

Enabling Security for Small and Medium Sized Government Agencies

Government agencies, whether they are federal, state, or local, have a high standard to uphold when it comes to security. While some government agencies enjoy the manpower, budget, and logistics of a large enterprise, there are many government agencies that operate a lot like SMBs (small and medium sized businesses). They have limited IT staff and budget, and are chronically pressed for time. However, they have the same public duty to protect citizen's data and confidentiality, so the security implications are just as important to the SMGA (small and medium sized government agencies) as they are to large government organizations.

While the SMGA faces many of the same technical, logistical, and staffing challenges as traditional SMBs, security is one of the toughest aspects for the SMGAs to handle. Issues such as resources, budget constraints and a lack of specialized security talent are an ever-present problem.

Staff charged with developing, deploying, and maintaining an IT infrastructure for an SMGA, need to be pretty good generalists. Larger organizations can afford to have staff that are highly specialized in certain fields, but SMGAs require a small IT staff to be familiar with networking, server maintenance, storage, backups, web development, and more. The SMGA has to put together a small, talented staff that can handle a bit of everything.

To many an SMGA, security remains a dark art. Not only is it a field that involves a high degree of specialization, but is one perceived to be out of reach for many. Information within the field of security often seems to be available to only those in the know. It can often be perceived (sometimes deserved) that security technicians and engineers are secretive, difficult, and unapproachable. However, the only thing more terrifying than dealing with a security expert who seems only too happy to show you your ignorance, is having client data pilfered by hackers. So, the SMGA needs to make security a priority, no matter how difficult the field may be perceived.

A Brief History of Security

SMGAs, and indeed large enterprises, have long relied upon traditional firewalls for security. Traditional firewalls work on Layer 4 of the OSI stack, which is the layer that corresponds to an IP address and a TCP or UDP port. The IP address denotes the server, and the port typically denotes the type of service (web server, mail server, etc.).

A traditional firewall protects a site in a couple of ways. First, it only allows connections for services that are authorized. A server might be running a web service as well as a remote login service, but you only want the web service to be accessible to the rest of the world. A firewall would prevent public access to other unauthorized services.

This protection is important as there was a time when many of the site attacks were through this unintended access. Superfluous services, such as the old Finger protocol, provided an attack vector allowing intruders to obtain administrator access to a device without even logging in locally. When an attacker is able to gain administrator access from outside your network without local access, it is known as the dreaded "remote exploit".

While traditional firewalls have helped prevent this by reducing the number of attack vectors available, would-be attackers have adjusted their tactics. More exploits have been discovered in available services, and services have been used in ways that were never intended, in order to gain access to local networks and systems. It has become necessary not just to lock the doors, but to patrol the hallways of the network, so to speak.

Keeping the web server patches updated has helped mitigate this issue, as have dutiful code audits by the various web server vendors (Apache, Microsoft, etc.). However, the landscape is always changing.

While the SMGA faces many of the same technical, logistical, and staffing challenges as traditional SMBs, security is one of the toughest aspects for the SMGAs to handle.

It used to be that the simple act of serving up a file was the primary function of the web. Web pages were largely static, and there wasn't a great deal of user interaction, but the web has dramatically evolved.

The web application has emerged. Instead of static pages with limited uses, developers have created dynamic pages with nearly unlimited uses. This has greatly improved the ability for the web to provide useful and, indeed, necessary services. CNN.com, Facebook, and webmail are all examples of web applications. Virtually all of the Internet is dynamic now, as just about anything useful on the Internet right now is a web application, instead of a simple web page. The term web page is rarely used now.

While web applications have brought untold benefits to the Internet, web applications have also introduced another attack vector: The web application itself. These potential exploits typically don't even involve the web server, they often involve the web application code with the underlying development platform (such as ASP, PHP, Java, .NET, etc.), or the underlying database layer (Microsoft SQL, MySQL, Oracle, etc.).

To combat this new threat, a new type of security device is required, one that looks at the actual requests by operating at the application layer (Layer 7) rather than the IP address and port information, as with a traditional firewall. The solution to the Layer 7 problem came first in the form of IDS systems (Intrusion Detection Systems). These are devices that actively sniff the network, searching for the tell-tale signs of a network attack, hacking attempt, or successful intrusion. If such an event were detected, they would report it either directly to staff, or report it to a secondary reporting system that would alert the staff.

The problem with this solution, of course, is that it is only event detection. IDS systems, by definition, don't do anything about an event except report it. Hopefully, someone receives the report, and is able to determine if a security event has occurred, and knows what to do about it. There are automation systems that can automatically react based

on what the IDS system reported, but the integration is typically not ideal.

The next stage in this evolution is the IPS, which can actively block attacks as they occur with no manual intervention required. An IPS still reports the event, but it could often take appropriate action on its own.

IPS devices are typically generic to the entire network in that they work with all types of network traffic that may traverse an organization's network infrastructure. Such traffic might include web traffic, FTP, remote desktop, streaming media, peer-to-peer, or other protocols.

As an added challenge for the SMGA, IPS systems need to be deployed in a very deliberate manner, to ensure that ideal network access points are monitored. It doesn't do any good to put a guard at the front door if intruders come in through the side door.

In response to this need for an easy-to-deploy IPS device, and one that specializes in web traffic, a new breed of firewalls have emerged to add a new layer of protection: The Web Application Firewall (WAF). A WAF will look at not only the Layer 4 information (IP address and port), but also unpack the network payload and look at the actual traffic coming into a website.

Web Application Firewalls

Web application firewalls are fluent in the language of HTTP. By looking into the HTTP payload itself, they can match traffic against a database of known exploits. Attacks can then be intercepted, and stopped at the WAF, without ever reaching the web application.


WAFs have the benefit of typically being much simpler to implement and maintain than IPS, for there is really only one place to put them in a network. With an IPS deployment, you need to determine what the best monitoring points are on your network, as well as arrange the appropriate mirrored data ports. With a WAF, it's as simple as putting the device in-line with traffic.

Application Delivery Controllers – A sensible approach

WAFs have recently been implemented within ADCs (application delivery controllers), or load balancers as they are commonly referred. All traffic for a site typically traverses the ADC, so it's the perfect place to perform this type of security checking. By putting a WAF in an ADC, it's a simple matter of just turning the feature on to get the added security benefit of an application-aware firewall.

Acting as the Swiss-army knife of web infrastructure deployment (load balancing, server health monitoring, WAFs, SSL acceleration/termination, etc.), ADCs allow a site to quickly and easily deploy a web infrastructure while providing PCI-DSS compliant security.

The duty of SMGAs to protect citizens confidential data often go beyond those of a traditional SMB. The SMGA has all the responsibility of a large government organization, with far fewer resources. By using an ADC equipped with a WAF, the SMGA can more easily and effectively manage security using fewer financial and staffing resources than would otherwise be possible.

High Availability	Layer 7 Persistence	Layer 4 Persistence	Resource Load Balancing
Layer 4 Load Balancing	 Application Delivery Controller		HTTP Compression
Layer 7 Content Switching			HTTP Caching
SSL Offload	SSL Acceleration	Persistence with SSL	Intrusion Prevention

Complete ADC Solution