Persistent OSPF Attacks

Gabi Nakibly[1,2] Alex Kirshon[2] Dima Gonikman[2] Dan Boneh[3]

[1]National EW Research & Simulation Center, Israel
[2]CS department, Technion, Israel
[3]CS department, Stanford

Overview

- They allow to remotely control a router’s routing table without having to control the router itself.
- A single compromised router inside an AS can compromise the routing of the whole AS.
- Potentially every OSPF implementation is vulnerable.
  - The attacks were verified against Cisco’s IOS.

Who is vulnerable?

- Potentially all commercial routers are vulnerable!
- The vulnerabilities were found in the spec of the OSPF protocol [RFC 2328].
- The attacks have been verified against Cisco IOS 15.0(1)M.
  - IOS's latest stable release

Outline

- OSPF primer
- OSPF security strengths
- The newly found vulnerabilities and attacks
- Attacks' effectiveness

Internet Routing – The Big Picture

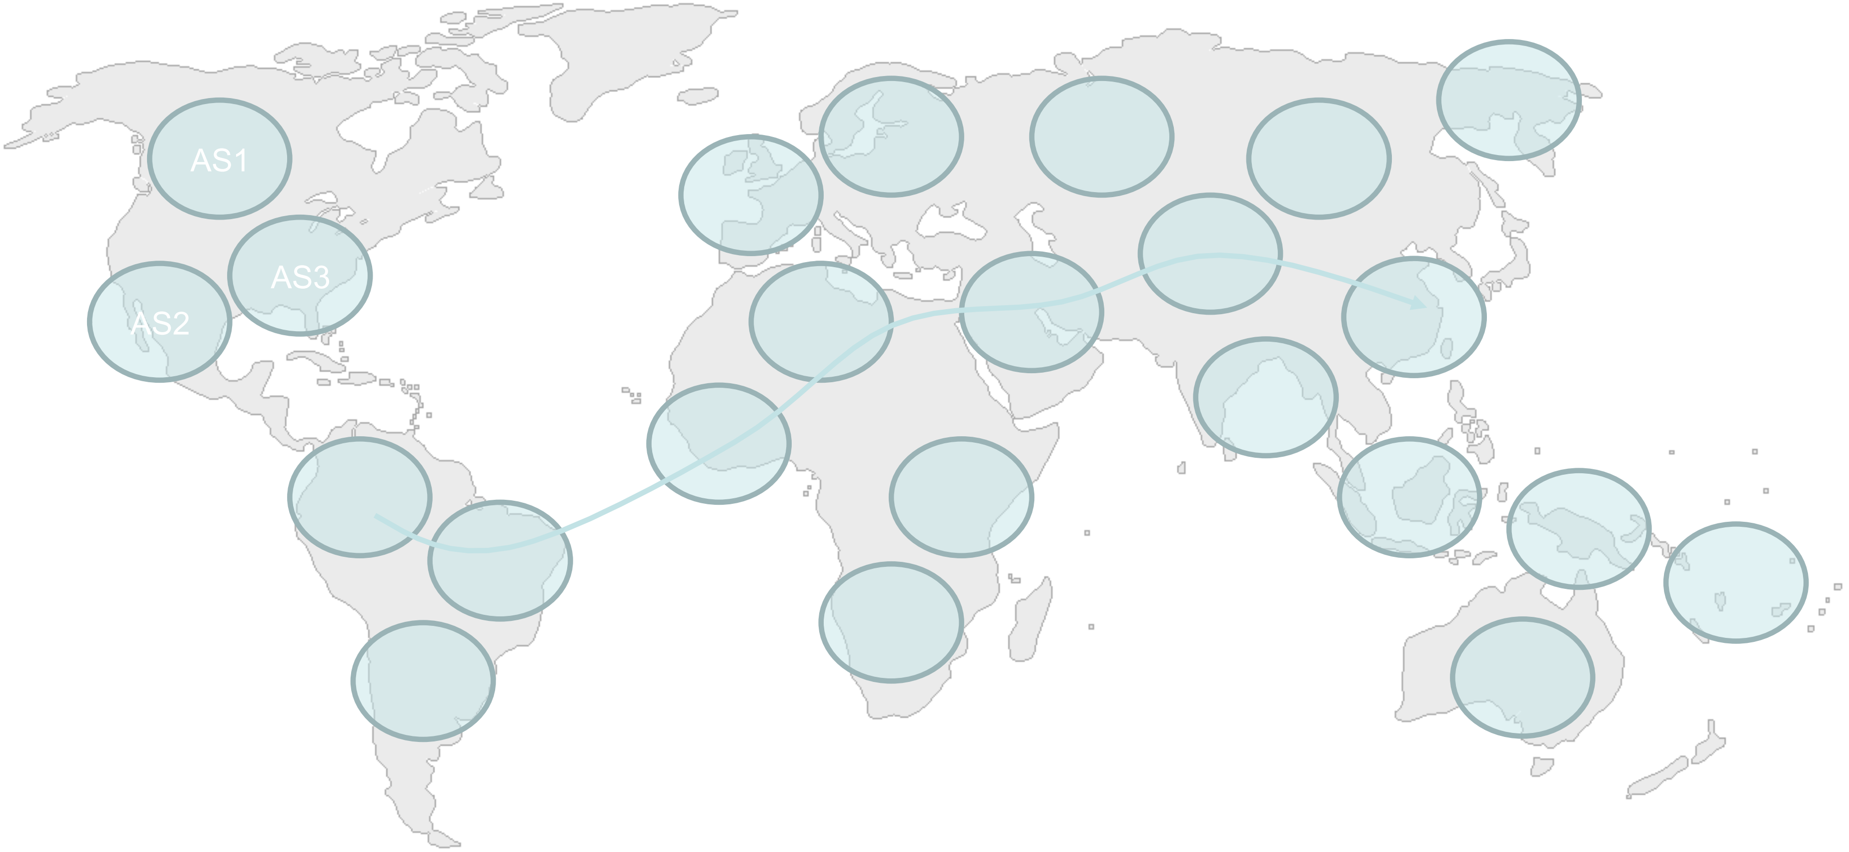

Inter-AS routing – BGP
Intra-AS routing – OSPF, RIP, IS-IS

How OSPF works?

