Threat analysis for routing-bridges

marcelo bagnulo

IETF62

Security goals

- minimum security expected from rbridges is to provide the same level of protection than regular bridges
  - i.e. that the introduction of rbridges in a bridged network does not introduce any new vulnerability.
- new features provided by rbridges may enable the usage of rbridges beyond current bridge capabilities.
  - security considerations may (and probably will) limit the recommended scope of application of rbridges.

Overview

- identify possible attacks to current bridges.
- threats related to the End-node Location Discovery Mechanism of rbridges.
- threats related to the Link- State Protocol
- security aspects that limit the usage of the rbridges beyond the scope of application of current bridges.

Overview

- **identify possible attacks to current bridges**.
- threats related to the End-node Location Discovery Mechanism of rbridges.
- threats related to the Link- State Protocol
- security aspects that limit the usage of the rbridges beyond the scope of application of current bridges.

Vulnerabilities of current bridges

- sending packets with spoofed link layer addresses
- Attacks to the STP

Scenario

- The attacker X has IP address IPX and link layer address LLX.
- Two nodes A and B have IP addresses IPA and IPB and link layer addresses LLA and LLB respectively.
- Assumption: attacker X, node A and node B are all in different links of the same bridged network, since the presented attacks are aimed to the bridging system.

Attack B.1

- The attacker X wants to establish a new communication with a node B pretending to be node A

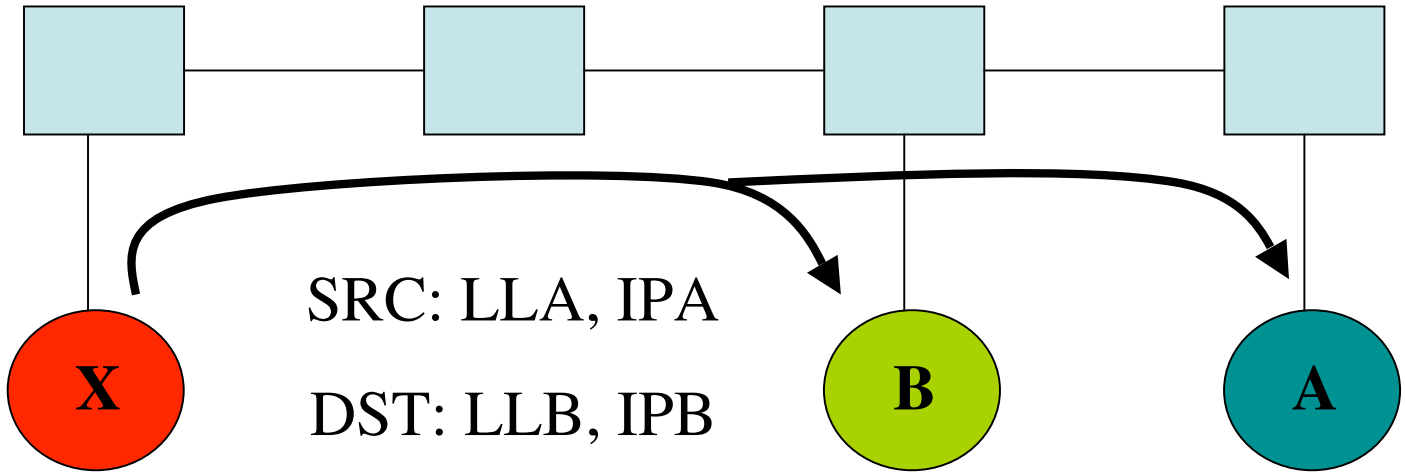

Attack B.1

- The attacker X wants to establish a new communication with a node B pretending to be node A

