May 22, 2024

DNSBomb Pulsing DoS Attack

Researchers from the Tsinghua University in Beijing, China, disclosed a new method for launching pulse wave denial-of-service (DoS) attacks. The new DoS attack vector, dubbed DNSBomb, abuses multiple widely deployed mechanisms to enhance the reliability and availability of the Domain Name System (DNS) to accumulate DNS queries sent at a low rate and concentrate them into short, high-intensity bursts of volumetric traffic. This overwhelms and disturbs TCP traffic flows on target systems and services. By combining both DNS queries and responses with multiple operational DNS resolution mechanisms, the Pulsing DoS attack technique obtains a bandwidth amplification factor (BAF) of over 20,000.

Pulsing DoS Attacks

Pulse Wave or Pulsing DoS attacks use repeated, short bursts of high-volume traffic to impact a target system or service. Pulses or bursts can last up to a few hundreds of milliseconds with a periodicity of a couple of seconds while the attack campaign can span hours or even days. The high-volume pulses in Pulsing DoS attacks only last milliseconds, unlike the more common Burst Attacks, which typically last several minutes. Due to their low average traffic bandwidth, Pulsing DoS attacks are harder to detect than traditional flooding attacks.

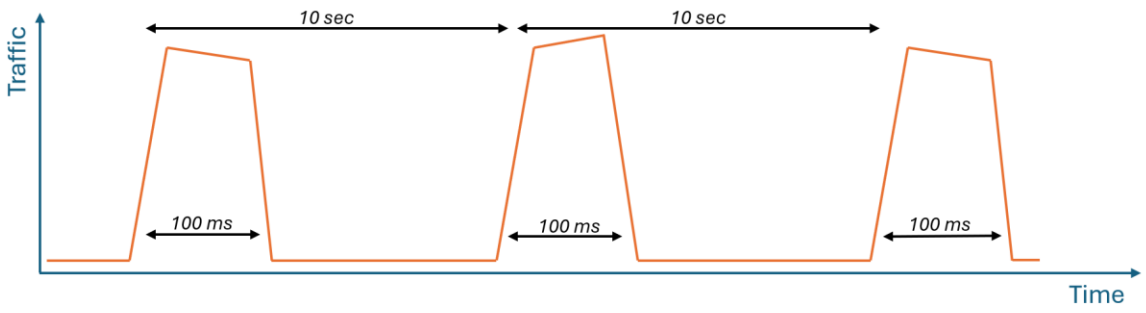

Figure 1: DNSBomb Pulsing DoS attack traffic pattern

Pulsing DoS attacks cause intermittent packet loss and degrade TCP connections by exploiting weaknesses in the TCP congestion control mechanisms that operate on timescales such as Round-Trip Time (RTT) and Retransmission Timeout (RTO). RTT is used to estimate the optimal amount of data that can be in flight in the network and adjust the sending rate (the congestion window) under normal conditions. If packet loss occurs, TCP will wait for a period of RTO after a packet is resent until receiving a valid packet. Upon further loss, TCP dynamically adjusts RTO based on RTT and its variation and continues the retransmission strategy. If the

total DoS traffic during an RTT-length pulse is sufficient to induce packet loss, the TCP flow will enter a timeout and resend a packet RTO seconds later. Moreover, if the DoS pulse period approximates the RTO, the TCP flow will continually incur a loss as it tries to exit the transmission state and eventually fail to exit resulting in a near-zero throughput.

Multiple studies demonstrated the potential of severe impact by Pulsing DoS attacks on various network components, including dynamic load balancers, wireless networks, VoIP networks, application servers, peer-to-peer networks, cloud data center networks, server-side sockets, cloud auto-scaling mechanisms, IoT protocols, SDN control channels, 4G/LTE networks, residential networks, low earth orbit satellite networks, Resource Public Key Infrastructure (RPKI) systems that are designed to secure the internet BGP routing infrastructure, and others. While those studies were mostly limited to model analysis and experimental simulation, they prove that any applications or services that provide adaptation or feedback-control mechanisms are susceptible to Pulsing DoS attacks.

DNSBomb Attack

Prior studies show it is challenging to tightly synchronize attack traffic from different bots as a bursting pulse at target servers, reducing the effectiveness of botnet-based Pulsing DoS attacks. The alternative, centrally generated, reflector-based Pulsing DoS attacks, either yield a small amplification factor or require a large pulse period (1,800 seconds) and have negligible impact on normal traffic during that period. In contrast, the DNSBomb Pulse DoS attack can be initiated with an arbitrary pulse period ranging from thousands of milliseconds to any duration of time and can reach an amplification factor greater than tens of thousands of times.

