



Application Delivery Optimization

Fueled by users who are hyper-sensitive to delays in application delivery, organizations are seeking solutions that can help them root out and eliminate delivery latency. Application Delivery Optimization promises to improve application performance in the data center and beyond.

White Paper
by Lori MacVittie



Introduction

Conversions. Revenue. Brand reputation. Return visits. Click depth. Loyalty. Satisfaction. Productivity. Abandonment rates.

No matter how you measure the impact of application performance, in the age of instant-everything, seconds count. And when you only have milliseconds to deliver content to a consumer or employee with a twitchy finger, that means every component in the delivery chain has merely microseconds to perform appointed tasks, or risk becoming a bottleneck.

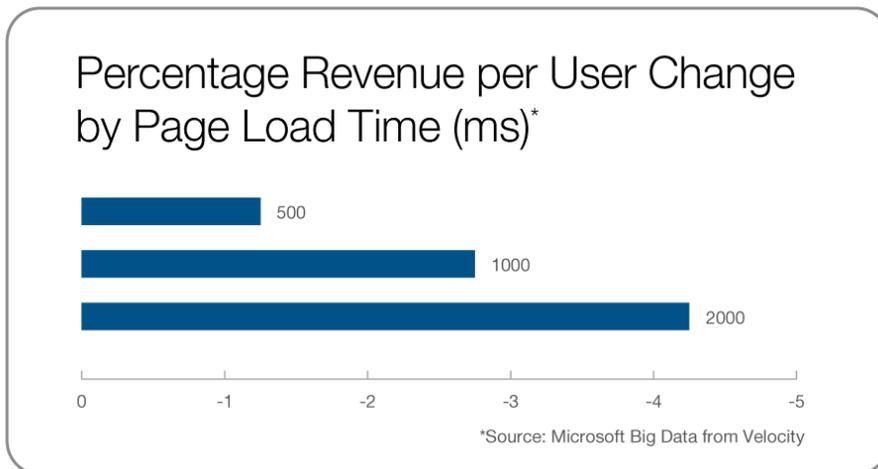


Figure 1: The relationship between delay in application delivery and revenue loss is significant.

Though the consensus may be that application performance is paramount to success, the debate about how to achieve that goal remains. Application performance is affected by such a wide variety of factors that from the lofty goal of a "fast application" has arisen a variety of niche solutions, all focused on leveraging different technologies and strategies.



WHITE PAPER

Application Delivery Optimization

Each argument for a specific niche solution bears consideration, as all address some facet of the larger problem. But no single niche solution encompasses the entire delivery chain, and thus no single niche solution can completely resolve the problem of poor application performance.

Resolving this problem requires taking a more holistic approach to optimizing the entire application delivery chain, and in doing so, eliminating the various bottlenecks that ultimately cause poor application performance.

The Bottlenecks

It's easy to say there is a need to maintain consistently high-performing applications. But it's quite another thing to realize that goal when faced with the myriad factors that can negatively affect performance. Traditional acceleration solutions focus on optimizing content: caching and geolocation to move content closer to the user; compression and transformation-based technologies to reduce the size of content; and protocol optimizations to reduce the amount of network traffic overhead associated with delivery that can affect network transfer times.



Nowadays, the [New York] Times claims, users drop off after a mere 400 milliseconds, and a difference in page load time of just 250 milliseconds is enough to convey a distinct advantage over your competitors.¹

