F5 White Paper

Providing Security and Acceleration for Remote Users

Delivering applications to remote users is a significant undertaking. Applications need to be available, and they must be delivered securely and quickly.

by Peter Silva
Technical Marketing Manager - Security

Contents

Introduction

The traditional IT model evolved from a world where resources, users, and access methods were under our control. Relationships among users, applications, and data were static and tightly bound; and applications were written with specific display layouts in mind. As remote and then mobile users were added along with partners, contractors, and guests, and as IT was distributed globally, the old model broke down. IT leaders are frustrated with their ability to respond to business needs because the underlying IT infrastructure constrains choice, slows response, and limits their ability to manage change. Costs increase, resources aren't shared, and management is complex.

These challenges are amplified as more application traffic moves across the web: While user access and sensitive corporate data needs to be managed, there's often only a limited view into and control over where a user navigates and what is accessed. Organizations want to provide basic web access to corporate networks; at the same time, they need to adhere to strict regulations concerning data access. In addition, as more workers trade their onsite office cubes for telecommuting from remote locations, workforce collaboration becomes even more critical.

Delivering the Data

Many organizations might be familiar with long waiting times while a web page loads or a large file download. In addition, there is a need to secure sensitive data on the Internet, especially due to the increasing number of telecommuters requiring a secure connection to email and corporate resources. Delivering IT applications is not a simple exercise. There are many issues to address: How do I make sure applications are always available? How do I secure applications? How can I make sure I'm using my resources (for example, servers, bandwidth) efficiently? How do I make sure delivery is optimized for the best user experience?

You can solve these problems one at a time, by buying simple load balancers and other single-purpose devices, by modifying the applications, and by paying for more resources. However, this strategy is costly to build and manage. Security is another concern when delivering applications. Here too, you may have some options such as modifying the application or installing point solutions. However, this too is costly and could still leave you vulnerable.

A typical network might only have visibility at the Internet Protocol (IP) or media access control (MAC) address level, and it cannot determine whether a user is a contractor, guest, or employee. This level of security has been acceptable to many organizations, because applications and other critical resources are protected via authorization and authentication processes (usually user ID and password) and identity access management (IAM) solutions. However, because networks are blind to a user's identity, the risk is that users "see" applications that they are not authorized to access. For example, a contractor who has been granted network access could "go exploring" (undetected) and attempt to access sensitive information.

Security and identity is an integral part of any application infrastructure that delivers authentication, authorization and accounting (AAA) services. Integrated application access control with scalable web security can drive user and group identity into the network for policy-based control to applications.

Application Delivery Made Easy

The Application Delivery Controller (ADCs) platform is specifically designed to address these issues by serving as a strategic point of control in your network, ensuring that applications are fast and available, and that applications and data are secure. The F5® BIG-IP® Local Traffic Manager™ (LTM) is the market-leading ADC. It load balances, secures, and optimizes application traffic, giving you the ability to add servers easily, eliminate downtime, improve application performance, and meet your security requirements. BIG-IP LTM provides the advanced features you need to direct users to the best possible resources at the application level.

Availability

Let's start with the first issue for application delivery: How can you make sure applications are always available? In the past, you could address this challenge with a simple load balancer. Spread the traffic among several servers and you're done. But as applications get more complex, your method for load balancing has to keep up. You can't just spread traffic around; the load balancer needs to understand the application to distribute the traffic appropriately. BIG-IP LTM ensures the best resources are always selected, has deep visibility into application health, and proactively inspects and responds to errors. While BIG-IP LTM manages the local resources in each data center, F5 BIG-IP® Global Traffic Manager™ (GTM) can automatically direct users to the closest or best-performing data center. When you couple the features of BIG-IP LTM with those of BIG-IP GTM, you can realize the full

potential of multiple data centers and provide seamless disaster recovery and routing based on quality of service or business criteria.