Although DNS has been demonstrated to be exploitable by DoS attacks, only DNS queries or responses were leveraged to establish traffic amplification or concentrate pulsing traffic. The researchers, however, discovered that a combination of DNS queries and responses opens a long-overlooked attack surface to launch new powerful Pulsing DoS attacks. They found that prevalent DNS mechanisms provide sufficient time to accumulate DNS queries and enable the rapid packet transmission to concentrate DNS responses with a low cost on the attacker's infrastructure. By combining both DNS queries and responses with multiple operational DNS resolution mechanisms, the researchers were able to achieve a BAF of over 20,000.

Attack Overview

The DNSBomb attack consists of three steps. In the first step, the attacker will generate a stream of DNS queries at a very low rate, directed at an exploitable DNS resolver. During the first step, DNS queries for the same domain are aggregated by the resolver. In the second step, the resolver will direct a single, aggregated query for all the DNS queries in step one to the attacker-controlled authoritative name server for the domain. The authoritative name server will respond to the query with a large-sized response. By holding this response until nearing the

timeout of the resolver, the resolver can be manipulated to simultaneously return all responses for the queries in step one to the target server resulting in a short, powerful, amplified traffic pulse.

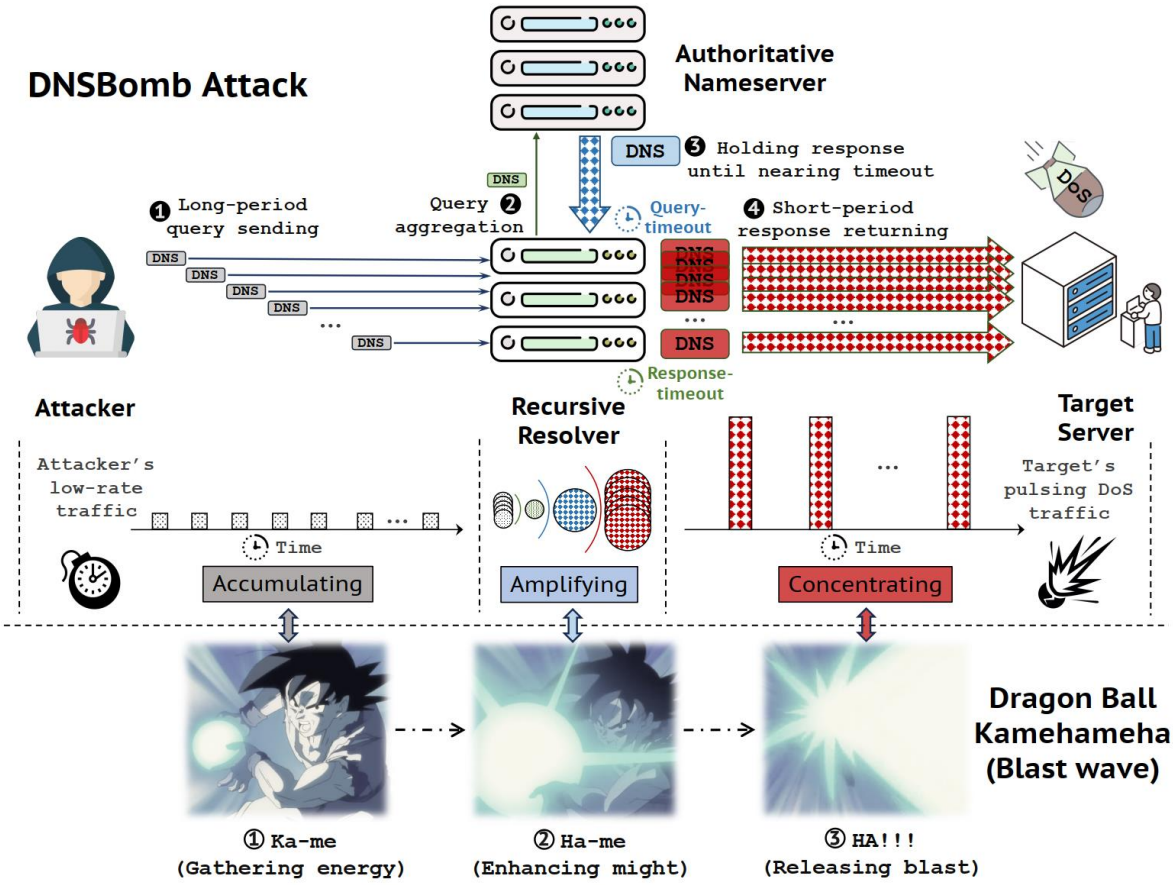

Figure 2: DNSBomb attack (Source: DNSBomb paper)

The researchers compare the three-step process to Goku's signature Kamehameha technique, a blast wave attack featured in the Dragon Ball anime series. The Kamehameha is performed with cupped hands to gather energy (step one: Ka-me), the energy is accumulated into a large force ball between the cupped hands (step two: Ha-me) until the hands are thrust forward to release a streaming, powerful beam of energy (step three: HA!!!). Analogous to the Kamehameha attack, a DNSBomb first accumulates sufficient DNS queries, then amplifies each query, and finally concentrates the responses into a powerful attack flood directed at the target server.

