UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Offensive technologies
Fall 2017

Denial of Service
Fabio Massacci
(Some slides courtesy of Jelena Mirkovic)

https://securitylab.disi.unitn.it/doku.php?id=course_on_offensive_technologies

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Denial of Service Attacks

- ***The goal is to stop the service from operating***
  - To deny service to legitimate users
  - Slowing down may be good enough
- ***Temporary effect → passes as soon as attack stops***
  - If this was a controller of a physical device this might be an extremely damaging effect
- ***How it works?***
  - Machine-based
    - Crash, put infinite loop, use lots of resources
  - Network-based
    - Crash routers on path to it, deny another service needed by it (e.g. DNS), use lots of network resources
- ***Typically use lots of resources and the cooperation of many machines (wittingly or unwittingly)***

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Is DoS a Real Problem?

- ***Yes, attacks happen every day***
  - One 2002 study reported ~4,000 per week[1]
- ***On a wide variety of targets***
- ***Tend to be highly successful***
  - 2009 Twitter
  - 2010 Visa & Mastercard
    - Anonymous as they stopped accepting payment on Wikileakes
  - 2012 BofA, Chase and Wells Fargo
    - Muslim group Izz ad-Din al Qassam Cyber Fighters
  - 2013 SpamHaus (and service provider CloudFlare)
  - 2015 BBC News (as well as Trump Web Site)
- ***Basically only stoppable by cooperation of ISPs***

[1]"Inferring Internet Denial of Service Activity," Moore, Voelker, and Savage, Usenix Security Symposium, 2002

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

DoS as the "Normality" of Internet

- ***The Internet was designed to deliver → lots of traffic from lots of places to lots of places***
  - Best effort service means routers don't do much processing per packet and store no state – they will let anything through
  - End to end paradigm means routers will enforce little security or authentication – they will let anything through
- ***It works well when all parties play fair***
- ***DoS is just one party who want to deliver lots of traffic from lots of places to one place***
  - Any individual packet can look proper to the Internet
  - Without sophisticated analysis, even entire flow looks legit
    - Example crash of the French Tax Web Server during the closing day of the first online Tax Submission

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Who Is Vulnerable?

- ***Everyone connected to the Internet can be attacked and can suffer damages***
- ***Your Machines are Secured***
  - yes but the bots are on somebody's else machine
    - Example of the Tragedy of the Commons
- ***You Have a firewall***
  - Attackers attack the firewall
- ***You have a VPN***
  - They fill the VPN with garbage (you'll eventually discard but have no resource for doing anything else)
- ***You are highly provisioned***
  - See Krebs attack from Mirai who brought down Akamai