Security

So, how do you secure applications? Once your applications are available to users, you still need to make sure only authenticated users gain access, and that they only access the resources they are authorized to see. Management of AAA in a web application deployment can be costly in an enterprise infrastructure. There are a number of ways to authenticate web users today. You can code authentication into the application during development, but that can be costly, difficult to change, and may not be as secure as you need. You can install agents on the servers, but that is difficult to manage, and not particularly interoperable. Additionally, this design can become costly in regard to both deployment and management since it's decentralized and every single server will need attention. Authorization and accounting, which might be required for regulatory compliance, can be overlooked or completely missing with these strategies. Another choice is to install additional specialized access proxies; however, they are unreliable, costly, and are not scalable. Plus, you'll add units to your infrastructure, defeating any consolidation efforts.

Adding F5 BIG-IP® Access Policy Manager™ (APM) to BIG-IP LTM brings identity, authentication, and access control to any BIG-IP environment, giving IT the ability to consolidate infrastructure, reduce AAA management costs, and drive user identity into the network. BIG-IP APM centralizes web single sign-on (SSO) and access control services, offers a full proxy L4-L7 access control (at wire speeds), adds endpoint inspection to the access policy, and includes the Visual Policy Editor (VPE) feature, which provides policy-based access control along with VPE Rules, a programmatic interface for custom access policies. VPE gives administrators control to create and manage security policies and resources with ease and the flowchart design shows exactly what types of inspections are enabled.

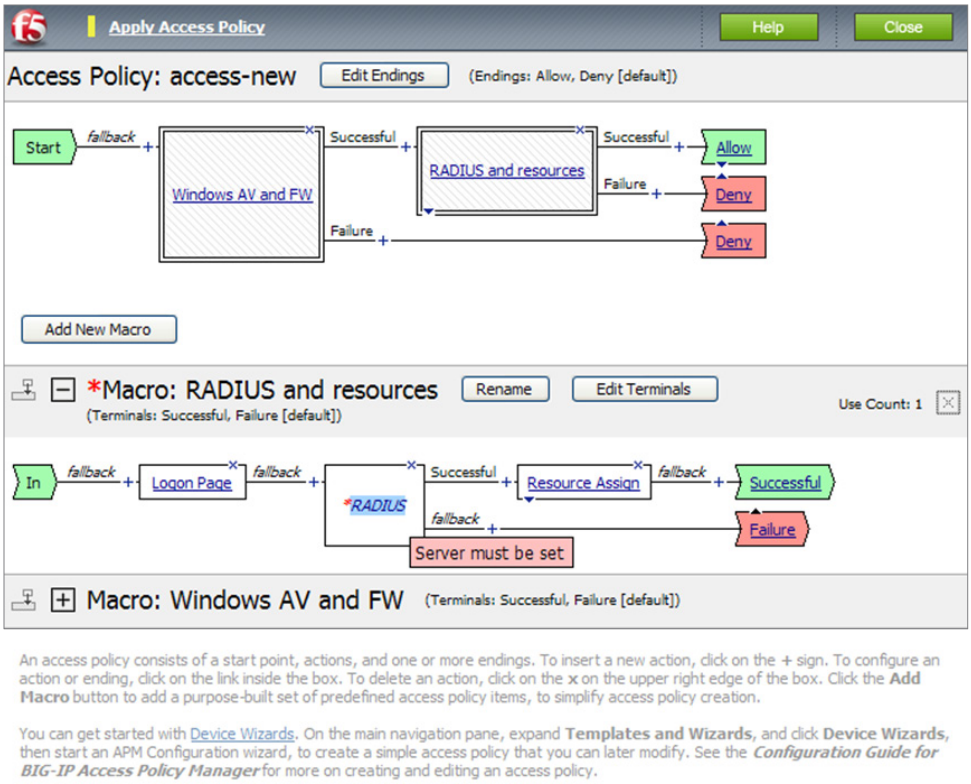

Figure 1: Create and manage security policies with ease using VPE in BIG-IP APM.

In the case of web application authentication, you can use BIG-IP APM to replace specialized access proxies or agents and gain superior scalability and high availability. You can also initiate an endpoint host inspection for any client requesting access to your web application, whether it be public- or internal-facing, to ensure a minimum security posture and enforce stronger authentication—such as HTTP or forms—than is typically available for web applications.