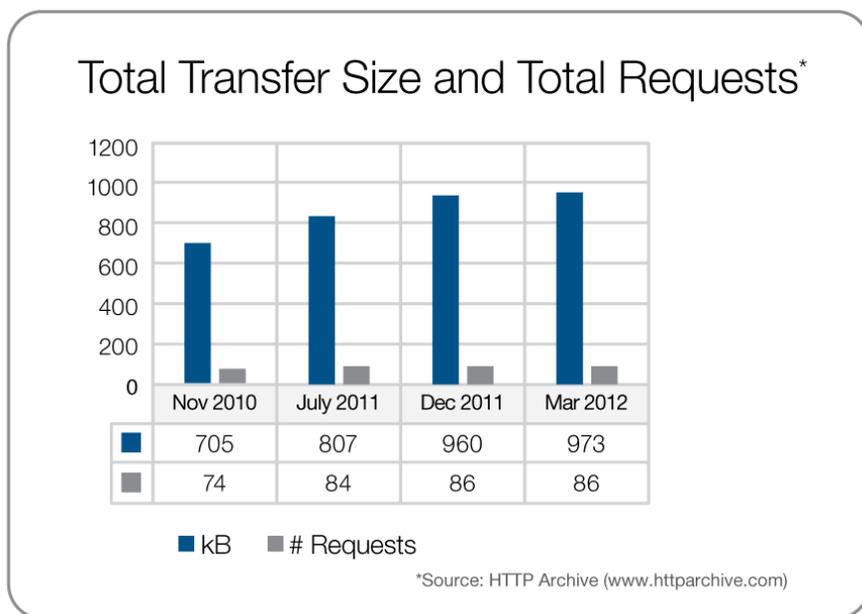


Figure 2: Everything about the web has experienced high rates of growth and shows no signs of stopping.

WHITE PAPER

Application Delivery Optimization

These technologies are still relevant when addressing performance challenges associated with modern web applications, but they often do not adequately address the challenges that have emerged with new web application architectures, deployment models, and end-user capabilities.

Everything about the web and the applications it comprises has gotten bigger- the number of users, the amount of time those users spend interacting with applications, and the devices from which web content is accessed.

The mix of media being delivered has also substantially shifted from primarily text-based content to rich, highly interactive content laden with video and images, and heavily integrated with other services.




Figure 3: Predictions of traffic and usage growth through 2015.

Growth in Web Page Size and Complexity

The size of web pages has tripled in eight years. But just as important as total content size is the complexity of the page. The number of objects comprising a web application today plays a significant role in response times and is increasingly the focus of optimization efforts.

Object complexity imposes overhead at both the transport (TCP) and application (HTTP) layers, requiring more network, delivery, and compute resources than websites' less complex predecessors. This growth impairs performance by requiring more connections, each of which are sustained for longer periods of time-straining capacity limits of the entire application delivery infrastructure.

Mobile Explosion

The explosion of mobile device usage among both consumers and employees is well known. Given the growing congestion of core IP networks, caused by additional traffic to the limited bandwidth and processing power of mobile devices, the impact of this explosion on performance is significant.



WHITE PAPER

Application Delivery Optimization

The ability of mobile devices to roam effortlessly between networks with very different performance characteristics introduces additional challenges for performance-minded IT. The dynamic nature of connectivity inherent in mobile device usage impairs IT's ability to realize consistent performance across a single user session.

Protocol Behavior

Protocols and performance have a curious relationship, particularly when network issues are present. Packet loss affects both TCP and UDP-based applications, but in different ways. TCP-based applications such as those delivered via HTTP generally cope with packet loss by retransmitting lost packets. The size of content transferred via HTTP means more packets and more opportunity to lose them. Additionally, constraints on capacity and the speed of WAN links and consumer network connections impose limitations on transfer rates.

For UDP-based application delivery—primarily video and voice services—the impact is directly on user experience. Packet loss translates into jittery, halting communications that interrupt and delay transmission, ultimately causing user frustration and abandonment.

Security

Performance and security vie continually for top priority among both operations and business stakeholders. Too often, security is sacrificed to improve performance. Ninety percent of respondents in a Crossbeam survey admitted to "making a trade-off between security and throughput performance," and 81 percent admitted to "shutting off functionality in a security product because it was slowing down their network."²

The need to address emerging application-layer threats that often sneak past security infrastructure, the increasingly disconnected authentication and authorization systems required by various end-point devices, and traditional network security requirements put additional pressure on IT to balance security with performance.

Cloud Computing

Cloud computing benefits operations financially with more affordable computing and improved agility within the application deployment lifecycle; but it introduces barriers to realizing consistently high-performing applications.