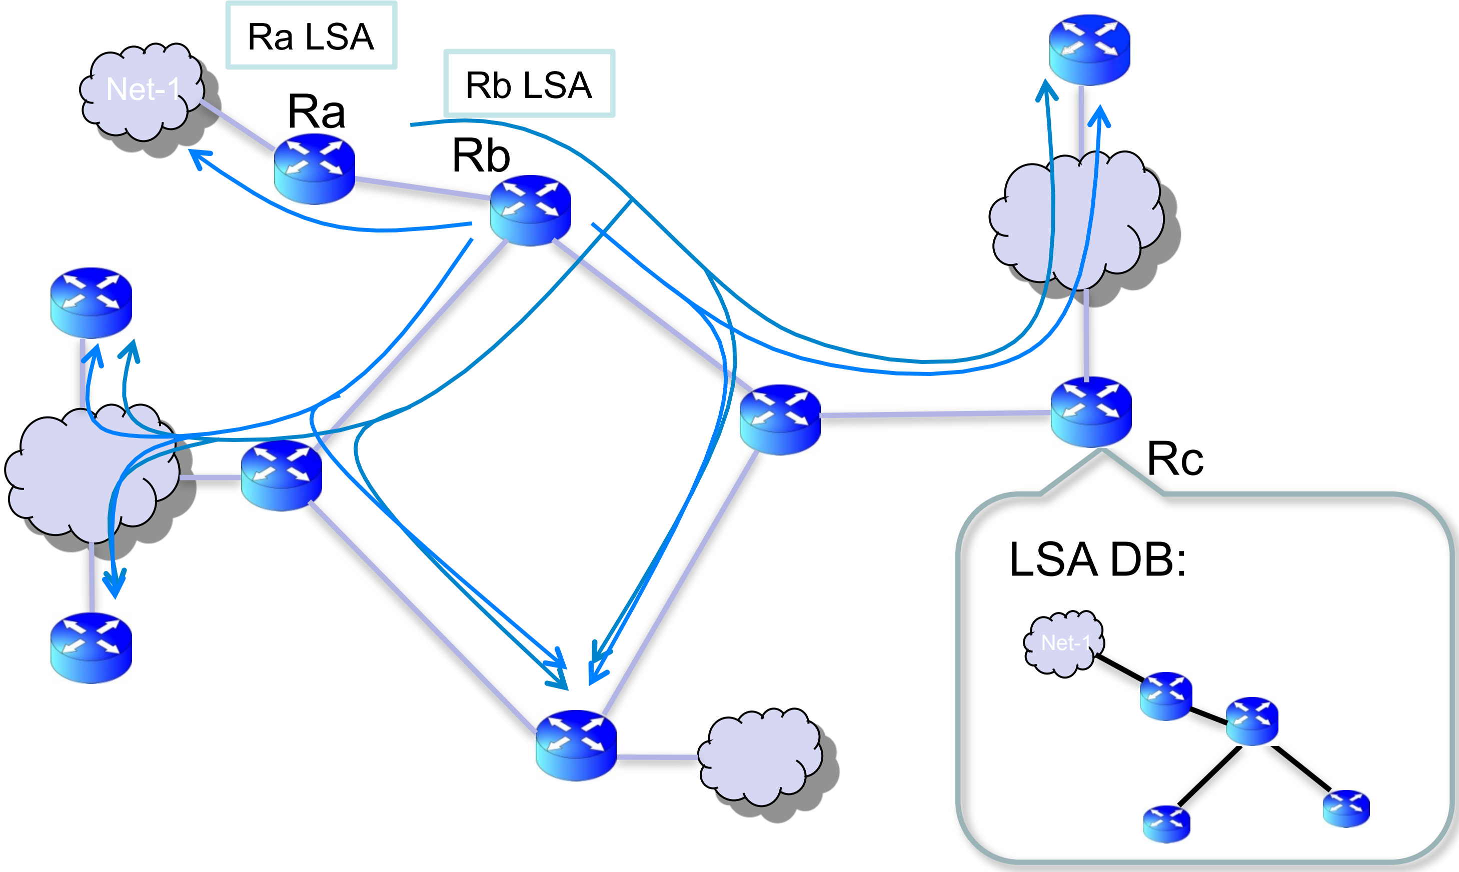

How OSPF works?

