Cloud WAF in Hybrid Environments

Whitepaper

SHARE THIS WHITEPAPER

Table of Contents

Why is Hybrid Cloud So Compelling?

Historically, applications are the reason why businesses turned to automated systems and new market penetrating ideas. These ‘apps’, the cornerstone in which all business gets accomplished, are migrating to non-premise or off-premise cloud providers. If you believe the likes of IBM and others who declared at the InterConnect Conference, “cloud” can no longer be categorized as private or public – but has shifted to a hybrid state.

To remain competitive and relevant, every business must transform and adapt. This paper examines three major reasons behind the idea of “cloud” being synonymous with “hybrid”.

Reason #1 Most Companies Will Retain Some Internal Application Delivery Infrastructure

Because of legacy business processes; legal, compliance, or resiliency reasons, complications from management and loss of real-time visibility, most businesses will not be able to completely eliminate IT infrastructure and rely solely on the cloud. Most businesses simply aren’t positioned to move all legacy applications to the cloud because of reasons such as:

- Applications which cannot operate in the cloud (e.g. manufacturing or Programmable Logic Device applications)
- Application inherently needs to be close to the customer
- Internet access is not yet resilient enough or robust enough for cloud service delivery models (e.g. some regions around the world where bandwidth or QoS is not yet world class)
- Compliance or legal rules will dictate that data must remain within a certain geography or domiciled in
- Internally accessed applications (not publically accessed)
- Data-intensive applications which access sensitive information (for applications that do not access sensitive data cloud storage and offer cost effective solutions, and for non-data intensive applications in which the overhead of accessing premise-based stored data is not significant)

Many providers define a hybrid cloud as the combination of any of the following models: premise-based tools, private clouds, and public clouds. However, hybrid is no longer relevant where IT service delivery is operated from and who owns the underlying infrastructure — the choreography and orchestration of these services defines it as hybrid.

Questions to Ask Yourself When Migrating Applications to the Cloud

1. Do I have both internally accessed application and publically accessed ones?
2. Will I have a transition period where I have application both on premise and in the cloud?
3. Are some applications planned to remain on premise (mission-critical, require access to sensitive internally stored data) for the longer term?
4. How will I secure both my premise-based and cloud-based applications?
5. How will I secure my apps during the transition to the cloud?
6. Will I migrate apps to a single cloud infrastructure or will I adopt a multi-vendor strategy?
7. How will I get a controlled and managed security policy for my applications in the various relevant scenarios listed above? How will I have a centralized reporting and security policy visibility?
8. Do I have visibility into the SDLC and the web app release process? Do I know what was modified? Do I know what the potential impact on the security policy is?
9. Do I have visibility into the different web frameworks and technologies used to develop and run your web apps? Am I aware of the various vulnerabilities of each such platform and technology?

> *A hybrid cloud approach does not require a complete migration of traditional IT infrastructure to a public or private cloud.*

Reason #2: Dedicated Infrastructures are a Luxury which Make Most Companies Uncompetitive vis-à-vis Hybrid Competitors

The verdict is in about the merits of virtualization and cloud in that it unleashes hidden efficiencies which were often elusive to classic data centers in the past. For example, when building a legacy data center, typical IT infrastructure must be designed and built for peak utilization. As demonstrated through server virtualization projects, the average utilization of physical servers tends to be no more than 12 to 18 percent. Businesses have started to virtualize physical servers to potentially increase this utilization rate to 30 to 40 percent or more.

