
# Radware DefensePro
# DDoS Mitigation Technology Audit

## Prepared for Radware Ltd.

September 27, 2016

Audited by ICSA Labs
1000 Bent Creek Blvd., Suite 200
Mechanicsburg, PA 17050
www.icsalabs.com

## Table of Contents

## Executive Summary

ICSA Labs visited Radware between August 1, 2016 and August 2, 2016 to perform an on-site audit. Per Radware's request, ICSA Labs audited the internal testing methodology and procedures that Radware uses to test the effectiveness of its DefensePro DDoS mitigation technology. A follow-up audit was performed remotely using WebEx on August 23, 2016 to verify that improvements to the test plan methodology and test bed had addressed minor issues observed during the initial visit.

After careful review of the testing environment, setup, tools and test procedures used by Radware to validate the efficacy of their DefensePro DDoS Mitigation Technology, ICSA Labs concludes that the test plan adequately demonstrates the protection technology that Radware's testing was designed to highlight. In addition, the attack tools and techniques included in the test plan represent the misuse of resources, both at the IP network layer and application layer, that has been reported in "in-the-wild" attack campaigns in recent years.

## Radware DefensePro DDoS Mitigation Technology Overview

Radware's DefensePro DDoS mitigation technology uses adaptive behavioral analysis to autonomously identify and mitigate both high rate and low & slow DDoS attacks in the network and application layers.

Radware's DefensePro operates on dedicated, high-performance hardware. The hardware architecture of DefensePro includes dedicated processors for layers 3 & 4 filtering of flood attacks (i.e., its DoS Mitigation Engine or DME), a custom ASIC for layer 7 regular expression acceleration engine, and multi-purpose CPUs for stateful and statistical analysis.

## Test Environments

During the audit performed by ICSA Labs, Radware demonstrated each test scenario using one of the following three test environments:

- First test environment: To demonstrate the DefensePro's features and functionality.

  The environment included a virtualized client, attacker and target server hosted in a VMWare ESXi system. The equipment for the test was located in Mahwah, NJ and San Mateo, CA.

- Second test environment: To demonstrate up to 300 Gbps throughput.

  The environment included clients and servers simulated by Spirent Avalanche, malicious traffic generated by Spirent Test Center, Cisco Nexus 9000 and DefensePro HTQE. The equipment for the test was hosted in Tel Aviv, Israel.

- Third test environment: To demonstrate integrated messaging and multi-tier mitigation

  The environment was intended to represent a typical ISP network that would include a BGP edge router, a scrubbing center, and a DefenseFlow virtual appliance or 3rd party flow collector. The equipment for the test was hosted in Ft. Lauderdale, FL.

## Test Plan Overview

Prior to the audit, Radware created and delivered to ICSA Labs its DDoS testing test plan that was designed to demonstrate the types of DDoS protection offered by DefensePro. ICSA Labs reviewed the test plan ahead of the onsite audit dates to make sure it was both adequate and sufficiently detailed.

Between the initial audit performed by ICSA Labs in early August and the subsequent audit later in the month, Radware implemented additional improvements to the test plan which ICSA Labs again examined. In addition to fixing some syntax errors, the most notable difference was the improved test plan's use of `jmeter` (https://jmeter.apache.org/) for legitimate traffic generation instead of the command line tool `curl`. The advantage of `jmeter` is that it includes a graphical interface for configuring network traffic parameters. The `jmeter` tool is also capable of graphically reporting on some protocol values, like HTTP server response codes.

The Radware DDoS product testing that ICSA Labs audited followed the steps outlined in Radware's test plan document. The test plan consists of eight sections, each containing one or more test scenarios. In total, there were thirty test scenarios. Each section contained a set of success criteria. Armed with the success criteria, an observer can verify whether or not the Radware DDoS protection solution had performed properly when faced with a DDoS or other related attack. To summarize the eight sections and success criteria were as follows:

1. *Volumetric IP layer 3 floods*. These are used by attackers to consume both Internet bandwidth and resources. Radware demonstrated that these attack scenarios are simple to launch. In addition they showed that the source IP addresses of traffic for the attacks are easily "spoofed".