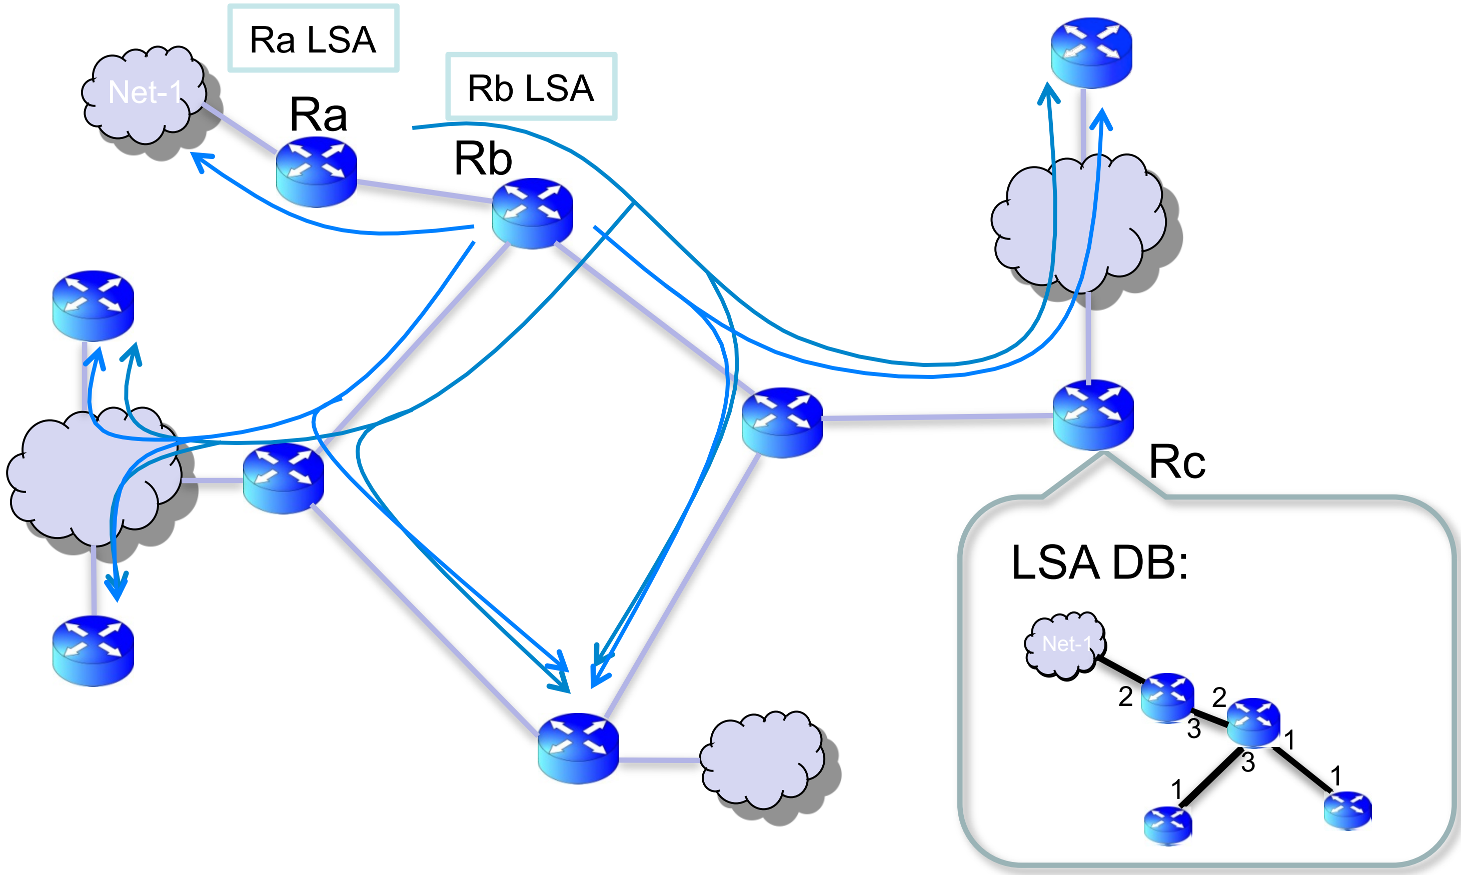

LSAs

- Each LSA is advertised periodically
  - Sequence number
    - To differentiate between instances of the same LSA
  - Age
    - To allow a specific instance of an LSA to expire

The Attacker

- Location: Inside the AS
  - Controls a single router
  - Arbitrary location
- Goal:
  - Persistent control over the routing tables of other routers in the AS

OSPF Security Strengths

- Per-link authentication
  - Every link has its own shared secret
- Every LSA is flooded throughout the AS
- The “fight back” mechanism

Known Attacks

- Falsify LSAs of:
  - The attacker’s router
    - Very limited
  - other routers
    - Known examples: Seq++, MaxSeq,…
    - Trigger immediate fight back
      - A non-persistent attack
  - phantom routers
    - Does not have an affect on the routing table

Known Attacks

- In summary,
  - The common conception is that even if the attacker is an insider it can not persistently falsify the LSA of a router it does not control.
    - Hence, it can not significantly poison the routing tables of other routers.

The New Attacks

- Attack #1 – Remote False Adjacency
    - Make a remote router include a non-existing link in its LSA
- Attack #2 – Disguised LSA
    - Falsify the entire LSA of remote router

Attack #2 – Disguised LSA

- The vulnerability
  - Two different instances of an LSA are considered identical if they have the same [RFC 2328 Sec. 13.1]:
    - Sequence number
    - Checksum
    - Age (+/- 15 minutes)
  - The actual payload of the LSAs are not considered!
- The attack
  - Advertise a false LSA having the same values for these three fields as a valid LSA.
    - The benefit: no fight back is triggered since the victim views the false LSA as a duplicate of the LSA it just advertised.

Attack #2 – Disguised LSA (cont.)

- The attack (cont.)
  - But, there is a problem: all other routers in the AS will also consider the false LSA as a duplicate
    - therefore, they will not install it in their LSA DB.
  - Solution: Disguise the LSA to the next valid instance of the LSA
    - While at the same time the victim originate this next valid instance
      - The trigger is done using the fight-back mechanism

Application

- The attacker floods consecutively the trigger and then the disguised LSA.