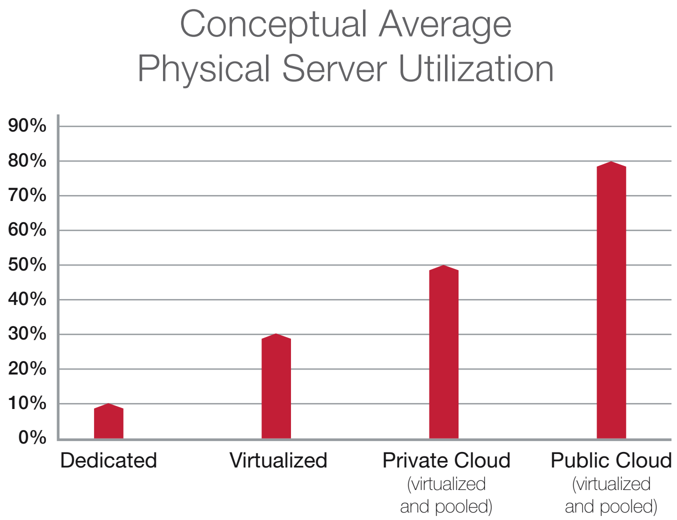

Figure 1: Physical Server Utilization

Cloud, in many ways was another attempt to build upon virtualization and increase utilization and efficiency. Cloud was a group of virtualized servers in resource pools with self-service portals that allow new server workloads to easily be created, modified and deleted. At its core, cloud was designed to take the complexity of virtualization away from the end user and fully enable self-provisioning and speed to service delivery. For example, to drive physical servers up 80 percent efficiency, multi-tenancy is required — which often drives the need for clouds with a mix of businesses that have varying peak demand time periods. Many consider this to be an oversubscription of servers for individual customers, but this model can be successful since customers will have peak needs at different times.

Additionally, these cloud systems need to work seamlessly with legacy applications and business systems often integrated technically and procedurally, leading to hybridization of the service delivery model.

Reason #3: Information Security & Compliance

Cloud security remains elusive and is considered standard practice that although information technology may be outsourced or provided by a third-party provider, the overall governance of security cannot be outsourced. In other words, security always remains the responsibility of the contracting company and has always been a 'hybrid' solution from inception.

From the very inception of cloud delivery models, security has been the main concern of inadequacies and other issues. In broad terms, cloud has given a rise to the need to provide information security in three categories.

Cloud Security Requirement Application Security: Web Based Attacks

The Open Web Application Security Project (OWASP) Top 10 provides the common threat classification in the web application security space. This category of attacks is generally not covered by traditional security technologies such as network firewalls and intrusion detection systems (IDS) which are focused on network-oriented threats and attack vectors. Among other types of security issues, companies needed to make certain threats:

- SQL Injections
- XSS (Cross-Site Scripting)
- CSRF (Cross-Site Request Forgery)
- Session Management attacks

However, there are many other attack vectors targeting web applications beyond the 10 listed on OWASP. While the severity of these attacks may be application dependent, in some cases these attacks may be harmful. Some of these additional attack vectors are listed on WASC threat classification.

- Brute Force
- Predictable Resource Locations
- Path Traversal
- HTTP Response Splitting
- Abuse of Functionality

Questions to Ask a Cloud Security Provider

1. What type of a solution can you offer to secure my premise and cloud based apps?
2. What type of a solution can you to offer secure internally and publically accessed apps?
3. How will you avoid generating false positives when the application changes?
4. What types of attacks are included in your protection SLA?
5. How does your solution secure against various web attacks?
   a. Upload of malicious content to the application
   b. HTTP protocol manipulation such as HTTP response splitting
   c. Parameter pollution
   d. Session management attacks and cookie poisoning
6. How many man hours will I (the customer) need to invest per month on maintaining the policy to properly secure my apps and not to generate false positives?
7. Do you offer Professional Services for policy tuning and configuration? At what additional cost?

In the end, most companies who are migrating applications from more internal deployments to a data center found fewer options and features in which to secure applications in the cloud.

Cloud Security Requirement: Compliance

Various compliance requirements caused a delay in cloud adoption as these requirements were often unknown. Moreover, the compliance of cloud companies need to be in conjunction with efforts and claims made by client companies within their premise-based or private cloud infrastructures. Some examples of huge compliance requirements are as follows:

- Payment Card Industry Data Security Standard (PCI-DSS)
- Health Insurance Portability & Accountability Act (HIPAA)
- Sarbanes-Oxley Act
- Patriot Act
- Various National & Regional Privacy Laws

In the end, standard frameworks have been largely adopted to address many of the standard approaches in information security and tests against those frameworks. These frameworks include SAE16 (replaces SAS70), ISO 27001 and HIPAA HiTrust.

