WHITE PAPER

Surviving DNS DDoS Attacks

Introducing self-protecting servers

SECURE64

Background

The current DNS environment is subject to a variety of distributed denial of service (DDoS) attacks, including reflected floods, amplification attacks, TCP and UDP data floods, malformed TCP/IP packets, and improper DNS requests. The intent of a DDoS attack is interrupt service by either overloading a target or forcing it to reset. Recent information on DDoS attacks include:

- On August 16, 2012, a DDoS attack focused on AT&T's DNS (Domain Name System) servers disrupted data traffic for business customers. The attack was focused specifically on AT&T DNS servers.
- In September 2012, Bank of America, NYSE, JPMorgan Chase and Wells Fargo were the targets of DDoS attacks affecting web services to customers.
- The number of DDoS attacks has increase by 50% and the volume of traffic used in the attacks has increased by 63%, from 2011 to 2012.
- 31% of mid to large sized enterprises in the US and UK suffered at least one DDoS attack in the last 12 months.

DDoS attacks overwhelm network resources in an attempt to stop valid traffic. Increasing armies of botnets are making these attacks ever larger. The results can include lost revenues, lost customers, and lost brand reputation. For today's always-on organizations, such as service providers, carriers, domain operators, and e-businesses, remaining highly available to valid Internet traffic during an attack is critical.

This white paper discusses conventional approaches to DDoS attack detection and mitigation, and introduces a new concept—a self-protecting DNS server that incorporates DDoS countermeasures directly into the operating system, allowing it to remain responsive to legitimate DNS traffic while defending against high-volume DDoS attacks.

Types of DDoS Attacks

Although there are many different types of DDoS attacks on DNS servers, they can be categorized as either protocol exploit attacks or flood attacks. Flood attacks can be further classified as:

- TCP SYN floods
- Illegal DNS packet floods
- UDP floods
- TCP floods

Protocol exploits

Protocol exploit attacks attempt to cripple the target DNS server by sending illegal IP or TCP traffic that triggers a failure in the I/O driver that is processing the network traffic. These attacks, while potentially crippling, are easily defended against by simply testing the I/O driver for all combinations of invalid traffic.

TCP SYN floods

Flood attacks, however, are much harder to defend against, as they send legal traffic (at least from a TCP/IP protocol perspective) at the victim, in an attempt to exhaust either its CPU or memory resources, causing a denial-of-service condition.

A TCP SYN flood is a common DDoS attack that takes down an unprotected server by consuming resources in the network stack. The attack mechanism abuses the three-way connection handshake (SYN, SYN-ACK, ACK) to crash the target as follows:

- The attacking machines send a SYN.
- The receiver allocates the resources to establish a connection in the target machine and responds with a corresponding SYN-ACK.
- The attacker never responds to the SYN-ACK with a final ACK. The allocated resources remain consumed for a substantial time, exhausting the resource pool.

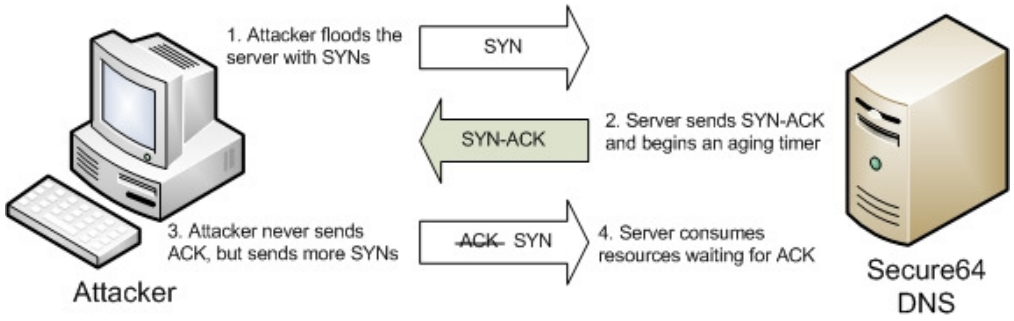

Figure 1. TCP SYN Flood Attack

Various methods have been invented to defend against this type of attack, including the use of SYN cookies. If deployed correctly, SYN cookies can be effective, but this method is computationally expensive for the defending machine. As a result, the target has limited ability to respond to valid traffic during a SYN flood attack.

