

Security Topics in Networking and Distributed Systems

CS 590D

Lecture 5

Department of Computer Sciences
Purdue University

Outline

- Practical Network Support for IP Traceback. Stefan Savage, David Wetherall, Anna Karlin and Tom Anderson. SIGCOMM 2000.
- Advanced and Authenticated Marking Schemes for IP Traceback. Dawn X. Song, Adrian Perrig. Proceedings IEEE Infocomm 2001



Incidents Reported WWW.CERT.ORG

1988-1989

1988 1989

6 132

1990-1999

1990	1991	1992	1993	1994	1995	1996	1997	1998	1999
252	406	773	1,334	2,340	2,412	2,573	2,134	3,734	9,859

2000-2003

2000 2001 2002 2003

21,756 52,658 82,094 137,529

Total incidents reported (1988-2003): 319,992

Please note that an incident may involve one site or hundreds (or even thousands) of sites. Also, some incidents may involve ongoing activity for long periods of time.

How to Defend Against Dos?

- Victim network
 - Intermediate network
 - Source network
 - Attack is **observed close to victim**
 - Attack **must be stopped close to the source**
- Intermediate network used to **“traceback”** the attack
- **Reactive** (after the attack) and **proactive** (prevents the attack) methods



Fighting DDoS (at 2000.)

- Ingress filtering
- Link testing
 - Input debugging
 - Controlled flooding
- Logging
- ICMP Traceback



Ingress Filtering

- Packets coming from “bad” IP addresses are stopped by the routers
- Requires to distinguish between legitimate and illegitimate addresses
- Improves robustness to Dos/DDos
- To be effective, it must be widely deployed



Link Testing

- **Technique**: start from the closest router to the victim and test the upstream links to figure out in which one is the attack coming
- **Requires**: the attack must be active; can not be done post-mortem
- **Input debugging**
- **Controlled flooding**



Input Debugging

- **Technique:**
 - victim realizes that it is attacked and identifies a common feature in the attack packets (attack signature)
 - Victim talks with network operator over the phone, the operator installs a filter on the victim's upstream egress port
 - repeat
- Trace performed manually or automatically
- **Requires:**
 - multiple ISP collaboration
 - Management overhead high
 - ISP sometimes not motivated (no direct economic incentive) in providing assistance

Controlled Flooding

- **Technique:** traceback without requiring assistance from network operator
- Tests the link by flooding them with traffic and observing changes in the rate of the packets received from the attacker; repeat...
- **Requires:** map of Internet topology
- The scheme is itself a denial of service
- Good topological map of the Internet
- How to distinguish between multiple links, if attack comes on multiple links

Logging

- **Technique:** log packets at key routers and use data mining to determine the path the packets traveled
- **Requires:** high resources and integration of large scale inter-provider database
- **Can be performed post-mortem**

