ΟΙΚΟΝΟΜΙΚΟ ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ

ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

Information-Centric Networks

Section # 7.3: Evolved Addressing & Forwarding

Instructor: George Xylomenos

Department: Informatics

European Union
European Social Fund

OPERATIONAL PROGRAMME
EDUCATION AND LIFELONG LEARNING
investing in knowledge society
MINISTRY OF EDUCATION, LIFELONG LEARNING AND RELIGIOUS AFFAIRS
MANAGING AUTHORITY

NSRF 2007-2013
programme for development
EUROPEAN SOCIAL FUND

Co- financed by Greece and the European Union

Funding

- These educational materials have been developed as part of the instructors educational tasks.
- The **“Athens University of Economics and Business Open Courses”** project only funded the reformatting of these educational materials.
- The project is being implemented as part of the Operational Program “Instruction and Lifelong Learning” and is co-financed by the European Union (European Social Fund) and national funds.

European Union
European Social Fund
OPERATIONAL PROGRAMME
EDUCATION AND LIFELONG LEARNING
investing in knowledge society
MINISTRY OF EDUCATION, LIFELONG LEARNING AND RELIGIOUS AFFAIRS
MANAGING AUTHORITY
Co- financed by Greece and the European Union
NSRF 2007-2013
programme for development
EUROPEAN SOCIAL FUND

Licencing

Week 7 / Paper 3

- Accountable Internet Protocol (AIP)
  - Michael Walfish, Hari Balakrishnan and Scott Shenker David G. Andersen, Hari Balakrishnan, Nick Feamster, Teemu Koponen, Daekyeong Moon, Scott Shenker
  - ACM SIGCOMM 2008
- Main point
  - Accountability at the forefront of the Internet
    - Prevention of source spoofing, DoS, route hijacking, route forgery
  - AIP uses a hierarchy of self-certifying addresses
  - Each component is derived from a public key

Introduction

- The Internet is rife with IP level vulnerabilities
  - Misconfigured routers wreak havoc on packet delivery
  - Hijacked routes used to send untraceable spam
  - Hijacked hosts spoof source addresses
  - DoS attacks occur on a daily basis
- Many solutions proposed, but all have shortcomings
  - Complicated mechanisms that change the Internet model
  - External sources of trust to certify BGP updates
  - Operator vigilance to keep updating filters
- Maybe the fundamental architecture is at fault
  - AIP uses self-certifying flat addresses
    - Hosts can prove they own an address without a PKI
    - But, flat addressing is a scalability challenge

AIP design

- Basic structure and function
  - Hierarchical addresses with two or more components
  - Each network is divided into Accountability Domains (ADs)
  - Each host is an Endpoint IDentifier (EID)
    - Both AD and EID are globally unique
    - Addresses have the form AD:EID
  - ADs may be subdivided into units
    - In general addresses are AD1:AD2:…:ADk:EID
  - The AD and EID is the hash of the public key of an entity
    - Each component is 160 bits long (8+144+8)
    - Direct link between identity (public key) and name (AD/EID)
  - Each AD/EID has the form Version:Key Hash:Interface
    - Version indicates the scheme used to generate the AD/EID
    - Interface indicates one of the interfaces of a host

AIP design

- Forwarding and routing
  - Each packet contains source and destination AD:EID
  - Multiple AD levels are treated as a stack
    - Forwarding proceeds towards the current AD
    - Border routers switch to the next AD in the stack
  - Interdomain routing can use BGP or any other protocol
    - Routers advertise reachability to ADs, not prefixes
    - ADs can be grouped into ASes if needed
- DNS and mobility
  - A multihomed host will have different ADs but the same EID
  - Transport protocols bind to the EID, not the AD
  - Mobile hosts need to change their AD only
  - Changes can be authenticated since EIDs are bound to keys

Uses of accountability

- Source accountability (no source spoofing)
  - Common source spoofing variants
    - Pretending to be a host at another network
    - Pretending to be another host at your network
    - Creating large numbers of unused addresses
  - EID verification: at first hop router
    - On reception of data from unverified host
      - Drop packet and return verification packet V
      - V contains source/destination addresses, packet hash and interface
      - V is signed using a secret that changes regularly
      - The sender returns V signed with its private key
      - The router verifies the signature and if correct passes next packets
      - Replay attacks at the router prevented by the secret
      - Replay attacks at the sender prevented by inserting random packet ID
      - Only the sender needs to cache hashes of recent packets

