Pulse Secure®

WHITE PAPER

Securing DOD Web Applications with a Scalable Web Application Firewall Architecture

Introduction

In the Department of Defense (DOD) dramatic growth of web applications, and the move to cloud technologies, has been accompanied by the rise of new security threats. Advanced persistent threats (APT) and insider threats are the two most acute security threats. These threats not only cause significant damage to national security, but also impact organizational momentum and have enormous cleanup costs. For example, the widely-publicized US Navy Iran APT incident impacted operations for over four months following a SQL injection attack, and required around $10m to clean up.[1]

When we consider the massive scale and breadth of DOD web applications, dealing with the web application-layer attack surface in an efficient, effective, and scalable manner is daunting. There is growing recognition within the DOD that better data and application security is required to adequately manage these new risks. DOD IT organizations also need to adapt to new challenges, such as:

- Security of distributed applications across the DOD Information Network (DODIN)
- Integration with cloud technologies
- Extreme scale
- Centralized control with delegated administration
- Custom security policies for multiple applications with differing risk profiles
- Variance in contractor secure coding practices
- Rapidly changing threats and application attacks
- Interworking with existing security architectures and policies

With program offices under pressure to meet deadlines, operate on tighter budgets and deal with lowest priced technically acceptable (LPTA) acquisition policy the risk of application security vulnerabilities is very high.

Software-based distributed web application firewalls are a game changing approach that arms DOD with the means to get ahead of these new, rapidly evolving risks.

Securing DoD Web Applications

Securing today's complex mission-critical applications requires a combination of both network and host-based tools. Pulse Secure® Virtual Web App Firewall (Pulse Secure vWAF) has a distributed, three tier, massively scalable security capability, which can be deployed across the full range of hosting environments (physical, virtual and cloud). In addition, Pulse Secure offers virtual and cloud optimized licensing options that allow you to manage costs based on the volume of data protected (gigabits per second) rather than traditional per-instance or per-appliance licensing models. In high-security environments where server certificate sharing exposes an unacceptable risk, Pulse Secure vWAF can be installed as a plug-in on popular web server platforms, giving massive scalability for distributed applications. Designed for massive scalability, security policies can be managed centrally, while enforcement is distributed across your applications.

Flexible Deployment

Pulse Secure Virtual supports the full range of deployment options enabling you to choose the best fit for your architecture and application risk profile. Pulse Secure vWAF can be deployed physically on the web server, physically as a network appliance or virtually as a virtual appliance in a customer data center or cloud provider—or even as an integrated package with Pulse Secure Virtual Traffic Manager for enhanced security and control of complex applications.

Policy Version History and Change Control

Pulse Secure vWAF stores policies and change history in an integrated version control system, allowing roll-back to previous rulesets with just one click, making it easy to revert to known policy sets if a new ruleset interferes with an application.

Dual Active Protect/ Detect Policies

Continuous improvement and iteration of policies is made easy by running two active parallel rulesets: one ruleset runs in Protect mode for blocking malicious traffic and protecting your application; the other ruleset runs in Detect mode for evaluating new policies.

[1] http://online.wsj.com/news/articles/SB10001424052702304732804579423611224344876

No Need to Store Copies of SSL Keys

You can choose whether Pulse Secure vWAF manages your SSL termination: in a distributed architecture, it will not store copies of your SSL keys, and needs no changes to your PKI architecture, avoiding a potential target for man-in-the-middle attacks.

Pure Software Form Factor

Pulse Secure vWAF gives you the choice of either a virtual appliance or a pure software form factor. With true portability across clouds, and flexible integration with orchestration platforms, you can run on Windows and Linux systems, secure clouds, or traditional data centers.

Powerful Scripting Engine

The Python-based scripting engine allows your security professionals to change and secure every part of an application's HTTP request/response flow. Complex flaws in applications like broken authentication models can be fixed inside Pulse Secure vWAF, without needing to change the application itself.

Comprehensive User/ Rights Management

Users and Groups can be assigned access rights to view or modify specific elements of the security configuration, so that policies can be restricted to closely-defined security groups.

Start Small, But Scale to Meet Demand

Pulse Secure vWAF can be scaled to meet transactions demands with a simple license key upgrade. Start small today, without over-provisioning, and then incrementally add more processing power and license keys as demand grows for transactions and throughput.