Cloud Security Requirement: Network and Application Availability Concerns

The task of keeping a business up and available while orchestrating various cloud delivered service models is challenging. Similar to the change of ‘just-in-time’ inventory in the manufacturing models, the cloud, with all of its benefits of cost and agility ushered in a whole new era of requiring high degree of up-time. Moreover, up-time is not simply just a contract issue based on committed SLAs with punishment remedies for violations. Up-time is becoming core to a business’ survival and operational processes and protections need to be assured. The issue of up-time is multi-faceted and needs to cover numerous categories of security threats such as:

- Volumetric attacks (e.g. DDoS) vs. non-volumetric (e.g. SQL injections leading to outage)

- Bots vs. human. It used to be the domain of premise-based technologies to determine the difference between legitimate humans and illegitimate humans (e.g. hackers). Now this area has exploded and includes differentiating humans (legitimate and illegitimate) from bots (or automated programs) which can also be legitimate and illegitimate

- Network Resource Exhaustion
  - Pipe saturation – often called DDoS typically has benefits in moving to the cloud as cloud companies generally have bigger bandwidth and larger scale. However, this does not make the cloud company immune to Internet pipe saturation and often increases risks of multi-tenanted attacks having a domino effect on all hosted companies.

  - Data center infrastructure – Within the data center itself, there are many opportunities to exhaust the network devices rendering transit of technology unavailable. Examples of such attacks are exceeding limits on sessions on firewalls, exceeding connection-per-second limits on applications and generally exhausting TCP stack resolution services.

  - Web Exhaustion Techniques - One of the least understood concepts is the contribution of web vulnerabilities in explaining the origin of business outage. According to Radware’s 2014-2015 Global Application & Network Security Report, 27% of all business outages are explained via exploited web application vulnerabilities. Although in theory any one of the standard OWASP categories can lead to a resiliency problem, outages at the web level revolve mainly around the following five vectors:

    - Basic HTTP floods (GET/POST) – either encrypted or otherwise
    - CDN: HTTP dynamic floods
    - Brute Force Attacks
    - Input Validation Attacks
    - SQL / JSON / XML / LDAP Injections

  Cloud service providers often don’t have detection tools or mitigation wherewithal to detect, block or report on these threats.

- Multi-vector attack campaigns would be bad enough if the attacks listed above were all an organization had to deal with, however these attacks are almost never unaccompanied. These attacks are known today as ‘multi-vectored’ attacks. Multi-vectored attacks come in many forms, however according to Radware’s 2014-2015 Global Application & Network Security Report, seldom have single vectors and on average have 7 distinct vectors within often containing at least one within each of the following categories:

- SYN Flood (volume attack)
- UDP Flood (volume attack)
- Crafted application attacks (e.g. injection or brute force attack)
- Encrypted attacks

Multi-vector attacks are an effective strategy for attackers because they:

- Complicate detection as they require a great amount of processing power on behalf of the victim to pick out and compare the attacking vectors.
- Provide “smoke-screens” for more nefarious, often piercing attacks.
- Complicate forensics and naturally triages most ‘soft’ spot in defense infrastructure.

Current Challenges with Cloud WAF in Hybrid Environments

There are three megatrends within technology which are contributing to decay in current information security deployments and technologies. These trends may be colliding and make the work of an information security practitioner very difficult.

The first trend is cloud adoption itself. Second, Internet of Things (IoT) is changing the number of end points and the type of access requirements immeasurably. Lastly, there is a revolution in the way routing is being conducted with new protocols and standards which are laying waste to decades of assumptions.

Current Disruptive Trends

- Cloud Migration
- Internet of Things (IoT)
- Virtualization of the Network Features (NFV/SDN)

These technology shifts have vastly changed business leaders’ expectations of IT and disrupted many of the security models we’ve come to expect. These changes have resulted in the following complications for security professionals:

- Different operating environments (e.g. premise, cloud, hosting, managed, collocated, etc.)
- Loss of visibility to the overall business picture
- High expectations: Businesses are now looking for IT to respond in hours or even minutes compared to what used to be days or weeks. Need to bring a new application online? The expectation is that it should be done today, not in a couple of weeks. Security professionals are often forced to make a tradeoff between putting in all the controls necessary to protect the business, and letting the business run at the pace it wants with security controls that may have weaknesses or gaps.
- Ability to detect threats with high quality in one location and react to these revelations in all operating environments in real time.
- Craft the right security rules in one location and automate these policies throughout the entire IT and application infrastructure regardless if internally owned or operated.
- Orchestrate changes to the affected systems quickly and universally. Making changes manually to all the necessary devices can take some time and be prone to mistakes.

