



FloodGuard™:

A Distributed Solution for Detecting and
Mitigating Flooding Attacks Leading to
Denial of Service

A White Paper
from
Reactive Network Solutions, Inc.

+1-650-298-9481

info@reactivenetwork.com

www.reactivenetwork.com



Introduction

This paper discusses some fundamental design limitations of the Internet and how these limitations create Denial of Service (DoS) vulnerabilities. We also explain the architecture of the FloodGuard system, a product specifically designed to defend enterprises from flooding attacks leading to denial of service.

The simplicity of the design of the Internet has led to both its explosive growth and its limitations, which include vulnerability to flooding attacks. The Internet protocol (IP) does not guarantee packet delivery as some of the routers may temporarily run out of resources. The transport control protocol (TCP) is used at the source and destination endpoints to make sure that the original data is received in its entirety, requesting retransmissions of lost packets from the source endpoint if necessary. TCP does not expect much from the network nodes (routers) and implements a complex and very adaptable communication process to best utilize the available communication resources.

This model enables routers to be as simple as possible, responsible only for forwarding packets, and pushes much of the complexity out to the edge of the network. Routers do not differentiate classes of packets, record any information about the packets going by or have any notion of which source endpoint is transmitting to which destination endpoint. Routers simply forward packets from one input port to another, making them relatively simple and easy to manage.

This simplicity is what makes the Internet scalable and fault-tolerant, with no single point of control or failure. It also makes the Internet cost effective and has enabled service providers and telecommunication companies to greatly reduce data communication costs and to extend Internet connectivity offerings to hundreds of millions of users worldwide.

At the same time, this simplicity creates limitations that have restricted the Internet's ability to offer best-effort, single-tier services. Most notably, service providers today lack the ability to offer tiered quality-of-service guarantees. Such guarantees are necessary for many real-time media applications such as telephony, video conferencing and video broadcasting and for applications with high availability requirements.

Given the pervasiveness of the Internet, any technology that can cost-effectively improve the ability of service providers to offer tiered quality of service guarantees is very attractive as it would allow the reuse of an existing infrastructure (the Internet) to support and extend new applications for millions of potential customers.

Reactive Network Solutions offers such a technology in its FloodGuard system, a distributed detection and response system that automatically detects and mitigates flooding attacks leading to denial of service. FloodGuard non-invasively augments network infrastructures with



distributed intelligence, ensuring scalability and no single point of failure; uses anomaly-based detection and a method of mitigation that leads to continuous refinement; can be deployed both locally in the enterprise network and upstream in a service provider's network; and has no adverse effect on network performance or reliability.

Internet DoS Attacks

Internet DoS attackers use the Internet to remotely and anonymously send packets that have adverse effects on end-systems and/or their needed communication resources, thereby denying access to legitimate users.

This is possible because the Internet infrastructure does not use the packet's source address for forwarding. In general, the source addresses are only used by endpoints to reply back. Most Internet DoS attacks exploit this fact and insert fake source addresses in the packets (spoofing) to avoid traceability. The spoofed packets are delivered to the destination as if they were authentic, thus fooling the receiver into accepting bogus packets and/or replying to bogus source addresses. Even if the receiver discovers the attack, because routers do not record any information about the packets going by, packets with spoofed source addresses cannot be traced.

Although all Internet DoS attacks exploit the intrinsic anonymity of the Internet, the individual techniques used are quite different and therefore require different mitigation approaches. For example, DoS attacks that target deficiencies of certain software implementations are easily perpetrated but can also be easily fixed by patching the software. In contrast, DoS attacks that misuse specific communication protocols may require changes or amendments to standards to fix, a time consuming and expensive process. Fortunately, these attacks are difficult to devise and therefore rare. +

Flooding attacks are the most popular DoS attack threat because they are easy to perpetrate, difficult to stop and have destructive outcomes. In a flooding attack, the perpetrator sends large amounts of bogus traffic toward the victim's network to consume its bandwidth, exploiting the lack of bandwidth resource management in IP networks. Flooding attacks are today further classified into two types, distributed denial of service (DDoS) attacks and distributed reflective denial of service (DRDoS) attacks, although perpetrators are continually devising new ways of initiating flooding attacks.

In the next sections, we discuss how to mitigate flooding attacks in general and describe how the FloodGuard system implements a practical and comprehensive solution.

+In this category, TCP SYN attacks are an exception, being relatively common. The FloodGuard system mitigates TCP SYN attacks in addition to flooding attacks.