Illegal DNS packet floods

Two of the most common attacks against DNS servers are UDP reflected and amplified UDP reflected floods. These two attacks utilize open, resolving DNS servers on the Internet to "reflect" a torrent of DNS answers to an authoritative DNS server victim, causing it to become unavailable, or its connection to the Internet to fill. A reflected, amplified attack is illustrated in Figure 2 below. It is amplified because the query from the botnet army is small, whereas the TXT record sent to the victim is much larger. A reflected, non-amplified attack works in a similar manner, but the reflected responses to the victim are smaller.

These types of attacks generate legitimate UDP traffic (DNS responses are certainly valid traffic on the Internet), but are characterized as illegal DNS packet floods because the victim DNS server never issued the original queries that generated the response (the botnet did).

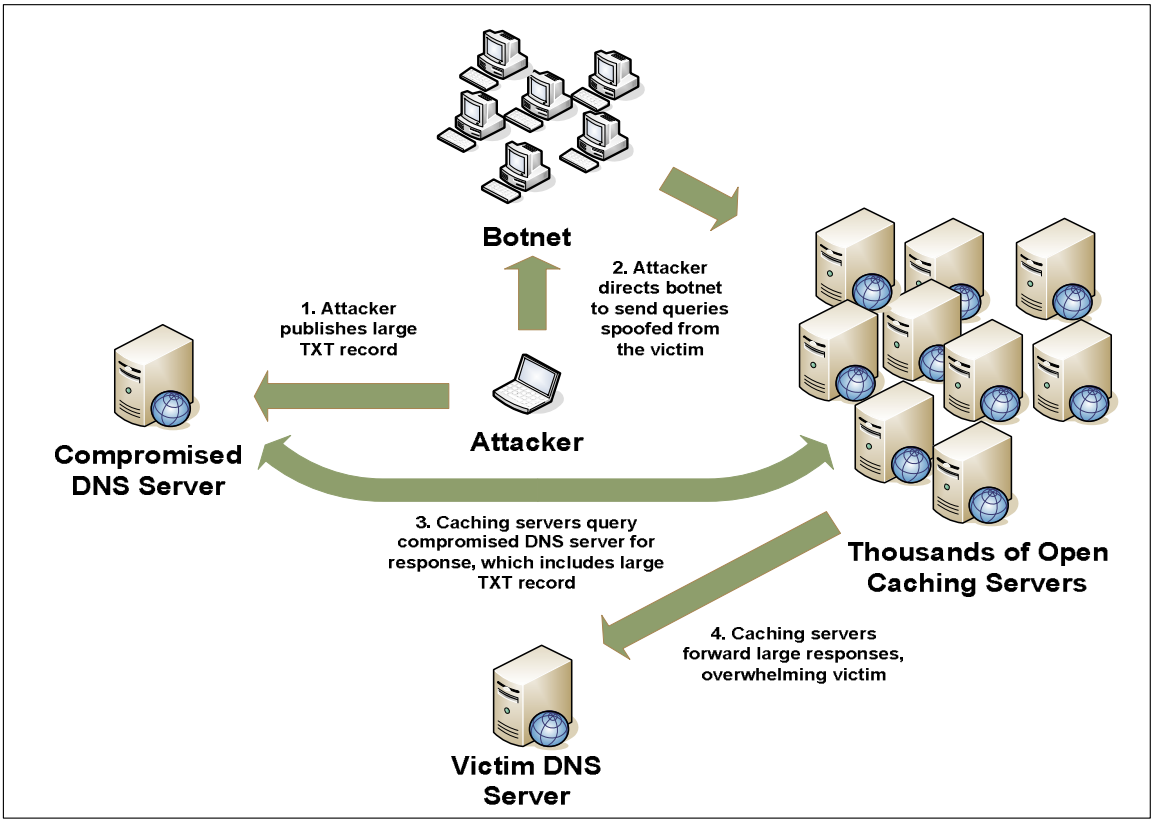

Figure 2. Reflected, Amplified Flood Attack

UDP and TCP data floods

The final type of attack occurs when a large number of bots make more requests of the DNS server than it can handle. This causes the DNS server to drop inbound DNS requests (in the case of UDP), or refuse to establish new connections (in the case of TCP), thus achieving a denial-of-service condition for legitimate users. Although rarer than other types of attacks on the DNS, this type of attack can be created by a large botnet sending a high volume of DNS requests of an authoritative DNS server, either with spoofed or non-spoofed source IP addresses, as shown in the illustration below.