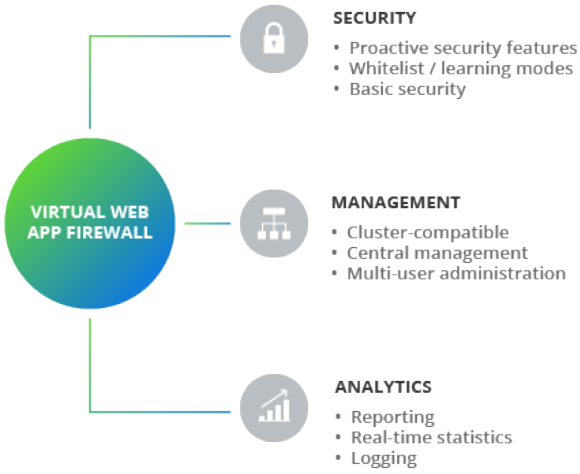

Figure 1: Defense in Depth: Pulse Secure Virtual Web Application Firewall provides multiple layers of protection.

Unmatched Scale and Performance

Pulse Secure vWAF provides Layer 7 application security at massive scale and resilience for the largest applications, by allowing your security infrastructure to scale up and cluster out. And because the software can run on commodity compute servers and cloud platforms, you are not restricted to a fixed appliance form factor.

Application Security—a Critical Component to Risk Management

Every year, thousands of new vulnerabilities are reported in web applications, and it is clear that it is becoming a major challenge to keep pace with the growing threat. With so many new vulnerabilities reported each year, many organizations find it difficult to secure, maintain and enhance applications due to the complexities of security analysis and testing.

Applications vendors themselves have a similar problem: not only do they need to remain alert to identify new vulnerabilities supported in their own applications; they need to determine whether custom integration projects might also compromise security of their own or other applications. Identifying and solving a security loophole can take significant time and resources in order to bring a patch or solution to market.

While some application vulnerabilities can be solved quickly with a patch in a third-party component or OS module, it is not unusual for logic flaws or data leaks to take months to solve in production systems—leaving doors open for hackers to enter through. Some off-the-shelf applications can go unpatched for a year or more, depending on the priorities of the application vendors, and the perception of risk.

Pulse Secure Virtual Web Application Firewall is a massively scalable solution for application-level security, both for off-the-shelf solutions, and complex custom applications including third-party frameworks. Pulse Secure vWAF can apply business rules to HTTP(S) traffic, inspecting and blocking attacks such as SQL injection and Cross-Site scripting,

while filtering outgoing traffic to mask personal identifiable information like social security numbers, and help maintain compliance with risk management frameworks such as NIST SP 800-53 rev 4, FedRAMP, DISA Security Technical Implementation Guides (STIG) and PCI-DSS.

Cloud Applications are Exposed to a Wider Range of Threats

Accessing cloud technologies requires a thin-client, typically a web browser, and security budgets have previously focused on network firewalls and antivirus solutions. However, more recently the primary target for attacks has shifted from the network layer to the application layer, because the operating systems and service interfaces available to cloud applications have been hardened to expose a reduced profile. As a result, it is now much easier to target the application logic or application framework than the actual server behind the hardened network perimeter. Most DOD applications are created in-house or by a contractor, but application security is rarely a core developer skill, which leads to a range of security problems throughout the application lifecycle.

Furthermore, wider adoption of cloud computing means that attack vectors continue to increase, as applications leverage external hosting and delivering infrastructure, platform and software. It is increasingly common for web applications to be built on open-source components or prebuilt web frameworks. This approach has clear benefits of interoperability and shorter development time, but patching to solve security vulnerabilities becomes more complex. A flaw in one component of code must be patched for each instance it is used throughout each application in which it is used. In Platform-as-a-Service (PaaS) scenarios, with dynamic infrastructure and application frameworks, this can become very difficult to manage or may be out of the client's direct control.

Applications developed specifically for a cloud are often very complex, designed for access speed, scalability and flexibility for third-party development through an open API. For example, Salesforce.com, Google Docs, Twitter, and Digital Common Ground Systems (DCGS) are good examples of APIs exposed to allow access from custom applications. These 'as a Service' applications are developed in two ways today: (1) by moving on-premise applications to a cloud (whether private, public or tactical), and (2) by developing and operating applications in an elastic dynamic cloud.

Applications that migrate out of the DODIN and into a public cloud infrastructure carry the risks of exposing protected software to external threats that they were not designed to combat. Common security threats include SQL injection attacks, cross-site scripting, and cross-site request forgery. These often have to be open to the general public and are a key entry point for an APT.