Where to Stop Flooding Attacks

Flooding attack mitigation requirements vary depending on the attack type and bandwidth and the type of networking hardware/software present on the path of the attack. In general, any mitigation technology needs to be able to combat flooding attacks in the network to overcome the worst-case flooding scenarios in which the flood overtakes the upstream service provider's access network. If network-based mitigation techniques are not available and the flooding attacks have a low intensity, it may be sufficient to handle the mitigation downstream near the customer's equipment.

Attack Volume (Mb/s)	Attack Types	Typical Bottlenecks	Possible Mitigation Points	Mitigation Approaches
0.1-1	Randomly spoofed SYN attack	Some types of operating systems	Enterprise	Upgrade operating system
1-10	Any SYN attack	Firewall or edge router	Enterprise NSP AL NSP DL	Filter attack packets
10-50	Any packet type	Access router	NSP AL NSP DL Upstream NSP	Filter attack packets
50+	Any packet type	Distribution or border router	NSP DL Upstream NSP	Filter attack packets

AL = access layer

DL = distribution layer

Table 1: Mitigation Approaches Based on Attack Type

Table 1 summarizes the possible places in the network we believe mitigation should occur for different attack volumes assuming a burstable 100 Mb/s connection. As attack volume grows it saturates (fills up) the capacity of servers (if they are not properly hardened), the enterprise edge equipment, the Internet link itself and finally the NSP peering point. Once the entire capacity is exhausted by the flood, no legitimate traffic can reach the enterprise's servers.



Flooding attacks that cannot be solved by improving the operating system of the servers need to be mitigated by approaches that filter out attack packets, freeing up bandwidth to allow legitimate traffic to reach its destination. Table 1 also reports where the mitigation needs to take place to relieve the bottleneck. In all cases, it is necessary to apply the bandwidth resource control mechanism upstream from the bottleneck to free up bandwidth. For some attack bandwidth ranges it is possible to augment the enterprise with appropriate support to relieve the downstream bottlenecks caused by firewalls, servers or low-performance edge routers. In general, though, mitigating attacks at the service provider's peering points can stop a very wide range of flooding attacks. If the attack is so large that it floods the peering points themselves, thus effectively shutting down the connectivity of the Internet core, mitigation needs to be propagated across peering points, which requires cooperation between network service providers.

Obviously these considerations are illustrative and cannot be generalized since the relative capacities of the different parts of the system may differ widely depending on the particular installations and architecture. However, these variations would not invalidate the fact that the mismatch in bandwidth capacity among different segments of the network will cause similar flooding vulnerabilities and similar considerations for mitigation.

We have performed some quantitative analysis to give a sense of the relative capacities of some common networking equipment. For example, we have measured that a high-end PC-based firewall (1GHz Pentium, 512 MB memory, 2 Intel interfaces) running the latest CheckPoint firewall software becomes the bottleneck of a randomly spoofed source address SYN attack at 2.5 Mb/s. 2.5 Mb/s of attack traffic causes the firewall to drop enough legitimate traffic to cause 50% of the legitimate Web site accesses to fail. We have also found that a very popular Cisco edge router (7200VXR) becomes the bottleneck at around 10Mb/s.

These measurements point out that typical edge equipment is very susceptible to flooding attacks and that most flooding attacks can be mounted with relatively little effort and resources.

Our Solution

All flooding attacks are essentially equivalent in nature and can generally be handled by a single flooding protection system that can reduce the amount of bad traffic hitting the victim. This is true regardless of what type of flooding packet is used. Such a bandwidth protection mechanism should combat the distributed nature of flooding attacks and filter packets on multiple links leading to the victim. It should also implement filtering strategies that take into account the fact that malicious packets, in general, cannot be distinguished from normal packets solely on the contents of the packet headers. Finally, it should not heavily depend on network operators to perform quick and sophisticated network reconfiguration operations.

Reactive Network Solutions' FloodGuard system addresses these issues. It is based on the concept of adding intelligence into the network by using adjunct (not in line), cost-effective devices. FloodGuard's ultimate goal is to detect and mitigate all types of flooding attacks directed at a victim network while the attacks are in progress and provide for continued access to the network services during attacks without adversely affecting the normal functioning, performance and security of the network at other times. FloodGuard deals with a diverse set of flooding attacks including:

- Centralized DoS, DDoS and DRDoS attacks
- Fixed and randomly spoofed sources
- Direct or indirect attacks
- TCP-, UDP-, ICMP-based attacks

The placement of FloodGuard's detection and mitigation devices depends on the particular network topology and mitigation requirements (as discussed in the previous sections). FloodGuard devices can cover a wide range of applications as they can be deployed both within the local area network (LAN) or intranet of organizations or the IP backbone networks as they feature a comprehensive set of network interfaces. This wide applicability of the technology also promotes incremental deployment where the devices are initially targeted to offer local mitigation (at the enterprise or datacenter level) and later reused for more global mitigation solutions which include upstream peering points.