Imagine, for example, going to a banking website or an internal corporate customer relationship management (CRM)/enterprise resource planning (ERP) system; and before you even get a chance to log on, there is a check to ensure you have antivirus protection and/or a local firewall enabled, or that you possess a certain certificate that verifies your identity. You might feel a greater sense of security and look to only do business with web sites that offer that level of protection. From an administrator point of view, this can have a two-fold effect. First, you'll be educating the user on the importance of security and second, you can ensure that the user only has access to their information with advanced authentication and access control along with detailed reports on every user session activity helping you maintain regulatory compliance. The BIG-IP APM dashboard gives an overview of active sessions, throughput, new sessions, and connections in a readily viewable and customizable reporting pane.

There may be situations in which one needs to provide authentication and client validation prior to giving access to unsecured applications. Secure authentication can be added for this purpose. This authenticating method is used first to create a secure session and then to give access to the services behind it. For instance, you may not want to lock down your public-facing website but certain requests could require authentication (for example, if a user were requesting access to a restricted folder). With secure authentication, anyone can navigate the main page but as soon as the user clicks a "member" area, access control provides the gate and fence.

BIG-IP GTM also has unique security support for Domain Name System Security Extensions (DNSSEC), even in a global server load balancing (GSLB) situation. DNSSEC will protect your domain from such attacks as DNS cache poisoning and other DNS vulnerabilities. DNS attacks are major threat on the Internet, and businesses do not want their users redirected to rouge or fraudulent sites. Since BIG-IP GTM is built on the F5 TMOS® architecture, it can interact with other BIG-IP devices deployed in the infrastructure, providing insight into the entire system, not just the local racks. BIG-IP GTM has IP geolocation features that identify where users are coming from and, in turn, where to send them for the most appropriate and available data center. With IP geolocation, you can also block suspicious IPs that may be creating havoc for the site.

BIG-IP APM for web access management offers enterprises cost-effective, policy-based user access control, unified application access control, application security for complying with regulations (for example, PCI, HIPAA), secure connections with SSL, and integration with existing enterprise infrastructure and applications.

Value of Access

The mobile workforce is continuing to grow at a rapid pace. The IDC predicts that the mobile worker population will increase to 1.2 billion in 2013, accounting for 33 percent of the worldwide workforce. Paramount to business operations is the ability for mobile workers to access corporate resources in a secure manner. However, there are abundant challenges. There is a significant cost in scaling out a secure remote access solution since the performance of current SSL VPN solutions can suffer when the load gets too heavy. AAA management becomes even more critical—not only to ensure that valid users are requesting access, but also to make sure they are only gaining access to the resources they are authorized to view. Additionally, global teleworkers will be using a wide range of clients connected to various types of networks. The new model for remote access requires access security, acceleration services, and application availability.

Efficient Use of Resources

How can you ensure you're using IT resources as efficiently as possible? With support for up to 40,000 concurrent users on one device, enterprise-level scalability, application delivery that is twice as fast as traditional SSL VPN, the F5 BIG-IP® Edge Gateway™ advanced remote access solution provides a unique set of features to consolidate and unify all your access needs for half the price of competitive solutions. BIG-IP Edge Gateway integrates with the existing enterprise infrastructure, providing AAA access services to networks, applications, and portals. Access administrators can configure and load balance multiple authentication mechanisms including Active Directory, LDAP, RADIUS, HTTP, and RSA SecurID, along with SSO and credential caching to enhance the user experience. Corporate users may hit the Active Directory server for authentication while partners query LDAP, but they both go through the same endpoint scrutiny prior to gaining access. With dynamic, per session Layer 4 and Layer 7 access control lists, you can keep users within their functional, authorized locations at a fine-grain level—even down to a specific folder path within a web application. All this is done over SSL, so every connection is encrypted. BIG-IP Edge Gateway also includes the same VPE as BIG-IP APM, which gives you the flexibility to create a single access policy to cover all access requests or to design specialized policies for each access method or access group.