Uses of accountability

- Source accountability (no source spoofing)
  - AD verification (at AD boundaries)
    - On reception of packet from AD B, AD A checks that:
      - If B is trusted to check packets, forward it
      - Otherwise check if the packet is on the route to the source
      - Otherwise drop packet and verify the source as with the EID
      - The last step allows multihoming and asymmetric paths
  - Ensuring scalability at border routers
    - Only packets that arrive from an unexpected route are remembered
    - If the AD:EID pairs for the same AD are many, use a wildcard AD:*
    - This is dangerous if an attacker controls some hosts in the AD
      - It can force the border router to insert a wildcard
  - Limiting address minting
    - Routers can limit the number of new EIDs they accept per minute
    - Similarly for ADs in border routers

Uses of accountability

- Shut-off protocol
  - Requires a smart-NIC that rate-limits transmission if needed
    - The NIC records hashes of recently sent packets
    - It also accepts Shut-Off Packets (SOPs) independently
    - A SOP contains a packet hash and a TTL, signed by the destination
    - The NIC accepts the packet if it contains a valid hash and signature
    - Then it shuts-off traffic for the TTL
    - The random packet ID prevents replay attacks
    - Hashes of thousands of packet can be kept in a Bloom filter
- Securing BGP
  - BGP peering sessions are encrypted with AD public keys
  - BGP routers sign their routing announcements
  - Routers only need to know the public keys of other ADs

Routing scalability

- Routing growth estimates
  - The AS diameter increases slowly (less than 5 hops)
  - Routing tables grow 17% per year (about 1.6 million in 2020)
  - Routing traffic grows linearly with table size
    - Estimate 1.5 updates per day per prefix
- Effects of moving to AIP
  - RIB/FIB size increase due to 160 bit AD and 2048 bit public key
    - Doubles for two level ADs
  - BUT: lookups are flat, not longest prefix
    - Estimated 80% reduction in memory accesses
  - The AS diameter may grow by 2-3 hops due to two level ADs
  - CPU costs that same as those needed for S-BGP

Routing scalability

- Semiconductor growth trends
  - Assume the ITRS roadmap of density doubling every 3 years
- Resource requirements
  - RIB (DRAM): Roughly triple the amount of RAM needed by IP
    - No problem with current growth trends
  - FIB (DR/SR/CAM): SRAM will grow much faster than FIBs
  - CPU: Loading time of tables to memory
    - Estimated to be less than 30 seconds
  - Cryptography: 66 seconds for all tables
    - May need acceleration as it is slower than the loading time

Key management

- Key compromise: what if a private key is exposed?
  - Key revocation is tricky but feasible
  - The biggest problem is the false confidence in lost keys
    - Therefore mechanisms for rapid loss detection are needed
  - Public registries for ADs are needed
    - They only store self certifying data (no PKI needed)
    - Can be automatically populated (no human intervention)
    - Local (per domain) and global registries possible
  - Types of information stored in registries
    - Identity/public key pairs
    - Revoked public keys, signed by private keys (write once!)
    - Peering relationships, signed by both peers
    - Certificates binding EIDs to ADs (may be multiple)
    - Certificates binding EIDs to first hop routers (can be multiple)

Key management

- Maintaining domain registries
  - Domains should be forced to sign AD:EID to get it into DNS
  - Hosts check the registries of the domains hosting them
  - Domains check the registries of their peers
- Key discovery: what is the destination address?
  - Use DNS and verify keys against AD identifiers
  - Establish out-of-band trust between peering domains
- Cryptographic algorithm compromise
  - This is why version numbers are added to identities
  - New algorithms can co-exist with older ones
  - A version number indicates a hash/sign combination pair

Traffic engineering and AD size

- Traffic engineering maps offered load to a set of paths
  - Performed by DNS and selective prefix advertising
  - ADs cannot be split to smaller prefixes for control
    - But they can be split hierarchically
  - AD granularity: a set of hosts under common administration
  - Splitting ADs for traffic engineering can be done
    - It is aggregating these prefixes that is impossible!
  - Experiments show that such aggregation is quite rare
  - Another possibility is to use the interface bits
    - Each interface could be advertised separately
    - Allows aggregation by zeroing out the interface bits
  - DNS load balancing does not change
    - The interface bits can make it easier

ΟΙΚΟΝΟΜΙΚΟ ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ

ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

End of Section # 7.3

Course: Information-Centric Networks, **Section # 7.3: Evolved Addressing & Forwarding**

Instructor: George Xylomenos, **Department:** Informatics

European Union
European Social Fund
OPERATIONAL PROGRAMME
EDUCATION AND LIFELONG LEARNING
investing in knowledge society
MINISTRY OF EDUCATION, LIFELONG LEARNING AND RELIGIOUS AFFAIRS
MANAGING AUTHORITY
Co- financed by Greece and the European Union
NSRF 2007-2013
programme for development
EUROPEAN SOCIAL FUND