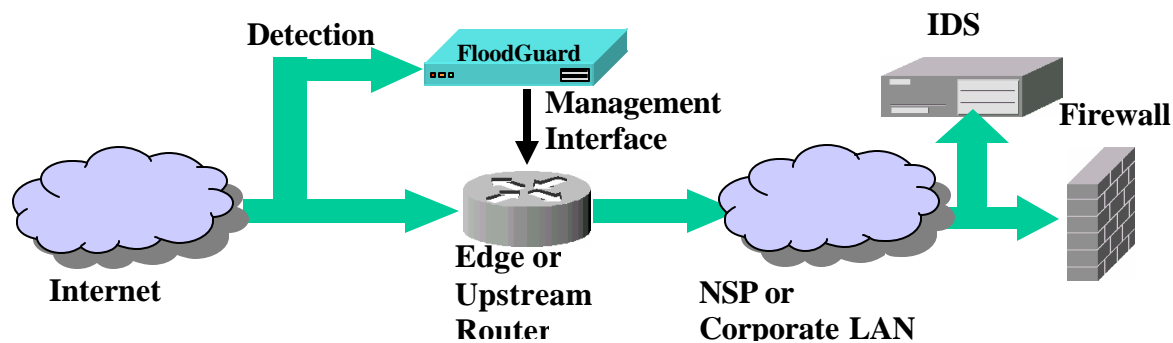


Figure 1: FloodGuard Device Deployment Architecture

As shown in Figure 1, FloodGuard devices feature at least one passive interface and one management interface. The passive interface is not addressable and is used solely for monitoring the production network's traffic. FloodGuard uses a Fast Ethernet management interface to transmit diagnosis and mitigation messages, control routers of the IP network, and inter-operate with the network management infrastructure. Figure 1 also shows the coexistence of FloodGuard with intrusion detection systems (IDS) and firewalls. FloodGuard does not overlap with the functionality of these devices but rather enhances their behavior and fault tolerance during flooding attacks.



Anomaly-based detection

FloodGuard detects anomalies in the traffic patterns by performing fine-grain analysis of the cross-product of packet inter-arrival times and IP address distributions, detecting significant changes to the temporal-spatial properties of the traffic stream. Mitigation is based on two algorithms that differentiate good and bad traffic sources through historical analysis of address patterns and measurement and enforcement of appropriate flow-control behavior (TCP back-off).

As with any anomaly-based detection, false positives are potentially a problem. Since FloodGuard implements a two-step analysis process to diagnose flooding attacks, false positives are extremely low even in extremely volatile environments. We have tested our detection accuracy on a very wide range of loading conditions reproduced using realistic behavior derived by replaying real HTTP logs and numerous field tests. We have tuned our system such that the impact to end-to-end performance and management overhead due to false positives is practically irrelevant.

Distributed FloodGuard Architecture

The FloodGuard system positions detectors and actuators near the network's most vulnerable points. In the event of an attack, detectors communicate with actuators immediately upstream, while actuators communicate with other upstream actuators. As a result, detectors and actuators are able to generate independent analyses and make autonomous decisions based on shared intelligence.

At the local level, FloodGuard protects enterprises against flooding attacks that may not affect upstream networks. For service providers, FloodGuard prevents malicious traffic from ever reaching their customers' servers, as shown in the following example of a FloodGuard deployment.