Comprehensive endpoint security keeps all devices that request access within your unique corporate compliance policy. As more personal computing devices are being used within the workplace, the concept of trusted vs. untrusted is becoming arcane. IT should treat every device, including IT-issued equipment, as untrustworthy until proven otherwise. Deep inspections can determine antivirus and local firewall status, whether there is a client certificate, if the system contains a particular identifier and the host's overall security posture. Administrators can also downgrade access if a certain device fails to meet one or more requirements. A user may receive full access when working from an IT-issued device that is up-to-date; however, if that same device fails some criteria, or if that same user is working from a personal device, rights may be restricted to intranet portals only.

For those employees with company equipment, the F5 BIG-IP® Edge Client™ can be installed as part of the overall corporate image package to enable always-connected application access. The BIG-IP Edge Client smart connection feature offers location awareness and zone determination for VPN support. It knows when a device is no longer connected to the corporate domain, and it can automatically initiate a secure SSL tunnel to allow the user to seamlessly and securely move from one network to another, improving client productivity and mobility. If a user loses their

VPN connection due to network conditions, BIG-IP Edge Client will automatically reconnect the device once connectivity is reestablished.

Optimization

Optimized delivery is one of the biggest challenges in providing an optimal user experience. Remote workers want an experience similar to being on the corporate LAN even if they are connected to a high-latency network. If an IT department wanted to deploy some sort of application acceleration, they would need to procure new specialized controllers and add them to the mix of equipment already being maintained by IT. This additional equipment adds both CapEx and OpEx to already tight IT budgets.

How can you ensure delivery is optimized for the best user experience? BIG-IP Edge Gateway comes equipped with application acceleration services to deliver LAN-like performance to any user around the globe. The WAN optimization features included in BIG-IP Edge Gateway employ adaptive TCP optimization to speed traffic, using techniques such as session-level application awareness, persistent tunnels, selective acknowledgements, error correction, and optimized TCP windows. Symmetric adaptive compression reduces data through the use of dictionary-based compression and advanced encoding schemes that make adjustments throughout the session to provide the best possible combination of speed and compression. The HTTP/HTTPS asymmetric acceleration rapidly increases the delivery of web applications whether they are being served in-house or from cloud infrastructures. Adding application acceleration through BIG-IP LTM enables traffic shaping and quality of service with adaptable compression for situations when latency is an issue. With these F5 solutions, web application deployment becomes a whole lot easier, safer, and more resilient—whether from a traditional data center or cloud environment.

Conclusion

Through a range of products, F5 has solved the security plus acceleration challenge for remote users starting with BIG-IP Edge Gateway and BIG-IP APM. With the powerful, easy-to-use management interfaces, IT administrators can create detailed access policies that are easy to understand and deploy. Scalability, performance, and reliability have been hampered by traditional SSL VPN solutions. However, F5 products directly address these key challenges. BIG-IP Edge Gateway (with WAN optimization included) touts 40,000 concurrent users, 8-gigabit throughput, and the F5 TMOS architecture, giving users BIG-IP LTM performance with added security

and optimization to provide a superior user experience. Supporting a breadth of clients, applications, and infrastructure, IT can consolidate its infrastructure, support global users, and offer security plus acceleration all on one BIG-IP Edge Gateway. BIG-IP APM on BIG-IP LTM brings identity, authentication, and access control with an easier, more powerful, and cost effective way to manage your network identity and web application infrastructure.

When you combine the solution with BIG-IP GTM, you add global intelligence to your data center infrastructure, providing seamless disaster recovery, IP geolocation, DNSSEC, and intelligent routing for the worldwide user base. All these BIG-IP solutions run on TMOS, which allows additional features and functionality to be deployed in a non-disruptive manner, enabling each to share context and real-time conditions on the network making application delivery fast, available, and secure.

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

F5 Networks, Inc. Corporate Headquarters	F5 Networks Asia-Pacific	F5 Networks Ltd. Europe/Middle-East/Africa	F5 Networks Japan K.K.
info@f5.com	apacinfo@f5.com	emeainfo@f5.com	f5j-info@f5.com

CS01-00012 0410