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Simple Denial of Service

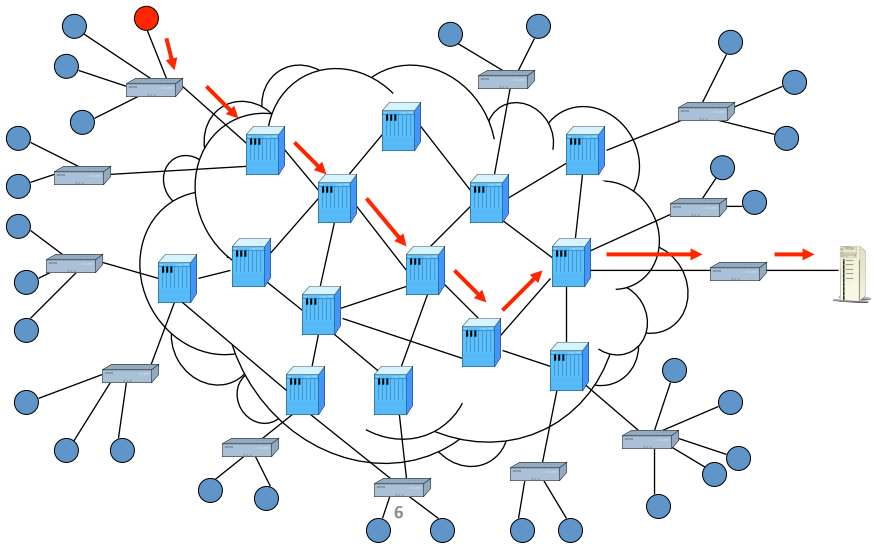

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

1-on-1 Denial of Service

- ***One machine tries to bring down another machine***
- ***Can it work?:***
  - Attacker must be "more powerful" than victim
  - OR there must be some asymmetry in the communication
- ***Asymmetry is key → Amplification Effects***
  - crafting a request is cheaper than composing the response
    - e.g. sending a bogus packet vs decrypting the packet and checking whether t it's bogus
  - Formulating a response requires keeping track of history
    - E.g. sending many bogus packets vs keeping track of all received packets from allegeddly many senders
- ***If so, one attack machine can generate a lot of requests, and effectively multiply its power if the responses are disproportionate***

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

DDoS – Distributed DoS

- ***Use multiple machines to generate the workload***
  - For any server of given power, enough attackers working together can overload it
  - Enlist lots of machines and coordinate their attack on victim
- ***Wittingly → lots of vulnerable machines***
  - Few gazillions machines (typically compromised) send data to victim → called a bot net
- ***Unwittingly → flawed protocol***
  - Send few bogus requests to some machines, protocol respond back with back gazillions of data to victim

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Distributed Denial-of-Service

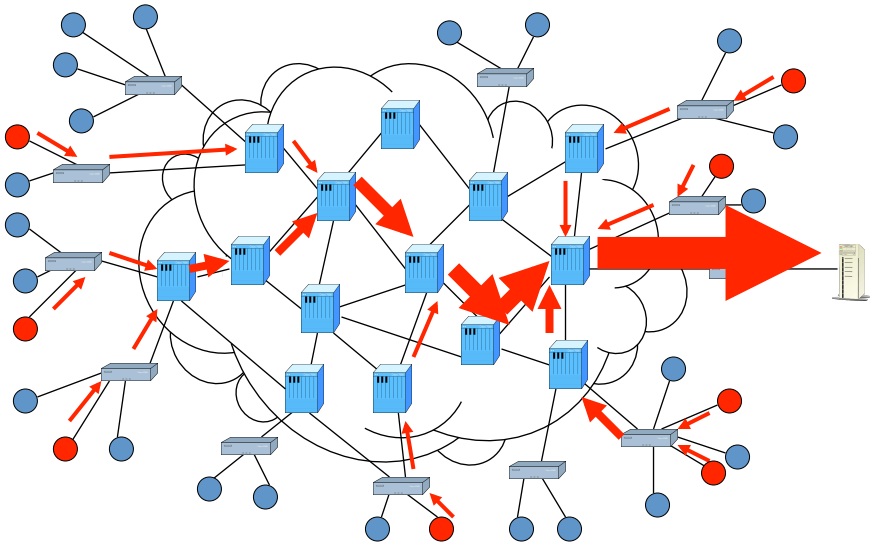

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

DoS - TCP SYN Flood

- ***Attacker sends lots of TCP SYN packets***
  - Victim sends an ack, allocates space in memory
  - Attacker never replies
  - Goal is to fill up memory before entries time out and get deleted
- ***Usually spoofed traffic***
  - Otherwise patterns may be used for filtering
  - OS at the attacker or spoofed address may send RST and free up memory

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

TCP SYN Cookies

- ***Effective defense against TCP SYN flood***
  - Victim encodes connection information and time in SEQ number for the server
  - Must be hard to craft values that get encoded into the same SEQ number – use crypto for encoding
  - Memory is only reserved when final ACK comes
- ***Only the server must change***
  - But TCP options are not supported
  - And lost SYN ACKs are not repeated

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

What Is IP Spoofing

- ***Putting a fake IP address in the IP header field for source address (requires root)***