There are a variety of FedRAMP approved public cloud frameworks, such as Amazon Web Services, Microsoft Azure, Autonomic Resources, as well as private DOD clouds such as DISA's MilCloud. There are many security challenges involved in developing web applications in a cloud, such as parameter validation, session management, and access control, which are key "hotspots" for attackers. Developers without experience of application security are more likely to create/develop applications that have security problems.

Mitigating web application-level attacks requires a WAF. Given the complexity of today's applications, these WAFs are often aligned to specific applications and configured to mitigate the specific threats to these applications. In large scale data centers or cloud hosting services, there will be hundreds of applications, each potentially with a dedicated WAF. The costs and time to procure and manage large numbers of physical appliance-based WAFs or WAFs licensed per virtual appliance can become a significant portion of an applications budget. Virtual can be licensed based on the volume of traffic it is protecting, allowing for any number of virtual WAF appliances or on-host WAF implementations as needed without generating additional costs.

Pulse Secure vWAF provides you the most cost-effective solution to arm your application with exceptional application-layer protection methods that can keep up with an APT's attack velocity, and naturally scale with elasticity in your cloud.

OWASP Top Ten—Managing Risks in Web Applications

OWASP, the Open Web Application Security Project, is the leading open source community group in web application security, and regularly publishes an annual "Top Ten" report showing the most common security challenges to web applications, ranging from business policy to application vulnerabilities.

For example, in their report of July 2013, OWASP noted that "Injection Flaws" had risen to the highest risk: because of the risk of high-value data leakage, this kind of vulnerability can give rapid access to sensitive enterprise information, hence the need for a solution which is able to validate inputs against business policies, and screen outgoing data for suspected data leakage, in a way which is independent of the underlying application architecture. In the US Navy's Iran incident, a SQL injection was the specific method initially used to compromise the environment, an attack vector easily addressed by a web application firewall.

All of these risks can compromise data and application security, but each year new tools and techniques are used by would-be attackers to identify new gaps and loopholes that could be exploited, so a solution needs to be agile enough to adapt to changing risk priorities, as well as robust enough to handle large-scale attacks on high-throughput web applications processing millions of concurrent transactions across many locations with tens of thousands of servers.

Why a Traditional Web Application Firewall Is Not Enough

Previously, web application firewalls (WAF) and other security solutions were limited to hardware appliances or virtual machines, which created a serious bottleneck for cloud environments and massively scalable DOD core datacenters. These legacy WAFs were often aligned to specific applications due to concerns about the risks associated with delegated administration within a single WAF appliance. As a result, cloud service providers or DOD datacenters were often forced to have racks full of dedicated WAF appliances—one per client or application—that take up unnecessary space and resources. The high costs associated with these dedicated WAF services often become a barrier to the efficiency of a fully virtualized cloud environments driving higher hosting costs and decreased usage.

In a cloud, the infrastructure and the services are shared between clients, meaning many organizations hardware. Each of these cloud operator clients adds a unique layer of policy settings, use-cases and administrative enforcement requirements. For the cloud or DOD datacenter, security quickly becomes very complex. The average provider may have hundreds or thousands of clients relying on its service, each with varied policy settings for individual groups within the client. The service provider now has to manage application filter settings for each and every client application. When you consider robust end-to-end encryption and authentication policies in DOD successfully injecting a web application firewall appliance as a man-in-the-middle to inspect traffic is difficult to impossible.

In an ideal world, applications would be designed from the ground up to meet the challenges of virtual and cloud architectures. However, it is critically important to implement affordable, scalable solutions to mitigate impacts of application based attacks. Pulse Secure vWAF is the industry leader in delivering to these complex, virtual and cloud based hosting environments.

Figure 2: OWASP.org "Top Ten" Report 2013 shows how attack vectors change over time.

Defining the Distributed Web Application Firewall (dWAF) for Cloud Security

A solution for web application security in DOD has to be massively scalable, flexible, adaptable and efficient to manage:

Massive Scalability

A distributed WAF (dWAF) must be able to dynamically scale across CPU, computer, server rack and location boundaries, customized to the demands of individual customers, without being limited to discrete hardware form factors. Resource consumption of this dWAF must be minimal, and scale in line with detection / prevention throughput rather than consuming increasingly high levels of dedicated CPU resources. Clouds come in all sizes and shapes, and the WAF needs to be flexible as well. Scale from mini tactical footprints forward deployed to massive public clouds is critical.