IT security personnel within large corporations are shutting off critical functionality in security applications to meet network performance demands for business applications.³



WHITE PAPER

Application Delivery Optimization

Symmetric optimization technologies require symmetric end points. These traditionally hardware-based solutions are not amenable to deployment on mobile computing devices. Software and virtual versions of such solutions do not completely mitigate this issue, as supporting today's myriad operating system and platforms on mobile devices can be challenging for vendors, especially when those systems are not supportive of solutions that need to control the network stack to provide value.

The dizzying array of user device options, deployment challenges, and support makes addressing the performance of cloud-hosted applications a particularly tough nut to crack.

On the other hand, cloud computing offers a mitigating solution to performance troubles when the root cause is capacity (load) or proximity to content. But this also introduces other challenges that may simply shift the potential cost of poor performance to other areas of IT, effectively wiping out the revenue gains with operational costs associated with disconnected policies and process management.

The next generation of application delivery solutions will need to address all these challenges-often simultaneously. As organizations extend the data center to incorporate cloud resources, they must embrace what are often high-latency, low-bandwidth networks as critical to that integration. The explosive growth of both mobile devices and the demand for them by employees must be met head-on; with productivity and brand reputation on the line, it's critical that organizations swiftly and effectively deal with the resulting impact on the network and inevitable performance-related issues.

Resolution: Application Delivery Optimization

Consumers are unmoved by the litany of excuses for poor performance. They don't care about the why, they simply want it resolved. The challenge for IT, then, is to be able to address all the various causes of poor performance. If IT can't achieve this with the tools they have, then they must try to offset those causes by leveraging other technology, for example, content reduction technology that can offset poorly performing networks. If one cannot make the car go faster to arrive on time, perhaps one can take a shorter route.

Application Delivery Optimization technologies and topologies comprise several solution sets, each focused on addressing a specific subset of causes of bottlenecks.



I'm not sure the cloud will ever be as big as anyone hopes. Not a single customer wants to go slower. Ever. Go ask any computer user if slower is OK and see what they say. Opening a large PDF on the cloud can be agony.⁴



Topologies

Symmetric acceleration. The most common implementation of symmetric acceleration is a WAN optimization solution. WAN optimization technologies address performance by optimizing higher-latency, lower-bandwidth links between clients and services. WAN optimization solutions focus on network optimization and data reduction as a means to improve performance; they are particularly adept at meeting the big data migration challenges. WAN optimization significantly reduces time on the wire, which enables replication and live migration processes to complete faster and reduces latency associated with a variety of IT concerns surrounding emerging technology like VDI.

Though WAN optimization is traditionally viewed as synonymous with symmetric acceleration, symmetric web acceleration affords the same benefits that WAN optimization offers at the network, but at the application layer.

Front-end optimization (asymmetric acceleration). Technologies that address bottlenecks caused by the size and complexity of content have long been associated with application delivery. Front-end optimization (FEO) encompasses image optimization, cache manipulation, reductive techniques (compression and minification), and re-ordering of content to enable faster rendering by the client. By reducing the size and/or location of content, FEO solutions can obviate many performance impediments.

An informal survey of 1,197 VMworld 2010 attendees found latency to be the top VDI network consideration (62%).⁵

Technologies

Protocol optimization. Protocol optimization focuses on improving the behavior of protocols to avert performance-impeding consequences. It includes dynamic adjustment of TCP parameters to maintain consistent delivery in the face of network congestion, and it attempts to circumvent jitter, inducing latency. Protocol optimization also leverages a number of application-layer technologies to address chattiness in protocols such as CIFS and HTTP.

Load balancing. Load balancing is not often directly associated with acceleration or optimization solutions, but its ability to non-disruptively add capacity resolves performance issues where load is the root cause. As cloud becomes more integrated into data center architectures, load balancing at the global (site) level is a critical component of Application Delivery Optimization, as it is the means by which users are directed to the site that can best meet performance requirements.