Until recently, no single web application firewall technology exists which addresses these problems above. Solutions offered by security vendors do not include a web application firewall that covers both on-premise and cloud protection. This causes the enterprise to integrate two different vendor solutions. Issues that arise from managing different vendors, including roadmap integration, management and support process, and blind spots can result in a security gap in protection coverage and quality.

This lack of integration between on-premise and cloud protection leads to limited visibility in the attacks and attackers on the network. Organizations cannot differentiate attacks that occur in the cloud from attacks on-premise. Was it the same vulnerability? Was it the same perpetrator in both attacks? These questions simply cannot be answered because the quality of detection is limited. Organizations need to be able to mitigate a security problem on-premise and in the cloud.

Securing applications on-premise, in the cloud and during the transition period (from the premise to the cloud) requires a hybrid solution that allows simple policy migration (from the premise to the cloud) to seamlessly migrate without exposing the newly migrated applications to web attacks.

Radware Cloud WAF

Radware's Cloud WAF Service provides a fully managed and always-on, cloud-based web application firewall service. Based on Radware's ICSA Labs certified market-leading web application firewall, Cloud WAF service provides full coverage from all OWASP top-10 attacks. The service implements both negative and positive security models by utilizing its unique ability to automatically adapt to the continuously changing threat landscape and on-line assets.

It's the industry's first hybrid-based cloud WAF service that integrates with Radware's on- premise devices to provide comprehensive coverage. The service provides full and unparalleled protection from web application-based attacks and is based on Radware's Attack Mitigation solution that is comprised of a web application firewall, perimeter attack mitigation device and cloud scrubbing service. These technologies operate in a distributed and scalable cloud-based configuration and provide unified protection with no security gaps between on-premise and cloud-based devices.

Radware's Cloud WAF Service is the only solution with integrated CPE and cloud WAF technologies where the same market-leading WAF technology is used by both on-premise appliances and as a cloud service. It provides unified protection with no security gaps between on-premise and cloud-based applications and facilitates quick and easy migration of applications to the cloud.

Unmatched Web Application Protection

Built with state-of-the-art machine learning technologies, Radware's Cloud WAF Service automatically detects application domains, analyzes potential vulnerabilities, and assigns optimal protection policies. The service continuously monitors and analyzes application usage patterns, and generates granular baselines for legitimate traffic. These allow the rapid detection and mitigation of zero-day attacks, and the continuous fine-tuning of security policies due to changing application usage patterns. Radware's Device Fingerprinting technology allows the automatic IP-agnostic tracking of malicious sources trying to obscure themselves behind dynamic IP changes.

Web assets are always kept protected, even while applications constantly change and threats rapidly evolve, assuring web security is future-proof.

Fully Managed Security Service

Includes 24x7 support, proactive log review and analysis, system monitoring and auto policy generation. It gives organizations full support and service before, during and after attacks. The service is backed by Radware's Emergency Response Team (ERT) – a dedicated group of security experts that actively monitor and mitigate attacks in real time.

Preemptive Attack Intelligence

Continuously mining data in numerous Web and Darknet resources, the Radware's Cloud WAF service proactively provides attack intelligence to its customers, and automatically raises defense levels before suspected attacks are launched.

Easy, Flexible Model

The service is offered in a simple, OPEX-based model with 3 packages to choose from (Silver, Gold & Platinum). It's simple to setup with no deployment process or download/install of items needed. Once set-up, Radware's security experts have immediate access and require no customer interaction or resources to get started. Cloud-based customer portal gives visibility and insights into application security.