Cross Platform Portability

The dWAF must be able to live in a wide variety of components to be effective, without adding undue complexity for DOD. Today's cloud service providers, DOD datacenters and tactical sites use a variety of traditional and virtual technologies to operate their web applications, so the ideal dWAF should accommodate a mixed environment and be available as a virtual software appliance, a plug-in, SaaS or be able to integrate with existing bare-metal hardware. Flexibility with minimal integration with the existing environment is essential for broad adoption and low total cost of ownership.

Distributed and Delegated Management

A web-based user interface must allow DOD users to easily administrate their applications. Configuration should be based on the applications under protection, not defined by a singular host, allowing far more granular settings for each application. Ruleset configuration must be simple, and supported by setup wizards and DISA STIG compliant templates so best practices, as well as compliance, are inherited immediately. Statistics, logging, and reporting must to be intuitive, easy-to-use and must integrate seamlessly into other systems. Most importantly for a dWAF, robust multi-administrator privileges must be available and flexible enough to effectively manage diverse policy enforcement schemes required by global organizations. DOD must deliver a set of core protection templates, with the ability to delegate policy for individual clients and applications to operate and refine. This approach leverages the velocity of a DEVOPS culture while ensuring central policy visibility and collaboration.

Detection and Protection

Foundation security using black, white and grey listings for application requests and responses must be possible. To make sure pre-set policy enforcements are not activated or deactivated without approval from an administrator, deployment and policy refinement through establishing rulesets must be possible in a shadow monitoring or detection-only mode. Once the shadow monitoring ruleset is stable, only then should it be allowed to deploy in an enforcement mode on the dWAF. This allows complete transparency for the administrator into the real-world effect of this ruleset, while at the same time allowing layered rulesets to be tested without compromising existing policy enforcement. Avoiding false positives and validation of changes in rulesets are essential for a real-world, usable dWAF in a hybrid cloud environment.

Automated learning and ruleset suggestions based on intelligent algorithms or recommendations from a static source code analyzer or web vulnerability scanner are also desirable from a manageability view. Again, this only holds true if the administrator retains full control over activation/ deactivation of each ruleset. Without this level of control, legitimate traffic may become blocked and policy settings would become compromised, which in turn will negatively impact the adoption of a web application firewall.

Application Shielding

Proactive security functions are highly recommended to reinforce any application in a cloud. Detection is simply not enough for today's web application security. Features like transparent secure session management, URL encryption and form-field virtualization provide strong deterrence to attack, while saving application development and deployment time. These features are effective because session management, URL encryption and form-field virtualization are done at the dWAF level and not in the application itself.

Integration with End-to-End Encryption and Authentication

DOD Public Key Infrastructure (PKI) policies and new JIE standards are rapidly moving towards private key storage in hardware security modules, coupled with end-to-end authentication and authorization using mutually authenticated TLS v1.2 with Message Authentication Codes (MAC). This prevents the private key from being shared with a WAF appliance, as well as preventing a man-in-the-middle from decrypting the traffic. This prevents a traditional WAF appliance from being able to decrypt the HTTPS traffic, keeping it from being able to conduct its primary role. As DOD layers in content-based and attribute-based authorization via Security Assertion Markup Language (SAML) injecting a WAF appliance as a man-in-the-middle is even more difficult. Where as a software-based dWAF plugin can natively running in the web server, after the encryption and authentication have already taken place, provides an easy to deploy alternative. This software-based dWAF plugin requires no changes to encryption or authentication, no private key sharing, no man-in-the-middle, no manpower to rack-and-stack hardware and low cost automation-based deployments make WAF adoption in DOD faster, lower risk, more secure, lower cost and centrally administered with delegated local control.

A Cost-saving GOTS-like Security Alternative

Application contracts often include complex and costly application security development requirements. A high degree of variation in quality results from thousands of different developers interpreting and implementing the requirements differently. Not to mention the additional develop effort and costs that add no real value to ultimate mission. If the DOD furnished the dWAF plug-in to the developer to integrate as the front-end application security asset significant costs savings would be captured, along with a repeatable high-quality security capability with no variation between applications. This would allow application security teams to ensure even early prototypes are extremely secure, easily meet DISA Application Security STIG requirements and are central administered and monitored at the enterprise-level.