Offload. One way in which the application delivery tier improves performance is by alleviating the burden placed on servers by compute-intensive processing incurred by cryptography (security) and TCP connection management. With the burden of these services removed to the application delivery tier, processing time at the application infrastructure layer improves, resulting in improved response time.

Location. A user's proximity to content has long been a factor in performance, primarily due to the latency incurred by the distance data must travel to reach the user. Location-focused solutions seek to improve performance by moving content closer to the user.

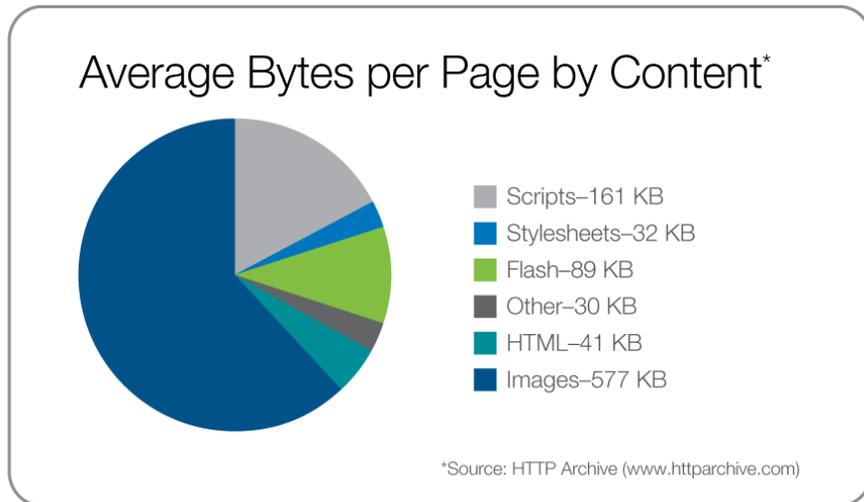


Figure 4: The distribution of content type indicates image optimization and minification of scripts will significantly boost web application performance.

The problem with today's focus on one solution or another is not that individual solutions do not solve real application delivery performance issues; it's that no single niche solution can address all of the bottlenecks that might impede performance. Complicating the matter is that the potential combinations of user devices, networks, resource locations, and resources are both volatile and numerous. Organizations need a more comprehensive strategy that embraces a holistic and dynamic approach.

Application Delivery Optimization is that approach: a holistic, integrated, and adaptable set of services with flexible deployment options and a consolidated operational platform that eliminates inefficiencies associated with gluing together multiple niche solutions across multiple locations.

The F5 ADO Approach

F5 Application Delivery Optimization (ADO) solutions offer a strategic, holistic approach to redressing the myriad bottlenecks that cause performance problems. The F5 ADO approach spans environments, networks, client devices, and application infrastructures to ensure organizations can meet a broad spectrum of challenges dynamically and while maintaining operational consistency.



WHITE PAPER

Application Delivery Optimization




Figure 5: The F5 ADO approach is a holistic, dynamic, and operationally consistent set of services focusing on eliminating application performance bottlenecks.

The F5 BIG-IP platform is purpose-built to enable a variety of services to be deployed in an operationally consistent environment. Its internal, integrated architecture ensures optimization solutions can share the context necessary to dynamically adapt to conditions that impact performance in real time. By sharing network and protocol optimizations across specific services, F5 ADO solutions combat common bottlenecks and allow operations to mitigate performance problems regardless of user, network, and application deployment location.

The BIG-IP system spans technologies and topologies that deliver a comprehensive ADO solution.

Offload (Technology)

A BIG-IP system deployment naturally places it in a strategic point of control in the network. As a mediating component between client and server, it provides a variety of services that offload computationally expensive processing from the application infrastructure.



WHITE PAPER

Application Delivery Optimization

Core to the BIG-IP platform is OneConnect, a TCP connection offloading technology that reduces the overhead associated with TCP management from servers. A single connection between a BIG-IP device and an application can efficiently serve 100 clients with better performance by eliminating the time required to open and close TCP connections. By virtue of the BIG-IP platform's core operating system, TMOS, all BIG-IP solutions are enabled with this capability, including WAN optimization and web acceleration.