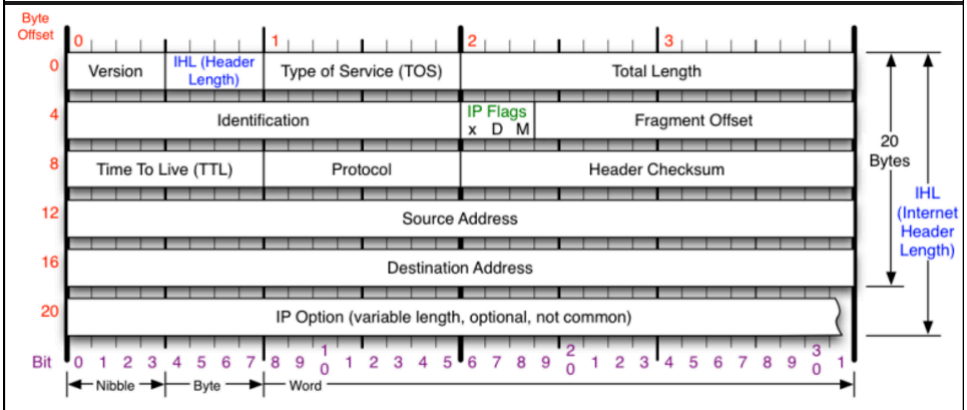

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Why Attackers Spoof?

- ***Hide their identity***
  - Put a blame on someone else
- ***Confuse the defense***
  - In DDoS, make traffic appear to come from many sources
- ***Acquire identity of a legitimate host***
  - Leverage some trust relationship (e.g., bypass a firewall)
  - Hijack a TCP connection
  - Perform DNS hijacking

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Why Defenders Spoofs

- ***Wait a minute why should defenders spoof?***
- ***Think of at least one thing that could be seen as "spoofing"***
  - Hint: you used it in a previous exercise

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

How Do You Detect/Foil Attacks?

- ***Have database of attack signatures***
- ***Detect anomalous behavior***
  - By measuring some parameters for a long time and setting a baseline
    - Detecting when their values are abnormally high
  - By defining which behavior must be obeyed starting from some protocol specification
  - It has some parameter values
  - It has certain behavior
- ***Filter Attack Streams***
  - Addresses the core of the problem by limiting data presented to target
  - Key question → what do you drop?
  - Good solutions drop all (and only) attack traffic
    - Drop everything but give priority to legitimate-looking traffic
  - Less good solutions drop some (or all) of everything
    - Devise filters that encompass most of anomalous traffic

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Filtering Vs. Rate Limiting

- ***Filtering drops packets with particular characteristics***
  - If you get the characteristics right, you do little collateral damage
  - At odds with the desire to drop all attack traffic
- ***Rate limiting drops packets on basis of amount of traffic***
  - Can thus assure target is not overwhelmed
  - But may drop some good traffic
- ***You can combine them (drop traffic for which you are sure is suspicious, rate-limit the rest) but you gain a little***

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Where Do You Filter?

In multiple places?

Near the source?

In the network core?

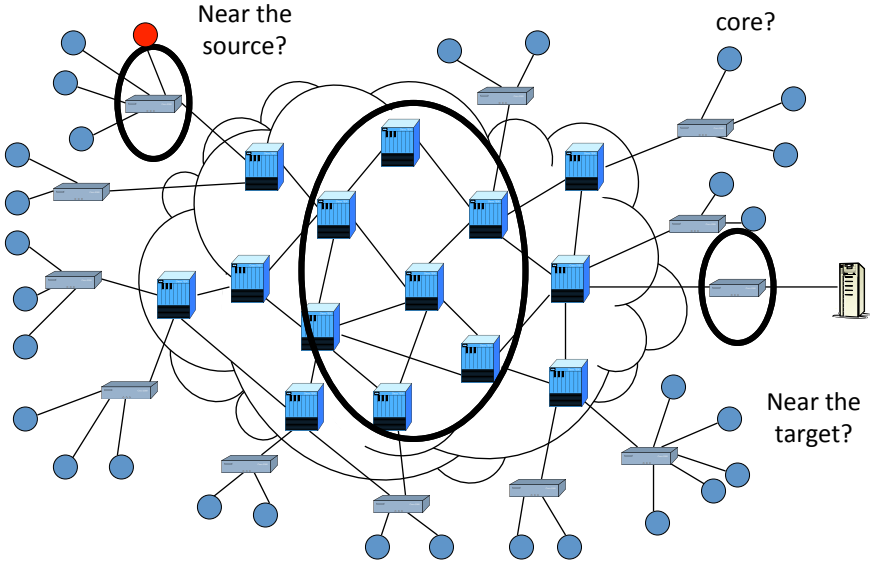

Near the target?