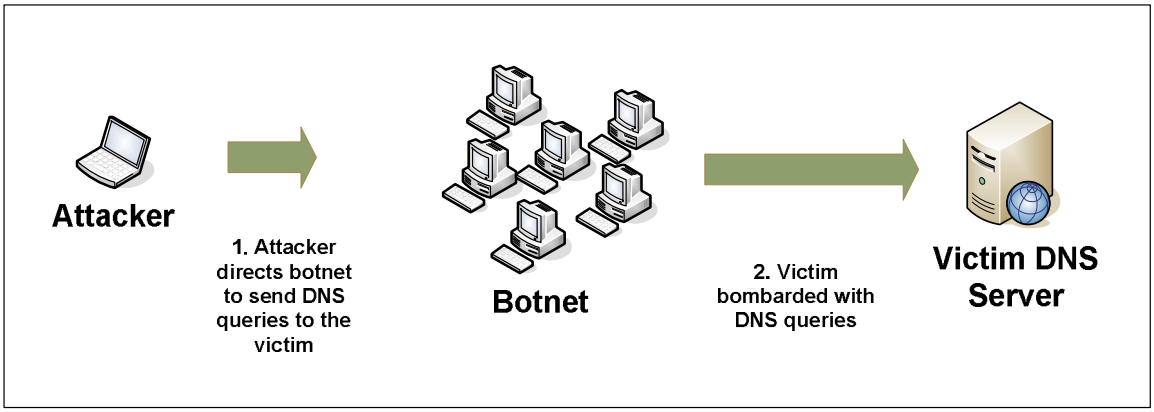

Figure 3. UDP Data Flood

Current Defenses Against DDoS Attacks

To defend against DDoS attacks, organizations have enlisted routers and firewalls, deployed security devices such as IPS systems, invested in dedicated DDoS equipment, and over-provisioned both network infrastructure and DNS servers for sufficient capacity to fend off attacks. These solutions are far from ideal, as they add cost, complexity and latency to the network, and are often only partially successful in defending against these types of attacks.

Router ACLs can be used to block attack traffic once it has been identified by other devices or systems in the network, but do not detect and mitigate DDoS attacks on their own. Furthermore, router ACLs cannot be used to defend against protocol exploit attacks, TCP SYN floods or flood attacks using spoofed IP addresses.

Traditional firewalls are used to pass or block certain types of traffic from ranges of IP addresses and ports, but are not usually designed to analyze the traffic that passes through them. This makes them unable to defend against any of the types of attacks described above. Certain newer firewalls incorporate limited capabilities to detect and block TCP SYN flood attacks, but are not designed to detect or mitigate protocol exploits or UDP or TCP flood attacks.