Offloading cryptographic processing to the BIG-IP system also affords a boost in application infrastructure performance. BIG-IP solutions take advantage of F5's integrated cryptographic acceleration hardware to dramatically improve the transaction and bulk encryption rates of secured services. This capability is paramount to enabling the implementation of the recommended best practice "SSL Everywhere" or "Always On SSL," as it reduces the impact of cryptographic processing across the entire application delivery chain. It is particularly useful when addressing long-distance transfers across the WAN to cloud computing and remote sites where a secured and fast transmission is desirable for security purposes.

Symmetric and Asymmetric Acceleration (Topology)

The advantages and disadvantages of symmetric and asymmetric optimization techniques have long been debated. The reality is that both are beneficial to optimization efforts. Each approach has varying benefits in specific scenarios, as each focuses on specific problem areas within the application delivery chain. Neither is necessarily appropriate for every situation, nor will either one necessarily resolve performance issues in which the root cause lies outside the approach's intended domain expertise. A successful ADO strategy is to leverage both techniques when appropriate.

This is particularly important when considering cloud computing environments as an extension of the data center, and in other scenarios in which symmetric optimization options may not be possible, such as unsupported mobile clients. Both asymmetric and symmetric acceleration solutions should be available and, if possible, able to dynamically provide services based on context to ensure optimum application performance. This includes the option of a "SoftWOC client"-a client-focused end point deployable on user devices that enables symmetric optimization when interacting with a WAN optimization-enabled server-side resource. BIG-IP Edge Client provides this functionality for a variety of operating systems, devices, and platforms.



Creating and tearing down HTTP connections is expensive, and adding SSL to the mix makes it even worse due to the necessary handshaking and encryption/decryption. Restrict SSL to data sensitive areas of your site such as the checkout only.⁶



WHITE PAPER

Application Delivery Optimization

Adding to the clear benefits of WAN optimization, the most common symmetric acceleration solution, a modern web acceleration solution like BIG-IP WebAccelerator can support both traditional symmetric and asymmetric deployments—even simultaneously. Deploying such a solution on a virtual platform—and thus within public cloud environments—further affords the ability to leverage cloud computing with less concern that it will negatively affect performance. When organizations combine web acceleration with WAN optimization in a comprehensive ADO solution, users experience significant improvement in performance due to the optimization at both the network and the application layers.

Perhaps most often associated with solutions designed to improve application performance are front-end optimization technologies, which are commonly associated with asymmetric acceleration technologies. Both focus on addressing performance issues that stem from the size and location of content, as well as limitations imposed by client browser and operating system environments. They require no special agent or client-side software to improve performance.

Both symmetric and asymmetric optimization techniques heavily focus on reducing content as a means to improve performance.

Content reduction techniques

Content reduction technology optimizes application delivery by reducing the size of data exchanges. This is particularly useful for mobile device clients, as mobile platforms are generally compute-constrained and their networks highly sensitive to latency. Reducing the size of content through compression and image optimization improves performance by speeding the transfer of data across such networks.

Similarly, data center-to-data center (or cloud center) transfers of large data sets, such as those required by file or database replication and highly interactive applications such as VDI, can be significantly slowed by high latency and low-bandwidth networks. WAN optimization solutions address these chokepoints by using deduplication and compression techniques to reduce the size of data.

Load Balancing (Technology)

When performance issues are ultimately caused by limited capacity and increasing load on application infrastructure, load balancing is the natural solution. By introducing a flexible tier of application delivery through load balancing, organizations can non-disruptively and dynamically achieve the elasticity required to meet demand efficiently and ensure top performance. Load balancing is the most critical component of an architecture that supports on-demand elasticity to eliminate capacity bottlenecks and the performance problems they cause.



Our BIG-IP Edge Client and BIG-IP Edge Portal applications give users secure, convenient, and fast access to company resources from the corporate and personal devices of their choice. With leading support for Android, Linux, and Apple iOS endpoints, F5 solutions provide flexible, optimized access to empower mobile users without sacrificing IT's ability to manage applications and infrastructures.