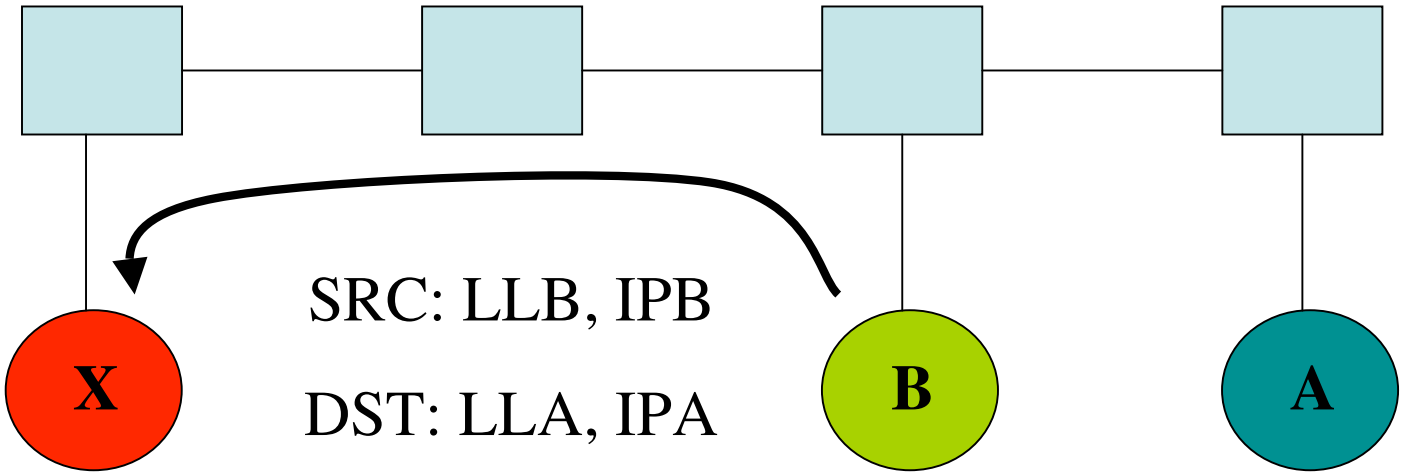

Attack B.1

- This is a masquerading attack, where node B is convinced that it is communicating with node A while it is actually communicating with the attacker X.

Attack B.2

- The attacker wants to impersonate node A in any new communication established by node B.

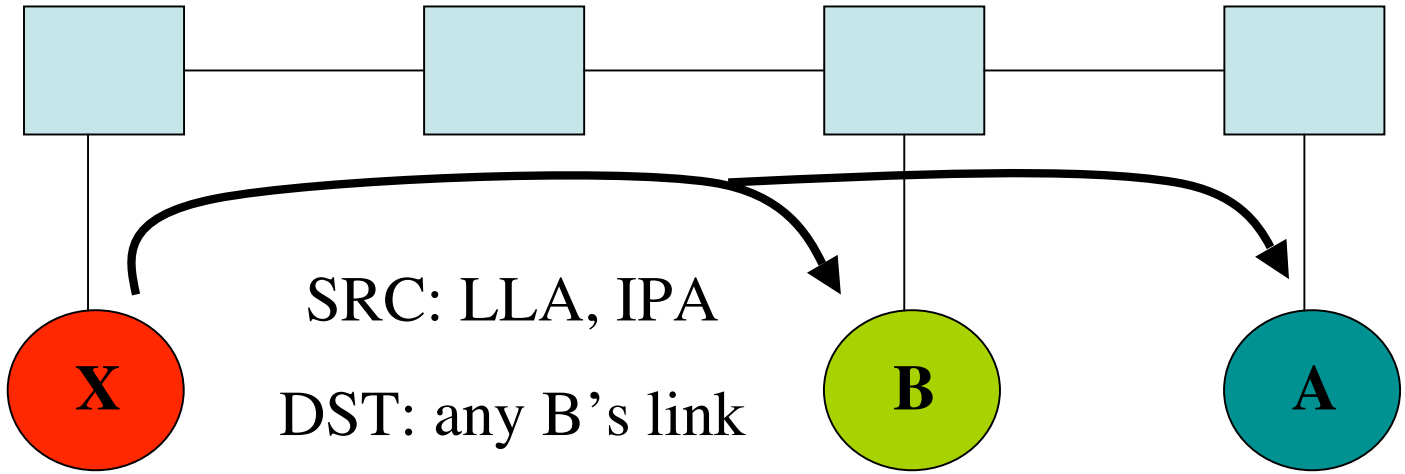

Attack B.2

- Repeat until B starts the communication
- What destination address? (only B or more)