Intrusion Prevention Systems (IPSs) and dedicated DDoS equipment rely on techniques such as packet inspection, stateful analysis engines, and traffic shaping to detect and mitigate DDoS attacks. While traditional IPS products do a good job of detecting invalid packets, they are prone to false positives and cannot detect attacks that employ valid packets (like UDP and TCP floods). DDoS mitigation solutions suffer from the reverse problem; they are designed to detect and mitigate UDP and TCP flood attacks, but can be defeated by exploit attacks and certain types of stealth attacks. Both types of systems can be difficult to configure and manage, requiring oversight by highly skilled administrators. In addition, these products are costly (solutions are approximately $50K per 1 Gbps of protection), and organizations normally install the devices in high availability pairs. Depending on the product and network architecture, the devices can also slow down the network by imposing more latency.

Over-provisioning DNS servers is an equally costly and ineffective solution to the problem, since over-provisioned DNS servers do not successfully fend off DNS DDoS attacks. Even a modest sized attack can consume the extra capacity of an over-provisioned DNS, eventually making the DNS servers unavailable. And attackers can simply increase the size and force of the attack to consume additional DNS resources, well before the network infrastructure is saturated.

Designing a Self-Protecting DNS Server

Secure64® DNS Authority is a dedicated, authoritative DNS name server that runs on the Secure64 SourceT® micro operating system on Itanium® 2-based servers. The SourceT network I/O stack is designed to be self-protecting, allowing it to identify and block attack traffic while continuing to respond to DNS queries from legitimate sources.

As shown in Figure 4 , SourceT includes four levels of protection to defeat these types of DDoS attacks: protocol exploit protection, TCP SYN flood protection, illegal DNS packet protection, and UDP/TCP flood protection.

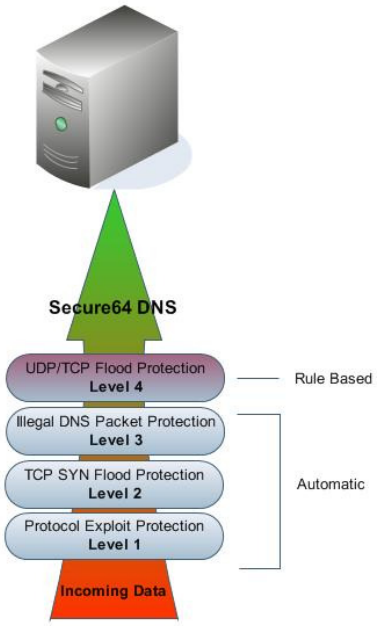

Figure 4. Four levels of DDoS Protection

Table 1 summarizes the countermeasures utilized by each of the four filters built into the system's I/O stack.

Filter	Attack Type	Mitigation Methods
Level 1	Protocol exploits	Automatic detection of malformed packets and invalid combinations of header bits, which are dropped automatically
Level 2	TCP SYN flood attacks	Automatic detection and mitigation via pre-connect/aging method
Level 3	Malformed DNS requests	Automatic detection and mitigation via DNS packet inspection
Level 4	UDP and TCP packet flood attacks	Automatic blacklisting, rules-based flood protection, and high-performance system that can operate while under attack

Table 1. Secure64 DNS Authority DDoS Protection

In addition, the SourceT micro operating system protects the name server from buffer overflow attacks and compromise by rootkits, viruses, Trojans, and other malware. These protections are described in another white paper, "Eliminating Malware and Rootkits: Six essential characteristics of a genuinely secure OS," available from the Secure64 web site.

Level 1: Protection from protocol exploits

Inbound packets are first processed by the NIC hardware and then by the I/O driver. Common exploits involving malformed packets or invalid combination of header bits are dropped immediately.

