

Defenses against Distributed Denial of Service Attacks

Adrian Perrig, Dawn Song, Avi Yaar
CMU



Carnegie Mellon CyLab
4615 HENRY STREET
PITTSBURGH, PA 15213
PH: (412) 268-7100
FX: (412) 268-7108
WWW.CYLAB.CMU.EDU


Internet Threat: DDoS Attacks

- Denial of Service (DoS) attack: consumption (exhaustion) of resources to deny access to others
- Distributed Denial of Service (DDoS) attack is a coordinated DoS with many attackers
- Homogeneity of computing systems enables an attacker to compromise tens of thousands of hosts
- **DDoS attacks pose a significant threat!**

Characteristic of a DDoS Attack

- Victim is flooded with packets
 - If attacker uses IP spoofing, victim does not know which packets to serve
 - If attacker uses link flooding, legitimate connections experience high loss rates and become unusable (similar effect to flash crowds)

- **Fact: during DDoS attacks, no legitimate clients can be served**



DDoS Challenges

- **Challenge 1: IP source address spoofing**
 - Attackers hide their origin by spoofing IP source address
 - Victim cannot filter out spoofed IP packets, wastes resources
 - Goal: identify spoofed packets
- **Challenge 2: Link flooding**
 - Attacker floods victim network
 - Victim's legitimate traffic gets dropped
 - Goal: receiver can control who can send packets to it

Outline


- **Challenge 1: IP source address spoofing**
 - **Pi**: first approach to identify IP-spoofing for every packet
 - *IEEE Security & Privacy Symposium 2003*
- **Challenge 2: Link flooding**
 - **SIFF**: first stateless approach to enable routers drop attack packets in network
 - *IEEE Security & Privacy Symposium 2004*

Other Defense Approaches

- **Traceback**
 - Reverse link flooding
 - Out-of-band tracing (iTrace)
 - In-band tracing (IP Traceback)
 - Router state (SPIE)
- **Traffic Filtering (Pushback)**
- **DNS Rerouting**
- **Ingress Filtering**
- **Detection of Spoofed IP address (Hopcount filter)**
- **Overprovisioning/replication (DNS root servers, Akamai)**
- **In-network detection and defense approaches (DWard)**

Pi: Packet Marking and Filtering Mechanisms for DDoS and IP Spoofing Defense

- **Basic Premise:**
 - Path “fingerprints”
 - Entire fingerprint in each packet
 - Incrementally constructed by routers along path
- **Detect spoofing by observing discrepancy between IP address and path fingerprint**



Pi: System Overview

- **Two phases**
 - Pi marking
 - Stack marking
 - Write-ahead marking
 - Pi filtering
 - $\langle \text{Pi}, \text{IP} \rangle$ filtering
 - Basic filter
 - Threshold filtering

Pi Marking: Description

- **Marking Scheme**

- Routers push n bits into the MSB position of the IP Identification field


- **Marking Function**

- Last n bits from the hash of the IP address of the current router concatenated with the IP address of the previous-hop router
[ie. $MD5(IP_i || IP_{i-1}) \bmod 2^n$]

- **Why not just $MD5(IP_i)$?**


- Probability of collision in “fan-in” scenario is better with previous hop information

Pi Marking



Pi Marking: Write-Ahead Improvement

- Problem: legacy routers do not mark



- Solution:

- Every router marks on-behalf of the next-hop router in the path. (Write-Ahead)
- Every router checks to see if its markings have already been added to the packet.

Pi-IP Filter

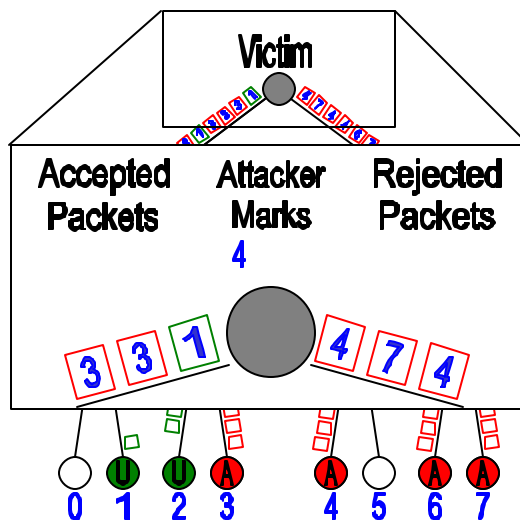
- Victim stores mapping of Pi mark and corresponding IP networks
- If victim is under attack, use Pi-IP table as indicator of which source addresses are not spoofed
- Experiment: over 99.9% of all spoofed source address packets rejected

