Analyzing Internet Routing Security Using Model Checking

Adi Sosnovich[1]; Orna Grumberg[1]; Gabi Nakibly[2]
[1]Computer Science Department, Technion, Haifa, Israel, [2]National EW Research and Simulation Center, Rafael, Haifa, Israel

CONTRIBUTIONS

- Defined and implemented useful BGP-based **aggressive reductions of the Internet** topology to enable an automatic analysis.
- Developed an **automatic analysis** that can **reveal** possible **traffic attraction scenarios** on the Internet and prove that certain scenarios are not possible.
- Identified **safe nodes** that are not amenable to attraction attacks and can be exploited to reduce vulnerability of other nodes in the internet.

BACKGROUND – INTERNET ROUTING

- The Internet is composed of Autonomous Systems (ASes).
- Each AS is administered by a single entity.
- **Inter-domain routing** determines through which ASes packets will traverse.
- Routing on the AS level throughout the **Internet** is handled by a **single** routing protocol called the **Border Gateway Protocol (BGP)**

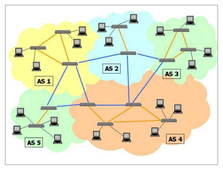

BGP VULNERABILITIES

- The Internet is vulnerable to **traffic attraction attacks**
- A malicious AS can manipulate BGP to attract traffic to, or through, its AS

Traffic diversion example
(Source:
http://research.dyn.com/2013/11/mitm-internet-hijacking)

MODELING BGP

- **Network topology:** a graph of AS nodes with edges of type **peer-peer** or **customer-provider.**
- **Destination node :** a **single** predefined **destination** AS in which the target network resides.
  - All ASes try to build routing paths to it.
- **Normal AS nodes**: a normal node sends valid routing advertisements based on normal policies.
- **Attacker node**: a **predefined** AS node that can send false routing advertisements.
  - Its **goal** is to achieve **traffic attraction**
  - It can send arbitrary routing advertisements and uses arbitrary policies

EXPORT AND PREFERENCE POLICIES

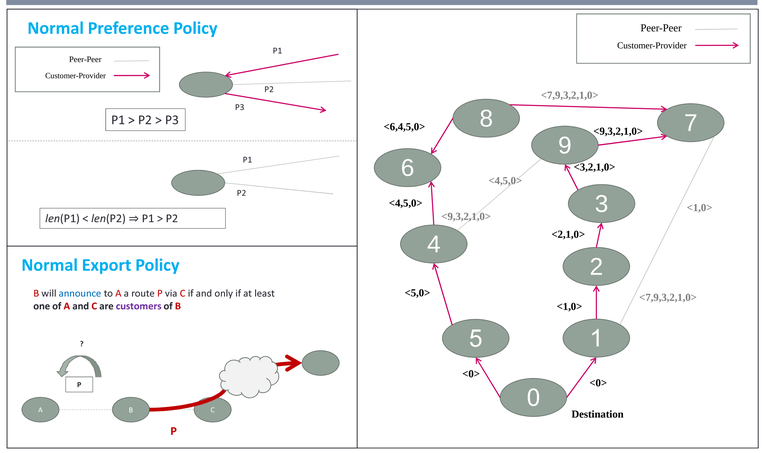

SPECIFICATION

- **Trivial attack strategy:** the attacker **sends** a **false** advertisement to **all its neighbors** that the **target** network is located within its **own AS**
- **Specification of a *non-trivial attraction/interception attack:***
  - **If the attacker can attract (or intercept) traffic from some victim, while it fails to do so in the normal run and in the trivial attack, the attraction (or interception) specification is satisfied**