The first filter protects against protocol exploits, including the following:

ARP malformed packet types

- Not Ethernet Type
- Not IPv4
- Packet Contains Bad OP Code

MAC malformed packet types

- Frame Too Short
- Frame Contains Bad CRC
- Frame Contains Bad Address

TCP malformed packet types

- Segment Duplicate Has Different Data
- Segment Overlap has Different Data

IP malformed packet types

- Frame Contains Bad Header Version
- Frame Contains Bad Header Length Info
- Frame Contains Bad Checksum
- Fragment Causes Overlay of L4 Header
- Fragment Is Last and Duplicate
- Fragment Offset is Beyond Last Expected
- Fragment is Duplicate
- Fragment is Overlapping
- Fragment Number Exceeds Maximum Allowed

The effectiveness of this filter was tested by ExtremeLabs, an independent test laboratory, who subjected Secure64 DNS Authority to a variety of exploit attacks. In one of these tests, Authority was placed under a nominal load of legitimate traffic, and then subjected to ARPFlood, in which a network segment is flooded with inaccurate ARP protocol-based information. Due to the pseudo-random addresses sent, some servers lose performance quickly and others become unavailable or crash.

> **Secure64 DNS Authority remained 100% responsive to legitimate queries while mitigating the SYN flood attack until the total data rate saturated the Gigabit connection at 830 Mbps of total traffic.**

Details of this and other Extreme Lab tests are available from the Secure64 web site, www.secure64.com.

Level 2: Protection from TCP SYN floods

Unlike conventional operating systems that allocate connection resources upon receipt of the initial SYN request, SourceT does not establish a connection until the three-way handshake is complete, using a pre-connect/aging method:

- When a SYN arrives, a small pre-connection entry is established in an OS data structure and begins to age.
- If the final ACK arrives before the timeout completes, the connection is legitimate and fully established.
- If the final ACK does not arrive before the entry times out, the entry is deleted.
- The platform establishes ongoing legitimate connections while avoiding the allocation of limited resources during SYN floods.

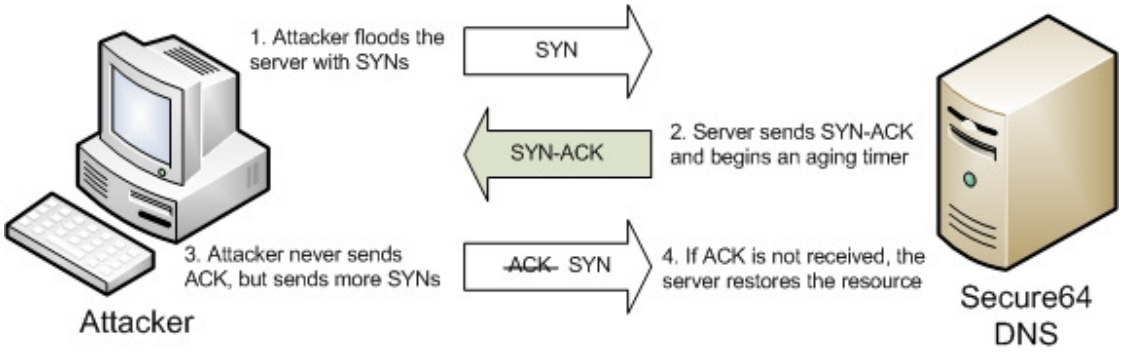

Figure 5. Secure64 DNS Authority TCP SYN Flood Protection

In addition to the pre-connect/aging method, the Secure64 platform automatically limits the number of incoming SYNs and the number of new pre-connections:

- Incoming SYNs are limited to a rate of 100,000 possible connections per second.
- Established connections continue to operate.

To validate the operation of this filter, the test laboratory placed the Authority server under a legitimate traffic load while TCP SYNs were generated from a number of different source IP/MAC pairs at an increasing rate until the saturation point of a Gigabit Ethernet connection.

> **Secure64 DNS Authority remained 100% responsive to legitimate queries while mitigating the SYN flood attack until the total data rate saturated the Gigabit connection at 830 Mbps of total traffic.**

Level 3: Protection from illegal DNS packets

To efficiently withstand attacks while processing valid DNS queries, the I/O driver performs DNS validation on all UDP packets sent to the server's configured DNS IP address(es) and port.