Using Pi to Filter DDoS Attacks

- Pi identifies path a packet traversed
- Observation: attack packets always arrive with same Pi mark
- Idea: only serve packets with a Pi mark that has a high probability to be from legitimate host
- Assumption: mechanism to detect malicious packets
- Pi mark offers a probability that incoming packet is a legitimate packet (similar: RFC 3514 Evil Bit)
- Goal: instead of serving <1% of clients, serve 60% of clients

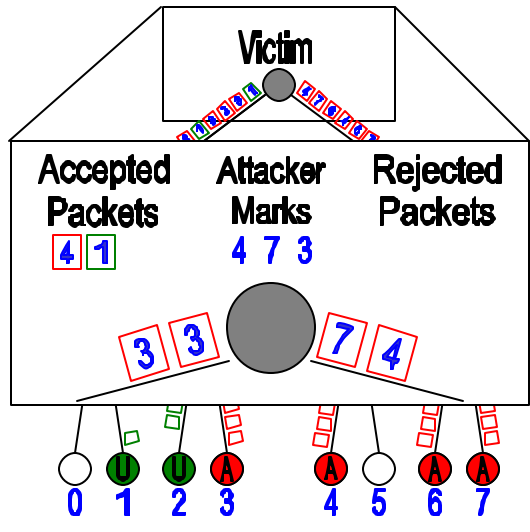
Basic Filter Example

- Pi mark used to identify traffic originating from malicious origins
- Victim rejects packets with Pi marks that carry attack traffic




Basic Filter Example

- Pi mark used to identify traffic originating from malicious origins
- Victim rejects packets with Pi marks that carry attack traffic

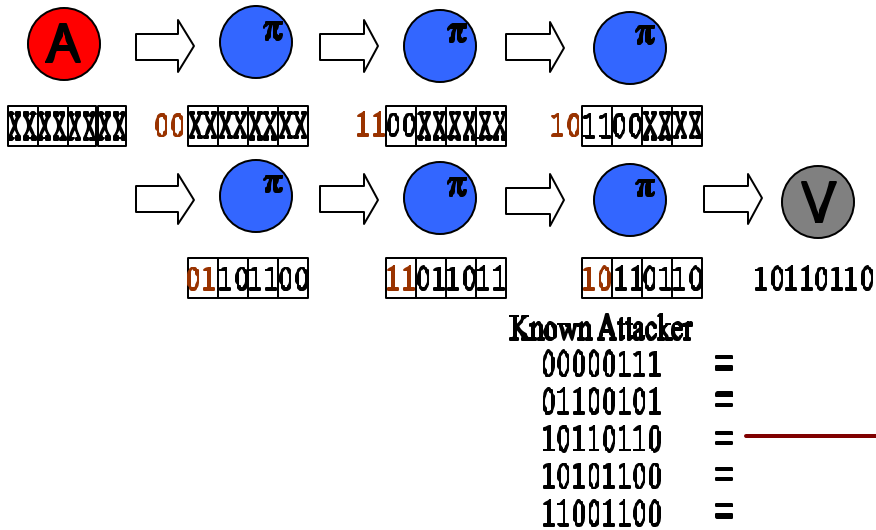


Basic Filter Example

- Pi mark used to identify traffic originating from malicious origins
- Victim rejects packets with Pi marks that carry attack traffic