By exploiting the DNS resolver's query aggregation security mechanism, the number of queries the attacker-controlled authoritative name server needs to resolve remains very limited. The

authoritative name server only needs to return a single response for the final query to amplify all the client's queries into multiple, large-sized responses directed at the target server.

The researchers demonstrated through controlled tests leveraging public internet resolvers that by sending a low-rate query stream of only a few hundred kilobits per second (Kbps), the resolvers produce a pulsing DoS attack with pulse peaks of several gigabits per second (Gbps) in magnitude. Depending on the DNS server software, the periodicity of the pulse waves, the accumulation timeout, ranged between 1.5 and 10 seconds. The pulse window, also referred to as the returning time, ranged from 8 to 240 milliseconds (Table 1 in DNSBomb paper).

Threat Model

The DNSBomb attack aims to generate short and periodic traffic bursts targeting a victim server using publicly exposed open DNS resolvers at a very low resource cost for the attacker. Similar to a traditional DNS-based DoS attack, the attacker needs the capability to spoof the source IP with the IP address of the victim. According to the May 2024 statistics from CAIDA, 21.7% of IPv4 Autonomous Systems (ASes) and 27.2% of IPv6 ASes allow IP spoofing. Attackers can leverage any bulletproof hosting service within one of these ASes for source IP address spoofing. Additionally, attackers need to initiate DNS queries for their own domain. That domain can be purchased through any domain registration platform and the authoritative nameserver hosted in any cloud platform and limited resource requirements.

Mitigation Solutions

The DNSBomb attack leverages DNS mechanisms that guarantee the availability, security and reliability of the domain name resolution system. To limit the impact of DNSBomb attacks, service providers and vendors must compromise by reducing resolution performance on their implementations. The researchers, through testing, provide guidelines for timeout, rate limiting, and response-returning timeout settings to reduce the BAF and limit the number of responses in a single pulse for different DNS implementations. The researchers have responsibly informed all affected parties and shared their findings with multiple DNS software vendors and DNS service providers.

Reasons for Concern

Given the limited resources required to perform the attack and the high potential for impacting services and applications, the DNSBomb attack technique should not be ignored. While the researchers have responsibly informed all affected vendors and service providers before publicly disclosing their findings, there is no silver bullet to stop DNSBomb attacks, merely a capability to limit its effects.

While high-traffic pulses only last a couple of hundreds of milliseconds, the impact on network components and TCP-based services can last minutes, allowing well-timed pulse waves to degrade or fully disrupt the availability of a service or application for a sustained period.

This new attack vector reemphasizes the importance of real-time, fully automated, advanced attack detection algorithms and mitigation capabilities that current DDoS protections should provide.

EFFECTIVE DDOS PROTECTION ESSENTIALS

Hybrid DDoS Protection – Use on-premises and **cloud DDoS protection** for real-time **DDoS attack prevention** that also addresses high-volume attacks and protects from pipe saturation

Behavioral-Based Detection – Quickly and accurately identify and block anomalies while allowing legitimate traffic through

Real-Time Signature Creation – Promptly protect against unknown threats and zero-day attacks

Web DDOS Tsunami Protection – Automated immediate detection and mitigation of Web DDOS encrypted high RPS and morphing attacks

A Cybersecurity Emergency Response Plan – Turn to a dedicated emergency team of experts who have experience with Internet of Things security and handling IoT outbreaks

Intelligence on Active Threat Actors – High fidelity, correlated and analyzed data for preemptive protection against currently active known attackers

For further **network and application protection** measures, Radware urges companies to inspect and patch their network to defend against risks and threats.

EFFECTIVE WEB APPLICATION SECURITY ESSENTIALS

Full OWASP Top-10 coverage against defacements, injections, etc.

Low false positive rate using negative and positive security models for maximum accuracy

Auto-policy generation capabilities for the widest coverage with the lowest operational effort

Bot protection and device fingerprinting capabilities to overcome dynamic IP attacks and achieve improved bot detection and blocking

Securing APIs by filtering paths, understanding XML and JSON schemas for enforcement, and using activity tracking mechanisms to trace bots and guard internal resources

Flexible deployment options including on-premises, out-of-path, virtual or cloud-based

LEARN MORE AT RADWARE'S SECURITY RESEARCH CENTER

To know more about today's attack vector landscape, understand the business impact of cyberattacks, or learn more about emerging attack types and tools, visit Radware's **Security Research Center**. Additionally, visit Radware's **Quarterly DDoS & Application Threat Analysis Center** for quarter-over-quarter analysis of DDoS and application attack activity based on data from Radware's cloud security services and threat intelligence.