ICMP Traceback

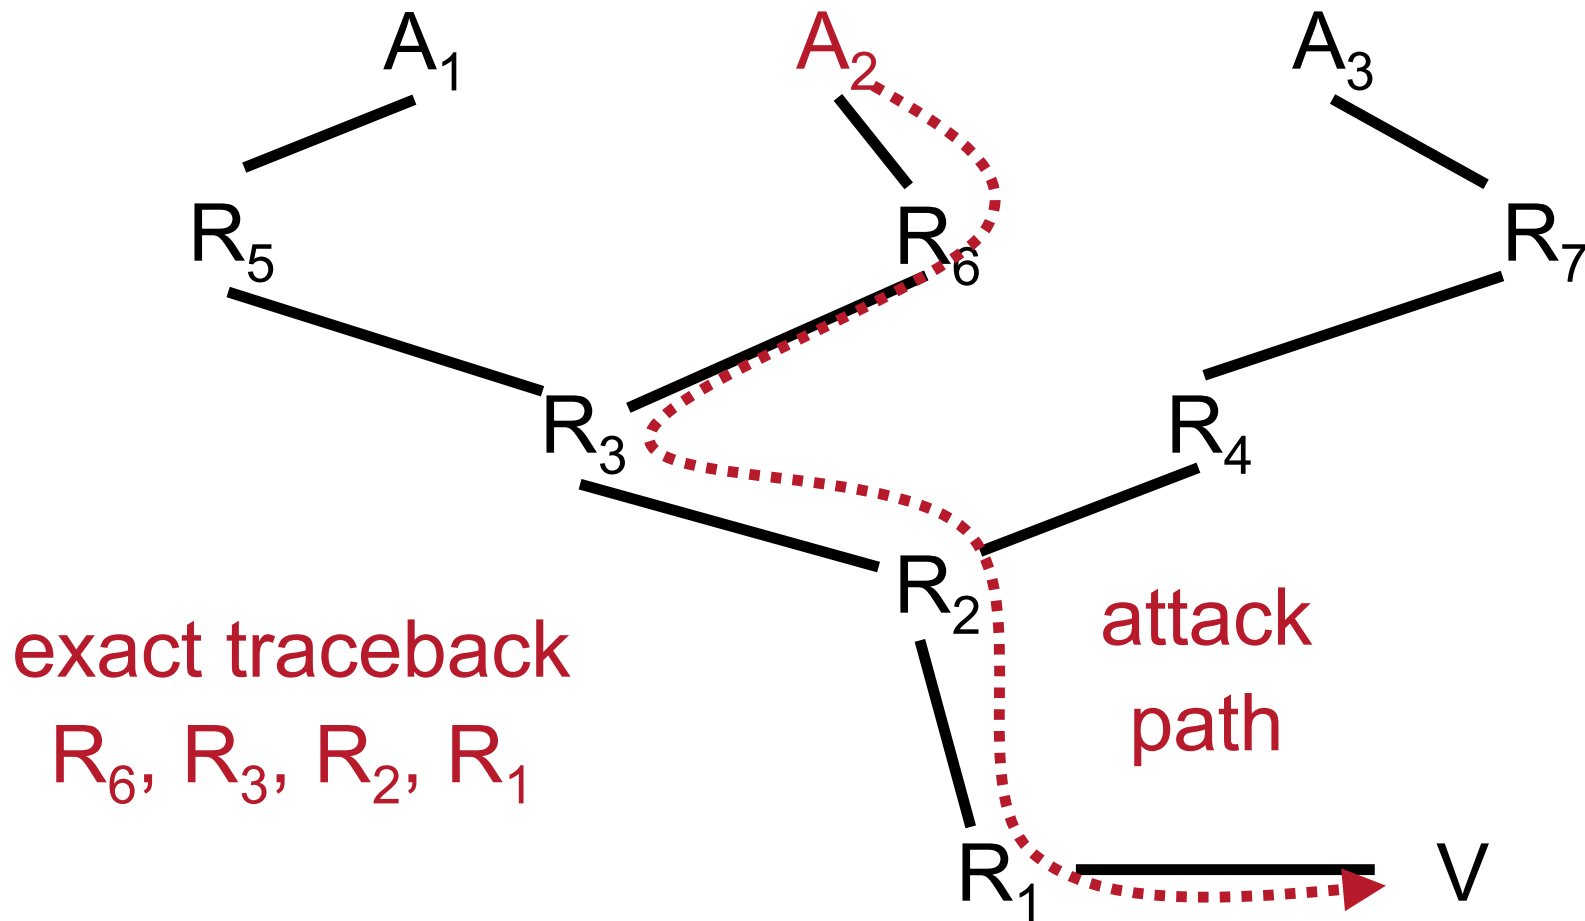
- **Technique:**
 - traceback the attack by using of router-generated ICMP messages;
 - with a small probability routers copy the headers of the packet it was forwarding to the header of a ICMP packet, including information about adjacent routers on the path
 - Victim can use these router-generated ICMP packets to reconstruct the path
- **Drawback:**
 - ICMP packets can be filtered out or rate-limited at routers
 - Relies on input debugging information not available in all architectures
 - Requires key distribution for authentication of ICMP packets

Probabilistic Marking


- **Technique:**
 - Routers mark the packets with the path, probabilistically
 - Victim uses information to reconstruct the path
 - Convergence time: # of packets to reconstruct the attack path
- **Advantages:**
 - Require no cooperation with ISPs
 - Does not cause heavy network overhead
 - Can trace attack “post mortem”