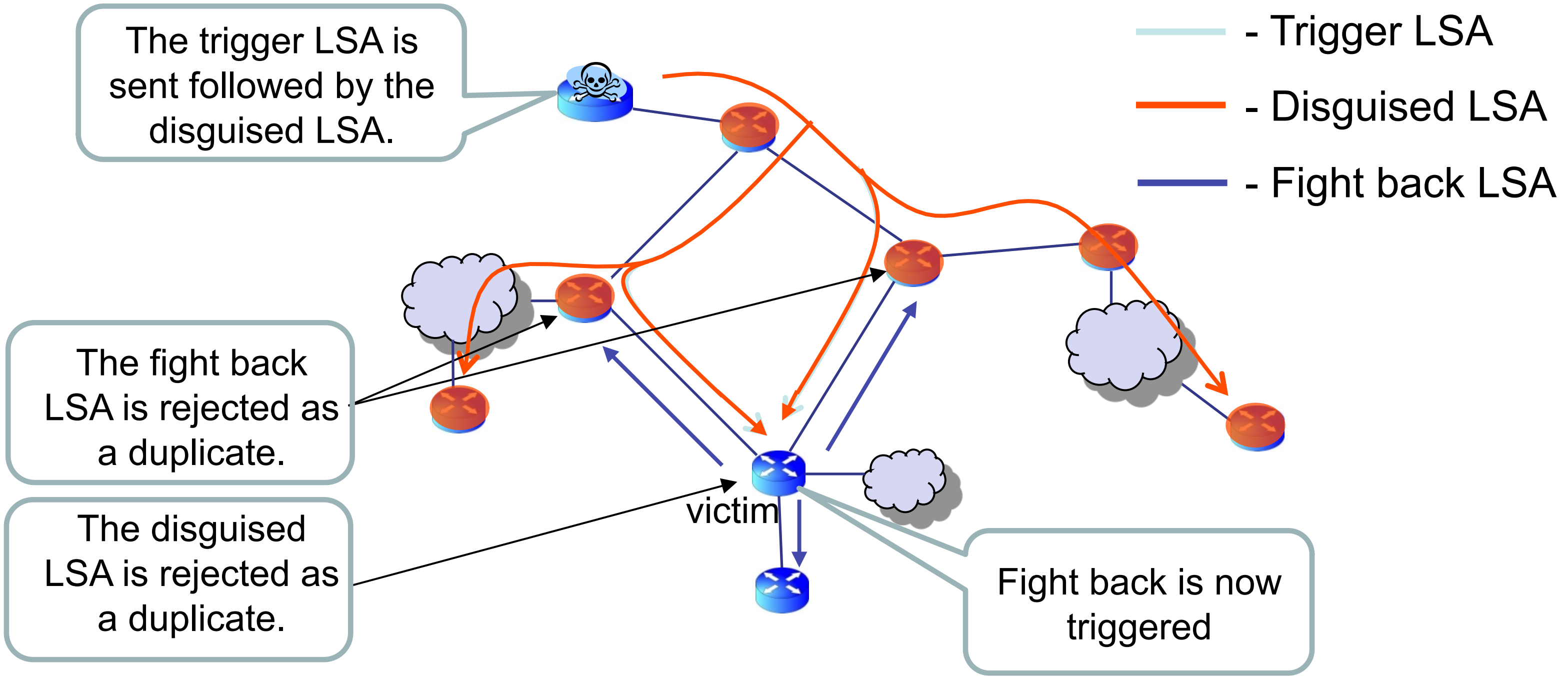

The trigger LSA is sent followed by the disguised LSA.

- Trigger LSA

- Disguised LSA

- Fight back LSA

The fight back LSA is rejected as a duplicate.

The disguised LSA is rejected as a duplicate.

Fight back is now triggered

How the disguised LSA can be crafted?

- Age: this is the easiest one.
  - The disguised LSA will be advertised within 15 minutes of the valid (fight back) LSA.
- Sequence: the value is always incremented by one.
  - The disguised LSA will have the sequence of the trigger LSA plus 1.
- Checksum: this is the hardest feat, but not that hard.
  - The content of the next valid LSA is deterministic and predictable, hence the checksum is also predictable.
  - A dummy Link entry in added to the payload of the LSA.
  - The value of this entry is calculated such that the entire LSA will have the desired checksum.
    - This can be done since a checksum is a 16-bit result of a linear calculation on the LSA octets.

Attack Effectiveness

- We simulated the attack on real ISP topologies
  - Inferred by the RocketFuel project
- We measured for every pair of attacker-victim locations what is the percentage of poisoned routers.