The system examines the DNS packets and rejects malformed queries, such as queries that are shorter than the minimum query length. It also rejects any DNS responses sent to the configured DNS port (which is listening for queries). Inspecting the DNS packets helps prevent attacks that flood the server with improper traffic to the DNS port.

DNS responses to queries sent by Authority (for example, a slave querying an SOA or a master sending a NOTIFY) are sent to a different UDP port and do not affect the validation process.

To test the effectiveness of this filter, Secure64 placed the Authority product under a maximum load of 102,000 legitimate queries per second and then subjected it to a reflected flood attack of approximately 150,000 packets per second of reflected attack traffic (approximately 82 Mbps of attack traffic). Because the inbound attack traffic is illegal (inbound query responses to port 53 are not legal DNS traffic), the illegal DNS packet filter engages to drop the attack traffic.

During testing, Secure64 DNS Authority was still able to answer 100% of the 102,000 legitimate queries per second while dropping the reflected flood attack traffic.

During testing, Secure64 DNS Authority was still able to answer 100% of the 102,000 legitimate queries per second while dropping the reflected flood attack traffic.

Level 4: UDP/TCP data flood protection

UDP DATA FLOOD PROTECTION

By the time traffic passes through the first three filters, illegal and invalid packets have been identified and dropped. What remains are valid queries to which the DNS server should respond. However, during high-volume flood attacks, there may simply be too many queries for the name server to handle properly. In this case, Authority supplies port-based protection from UDP packet floods through automatic IP-based rate limiting and configurable aggregate rate limiting as follows:

- Administrators can configure the software to limit the number of UDP packets accepted per second. This is an aggregate limit, regardless of the source IP address.
- Administrators can configure the software to limit the number of UDP packets accepted per second from each source IP address.
- The system continuously compares the limits with the average packet rate, on both an aggregate and a per-source-IP basis.
- If the UDP packet rate from a source IP address surpasses the source IP packets-per-second limit, the software automatically blacklists the IP address and drops its incoming packets. "Good" traffic continues to flow and new connections can occur, even though an attack is being blocked.
- If the flood from the blacklisted IP address backs off, traffic is accepted after a timeout period. This helps ensure that a spoofed IP address does not block a real user.
- Repeat offenses result in faster blacklisting than the first-time offense. The system retains blacklisted IP addresses for a period of time.
- If the configured aggregate UDP limit is reached, a rate-controller mechanism begins probabilistically dropping packets in order to maintain the incoming rate at the configured limit. This may result in some good traffic loss; however, the server can remain available to a quantity of good traffic and to administrators, when other DNS servers would become overwhelmed.

Figure 6 compares how Authority and Linux running BIND protect against a direct, non-spoofed UDP flood, in which bots direct a high rate of legitimate DNS queries at the victim server in an attempt to overwhelm it. As the figure shows, the Authority server blocks traffic from the attacking bots, maintaining 100% availability

to legitimate queries. The Linux/BIND server, which has no such rate limiting abilities, becomes overwhelmed by the attack, eventually becoming completely unavailable to respond to legitimate traffic.

