F5 White Paper

Simplifying Application Access Strategies

Unified, ubiquitous access has historically been out of reach and enterprises are clamoring for a single appliance to handle all their secure access needs. Typically companies needed multiple single purpose devices to handle their access needs…. until now.

By Peter Silva
Technical Marketing Manager

Contents

Introduction

For years, corporations were forced to deploy many different appliances, controllers, gateways, and other types of specialized equipment to control access to their internal network and ultimately, the applications and data that reside there. Initially, you had to be on the LAN to interact with internal applications. Soon, Remote Access Service was introduced to enable users to modem-dial into the network, so they could use those same applications when away from the LAN. Over the years, many vendors have offered a variety of both hardware and software remote access solutions, including today's IPSec and SSL VPN.

Other ways of connecting to a network also appeared, like wireless access. At first, Wireless (802.11) Access Point's SSIDs were broadcasting, open, and the transmissions were unencrypted. Evolving technologies introduced WEP/WPA, administrators changed default passwords, and access points became more secure, often requiring a key to gain access. However, these connection points also became another vulnerable entry point to the applications.

Lately, NAC (Network Access Control) garnered accolades as the way to protect LAN resources and control access. However, those reviews went from outstanding to disappointing in less than two years. With no real NAC standards in place and various vendors offering different point products, it's no coincidence that NAC popularity is declining.

No matter what method is used to gain admittance to internal resources, IT administrators usually have to control that access with AAA (Authentication, Authorization, and Accounting) servers, enable some endpoint inspection, manage resources based on certain criteria and potentially, secure the desktop—and that's just for the trusted employee devices. What about that guest user? Complex, costly deployments along with cracked security is not a good way to protect your applications. You can protect critical assets while keeping it simple for both user and administrator.

It's the Application

What is the most critical part of your IT infrastructure to protect? Depending on who you talk to, you'll get different answers. Some will argue it's the network. This makes sense if all you're going to do on the network is acquire an IP address, "get connected," and nothing else. But, if you want to exploit a specific vulnerability on an internal system, the network gateway and firewall aren't going to prevent that since you're already on the internal network .Also, these devices are not aware of the application (except for how to find it) after you pass their policy. Layer 3 devices are not really designed to be layer 7 aware.

Because of this, your data and applications are the most critical to secure. In fact, federal laws have been passed regarding the protection and integrity of data. Client devices, which are on the other side of your infrastructure, should be the second most critical part of your IT infrastructure to secure. You don't want an infected machine on the network, let alone interacting with your applications, nor do you want the potential data-leakage of any device connected to your infrastructure. The network is just the road—the connection and paths that lead to your data. Identity management and application access and delivery should be your primary focus.

Challenges

Wireless Challenges

Almost every company has some sort of wireless access in their offices, both for employees and guests. Historically, and still often today, the access point's SSID is broadcast as an open, available signal with no key required. Sometimes, it is secure but the password is so well known that it might as well be an open signal. It is better to have a good password and decent encryption to enable basic Internet access for both employees and guests. Often, companies offer just a VLAN for outbound traffic, which is perfect for your partners and contractors, but what about employees? Sometimes organizations provide two access points so employees can tune into another broadcast with a different pass code that segments them to their normal "internal" access. However, you'll need a WLAN controller to govern those access points and depending on your wireless infrastructure, this could mean many devices.

Remote Access Challenges

SSL VPN has emerged as the tool of choice for most new remote access deployments. IPSec is still used (along with site-to-site) in a lot of corporate environments and there are lingering philosophical debates about the merits of IPSec and SSL VPN, but for this discussion we'll focus on SSL VPN. SSL VPN has evolved from a niche remote access method to a mature technology. Nowadays, it can do customized pre-logon security checks, AAA anything including two-factor, VLAN, quarantine, port-map, reverse proxy, granular application control, and the list goes on. As the concurrent user count grows, however, SSL VPN presents challenges. All of the top market-leading SSL VPN vendors can only support a few thousand users on a single unit. Many IT departments understand the benefits of SSL VPN; however they'd need about five SSL VPN appliances to replace every one IPSec unit. Even with the advantages of SSL VPN, this alone makes the transition a tough sell. IPSec concentrators typically have a higher user capability, but not the endpoint security features. While SSL VPN is perfect for some organizations, it frequently won't suffice for global deployments.

NAC Challenges

The biggest hurdle to Network Access/Admission Control (NAC) is that it is still a concept. There are no standards for NAC, so vendors provide their own interpretations of the concept, leaving companies to sort through the confusion. Interoperability, complexity, and costs are major inhibitors of NAC along with the gap between expectations and reality. Throughout the last year, research firm TheInfoPro (TIP) noticed that the use of NAC is actually declining from about 35% to 25% in enterprises. They've also seen a 3% increase in the number of enterprises that do not plan on implementing NAC.