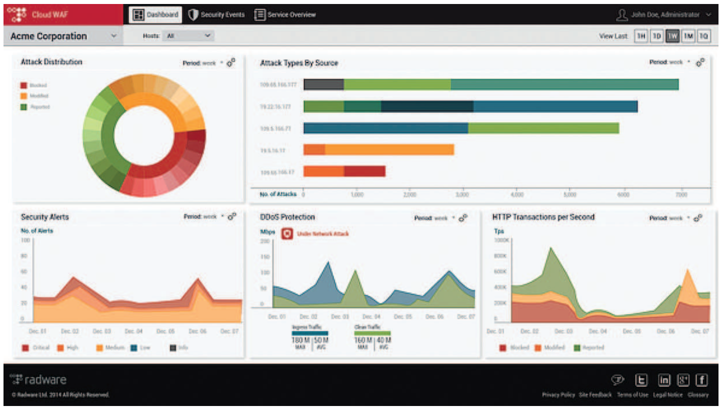

Figure 2: Radware Cloud WAF Customer Portal

Always-On DDoS Protection

With DefensePro, Radware's market-leading DDoS attack mitigation technology, the service includes anti-DDoS, NBA and IPS technologies to protect from network and application downtime, application vulnerability exploitation, malware spread, network anomalies, information theft and other emerging cyber-attacks. Radware's DDoS protection employs multiple detection and mitigation modules including adaptive behavioral analysis and challenge response technologies, in addition to signature detection to minimize false positives and impact on legitimate traffic.

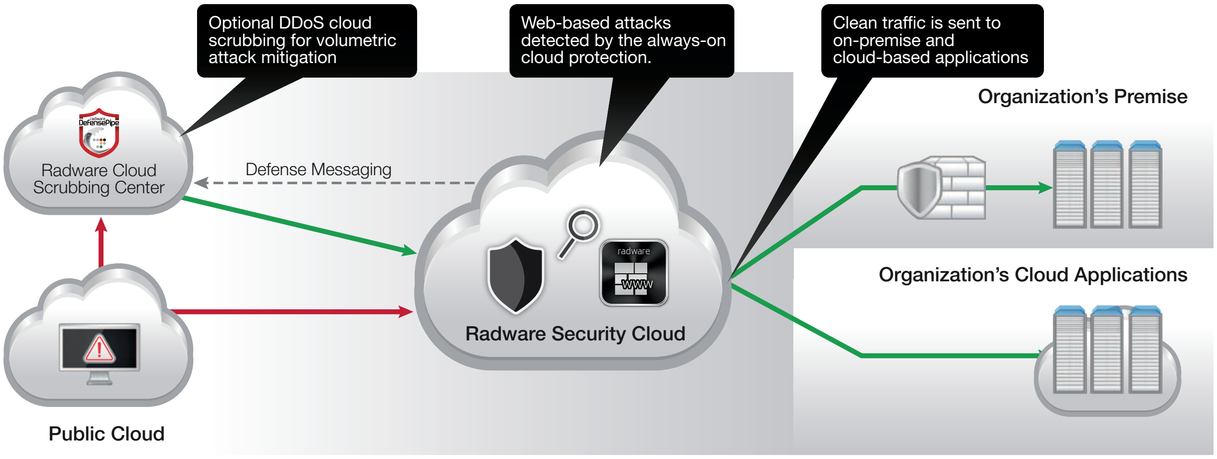
Optional DDoS cloud scrubbing for volumetric attack mitigation
Web-based attacks detected by the always-on cloud protection.
Clean traffic is sent to on-premise and cloud-based applications
Organization's Premise
DefensePipe
Radware Cloud Scrubbing Center
Defense Messaging
radware
WWW
Radware Security Cloud
Organization's Cloud Applications
Public Cloud

Figure 3: Hybrid Web Attack Mitigation in the Cloud

Radware's Cloud WAF Service offers an Enterprise-grade WAF in any easy to use, hassle free service. It provides unmatched, always-on protection from web-based attacks that includes full coverage of OWASP Top-10 threats and automatic learning and generation of policies, as well as behavioral-based DDoS attack protection.

This document is provided for information purposes only. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law. Radware specifically disclaims any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. The technologies, functionalities, services, or processes described herein are subject to change without notice.