   Success Criteria:
   - Verify APSolute Vision reports on the flood attack with a relevant alert
   - Verify low CPU utilization on target host
   - Verify attack traffic is not reaching the target host
   - Legitimate traffic is reaching the target host

2. *Very High Throughput attack mitigation*. On the Internet, these kinds of attacks would require a distributed network of attackers in order to reach DDoS bandwidths of 200-300 Gbps. Radware simulated 15 million packets per second (PPS) in these attacks with their Spirent gear.

   Success Criteria:
   - Verify APSolute Vision reports on the flood attack with a relevant alert
   - Verify low CPU utilization on DefensePro multi-purpose CPU
   - Verify baseline 30Gbps simulated legitimate HTTP traffic is still passed

3. *Application layer DoS/DDoS floods*. The layer 7 DDoS floods demonstrated by Radware are based on legitimate requests generated by a botnet targeting a victim application server (e.g., NTP, DNS or HTTP Server). The primary goal of this Denial of Service traffic is to utilize significant amounts of the server's processing resources, not to consume all the available connection bandwidth (although the latter also sometimes occurs).

   Success Criteria:
   - Verify APSolute Vision reports the corresponding layer 7 flood attack
   - Verify low CPU utilization on target server
   - Verify attack traffic is not reaching the target server
   - Verify target service responds to legitimate requests

4. *Low and slow plus exploits*. These scenarios utilized either low and slow DDoS tools or commonly known Internet attacks (e.g., Shellshock and Heartbleed).

   Success Criteria:
   - Verify APSolute Vision reports the attack as an intrusion
   - Verify attack traffic did not reach the target server
   - Verify legitimate requests from a valid client succeed

5. *Handling SSL*. In this scenario, Radware showed how an HTTPS GET Flood can be handled using DefensePro through careful observation of how the encrypted sessions are established.

   Success Criteria:
   - Verify APSolute Vision reports on HTTPS SYN flood attack
   - Legitimate HTTPS transactions succeed while under attack

6. *Mitigating brute-force and other application misuse*. Radware demonstrated application scanning, brute-force attacks on credentials and other application misuse of NTP, HTTP, and SIP.

   Success Criteria:
   - Verify APSolute Vision reports a Scan attack on the corresponding protocol
   - Verify scan traffic is not reaching target service
   - Verify legitimate queries from a valid client succeed

7. *Diverting traffic further away from the victim's core network*. A typical multi-tier network of Radware and 3rd party solutions were deployed to demonstrate how attack mitigation downstream of the border gateway combined with DefenseFlow workflow is used to detect and divert attack traffic.

   Success Criteria:
   - Verify APSolute Vision detects volumetric flood with bandwidth increasing over time
   - At predefined traffic levels, verify a pending notification appears in the UI
   - Verify diversion routes are advertised to BGP peers
   - Verify traffic and protection policy is pushed to scrubbing center
   - At pre-defined upper limit, verify traffic is "black-holed"

8. *WAF with Defense Messaging*. This scenario was a demonstration of Radware's Defense Messaging in which the AppWall WAF deployed out of path and configured to detect attacks originating from behind a Content Delivery Network (CDN) signals the DefensePro to block the attacker based on the *True-Client-IP* HTTP header.

   Success Criteria:
   - Verify the WAF attack is displayed with the packet source IP and true client IP
   - Verify the DefensePro protection policy contains true client IP from WAF alert
   - Verify attack client receives 403 Forbidden response
   - Verify legitimate request from same source IP, receives 200 OK response

## Testing

ICSA Labs observed Radware began each scenario by sending legitimate traffic to the protected server throughout the duration of the test. This legitimate traffic was monitored by the sending host, the inline device and the target server. Once the expected amount of legitimate traffic was verified, and ICSA Labs

agreed, then malicious traffic applicable to the scenario was initiated by an attacking host. This malicious traffic is also monitored by the sending host, the inline device and the target server. The expected result for each scenario is that the attack is detected and mitigated by the inline DefensePro, which it always was. Finally, all traffic observed at the target host was monitored by `tcpdump` and `slurm` running locally on the target host.

## Observations of Interest

Note that in order to isolate and test the robustness of the DefensePro's layer 7 DDoS attack mitigation capabilities, Radware disabled some of its other protection modules that would have triggered earlier. For example, the DefensePro Layer 3 SYN flood protection was disabled during an HTTP GET flood attack so that the attack would continue long enough to reach the application layer HTTP Flood protection module. ICSA Labs observed Radware removing these protection modules from the active protection profile such that only the specific protection being examined remained enabled.