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Filtering Location Choices

- ***Near target***
  - Easier to detect attack → if nothing works you'll notice
  - May be hard to prevent collateral damage
  - Can't handle large attack volume
- ***Near source***
  - How do you know it's a source?
  - Easier to prevent collateral damage
- ***In core***
  - How does it know it's an attack?
  - Sees everything (with sufficient deployment)
  - Easier to handle attack volume → this is what happens at the end of the day: ask your ISP

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Ingress Filtering

- ***RFC 2827, BCP 38***
  - Collect a list of your prefixes P
  - Filter out outgoing traffic whose source IP is not from P
  - Filter out incoming traffic whose source IP is from P
- ***Sounds simple?***
  - It took routers long time to put this kind of filtering on the fast path
  - Implementation may be tricky (multihoming)
  - It helps others, not you
  - It does not completely eliminate spoofing

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Implementing Ingress Filtering

- ***ACL: Manually collect a list of your prefixes***
  - Works for edge networks but not for ISPs
    - there are way fewer ISPs (~ 6 K) than edge networks (~ 33 K) so implementing something at ISPs is faster
  - If a network is multihomed and does not update its new ISP with its prefixes it may lose traffic
- ***Strict reverse path forwarding***
  - If my next hop to P is A then only A can send me traffic from P (however lots of routes are asymmetrical between ISPs)

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Implementing Ingress Filtering

- ***Feasible reverse path forwarding***
  - Remember all advertised next hops for P, one of them is a valid previous hop
  - Works correctly but lets some spoofed packets through
- ***Loose reverse path forwarding***
  - Only drop packets if their source IP is not routable
  - Only 1/3 of the IPv4 space is routable so randomly spoofed packets would be dropped 2/3 of the time

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Ingress Filtering w Multihoming

- ***Multihoming: having 2 or more upstream ISPs***
  - For backup (but use only some)
  - For good performance (use the fastest one at the moment) or load balancing (use them equally)
  - Changing providers (temporary multihoming)
- ***A network may announce its prefixes only to one ISP (for incoming traffic) but use both for outgoing traffic or vice versa***
  - This will lead to ingress filter drops at the ISP which is used only for outgoing traffic (ACL, strict RPF)

UNIVERSITY OF TRENTO - Italy

Digital MASTER SCHOOL

Ways Around Ingress Filtering?

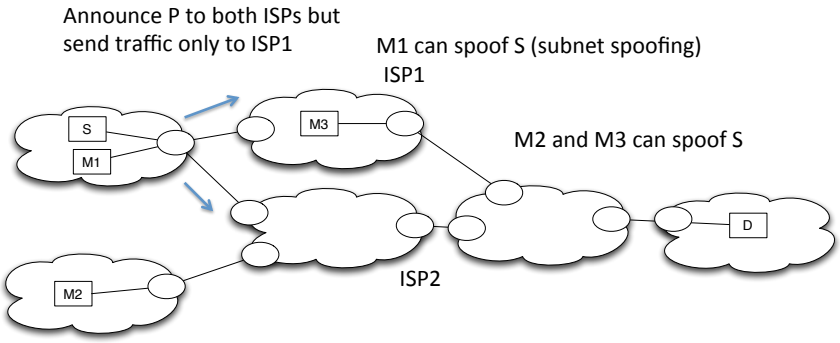

UNIVERSITY OF TRENTO - Italy

Digital MASTER SCHOOL

Poor Cooperation In the Internet

- ***It's hard to get anyone to help you stop or trace or prevent an attack***
  - Even your ISP might not be too cooperative
  - Anyone upstream of your ISP is even less likely to be cooperative
  - ISPs more likely to cooperate with each other, though
- ***Even if cooperation occurs, it occurs at human timescales***
  - The attack might be over by the time you figure out who to contact
  - Besides, how do you contact your ISP if you have been thrown off the internet?

