. However, the reality is plain to see. Whatever the actual costs and benefits of virtualization efforts are, adding server offload technology can demonstrably lower those costs and amplify those benefits.

White Paper

The ROI of Application Delivery Controllers in Traditional and Virtualized Environments

Conclusion

The compute resource cost of mathematically complex functions, such as SSL operations and compression, is a significant burden on web and application servers. These complex operations are CPU-bound and consume resources in a way that negatively impacts application performance and capacity. This is true whether the server in question is traditional or virtual. So, it's not magic or pixie dust, nor is it a trick with numbers: the ROI offered by an Application Delivery Controller employing SSL offload, compression offload, and TCP optimization is real. In addition, implementing these technologies in conjunction with existing virtualization and consolidation efforts can amplify an organization's cost savings and accelerate the overall ROI.

¹Energy Information Administration: see www.eia.doe.gov/cneaf/electricity/epm/table5_3.html for averages across industries 1995-2009; see www.eia.doe.gov/cneaf/electricity/epm/table5_6_b.html for specific industry and state guidance. Note that figures for 2008 and 2009 are estimated.

²Based on 1 hour/month/server at US \$24.

³Netcraft, <http://news.netcraft.com/SSL-Survey/>.

⁴Gartner, Worldwide Server Forecast Database, September 15, 2009.

⁵Assumes a high availability implementation comprising two ADCs.

⁶F5 Deployment Guide, [Tuning the OneConnect Feature on the BIG-IP Local Traffic Manager](#)

⁷Web Performance, Inc., [Measuring the Performance Effects of Dynamic Compression in IIS 7.0](#)

Web Performance, Inc., [Measuring the Performance Effects of mod_deflate in Apache 2.2](#)

Intel® Software Network, [HTTP Compression for Web Applications](#)

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

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
info.asia@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com