The DefensePro management software (APSolute Vision) displayed inbound and outbound traffic for the inline DefensePro appliance. The dashboard showed the traffic arriving at the device and the traffic that was discarded in each direction. The APSolute Vision console was used to display throughput and indicate if the attack had been detected and mitigated. Some of many console screenshots taken during the audit are included in Figure 1 below.

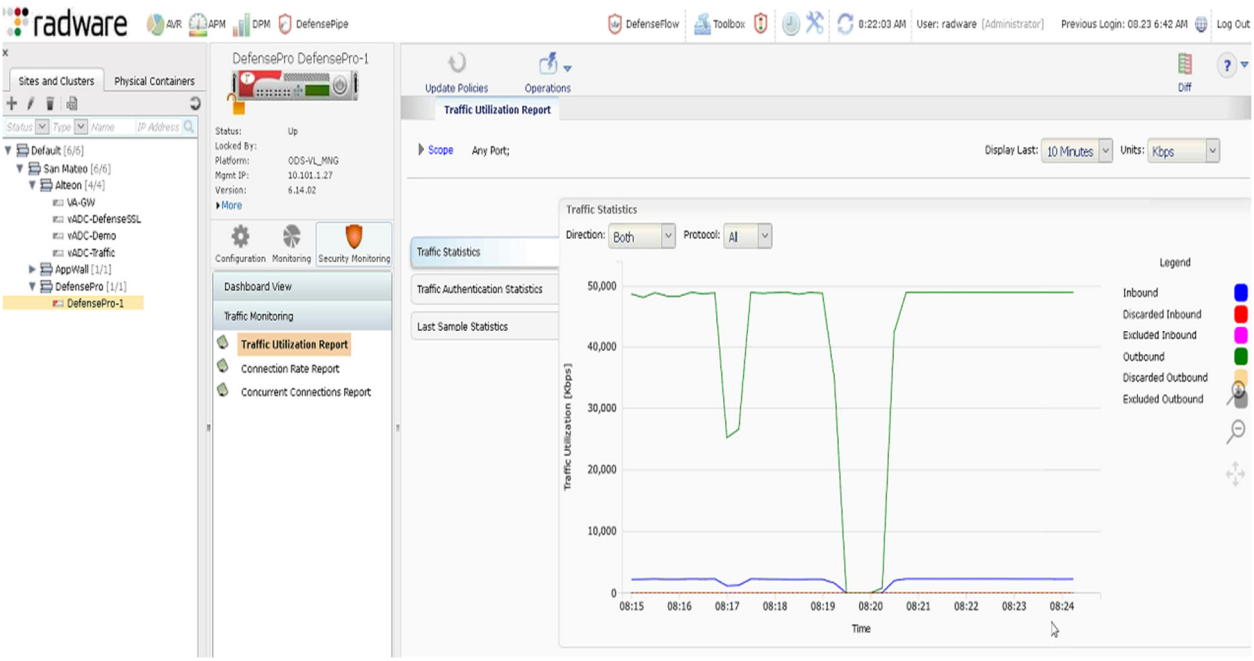


Fig. 1 - APSolute Vison Traffic Utilization Report showing inbound and outbound traffic

ICSA Labs observed several scenarios that demonstrated the DefensePro's capability of behavioral DoS protection (BDoS) by generating a real-time signature (referred to as a "footprint"). This ephemeral footprint uniquely matched characteristics of the current attack campaign (e.g., packet TTL or destination port). The result is that attack traffic is automatically pushed to the dedicated ASIC so that the attack flood is easily mitigated by the inline device. After the footprint is in place, general purpose CPU utilization on the DefensePro is very low and it easily handled forwarding legitimate traffic destined for the target server. ICSA Labs observed the DefensePro detecting the attack, creating a footprint and mitigating all the attack traffic within 12 seconds. An example of this is below in Figure 2.