There are a few different types of NAC, including Client, Edge Enforcement, and Inline/Switch based. Client-based NAC focuses exclusively on the endpoint device and/or security posture of the client. The vendor-specific control software is either pre-installed or "pushed" on demand via browser components/controls. If you're considering an Edge Enforcement-type NAC, consider this—they focus almost exclusively on the endpoint device and most of the time, you need them at every access point on your network.

The Switch/Network/Inline based NACs rely heavily on 802.1x. Simply put, 802.1x is a sentry-type protocol which means it either allows or doesn't allow traffic. It is port-specific and can implement VLAN/ACL restrictions, but it can also be expensive as all network devices need to support the protocol. Feature upgrades from typical vendors can range from $20 to $200 per seat and there is still no application awareness.

All these concepts can get expensive, expansive, and exhaustive to manage. Some can disrupt the user experience and none of the above mentioned solutions can give you granular access to your most critical asset, the application. None of them even completely solves the idea of NAC.

So where does this leave us? We have a bunch of equipment that does different things, managed by different people, potentially with different policies. In addition, most of the access blueprints do not have application awareness. Those that have application awareness can't sustain the performance needed for a global strategy. Lastly, the divisions between remote, wireless, and LAN continue to blur together; all of these lead to the same thing anyway—the application's data.

BIG-IP Secure Access Manager Solution

Protecting your intellectual property is the first, important goal, along with striving for unified secure access. Many seem to think security is about restricting access but it's really about allowing the right users access to the applications they need, depending on their location, device, permissions, and other factors.

F5® BIG-IP® Secure Access Manager™ leverages SSL technology to provide secure connectivity to any user regardless of the location or network the user is coming from (remote, wireless, internal LAN) while also enabling access to any authorized resource in the corporate network. BIG-IP Secure Access Manager is capable of securing the endpoint, the network, and the application by providing and enforcing granular access control, visibility, and auditing at unprecedented levels of performance. Unlike current point products that require IT administrators to deploy and define policies on multitude of systems to support individual access scenarios, BIG-IP Secure Access Manager offers a flexible, easy to manage, high-performance platform to provide secure access to any user (partners, employees, suppliers, contractors, and more), from any location and client device (managed and un-managed) to any corporate resource.

How can a single device achieve what many have not? BIG-IP Secure Access Manager is built on F5 TMOS™. TMOS is F5's purpose-built, custom architecture that has become the powerful and adaptable foundation for F5 products. TMOS is a full-proxy based architecture so it completely understands the communications and protocols passing through it.

Simplifying and Unifying

One of the keys to a successful application access strategy is easy and centralized policy creation and management. In addition, you need to control application access at a granular level and need to employ strong endpoint security and AAA. You need your users identified and secured, no matter how they are accessing the data and applications and you shouldn't have to add a bunch of devices or upgrade your hardware. With BIG-IP Secure Access Manager, you get the universal secure access features that ensure all of these capabilities.

With BIG-IP Secure Access Manager's Access Control, you get full command over your access policies and their accompanying profiles. Access control also governs your AAA identity servers along with granular application access control.

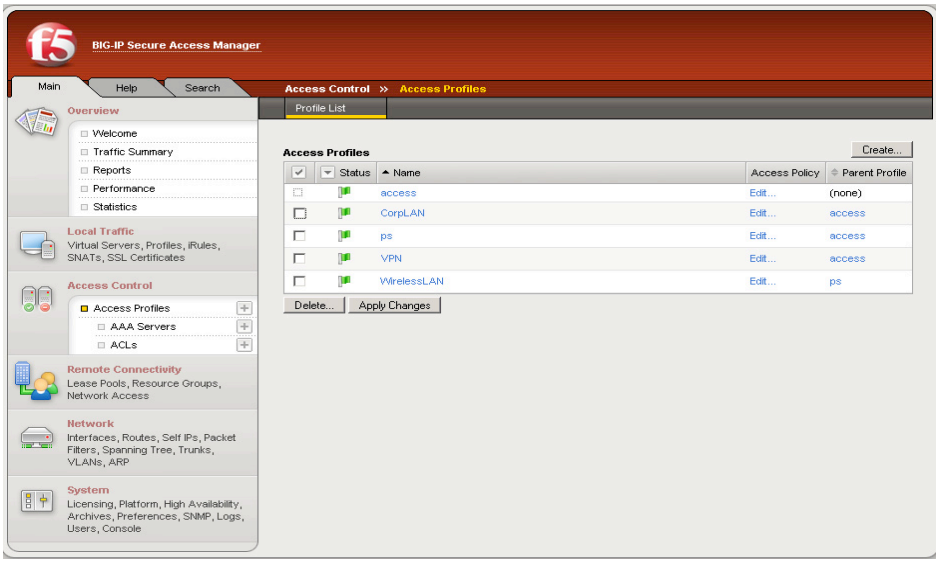

Figure 1: Screenshot of BIG-IP Secure Access Manager's Access Control

You can create multiple access profiles each with their own access policy. There is a corporate LAN profile—one for the VPN and another for the 802.11 connections. Many companies will have just one access profile that covers every connection from any device. Determining if there is a singular policy covering all access methods or if there is one for each separate method would depend on many factors such as regulatory obligations or compliance, current laws, corporate policies, and overall risk to the business. BIG-IP Secure Access Manager offers both options.