Integration with Existing Technology

Avoiding vendor-lock-in is a best-practice for both networking and application security. Any technology that is added to an infrastructure, platform or application itself must connect as seamlessly as possible with existing technology and DOD processes. Security technologies must be layered to create the best possible protection, so a dWAF must tightly integrate with security incident event management systems (SIEMs), ICAP-compliant malware detection systems, threat databases, automated mitigation systems and have APIs for custom integration efforts. It must also be open and extensible so it may be modified as the DOD mission requires.

Pulse Secure WAF Solutions

Pulse Secure Virtual Web Application Firewall is a pure software web application firewall designed to support these best practices. Due to its modular construction, it can be deployed very easily in a cloud-computing environment, making it a scalable solution for application-level security. Pulse Secure vWAF can apply business rules to HTTP(S) traffic, inspecting and blocking attacks such as SQL injection and Cross-Site scripting, while filtering outgoing traffic to mask sensitive data, and help meet compliance mandates for NIST RMF, DISA STIG, PCI-DSS and HIPAA. The product consists of three scalable components:

1. Enforcer
2. Decider
3. Administration Interface

These three elements can be configured either as a pre-packaged WAF solution, such as an add-on module for Pulse Secure Virtual Traffic Manager to manage a cluster of applications, or as a fully distributed solution using plugins across hundreds of web servers and many sites for maximum scalability and performance. The same distributed management interface can be used to protect both types of deployment, or even in a shared services environment.

1. Enforcer—Policy Enforcement Point

The Enforcer is a small plug-in, which can be installed into any kind of device. A device can be a web server or proxy (Apache, MS IIS, etc), a network firewall (like MS ISA) or integrated into a software load balancer or ADC (such as Pulse Secure Virtual Traffic Manager). The Enforcer sends request and response data to a component called the Decider and also modifies requests and responses if needed. The Enforcer is an adapter for the dWAF to get the data it needs to enforce the policy.

2. Decider—Policy Decision Point

The policy engine checks the data from the Enforcer module and decides what to do with the request/response. The Decider's unique architecture allows it to scale from one to many CPU cores and is also capable of running on multiple machines to scale horizontally. The Decider is the compute-intensive part of the solution, and the workload on the Decider depends on the load of the web infrastructure behind it. The more traffic generated by the end users, the more CPU resources are used in the Decider. In small environments the Enforcer and Decider may be deployed on the same host.

3. Administration Interface—Policy Management Service

The administration system can be deployed decentralized or as a single server. In a cluster installation every cluster node can be used to administrate applications and their policies. This decentralized architecture is very robust against node failures, and allows a large group of web application security administrators across many sites to work on individual application policies. In addition, it provides verbose central monitoring and alerting functions. Its open nature allows for it to be easily integrated with other security solutions.

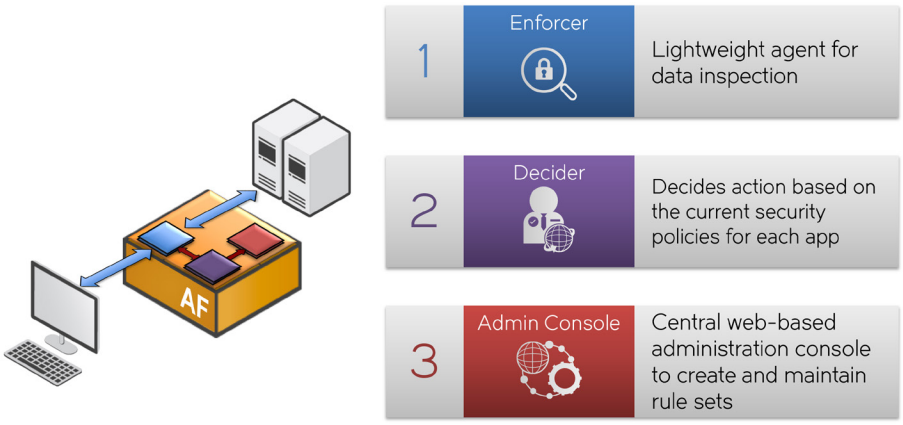

Figure 3: Pulse Secure vWAF consists of three scalable components.