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

DDoS on Twitter

- ***August 2009, hours-long service outage***
  - 44 million users affected
- ***At the same time Facebook, LiveJournal, YouTube and Blogger were under attack***
  - Only some users experienced an outage
- 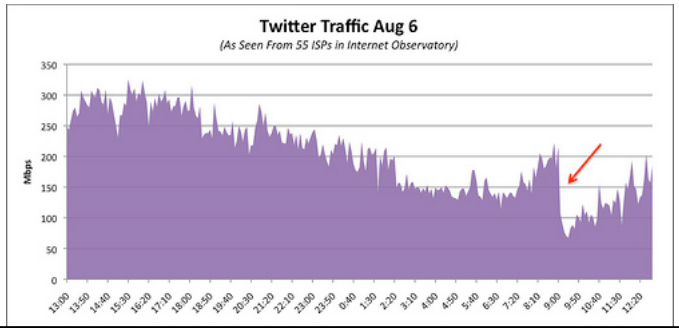

Image borrowed from Wired.com article. Originally provided by Arbor Networks

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

DDoS Attack Code

- ***Attacker can customize:***
  - Type of attack
    - UDP flood, ICMP flood, TCP SYN flood, Smurf attack (broadcast ping flood)
    - Web server request flood, authentication request flood, DNS flood
  - Victim IP address
  - Duration
  - Packet size
  - Source IP spoofing
  - Dynamics (constant rate or pulsing)
  - Communication between master and slaves

UNIVERSITY OF TRENTO - Italy

Digital MASTER SCHOOL

Typical Attack Modus Operandi

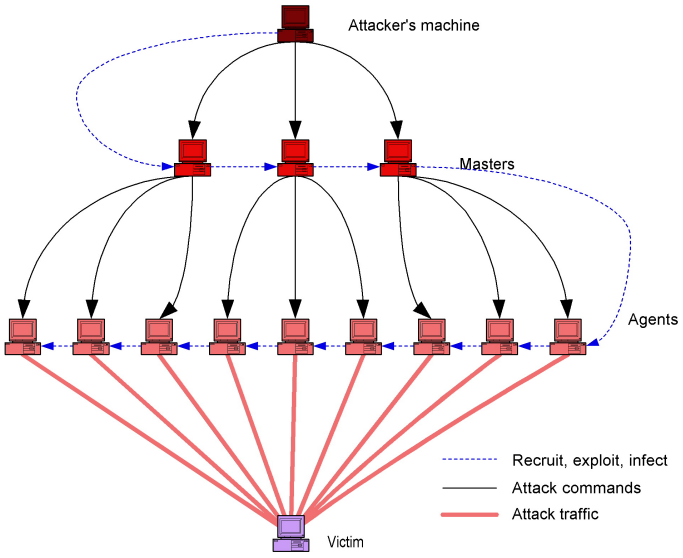

UNIVERSITY OF TRENTO - Italy

Digital MASTER SCHOOL

DDoS Attack Trends

- ***Attackers follow defense approaches, adjust their code to bypass defenses***
  - Use of subnet spoofing defeats ingress filtering
  - Use of encryption and decoy packets, IRC or P2P obscures master-slave communication
  - Encryption of attack packets defeats traffic analysis and signature detection
  - Pulsing attacks defeat slow defenses and traceback
  - Flash-crowd attacks generate legitimate (well-formed) application traffic
  - Social-network recruitment

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Small is Good (to decoy)

- ***Small-Packet Floods → Overwhelm routers***
  - Create a lot of pps
  - Exhaust CPU
  - Most routers can't handle full bandwidth's load of small packets
  - No real solution, must filter packets somehow to reduce router load
- ***Shrew Attack → Periodically slam the victim with short, high-volume pulses***
  - Lead to congestion drops on client's TCP traffic
  - TCP backs off
  - If loss is large back off to 1 MSS per RTT
  - Attacker slams again after a few RTTs
  - Solution requires TCP protocol changes
    - Tough to implement since clients must be changed