Pi Marking and Basic Filter

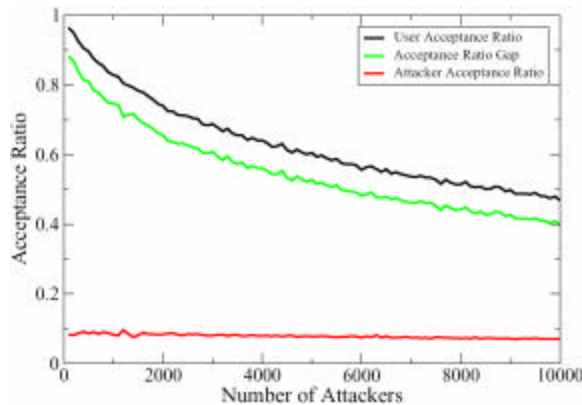


Pi Performance: Experimental Setup

- **Two Internet Topologies**
 - **Internet Map Project**
 - 81,953 unique endhosts
 - **CAIDA Skitter Map (f-root DNS server)**
 - 171,472 unique endhosts
- **5,000 legitimate users, 100-10,000 attackers**
- **$n = 2$ bits**
- **4 router non-marking ISP perimeter**
 - **Victim ISP marks unnecessary/undesirable**

Pi Performance: Basic Filter

- Pi Works!
- At 10,000 attackers
 - Over 45% legitimate users' packets accepted
 - Over 90% attackers' packets dropped



Pi Filtering - Thresholds

- Problem
 - Single attacker packet may cause multiple users' rejections
- Solution
 - Assume, for a particular Pi mark, i :
 - a_i = number of attack packets
 - u_i = number of legitimate users' packets
 - Victim chooses threshold, t , such that if:

$$t < \frac{u_i}{a_i + u_i}$$

then all packets with Pi mark i are dropped

Pi Performance: Threshold Filters

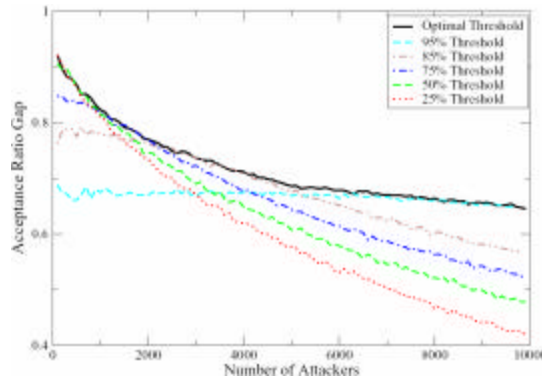
- Observations?
 - Increased attack severity requires increased threshold.

- Optimal threshold value

$$t_{opt} = \frac{P_U}{P_A + P_U}$$

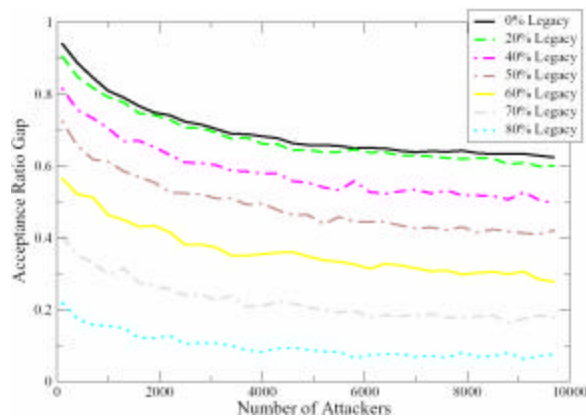
P_U – Total user pkts

P_A – Total attack pkts



Pi Performance: Legacy Routers

- Pi is robust to the presence of legacy routers
- Benefits even when only 20% of routers implement Pi



Pi Summary

- **Pi provides DDoS protection**
 - After first identified attack packet
 - Extremely low overhead at routers & endhosts
 - Does not interfere with IP Fragmentation
 - No need for inter-ISP cooperation
 - Great incremental deployment properties and incentives for deployment
 - First approach to enable per packet DDoS filtering