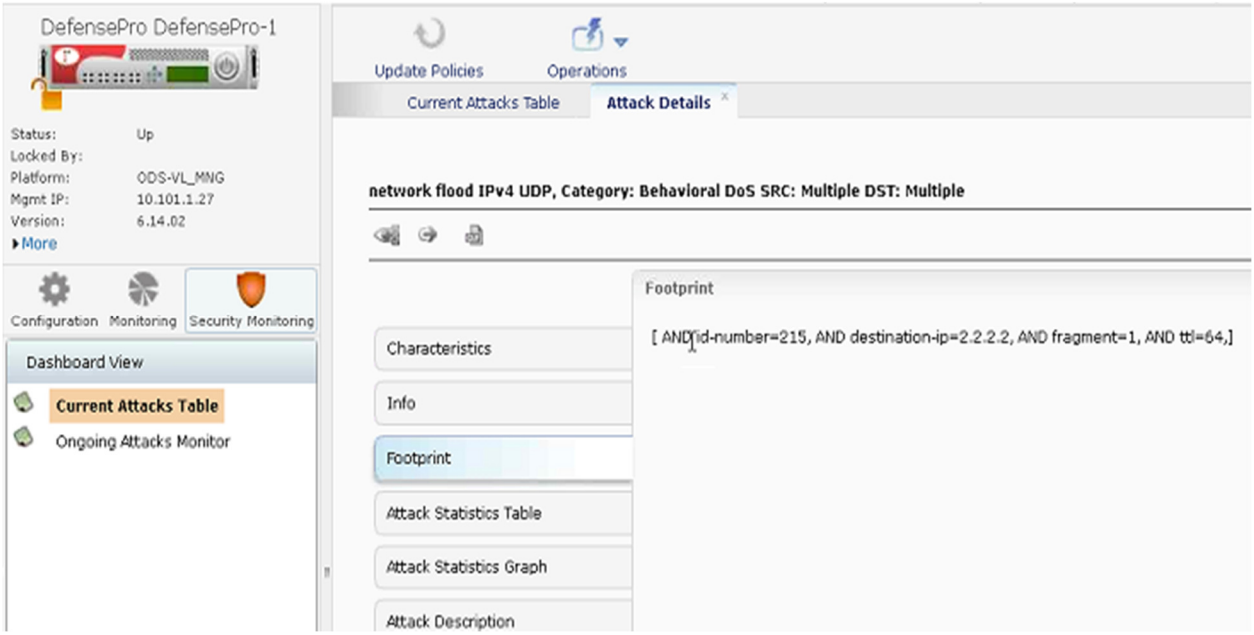


Fig. 2 - Automatically generated "footprint" (see middle right on figure) with four "AND" operators

In a UDP fragmented attack with random packet size and port, ICSA Labs observed several characteristics of the UDP flood attack change over time as the duration of the attack grew. Initially the modification to the attack allows it to evade the footprint, but ICSA Labs observed the DefensePro detecting that some of the flood traffic was no longer matching the footprint and automatically updating its parameters. When it did so, it began to keep pace with the attacker.

One concern with automatically generating a footprint is the risk of unintentionally blocking legitimate traffic that also matches the criteria resulting in self-inflicted DoS. ICSA Labs observed this during an ICMP flood with a spoofed source IP. The footprint generated matched the ICMP type, destination IP and packet TTL, all of which would also match legitimate ICMP echo requests. ICSA Labs observed Radware engineers resolving this issue by whitelisting the IP addresses of servers allowed to ping the target. Note that Radware did also demonstrate the DefensePro can be configured to bypass specific parameters to avoid the need for manual whitelisting.

A technology employed by Radware DDoS mitigation product that helps it to eliminate the occurrence of false positives is HTTP Challenge. When an HTTP flood is detected the DefensePro was configured to challenge the attacker with a 302 Response Code (HTTP redirect). A legitimate browser client has no trouble following the redirect, but most attack tools will not. When an IP address fails to follow the redirect it was blacklisted for a predefined time. ICSA Labs observed the 302 challenge being sent by the DefensePro using the Firefox browser with a plug-in that displayed HTTP headers. Radware also said it supports JavaScript challenges, but the capability was not observed during the audit.

DefensePro also includes a DNS Challenge. The DNS Challenge is implemented similarly to HTTP Challenge. When a DNS flood is detected, clients are requested to resend the request. Valid DNS clients happily oblige, while most test tools do not resend. When an IP address fails to resend it gets blacklisted for a predefined time. ICSA Labs observed the DNS resend using `tcpdump` on a legitimate client. Note that while DefensePro supports passive DNS challenge (waits to resend) as was demonstrated, Radware also said it supports active challenge (switches to TCP) methods but the latter capability was not observed during the audit.