—Mark Vondemkamp, Director of Product Management at F5

Together, BIG-IP WebAccelerator and BIG-IP WAN Optimization Manager (WOM) offer a variety of content reduction techniques:

- Minification
- Image optimization
- Data deduplication
- Adaptive compression



WHITE PAPER

Application Delivery Optimization

Both on- and off-premises cloud computing offers IT organizations the ability to dynamically and efficiently expand and contract capacity. F5 ADO solutions integrate with leading virtualization management platforms to ensure on-demand elasticity within and across environments. Combined with the ability to define thresholds based on both capacity and application performance, these integrations enable organizations to maintain the application capacity required to meet performance goals.

BIG-IP Local Traffic Manager (LTM), which can be deployed on both hardware and virtual platforms, offers organizations the flexibility to architect an ADO solution that can ensure optimal performance regardless of location or environment.

Location (Topology)

The speed of light is non-negotiable, and therefore there is a finite set of transfer-related improvements that can be made to address availability and performance issues. One proven way to combat network-related availability issues is to move the content closer to the user, often by leveraging multiple, geographically dispersed data centers, or increasingly today, by taking advantage of cloud computing. But to use location to improve performance requires a global application delivery service that can make contextual decisions about routing application requests. To achieve this, it's critical organizations have visibility into environments and integration with local application delivery services.

F5 ADO solutions can natively extract geographic location from clients. When BIG-IP Global Traffic Manager (GTM) is integrated with locally deployed BIG-IP LTM devices—each providing performance- and capacity-related data— it can combine that data with the geographic location to direct requests to the site that can best meet user, operational, and business performance requirements.

Location optimization

Caching technologies have long been used to "move" content closer to the edge of the network, thus eliminating transfer times across the infrastructure. More recently, caching technology has evolved to include a more intelligent approach that also takes the client into consideration. These techniques take advantage of client-side caches to "move" content even closer to the client, eliminating most of the delay caused by transfer times by obviating the need to transfer it.

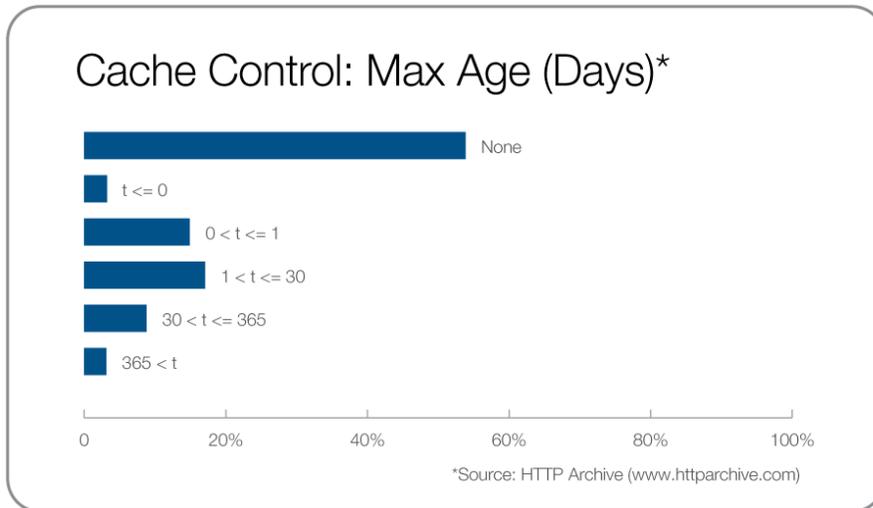


Figure 6: Over half of content has no cache control headers, indicating a good area for potentially improving performance.

Such techniques also use server-side cache management to improve the cachability of content. Inserting and/or manipulating cache headers to optimize cache strategies by leveraging both the client and the server-side caches dramatically improves performance by eliminating many transfers of data across multiple networks. BIG-IP GTM's geolocation capabilities in conjunction with BIG-IP WebAccelerator's device caching at remote locations can effectively replace a CDN, so organizations can eliminate reliance on external services to provide these capabilities.