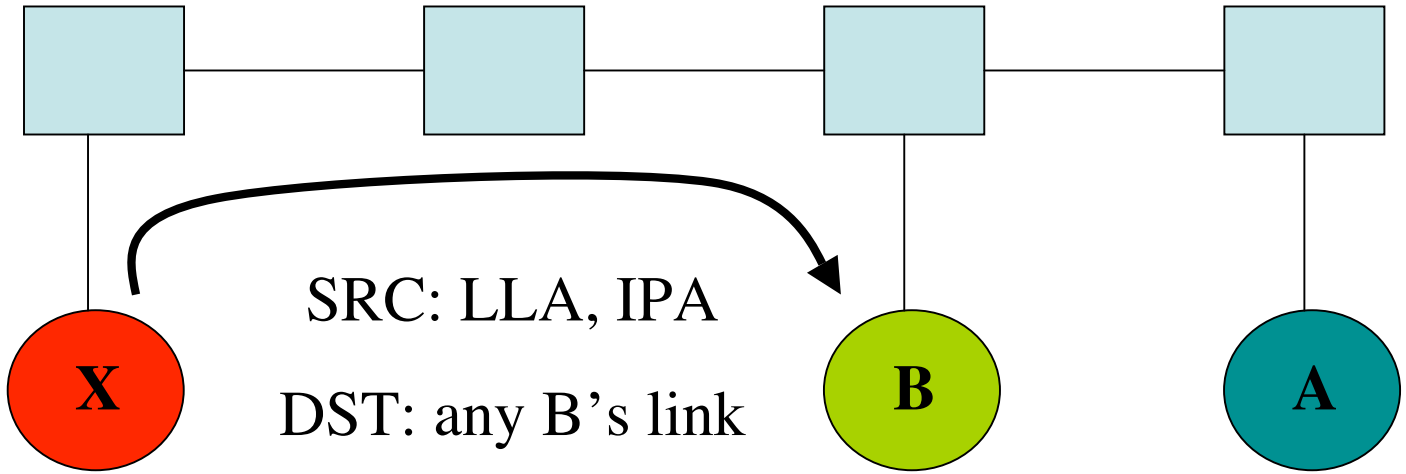

Attack B.2

- B starts the communication => ARP/ND

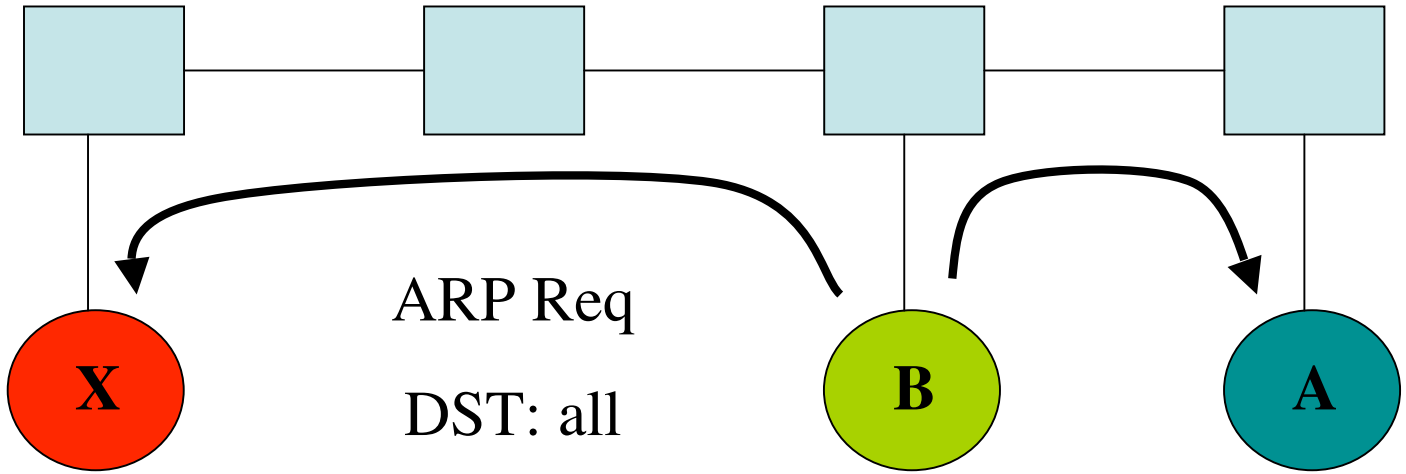

Attack B.2

- A Replies and the attack is suspended

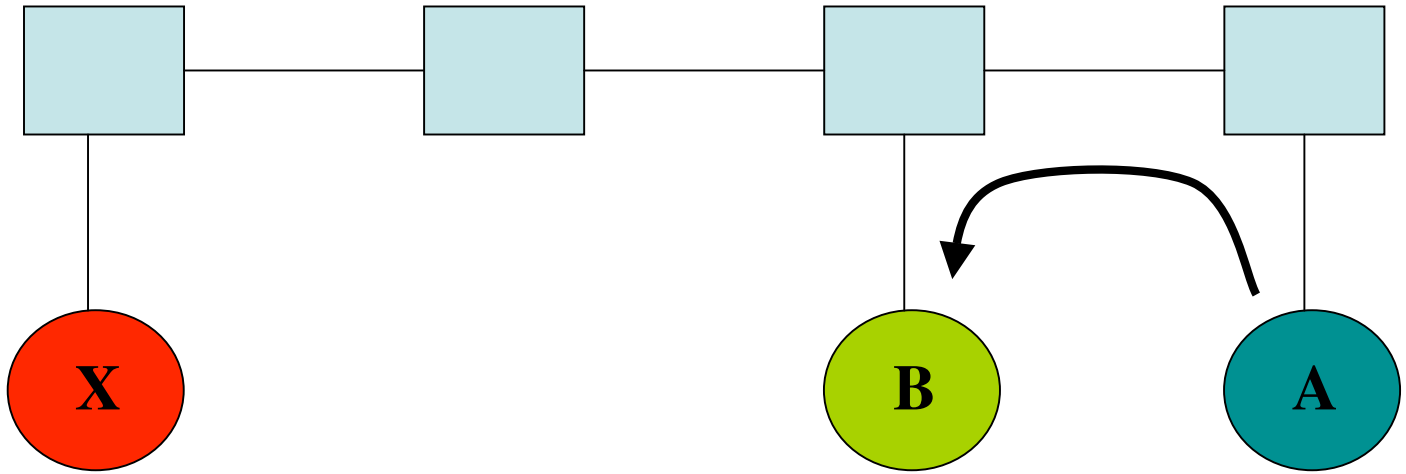

Attack B.2

- X sends a delayed reply, and the attack is restored