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Flash-Crowd and Reflectors Attack

- ***Crowd → wittingly generate legitimate application traffic to the victim***
  - E.g., DNS requests, Web requests
  - Usually not spoofed
  - If enough bots are used no client appears too aggressive
  - Really hard to filter since both traffic and client behavior seem identical between attackers and legitimate users
- ***Reflectors → unwittingly generate service requests to public servers spoofing the victim's IP***
  - Servers reply back to the victim overwhelming it
  - Usually done for UDP and ICMP traffic (TCP SYN flood would only overwhelm CPU if huge number of packets is generated)
  - Often exploit amplification effect – some service requests → huge replies; this lets attacker amplify his attack
    - try ping broadcast on a LAN with a spoofed

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Availability Of Attack Computers

- ***DDoS is feasible because attackers can enlist many vulnerable computers***
  - if you don't care which 1M computers
  - If many Computers don't look like Computers BUT they are so (See first lecture)
  - If people want to pay those computers with peanuts they are going to get a monkey job for the software
- ***Botnets numbering hundreds of thousands of hosts have been discovered***
  - Eg. Mirai botnet → "shame shame cameras come with a default password" → ask yourself "what is the cost of a Mirai Infected Camera?"

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Lack Of Enforcement Tools

- ***DDoS attackers have never been caught by tracing or observing attack***
- ***Only by old-fashioned detective work***
  - Really, only when they're dumb enough to boast about their success
- ***The Internet offers no help in tracing a single attack stream, much less multiple ones***
- ***Even if you trace them, a clever attacker leaves no clues of his identity on those machines***

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Defences: Resource Limitations

- ***Don't allow an individual attack machine to use many of a target's resources***
- ***Requires:***
  - Authentication, or
  - Making the sender do special work (puzzles)
- ***Authentication schemes are often expensive for the receiver***
- ***Existing legitimate senders largely not set up to handle doing special work***
  - Would you use a web site that requires you doing extra work?
- ***Can still be overcome with a large enough army of bots***

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Defences: Trace and Stop Attacks

- ***Figure out which machines attacks come from***
- ***Go to those machines (or near them) and stop the attacks***
- ***Tracing is trivial if IP source addresses aren't spoofed***
  - Tracing may be possible even if they are spoofed
- ***May not have ability/authority to do anything once you've found the attack machines***
- ***Not too helpful if attacker has a vast supply of machines***

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Traceback1

- ***Goal: locate the agent machines***
  - "Practical network support for IP Traceback," Savage, Wetherall, Karlin, Anderson, ACM SIGCOMM 2000
- ***Each packet header may carry a mark, containing:***
  - EdgeID (IP addresses of the routers) specifying an edge it has traversed
  - The distance from the edge
- ***Routers mark packets probabilistically***
- ***If a router detects half-marked packet (containing only one IP address) it will complete the mark***
- ***Victim under attack reconstructs the path from the marked packets***

UNIVERSITY OF TRENTO - Italy

eit Digital MASTER SCHOOL

Traceback and IP Spoofing

- ***Traceback does nothing to stop DDoS attacks***
  - It only identifies attackers' true locations
    - Comes to a vicinity of attacker
  - If IP spoofing were not possible in the Internet, traceback would not be necessary
- ***Incrementally deployable, a few disjoint routers can provide beneficial information***
- ***Moderate router overhead (packet modification)***
- ***A few thousand packets are needed even for long path reconstruction***
- ***Path reassembly is computationally demanding, and is not 100% accurate:***
  - Path information cannot be used for legal purposes
  - Routers close to the sources can efficiently block attack traffic, minimizing collateral damage
- ***Does not work well for highly distributed attacks***

What are the features of Internet?

- ***No need of***
  - validating IP source address
  - enforcing amount of resources requested
  - tracking traffic flows
    - Or tracking those controlling traffic flows
  - assigning responsibility for packets or packet streams
  - determining who accessed a machine
- ***BUT no need = no way***
  - In good and in bad fortune…