In the scenarios related to delivering network DDoS protection as a network service, Radware demonstrated DefenseFlow for ICSA Labs. According to Radware, DefenseFlow is an orchestration engine which manages the entire attack lifecycle from detection to diversion and mitigation. Workflows and operational procedures enable the ability to define operations to follow based on predefined criteria such as when an attack reaches a certain throughput level to redirect traffic via BGP or BGP FlowSpec to a scrubbing center, inject a blackhole rule or rate-limit the attack traffic at the perimeter. The test bed simulated a typical service provider environment where a lower capacity DefensePro would start the mitigation at the customer premise. Upon attack escalation, DefenseFlow extracted the traffic baselines from the CPE mitigation device, provisioned the security policy on the DefensePro at the scrubbing center, and injected a BGP diversion to re-route only the attack traffic to continue the mitigation. As the attack continued to grow, ICSA Labs observed DefenseFlow advertise a BGP "blackhole" route to all perimeter network devices in order to avoid any impact to the service provider infrastructure.

Radware demonstrated for ICSA how its Web Application Firewall (called AppWall) can be configured out-of-path to detect web attacks and signal an upstream DefensePro to mitigate the threat. Radware calls this capability "Defense Messaging" and it enables the solution to apply its detection technology for successful mitigation of complex and multi-vector attacks. The application layer signature supports adding HTTP headers such as *XFF* or *TRUE-Client-IP*, which enables the capability to block an attack originating from behind a content delivery network (CDN) while allowing all the legitimate requests accessing the site via the same CDN IP address. ICSA Labs observed a malicious IP address behind a CDN being blocked while other legitimate traffic from behind the same CDN IP address was allowed.

## Conclusions

Through a combination of onsite and WebEx-enabled sessions, and over the course of three days, ICSA Labs observed Radware engineers execute each of the thirty test scenarios in their test plan. ICSA Labs had reviewed the test plan even before the audit and found it was adequate for testing a DDoS mitigation product, like DefensePro. It is ICSA Labs' assessment that each test scenario was thoughtfully designed and properly executed such that each test case adequately demonstrated the protection technology which the test was designed to highlight.

Furthermore the attack tools and techniques used in the set of test scenarios adequately represent the misuse of IP network layer and application layer resources that has been reportedly observed in use by "in-the-wild" attack campaigns in recent years.

## Appendix A – Test Tools

During the audit performed by ICSA Labs, Radware used a variety of well-known and publically-available tools to execute the tests and monitor traffic flow on the systems in the test bed. Radware shared all the wrapper scripts and configuration files that were used in each test case. In that way, ICSA Labs was able to review and confirm that the tools when executed properly would produce the expected network traffic for the respective test case. Below is a list of the attack tools, monitoring tools, and applications used for legitimate clients and servers including the URL (where applicable):

- Attacker - Kali
  - Slowloris - http://ha.ckers.org/slowloris/slowloris.pl
  - LOIC - http://sourceforge.net/projects/loic/
  - Bonesi - https://github.com/markus-go/bonesi
  - slowHTTPtest – https://github.com/shekyan/slowhttptest
  - nikto – https://github.com/sullo/nikto
  - svwar – https://github.com/EnableSecurity/sipvicious
  - fierce – https://github.com/mschwager/fierce
  - Hping3 – http://www.hping.org
  - Apache JMeter: https://jmeter.apache.org

- Target – Kali
  - Apache2 / MySQL
  - OWASP Bricks - https://www.owasp.org/index.php/OWASP_Bricks
  - Bind9 – DNS Server
  - Openntpd – NTP Server

- Legitimate traffic generator
  - curl –v –L –b cookie.txt –c cookie.txt <target_IP>
  - Apache JMeter: https://jmeter.apache.org
  - ntpdate
  - nslookup / dig

- Legitimate traffic workstation – Windows
  - Firefox (using developer tools or Firebug)
  - nslookup / dig

- Result Monitoring (to validate results of mitigation)
  - APSolute Vision (alerts, dashboard, traffic monitoring)
  - Tcpdump or Wireshark on target
  - Slurm (traffic graph) on attacker and target

- High Throughput Traffic Generation and monitoring
  - Spirent Test Center: 4.64 (100G connectivity)
  - Spirent Avalance 4.58

## Test Facility Information

This audit report is issued by the authority of the Managing Director, ICSA Labs.

### Test Location

Radware, Inc.
575 Corporate Dr,
Mahwah, NJ 01730

### Lab Report Date

September 27, 2016