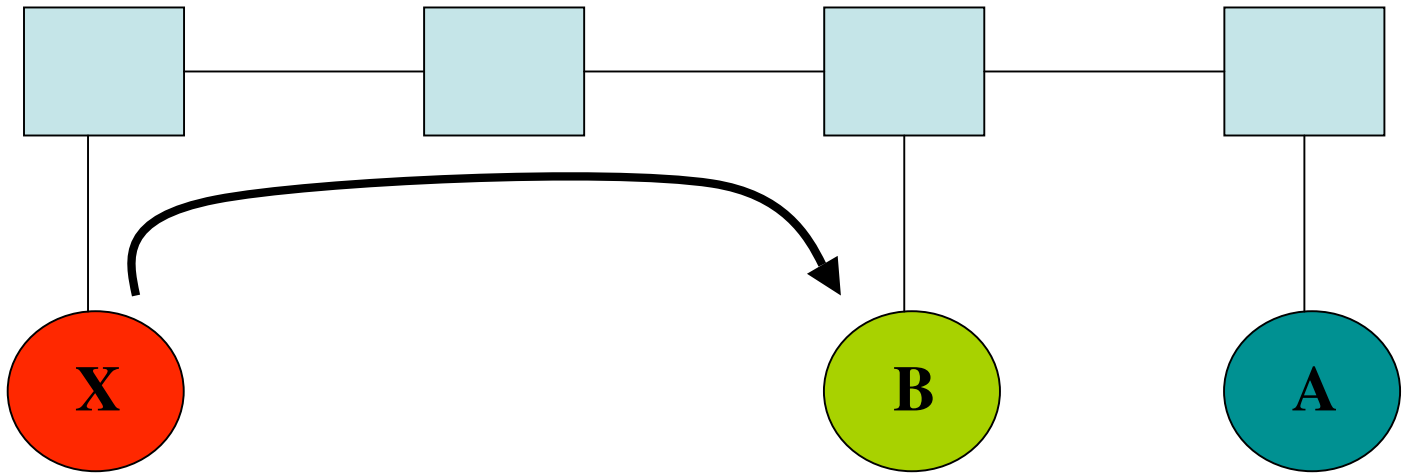

Attack B.2

- B start the communication with X

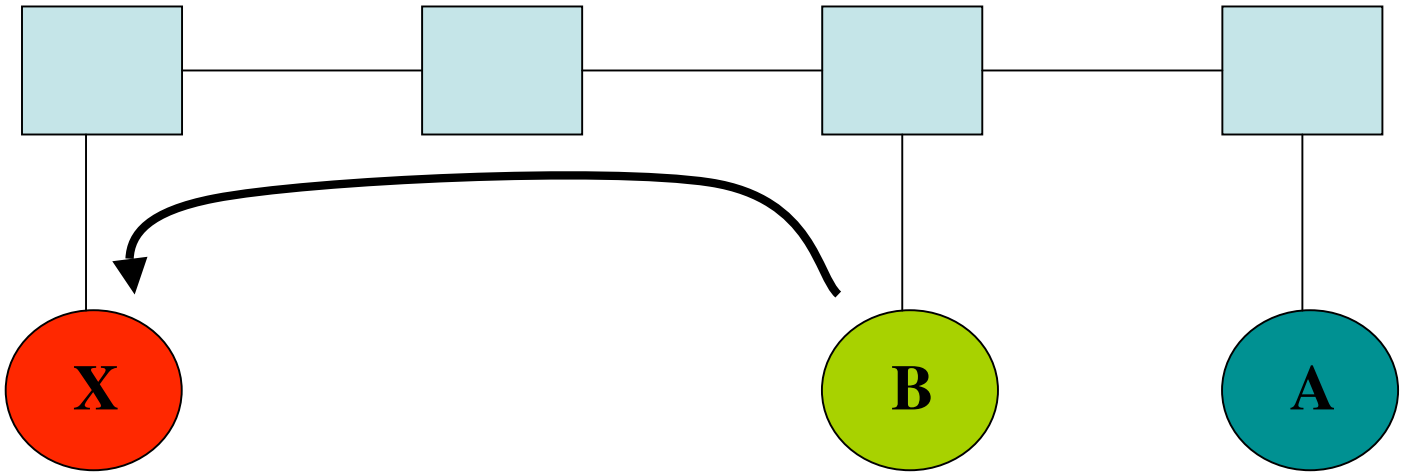

Attack B.2

- This is a masquerading attack to node B, since node B believes that it is communicating with A while it is actually communicating with the attacker X
- it is also a DoS attack to node A, since node A does not receive the traffic intended for him.
- this can be a DoS attack since the traffic generated by node B is flooding the path between node B and the attacker's link (especially if affects more than a single B)