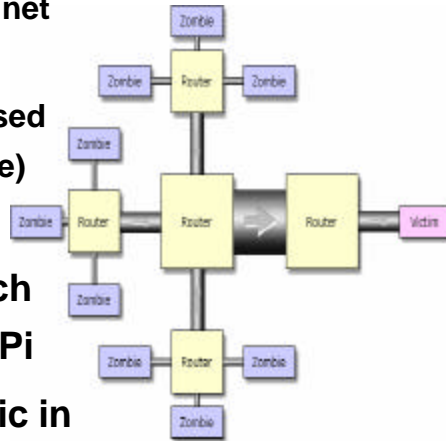
Outline

- **Challenge 1: IP source address spoofing**
 - **Pi**: first solution to identify IP-spoofing for every packet
 - *IEEE Security & Privacy Symposium 2003*
- **Challenge 2: Link flooding**
 - **SIFF**: first stateless solution to enable routers drop attack packets in network
 - *IEEE Security & Privacy Symposium 2004*

Problem: Bandwidth DoS

- What about DDoS attacks?
 - CERT reports 140,000+ bot net
 - Mydoom worm had 100,000-600,000 compromised hosts (depending on source)

- Too distributed for effective filtering, too much bandwidth disruption for Pi
- Goal: filter malicious traffic in core of network without per-flow state



Fundamental Problem

- DDoS attacks exploit fundamental problem: receiver has no control over who can send traffic to it
- We need to enable receiver to stop misbehaving senders
 - Adkins, Lakshminarayanan, Perrig, and Stoica: "Taming IP Packet Flooding Attacks", HotNets 2003
- Challenges
 - Need per flow state in network?
 - Where to filter?
 - Need trust relationships between ISPs?
 - Routers need to authenticate receiver requests to stop flows?

SIFF: A Stateless Internet Flow Filter to Mitigate DDoS Flooding Attacks

- **Goal: enable receiver to control its traffic**
- **Key ideas**
 - **Use Pi fingerprints as authorization to send traffic**
 - Pi fingerprint is used as a capability
 - Only clients who know their Pi mark get authorization
 - **Authorized or “privileged” packets get priority over non-privileged packets**
 - In bandwidth DoS, privileged packets are undisturbed by non-privileged packets

SIFF Overview


- **Sufficient space in packet header**
 - **Marking Field (128 bits)**
 - **Flags Field (3 bits)**
 - **[Optional] Update Field (128 bits)**
- **Packet classes**
 - **Unprivileged or best effort: signaling traffic / legacy traffic**
 - **Privileged: data traffic, displaces unprivileged traffic**

SIFF Properties

- **DoS-less client/server communication**
 - Packet receiver can stop flow if it consumes local network resources
- **Limited spoofing**
 - Equivalent to universal Ingress Filtering
- **Lightweight at routers**
 - Small constant state/processing per packet
- **Incremental deployment / backward compatible**
- **No trust required between ISPs, no authentication required at routers, only sender must trust receiver to receive correct capability**


SIFF Marking: Unprivileged Packets

- SIFF routers mark best effort packets traveling from sender to receiver
- Each router inserts 2 bits into marking field
- First router inserts 1-bit to signal length of marking to receiver



SIFF Marking: Privileged Packets

- SIFF routers check their marking before they forward packets
- Packets whose marking does not match are discarded
- Marking acts as a capability to forward packet from sender to receiver




Privileged Flow Setup

- Receiver receives marking with initial best effort packet, can send marking to client for use as a capability for privileged flows
- Advantage: only senders with capability can send privileged traffic
- Disadvantage: receiver cannot stop any privileged flows
- Approach: routers use changing keys to compute packet marking


SIFF Key-based Marking

- Each router has a secret key K , computes marking as $H_K(\text{currIP} \mid \text{prevIP} \mid \text{sourceIP} \mid \text{destIP})$
- On each router, matching marking enables packet to pass



SIFF Marking: Multiple Keys

- Each router has two keys: an old key and a new key
- On each router, either old marking or new marking enables packet to pass
- Router inserts marking with new key into packets



Receiver-controlled Flows

- As packet flow carries on, receiver receives updated markings
- If receiver wants to continue to enable sender to send privileged traffic, receiver sends updated marking as capability to sender
- If receiver wants to terminate malicious flow, receiver simply stops updating sender with new capability, and routers will soon stop the flow early in network