Simulation Results

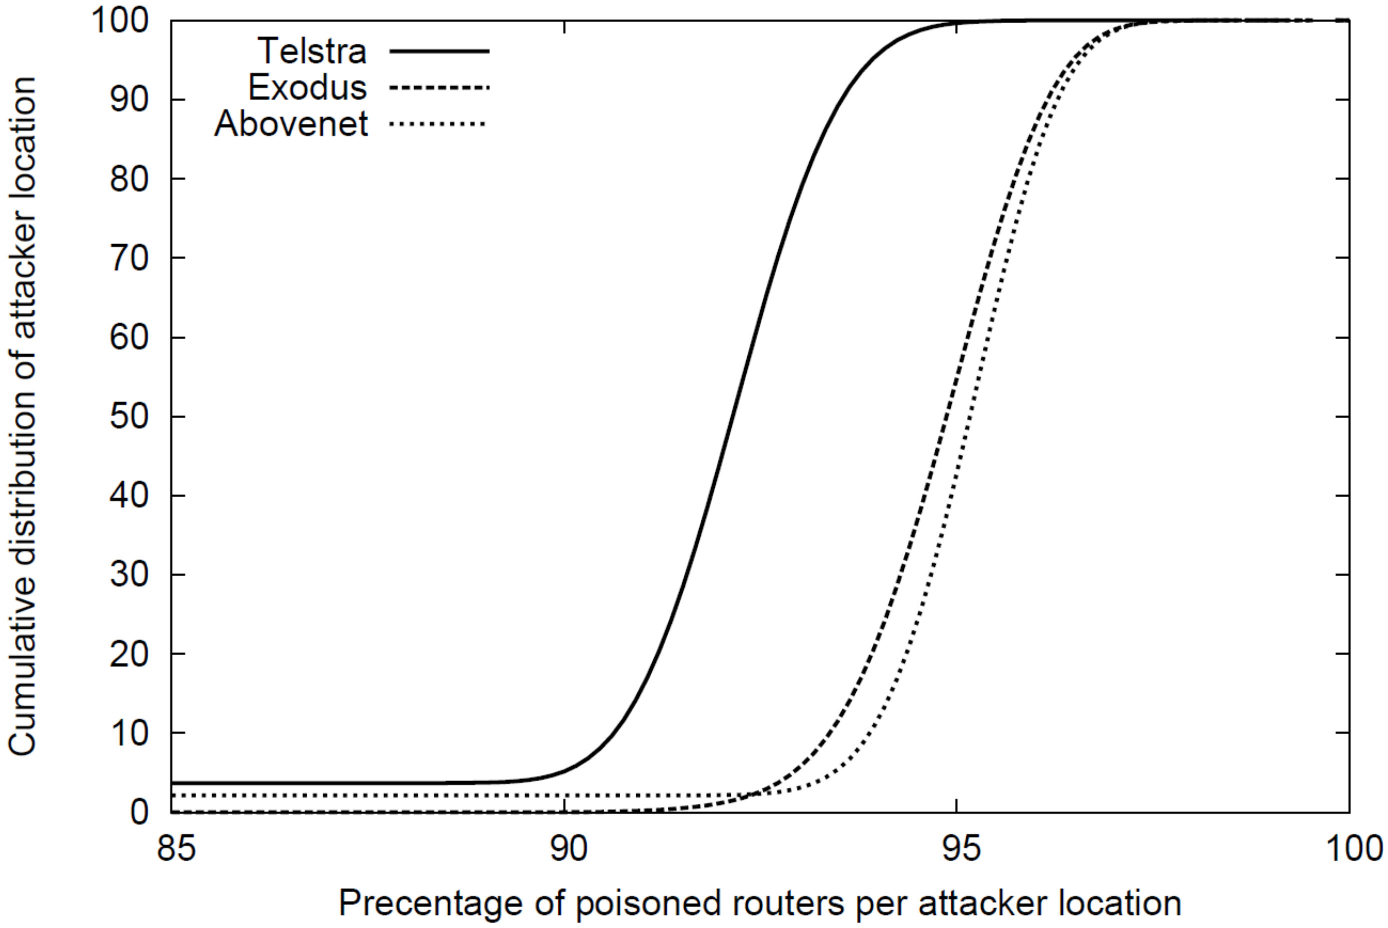

Conclusions

- Up until now the common conception was that even if the attacker is an insider it can not persistently poison the routing table of a router it does not control.
  - The new attacks shatter this misconception.
- **Using these attacks one can control the entire routing domain from a single router.**

In Summary …

Using these attacks one can control the entire routing domain from a single router.