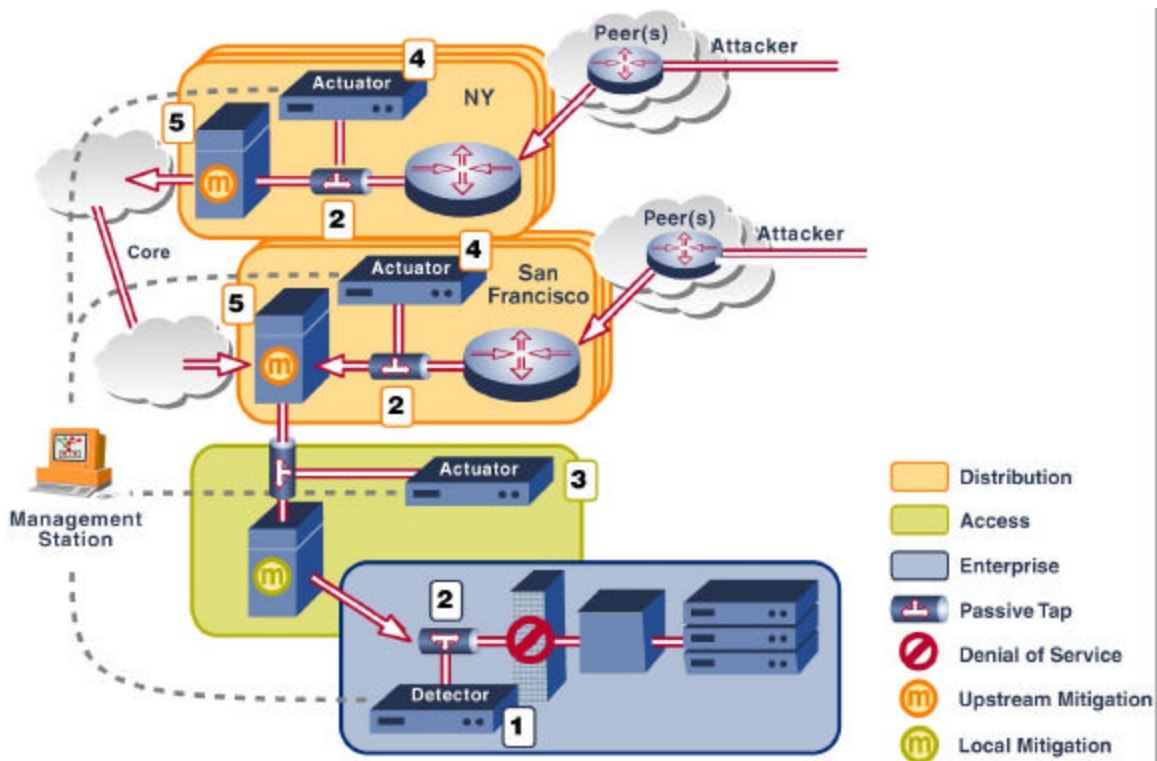


Figure 2: FloodGuard Architecture

Detectors reside near firewalls, edge routers, or routing switches, rather than inline where they might interfere with the network. Figure 2 shows a detector (1) located near the firewall of one of the service provider's enterprise customers.

The detector continually monitors traffic via a passive tap (2) or mirroring port, looking for anomalies in traffic behavior based on its traffic profile. The detector builds traffic profiles by observing characteristics of normal traffic, including packet volume, rate, type, source, and destination.



If a flooding attack is underway, the detector sends an alarm to the actuator positioned here near routers directly upstream of the enterprise (3), though actuators can also reside remotely in the distribution layer (4). The alarm includes information about the attack source, volume of traffic, and packet types. Detectors continuously send ALARMS upstream as long as the flooding condition persists. Loss of some of the ALARM messages is tolerable as long as some of the messages can reach the upstream actuators.

Like detectors, actuators continually monitor and analyze incoming traffic characteristics based on two patent-pending algorithms. When an actuator receives the attack alarm, it compares the alarm information to its own traffic profiles to determine if its segment of the network is threatened.

If the actuator's own segment of the network is threatened, it forwards the alarm to any upstream actuators and sends a trace back message to the detector confirming the threat and recommending a set of filtering rules to mitigate the attack upstream (5). FloodGuard offers three mitigation modes: automatic filtering, user-approved filtering, and user-controlled filtering.

Depending on configuration policies the actuators actions results in one or more of the following:

- Logging
- Forwarding of the ALARM to other upstream actuators
- Suggest to operators appropriate filtering actions
- Prompt operators to authorize specific filtering actions
- Install a rule set (deny rule followed by a set of zero or more permit rules) on the router associated with the actuator
- Send traceback messages to the detector from which the *ALARM* originated

The process is recursive and each actuator will either perform one or more of the above actions or will simply ignore the request depending on the texture and volume of traffic it sees on its interfaces.

FloodGuard is designed to tolerate failures of individual actuators and still provide partial functionality. Failure of the detector itself while deactivating FloodGuard will have no adverse effect on the protected services or the associated network as its placement is adjunct instead of on the critical path (like routers, firewalls and proxy caches). Due to its distributed architecture, FloodGuard can tolerate high failure rates in the communication between components during attack mitigation and when reporting information to the management system.

Conclusion

Given the fundamental characteristics of the Internet, DoS attacks can easily be perpetrated, causing dramatic reduction in availability of service. FloodGuard detects and mitigates even the most serious types of Internet DoS flooding attacks in a framework that is consistent with the scalability and fault tolerant architectural principles of the Internet.