Protocols and Protocol Behavior (Technology)

One of the most effective means of combating performance issues is to manipulate the protocols that transport and deliver applications. Transport and application layer protocols that attempt to provide reliable and fast delivery of applications have a variety of mechanisms designed to compensate for network layer conditions that can degrade performance and negatively affect business and operational expectations. Increasing demand on networks is putting pressure on organizations to explore not only new ways to mitigate performance problems, but to innovate and define new protocols that are based on the evolution of usage patterns and application development. Optimization techniques and support for emerging protocols are required of any successful ADO solution, and F5 recognizes this need by supporting both.

BIG-IP WebAccelerator offers a variety of location-based techniques:

- CSS caching
- Client-side cache control
- Cache-header management
- Small and large object caching
- SaaS optimization



Protocol optimization

There are a variety of TCP-related standards (RFCs) all specifically designed to address one thing: performance issues associated with TCP. F5's innovative approaches and techniques are designed to offset or completely avoid the effect of protocol-related problems.

For example, the BIG-IP system's TCP Express-core to all F5 ADO solutions- includes hundreds of custom features and improvements, as well as industry-standard solutions that redress performance problems regardless of their root cause. This enables ADO solutions to react dynamically to changing conditions and bottlenecks to ensure consistent and fast application performance.

TCP Express is a collection of TCP efficiency improvements, in the form of Internet standards (RFCs), and hundreds of custom F5 features based on extensive real-world experience. TMOS supports every modern TCP efficiency improvement, including:

- Delayed and selective acknowledgements, (RFC 2018)
- Explicit Congestion Notification (ECN), (RFC 3168)
- Limited and fast retransmits, (RFC 3042 and RFC 2582)
- Slow start with congestion avoidance, (RFC 2581)
- Adaptive initial congestion windows, (RFC 3390)
- Timestamps and window scaling, (RFC 1323)

For more information about TMOS, read [TMOS: Redefining the Solution](#).

Emerging protocol support

The technology landscape is constantly evolving, and there are new protocols that have both direct and indirect effects on performance. While some of these protocols, such as SPDY and HTML5 WebSockets and the recently proposed Speed+Mobility, seek to improve performance by changing core characteristics of both the transport and application layers, others such as PCoIP have emerged to offset the negative effect on performance of emerging technology like VDI.

These protocols, despite being developed as a response to growing performance concerns, can actually result in further impaired performance when the application delivery tier is not wholly prepared to support them.



WHITE PAPER

Application Delivery Optimization

Support for protocol-related techniques that address these issues is imperative to optimizing application delivery. In the face of the growing size and complexity of web content, SPDY has emerged as an increasingly popular protocol-especially for mobile devices and cloud computing. SPDY is based on HTTP, but is asynchronous, compresses most HTTP headers (unlike HTML5 WebSockets, which eliminates most HTTP headers), and prioritizes requests. In a dramatic divergence from HTTP, which is limited by browser and server configurations, SPDY further restricts clients to a single TCP connection per domain (but allows multiple asynchronous streams within that connection to improve overall performance).

The rather radical departure from the request-reply model of HTTP renders existing web and application servers-and most application delivery solutions-unable to take advantage of SPDY's considerable performance optimizations, despite the fact that Chrome and Amazon's Silk browser support SPDY, and it is anticipated to be supported by Firefox 11.

F5 resolves the issue of non-support by existing application infrastructure by including a SPDY gateway in its core application delivery platform, TMOS. With this gateway, organizations can take advantage of SPDY's optimizations to enhance application performance without requiring disruptive and potentially costly upgrades to its application infrastructure. This is also a boon for cloud-deployed applications constrained to specific infrastructure in off-premises deployments, as the BIG-IP system can provide SPDY gateway services to these applications via both virtualized and bridged cloud integration models.