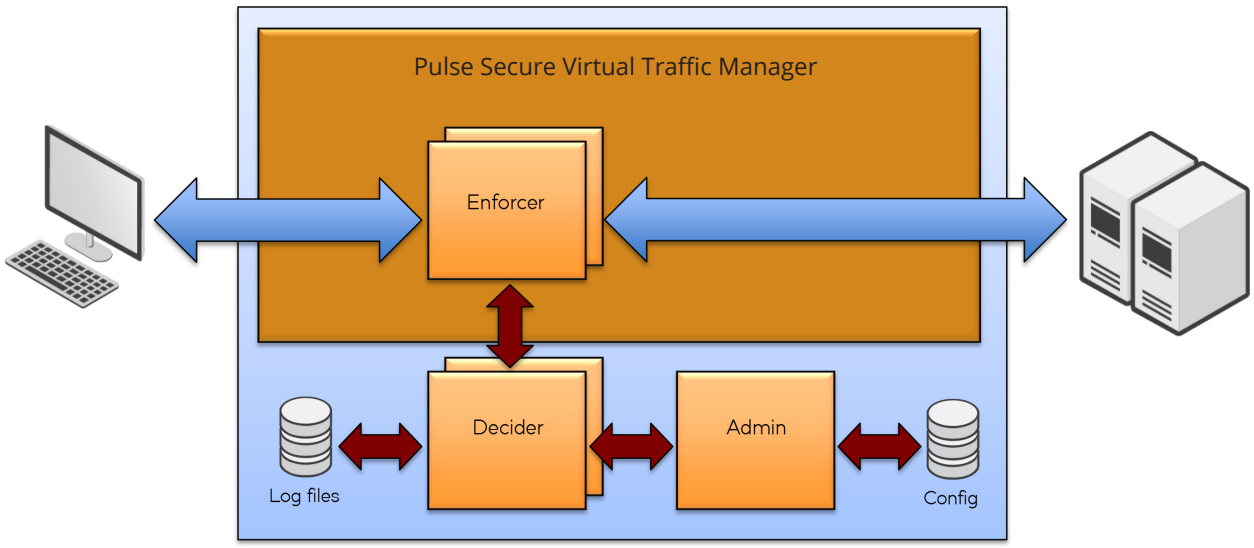

Figure 4: Pulse Secure vWAF as an add-on to Pulse Secure Virtual Traffic Manager.

Stand-alone Implementation

In stand-alone environments, Pulse Secure vWAF software is licensed as an add-on for Virtual Traffic Manager, and can be deployed either on a server appliance, or on a VM in virtual or cloud infrastructure. Enforcers and Deciders are co-resident inside the Traffic Manager package, administered as a single platform. The Admin GUI is accessed via the standard Pulse Secure Virtual Traffic Manager console.

PaaS and SaaS Implementation

Software as a Service (SaaS) offers user access to virtualized applications through a thin-client, usually any standard web browser. The benefit for users is access to software without any of the headaches of owning the programs—scaling and resources, and patching and upgrades are all taken care of. SaaS applications are usually billed on a per-user/per-use basis.

Platform as a Service (PaaS) provides organizations with virtual databases, storage and programming models to build custom applications. AWS, Microsoft Azure and DISA MilCloud are examples of PaaS environments available to DOD. This type of service provides scalable resources behind the platform and allows customers to grow throughout the lifetime of the application. It is an effective solution for organizations of all sizes, and is often billed on a usage model.

In a shared cloud (PaaS and SaaS) environment, many Enforcer plug-ins deployed in software load balancers or natively in web servers could communicate with an elastic Decider cluster. If the first virtual machine of this Decider service reaches 80 percent of the available CPU resources, a new virtual machine on a different instance of a cloud will automatically be provisioned, started and added to the Decider cluster. If the cluster-wide CPU usage of the Decider service drops below 40 percent, the Decider instances will automatically be removed from the cluster to release resources back to the cloud.

This ensures a DOD application, or cloud, has a truly elastic application security architecture with unlimited scale, zero-touch elasticity, and zero physical footprint. If an enemy attacks with a massive Distributed Denial of Service (DDoS) attack the dWAF will scale elastically with the web server infrastructure, requiring no physical intervention to achieve the scale benefits. This unique deployment architecture supports the JIE tenants of efficient, effective and secure—at lower total costs to legacy physical or virtual Application Firewall appliances.