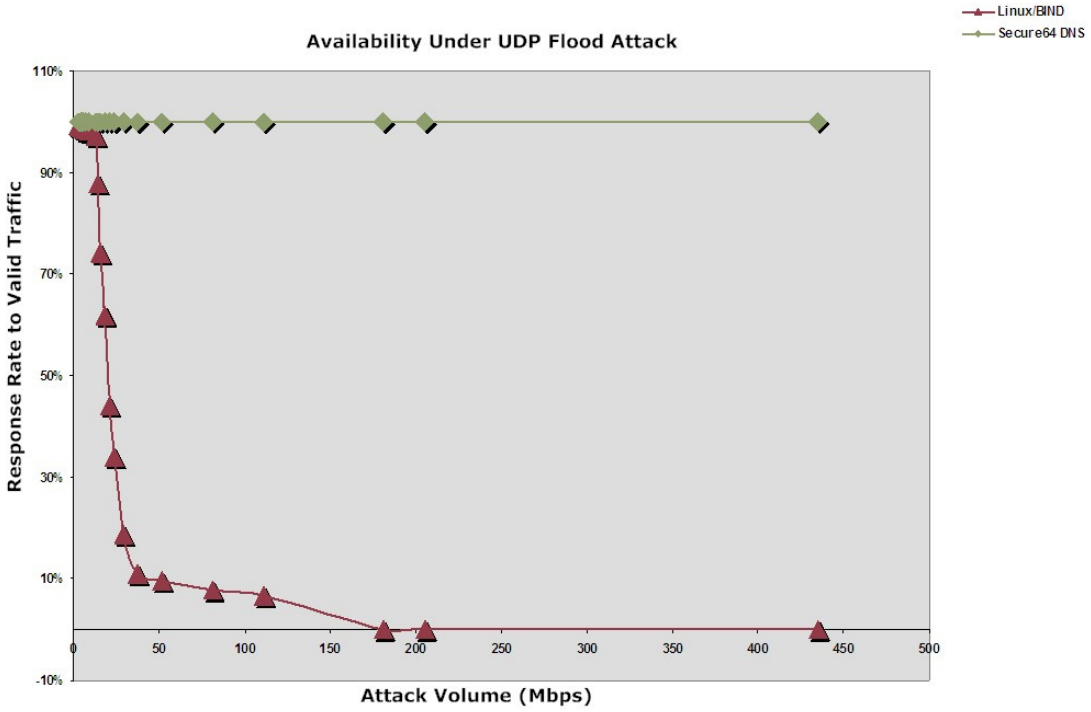

Figure 6. Availability Under UDP Flood Attack

TCP DATA FLOOD PROTECTION

For TCP traffic, Authority tracks resource usage to detect and block specific overload situations. The mitigation process looks for data flows that are consuming more than their fair share of system resources and then blocks those specific flows.

The measurement of resource consumption is a unique feature of the attack-detection mechanism. In contrast, other commonly deployed mechanisms examine data signatures, process firewall/SNORT rule sets, or simply measure traffic for unusual patterns. Measuring resource consumption detects both high-speed and stealth flood attacks.

Stealth attacks consume network resources slowly, without resorting to brute force, high-speed floods. Rate measurement does not typically detect this situation. But resource measurement detects the “back-pressure” of any flood and triggers detection.

For TCP traffic, Authority provides the following detection and mitigation functions:

- Administrators can configure the software to limit the number of packets in the network buffers. This is an aggregate limit, regardless of the source IP address.
- Administrators can configure the software to limit the number of TCP packets accepted per second from each source IP address.
- The system continuously compares the configured aggregate limit with resource consumption on an aggregate basis and compares the source IP packets-per-second limit on a per-source-IP basis.

- Administrators can specify that trusted TCP traffic, such as an internal DNS caching server, is not subject to the mitigation features.
- If an IP address surpasses the source IP packets-per-second limit, the software issues a TCP reset (RST) packet to block the specific attacker. (A RST is issued because blocking TCP traffic based on IP address can incorrectly stop valid TCP traffic.)
- If the configured aggregate limit is reached, a TCP flow controller mechanism begins to probabilistically drop packets, but in a manner that maintains distribution of resources (network buffers). This ensures that neither high-intensity floods nor stealth attacks can consume more resources than specified. As with UDP flood control, this may result in some good traffic loss; however, the server can remain available to a quantity of good traffic and to administrators, when other DNS servers would become overwhelmed.

The following figure illustrates a TCP flood scenario, as follows:

- Authority master and slave servers are configured to allow trusted TCP traffic for zone transfer data (note that zone transfers are also protected by ACLs, and TSIG is supported). This traffic is not mitigated.
- An attacker sends a TCP flood from a spoofed IP address to the master server.
- When the per IP-address limit is reached, Authority issues a RST packet to stop the attack.