The Role of Visibility in ADO

Visibility plays a unique role in enabling a successful Application Delivery Optimization strategy. While visibility is neither a technology nor topology that directly affects performance, administrators can leverage it to improve performance by using the statistics and metrics gathered and shared by performance monitoring capabilities to guide their choice of technologies and topologies. For example, it would be hard for load balancing services to distribute requests based on performance-related metrics without visibility into those metrics. Similarly, adaptive compression techniques can be detrimental to performance when applied to requests transmitted over high-speed, low-latency connections. Without extensive visibility, ADO solutions cannot be effectively applied to communications.



Initially, SPDY was a proprietary Google protocol implemented only in Google's Chrome browser. That's changing, however. Amazon's Silk browser includes SPDY support, and Firefox 11 will include preliminary SPDY support. Partially motivated by SPDY's uptake, the IETF's HTTPbis Working Group-the team of industry experts tasked with maintaining and developing the HTTP specification-is considering the development of a new specification, HTTP/2.0, with the goal of improving the performance of HTTP connections.⁷



WHITE PAPER

Application Delivery Optimization

A lack of visibility can exacerbate performance problems by making the process of finding and remediating bottlenecks longer and more difficult.

F5 ADO solutions are all application-fluent. They have the visibility of key performance and availability indicators that provide the critical data needed to make intelligent decisions regarding routing of requests, even across environments.

This critical data extends beyond the application to the end-user and both the server- and client-side networks, ultimately enabling the application delivery tier to make context-aware decisions, including the dynamic application of optimization policies to ensure per-session optimizations, rather than per-application optimizations that can negatively affect a subset of end-users.


Because of its deployment as a strategic point of control between clients and services, the BIG-IP system can provide the depth and breadth of statistics organizations need to efficiently and effectively use ADO solutions. F5 Analytics provides granular views of metrics gathered that power decisions ranging from when to apply compression to where to service requests. Visibility into these statistics enables the BIG-IP system to more accurately make decisions that improve performance.

F5 Enterprise Manager gives organizations a complete view of application performance with Centralized Analytics, which improves the user experience by enabling just-in-time provisioning of additional resources to ensure capacity demands are met.

Conclusion


F5 Application Delivery Optimization is not just about content acceleration or WAN optimization or implementing a caching solution. Optimizing application delivery in today's highly dynamic and dispersed environments requires a comprehensive, holistic approach that improves performance without sacrificing the agility of modern data center architectures.

With an integrated, overarching view of application delivery, F5 ADO solutions can collaborate, share critical visibility data, and dynamically apply the appropriate acceleration and optimization techniques to ensure a consistently fast end-user experience that mitigates the risk of losing users-and their associated revenue.



The iStats feature of the BIG-IP system can collect a wide variety of statistics that can be used to optimize application delivery by adjusting request routing and acceleration policies:

- Server latency
- Page load time
- Throughput
- Response codes
- Methods
- URLs
- Client IPs
- Client geographic
- User agent
- User sessions



¹ Gilbertson, Scott. "Users Expect Websites to Load in the Blink of an Eye," www.webmonkey.com. March 2, 2012.

² Billing & OSS World (B/OSS), "Survey: Security Sacrificed for Network Performance." July 19, 2011.



WHITE PAPER

Application Delivery Optimization

- ³ Billing & OSS World (B/OSS), "Survey: Security Sacrificed for Network Performance." July 19, 2011.
- ⁴ Barney, Doug. "The Great Cloud Bottleneck: How Capacity Issues Can Kill Your Cloud Project," Redmondmag.com. December 6, 2011.
- ⁵ Scarpati, Jessica. "VDI over the WAN: How latency affects virtual desktop performance." SearchEnterpriseWAN.com. October 6, 2010.
- ⁶ Bustos, Linda. "Every Second Counts: How Website Performance Impacts Shopper Behavior," Get Elastic. November 5, 2009.
- ⁷ Bright, Peter. "Google Works on Internet Standards with TCP Proposals, SPDY Standardization," www.webmonkey.com. January 25, 2012.

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