NON-TRIVIAL INTERCEPTION

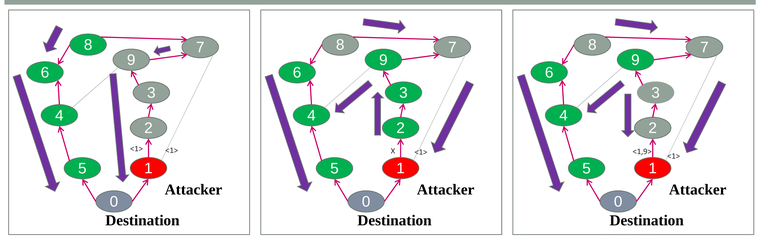

Normal outcome and trivial attack | Non-trivial interception | The newer strategy found by MC – non-trivial interception (with additional attraction)

THE METHOD

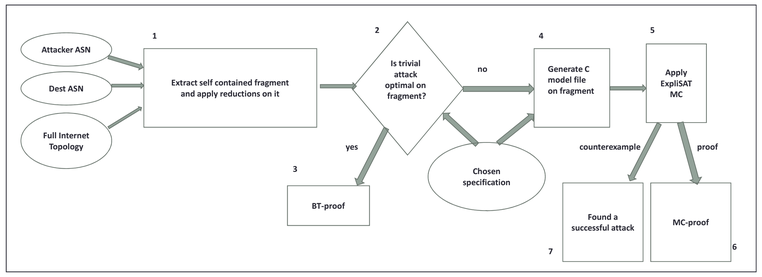

BGP NETWORK REDUCTION

- Applying model checking on the full Internet topology is **infeasible**
- We develop reductions to obtain a manageable sized **fragment** of the Internet

Network Reduction – First Attempt

Pick an **arbitrary** sub-network from the Internet

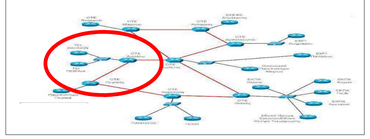

Problem:
If some attraction scenario is **found**, it is **not guaranteed to be preserved** in the context of the full Internet topology

Solution:
Find an isolated sub-network that is not affected by ASes outside, by using **valid paths**

- **Valid Paths**: **Export actions of regular nodes is performed only along valid paths**
- **Self-contained fragments**: A sub-network S of a BGP network is a self-contained fragment of a BGP network if for every node *n* outside S, there is no BGP run in which an export action from *n* to some *n'* in S is performed.
- **Lemma:**
  - Let N be a (large) BGP network and let S be a self-contained fragment of N
  - Then, any traffic attack found on S can occur on N as well
  - Moreover, if we obtain a proof that an attacker cannot attract traffic from some victim within S, then the proof applies for N as well

Extracting Self-contained Fragments

- Initially, {Dest,Attacker} and their neighbors are in *S*
- A node $c \notin S$ is added to *S* if:
  1. *c* is a neighbor of some $n \in S$
  2. *c* is on a valid path from some originator to *n*

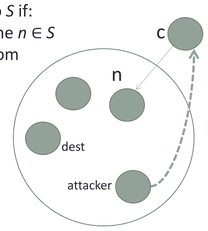

Definite Routing Choices

- Identifying nodes that **never route via the attacker**
- A node has a **definite routing choice** if its **chosen path is via the destination and not via the attacker for every possible run**, regardless of the attacker's actions

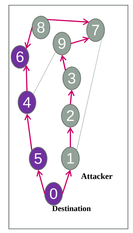

IDENTIFYING SAFE NODES

Partition of node types

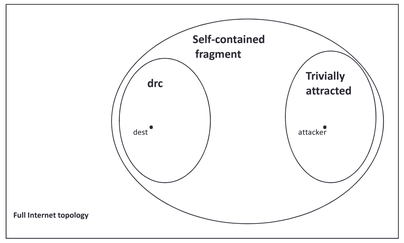

Best Trivial attraction proof scenario

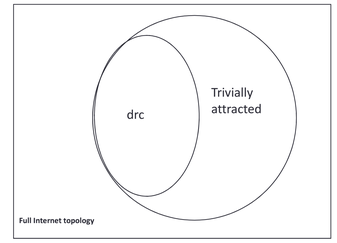

- We identify two types of safe nodes:
  - Nodes that have a definite routing choice
  - Nodes for which the model checker provides a proof that there is no attacker's strategy that can attract them