Exact Traceback Problem



Approximate Traceback Problem



approx. traceback

R_5, R_6, R_3, R_2, R_1

Assumptions

- DDoS Assumptions
 - Attacker may generate any packet
 - Multiple attackers may conspire
 - Attackers may be aware they are being traced
 - Packets may be lost or reordered
- Design assumptions
 - Attackers send numerous packets
 - Route between attacker and victim is fairly stable
 - Routers have limited CPU and memory
 - Routers are not widely compromised

Marking Algorithms

- Node Append
- Node Sampling
- Edge Sampling

Node Append

- **Technique:** Append address of each node to the end of the packet
- **Result:** when delivered, packet contains the complete, ordered list of routers attack path
- **Convergence very fast**
- **Requires:** high overhead, not-know how much space is required

Node Sampling

- Router writes its address in node field with a probability p
- Reconstruct path using a number of packets
- Low overhead, only one additional write
- Robust against single attacker when $p > 0.5$
- d is # hops and p identical at each router, then probability that victim receives a packet from the furthest router is

$$p(1-p)^{d-1}$$

Node Sampling

- What is the convergence?
- If $d = 15$ and $p = 0.51$, the receiver requires **42,000 packets to guarantee that it received a sample from the furthest router**
- To guarantee that the order is correct with 95%, $7 \cdot 42000$ packets required
- Technique is not robust against multiple attackers

Edge Sampling

- Store edges instead of nodes
 - start and end addresses of edge routers
 - distance from edge to victim
- Marking algorithm:
 - router writes its address in the *start* field, and 0 into the distance field
 - If distance field is zero, the packet is already marked, router writes its address in the end address field and increases the distance field by 1


Edge Sampling (cont.)

- Convergence: a smaller number of packets, converges quicker than node sampling
- Can distinguish multiple attacks
- Requires: additional space in the IP header
72 bits of in IP packet: 2×32 (IP address) + 8 bit(distance)

Encoding for Edge Sampling

- Requires 72 bits of in IP packet: 2 x 32 (IP address) +8 bit(distance)
- Address compatibility + reduce space
- Idea:
 - use the IP identification field
 - store the XOR of the edge addresses (edge-id)
- Increases reconstruction time

Store the XOR ...



Key observation: $d \text{ XOR } (d \text{ XOR } c) = c$
Reduces 64 bits to 32 bits
Still too expensive

Reducing the space ...

- Divide the XORed data into k non-overlapping packets
- Need offset of fragment
- Fragments are not unique, with multiple attackers, multiple edge fragments with the same offset and distance
- Store both XORed data and its hash, bit-interleaved

Other Issues

- Backwards compatibility
- Distributed attacks: probability of misattributing an attack and the amount of state needed increases with the fan-out of the attack: days of computation + false positives, for only 25 attackers
- Suffix validation: Attacker can spoof end edges
- Attack origin detection
- Traceback will solve the problem of finding the host used to attack by not the attackers

Advanced Marking Schemes

- If map of upstream routers are known (www.caida.org), all is needed is to mark the last router
- Less false positives
- Encoding: 11 bit for the hash of the IP address of the router and 5 bits for the distance
- Can distinguish 50 attackers
- Improve: use two sets of independent hash functions to minimize collision

Authenticated Marking Schemes

- Packets not authenticated: the attacker can forge the markings and thus mislead the victim.
- Digital signatures: very expensive: time (10ms/sign) + 128 bytes storage
- Use MACs: each router shares secret keys with the victim; Key management complex. Scheme impractical
- Use time-released keys;


Time-Release Keys

- Router selects a random key K_N then sing a cryptographic hash function (such as MD5) and K_N , the router creates chain of keys K_0, K_1, \dots, K_{N-1} .

$$\text{hash}(K_{j+1}) = K_j$$

- The hash ensures function that anybody can compute keys forward but nobody can compute keys backward
- Keys will be used in order K_0, K_1, \dots, K_{N-1} , and released in time

Next Lecture



- Router-Based Defense Against DDoS
 - On the Effectiveness of Route-Based Packet Filtering for Distributed DoS Attack Prevention in Power-Law Internets. Kihong Park, Heejo Lee. SIGCOMM 2001
 - Implementing Pushback: Router-Based Defense Against DDoS Attacks. John Ioannidis, Steven M. Bellovin. NDSS 2002.