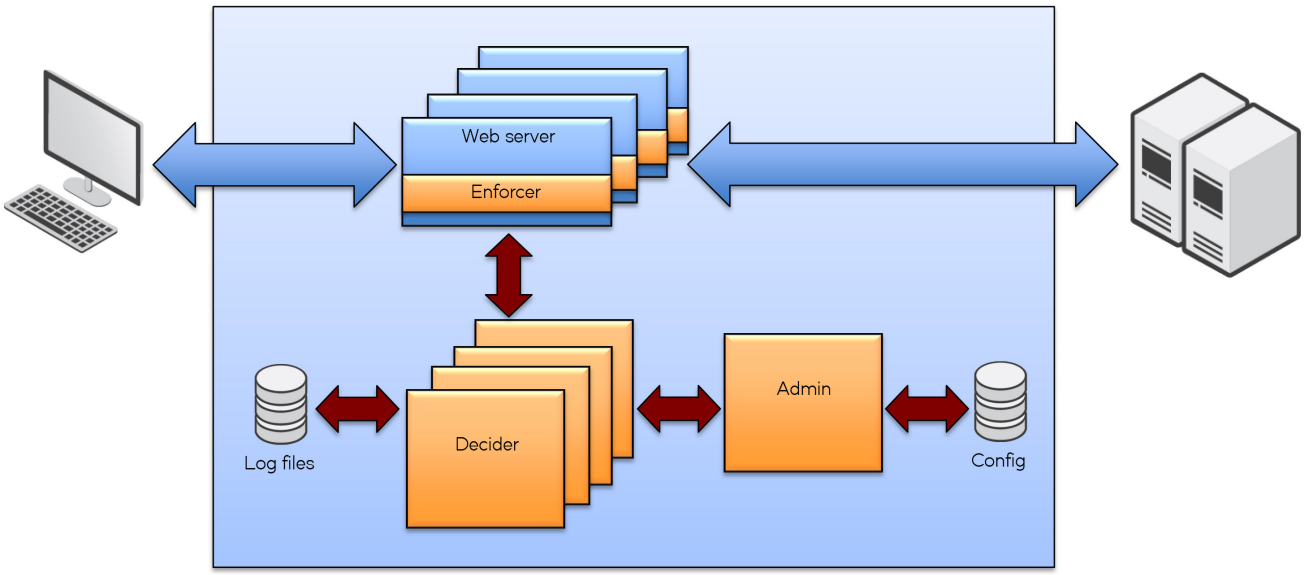

Figure 5: Pulse Secure vWAF in a distributed deployment for a single application.

IaaS Implementation

Infrastructure as a service (IaaS) allows access to large-scale resources to build and manage a complete virtual network. Organizations can commission and decommission virtual resources depending on their need, as application workloads change over time. Customers select from a range of infrastructure components and service level agreements to match their application requirements, matching their infrastructure usage costs to their service requirements.

Pulse Secure vWAF can be deployed similarly on an IaaS in a cloud. Each organization can have a private instance of all three modules, which are not shared between organizations, or even applications if desired. As before, many instances can be clustered together to scale up to meet DOD-scale traffic demands. Cloud providers can also provide automatic scaling within the cloud infrastructure to follow the lifecycle of individual applications.