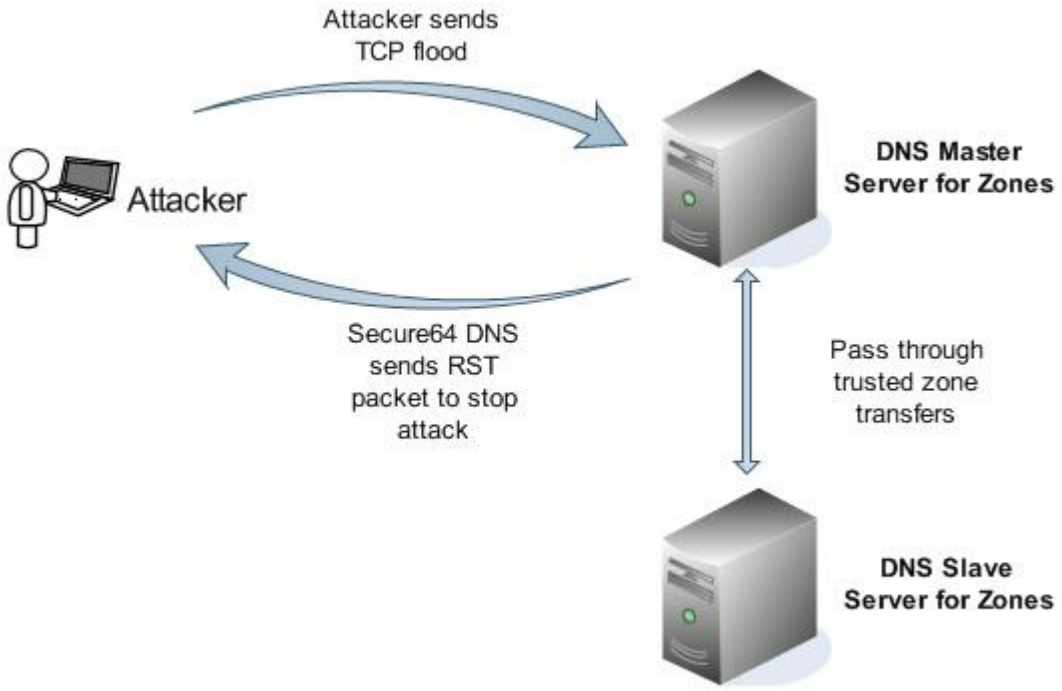

Figure 7. Secure64 DNS Authority TCP Flood Protection

Conclusion

Self protecting servers, such as Authority, detect and block a wide variety of common denial-of-service attacks with little, if any, degradation in server performance. Deploying servers with such an architecture can reduce

the need to overprovision server resources and eliminates the need to protect servers with network security devices, greatly decreasing network cost and complexity while increasing performance.

Even with the aggregate rate limiting and mitigation features in Authority, an amplification attack that employs a vast botnet can flood a victim's pipes to the extent that some good traffic is lost. However, Authority remains available during attacks that would overwhelm other DNS servers, allowing administrative access for attack response and recovery. And, depending on the flow of the attack, the server can continue to respond to requests throughout the event.

About Secure64

Headquartered in Greenwood Village, Colorado, Secure64 is a software company providing the most secure DNS products available to their customers in the government and telecommunications industry. Partially funded by a grant awarded by the Department of Homeland Security, Secure64's patented technology has been proven to be immune to compromise from rootkits and malware and resistant to network attacks that are the source of today's most serious security threats. The company offers a suite of trusted and secure DNS software appliances for caching, signing and authoritative use. Secure64's products are sold and serviced worldwide by both Hewlett-Packard and Secure64. Notable customers include the U.S. Departments of Commerce, Labor, and Interior and Qwest. For more information, visit http://www.secure64.com.

For More Information:

(303) 242-5890
www.secure64.com

Secure64 Software Corporation
5600 South Quebec Street, Suite 320D
Greenwood Village, CO 80111