Creating, editing, and managing policies should be simple. BIG-IP Secure Access Manager offers a Visual Policy Editor (VPE) for the entire process. Based on the award-winning F5 FirePass® VPE, BIG-IP Secure Access Manager takes it a step further. The VPE helps you create a policy to check endpoint client's security posture, like you would any other requesting device, but also enables the administrator to configure the authentication method and assign resources.

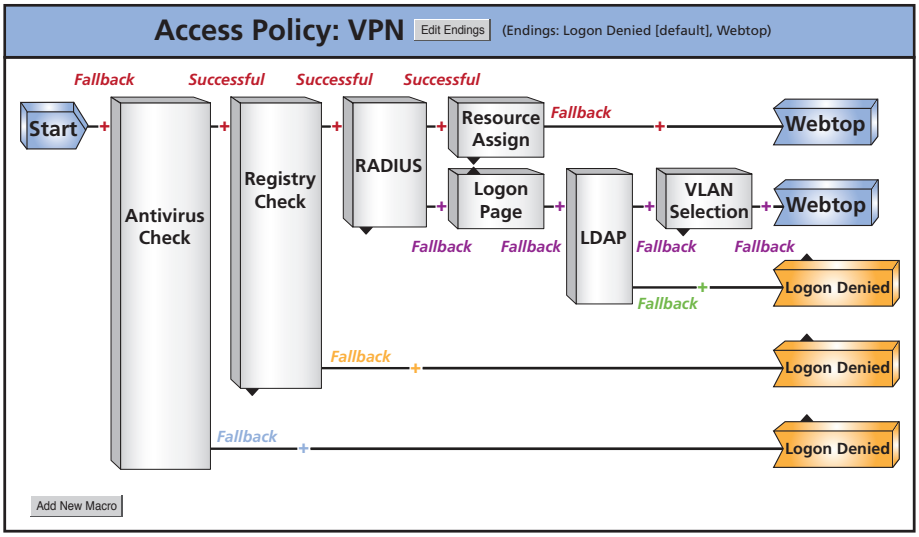

Figure 2: Example of a Simple VPN Policy

The access policy shown is basic, but this type of policy could also be used for VPN users, WiFi access points, and/or the local LAN access. It only takes a few clicks to give you a comprehensive, easy-to-understand policy. Actions are added with ease, and each action has specific rules attached. Consequent results and subsequent actions are based on the inspection, query, user, device, and many other attributes. To make changes to actions, simply click the box and configuration details become available. For example, the RADIUS authentication action could easily be changed to Active Directory or any of the other authentication methods. The same with the Antivirus Inspector—you are able to check for only the corporate-issued software on any of more than 100 antivirus vendors.

Administrators can also check the firewall, registry, OS, file, and perform many other client-side checks as part of the pre-logon inspection. In many ways this is just like a client NAC, but better and easier. After scrutinizing the device and determining if it meets the minimum security requirements, BIG-IP Secure Access Manager can authenticate the user, authorize access and assign appropriate resources all while virtualizing the application. Depending on the policy, those resources might be an internal web application, terminal server, or a layer 3 network access resource, if the user is remote.

Access Control Lists (ACLs) also reside under Access Control. ACLs are a bit stronger than IP-Filters in that an ACL can be set for both layer 4 and layer 7. Layer 4 ACLs are the usual network type, designating IP address(s), port(s), protocol(s) for both source and destination. Actions include Accept, Discard, Reject, and Continue. The continue action prompts a move to the next ACL. which may be even more granular for a specific application or VLAN/subnet. Application layer ACLs can be even more powerful. Administrators can define the same rules as in layer 4 but also add some application intelligence. An administrator, for instance, can restrict a user to a specific path or folder within the application. By doing this, BIG-IP Secure Access Manager restricts users and devices from mistakenly or purposely venturing to other locations. It also ensures users are getting the appropriate access based on their rights, security posture, location, or connecting device.

F5 Networks, Inc.
Corporate Headquarters
401 Elliott Avenue West
Seattle, WA 98119
+1-206-272-5555 Voice
(888) 88BIGIP Toll-free
+1-206-272-5556 Fax
www.f5.com
info@f5.com

F5 Networks
Asia-Pacific
+65-6533-6103 Voice
+65-6533-6106 Fax
info.asia@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
+44 (0) 1932 582 000 Voice
+44 (0) 1932 582 001 Fax
emeainfo@f5.com

F5 Networks
Japan K.K.
+81-3-5114-3200 Voice
+81-3-5114-3201 Fax
info@f5networks.co.jp

Simplifying Application Access Strategies 01/09