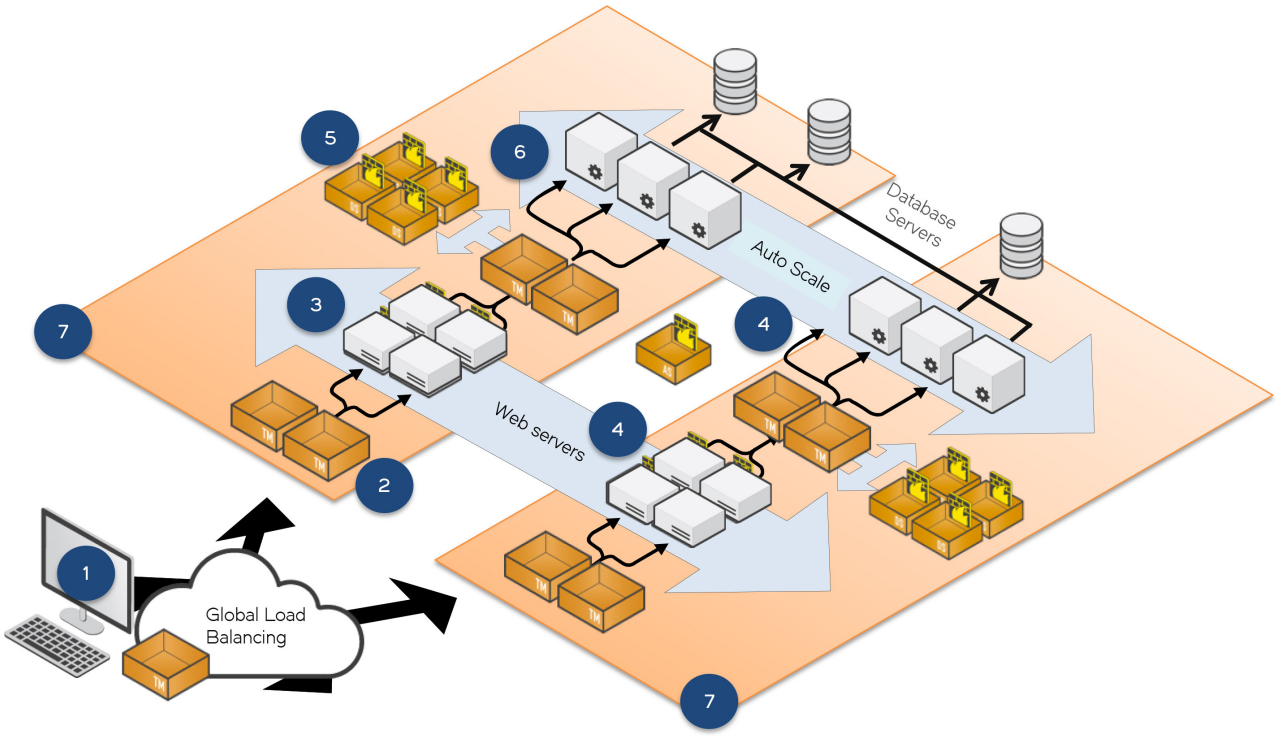

Very Large Scale Applications

For the very largest applications deployed in public clouds, a distributed application firewall needs to scale dynamically to the very highest traffic levels, and adapt to meet the changing transaction demands of public-facing applications. For example, Pulse Secure vWAF can be deployed in a distributed cloud architecture from a FedRAMP cloud service provider. Since Pulse Secure vWAF is pure software it clearly ports into any FedRAMP certified cloud service provider, or multiple hybrid clouds, with no re-tooling. Providing DOD true cloud portability, freeing them from provider lock-in and enabling seemly multi-cloud burstability without application-layer security compromises. Cloud-portability isn't available with physical or traditional WAF appliance approaches.

The components of a highly scalable web application firewall

1. **Global Load Balancing:** In this example, Traffic Manager software serves incoming DNS requests using Global Load Balancing, and routes network traffic to the closest cloud data center.
2. **HTTP Distribution:** Traffic Manager handles incoming HTTP requests, which are distributed across multiple web servers in different availability zones. In addition to fine-grained health checks in the application layer, Pulse Secure software can provide SSL Offload, Compression, Caching, and Layer 7 based request routing.
3. **Application Firewall—Enforcer:** In this example, the Pulse Secure vWAF Enforcers are installed as a plug-in on the web servers. The Enforcer traps HTTP requests and response data and sends it to the cluster of Pulse Secure vWAF Deciders for security processing.
4. **Auto-scaling of Web Servers:** Traffic Manager manages the automatic spin up and spin down of web servers as required to service your application effectively using Traffic Manager's Auto-scaling capability.
5. **Application Firewall—Decider:** A scalable group of Pulse Secure vWAF Deciders processes the HTTP requests and responses to ensure compliance with security policy. Pulse Secure vWAF provides a hybrid security policy model combining traffic signatures and positive security postures to prevent HTTP based security attacks.
6. **Load-Balancing and Health-Checks:** Traffic Manager provides granular application health checking and layer 7 load balancing into the application tier.
7. **Cross-Zone Availability:** Traffic Manager works seamlessly across multiple Availability Zones, multiple regions or across multiple clouds, for the ultimate in cloud high availability.

Discover Pulse Secure vADC with a Free Trial

It's easy to test all of the capabilities available in the Pulse Secure vADC product family. Download the Pulse Secure Virtual Traffic Manager Developer Edition today to find out how to realize greater ROI from data center consolidation and transformation programs.

The Pulse Secure vADC Developer Edition, available either as pure software, or as a virtual appliance, makes the complete ADC technology platform available to every application developer in your organization, enabling them to develop applications faster, test them in a production-identical environment, and bring them to market more quickly.

Find out More

To find out more about Pulse Secure Virtual Web Application Firewall, or to arrange a demonstration or product evaluation, please visit http://www.pulsesecure.net/vadc/vwaf

Corporate and Sales Headquarters
Pulse Secure LLC
2700 Zanker Rd. Suite 200
San Jose, CA 95134
www.pulsesecure.net

Pulse Secure®