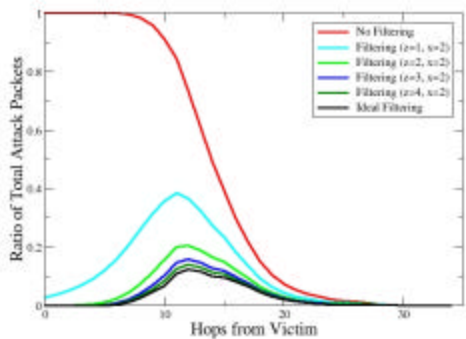
SIFF Performance I

- Three parameters for SIFF
 - z = number of bits per router mark
 - x = number of marks in router's window
 - T_K = time between router key changes
- DDoS 1: Attackers flood unprivileged traffic
 - Result: Privileged users unaffected, "1-packet" privileged users


SIFF Performance II

- DDoS 2: Attackers flood “forged” privileged traffic
 - Probability of fooling a SIFF router:

$$P(x,z) = 1 - (1 - 1/2^z)^x$$
 - Probability of fooling d SIFF routers: $P(x,z)^d$



Filter effect at distance x



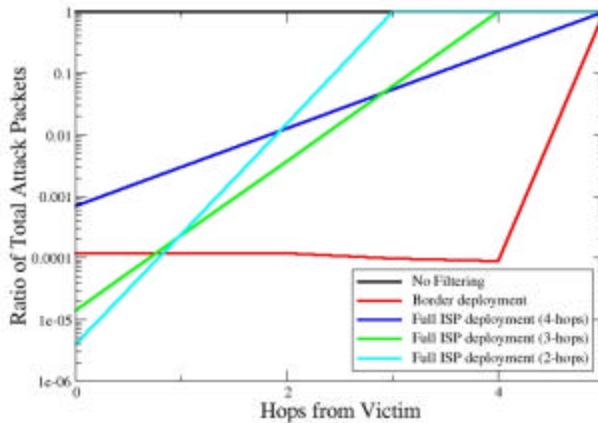
Attacker distribution

SIFF IPv4 Design

- Field Conversions:
 - Flags (3) → One reserved bit (set to 0 by legacy hosts)
 - Marking (128) → IPv4 ID field (16)? (Not Quite)
- Assumption
 - ISPs networks get DoS attacks (not backbone)
- Proposal
 - Assume x bits available for marking, half for capability, half for capability update
 - Only victim's ISP marks packets
 - Full ISP deployment: (approx) 3 hops @ $x/8$ bits each
 - Border ISP deployment: 1 hop @ $x/2$ bits

IPv4 Performance

- Packets drop immediately at ISPs “border” (3-hops)



SIFF Issues

- **ISP Border Deployment**
 - ICMP errors may reveal capability
 - Slight modification ICMP
 - Flooding may occur at non-SIFF router
- **Path Stability**
 - Route changes invalidate capabilities
 - Demote privileged to unprivileged when verification fails

Discussion: Deployment Incentives

- **Lack of incentive for ingress filtering**
- **Pi provides incentive for ISP**
 - Customers benefit from Pi marking
- **Attackers within ISP cause blocking of other ISP customers**
 - ISP has incentive to block attack
 - Incentives for ingress filtering
- **Market pressures drive Pi & SIFF deployment**
 - Large-scale Internet sites > ISP > router manufacturer

Related Work

- **Gligor: “A note on the denial of service problem”, IEEE S&P '83.**
- **Gligor: “Guaranteeing access in spite of service-flooding attacks”, SPW 2003.**
- **Adkins, Lakshminarayanan, Perrig, and Stoica: “Taming IP Packet Flooding Attacks”, HotNets 2003.**
- **Anderson, Roscoe, Wetherall: “Preventing Internet Denial-of-Service with Capabilities”, HotNets 2003.**

Conclusions

- Changing end-hosts and routers seems necessary to defend against DDoS threat
- Pi enables detection of spoofed IP address
- SIFF mitigates flooding attacks
 - End-host identifies attack flows
 - No per-packet state at routers!
- Strong deployment incentives for ISPs!