Attack B.3

- The attacker wants to hijack an ongoing communication

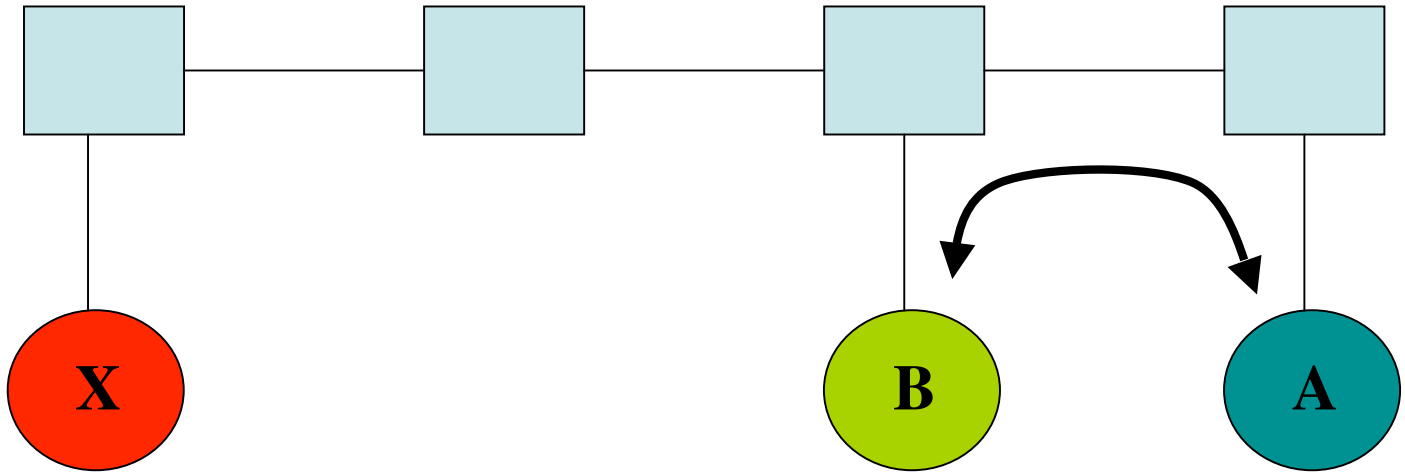

Attack B.3

- The attacker wants to hijack an ongoing communication

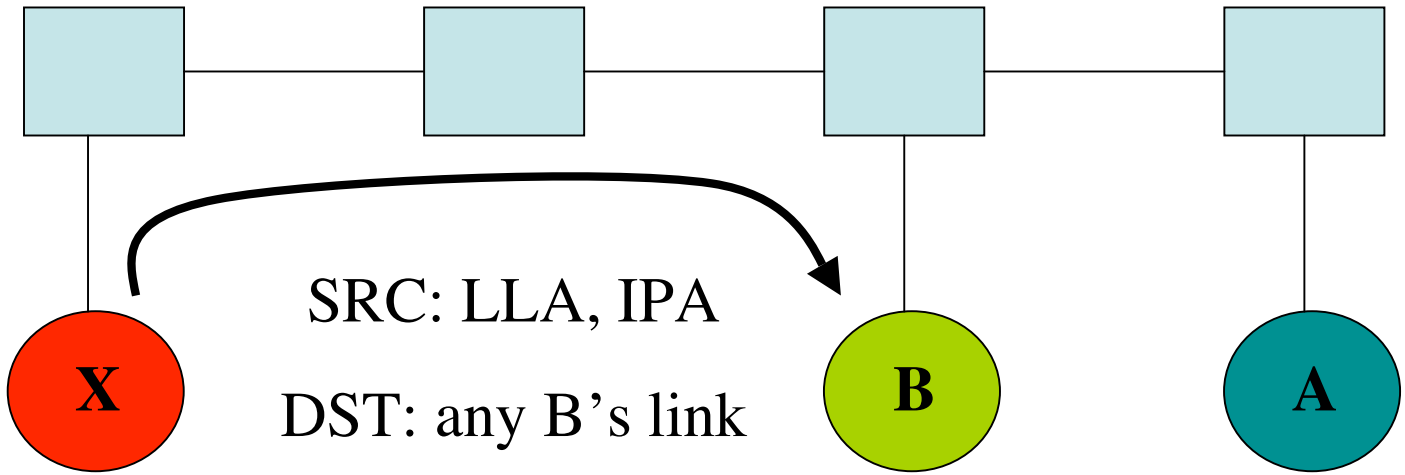

Attack B.3

- The attacker wants to hijack an ongoing communication

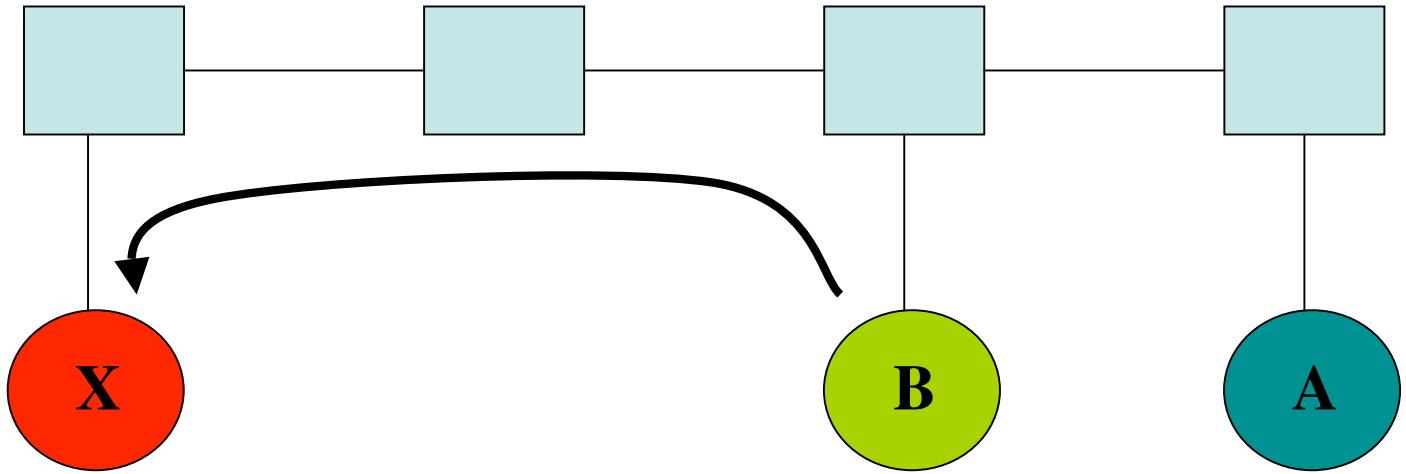

Attack B.3

- The attacker wants to hijack an ongoing communication

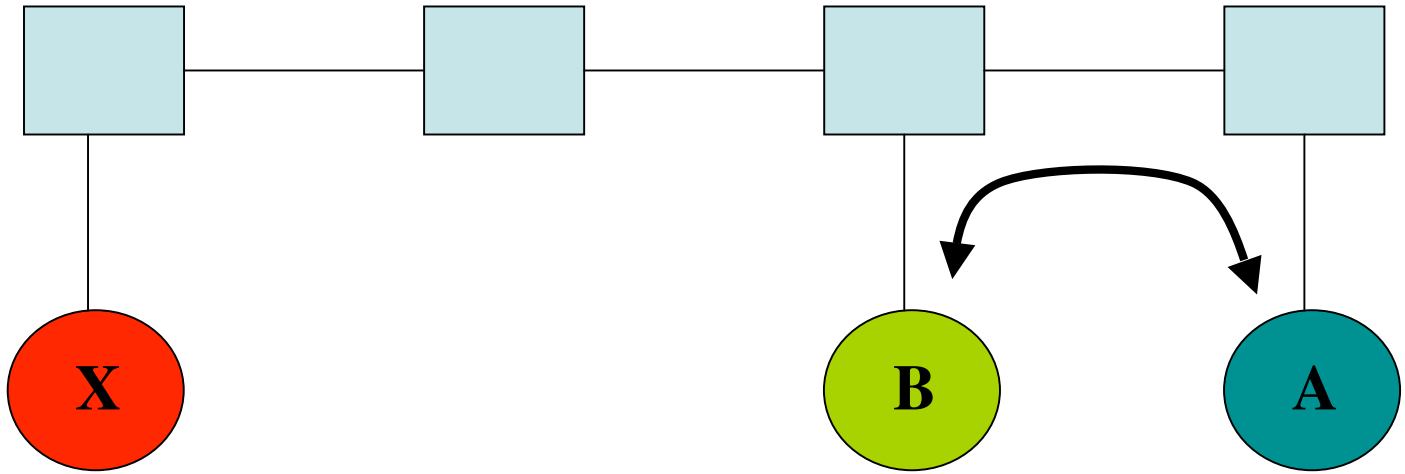

Attack B.3

- Unstable situation
- X can transmit with a high frequency, and managing to hijack
- Sending packets to different destinations, can affect all communications of A
- This is a masquerading attack to node B
- it is also a DoS attack to node A
- this can be a DoS attack since the traffic generate by node B is flooding the path between node B and the attacker's link (especially if affects more than a single B)

Attack B.4

- Attack to the spanning tree protocol
- X convince all the bridges in a link that he is the Designated Bridge on that link.
- This would imply that no bridge will act as DB in the bridge
- X can become the DB of a given link by advertising configuration message with the lowest cost to the root.
- This s DoS attack.

Attack B.5

- Attack to the STP
- X becomes the root of the spanning tree,
- This is achieved by advertising configuration messages with the lowest root ID.
- So far, not very harmless
- The attack is caused when the root is flicking
- This would cause spanning tree reconfiguration
- The effects are worse because of delayed port startup
- This is a DoS attack.

Attack B.6

- Cache overflow
- X sends packets with different (spoofed) source addresses,
- cause the cache of the bridges to overflow.
- following packets will be flooded, increasing the traffic of the network.
- This is a DoS attacks.

Assumption about the rbridges

- when an rbridge has multiple available paths to a given end-node, it only forwards packets using ONE of the available paths, probably the shorter one.

Overview

- identify possible attacks to current bridges.
- **threats related to the End-node Location Discovery Mechanism of rbridges.**
- threats related to the Link- State Protocol
- security aspects that limit the usage of the rbridges beyond the scope of application of current bridges.

Attack RB.1

- On-campus attacker X wants to establish a new communication with a node B pretending to be node A

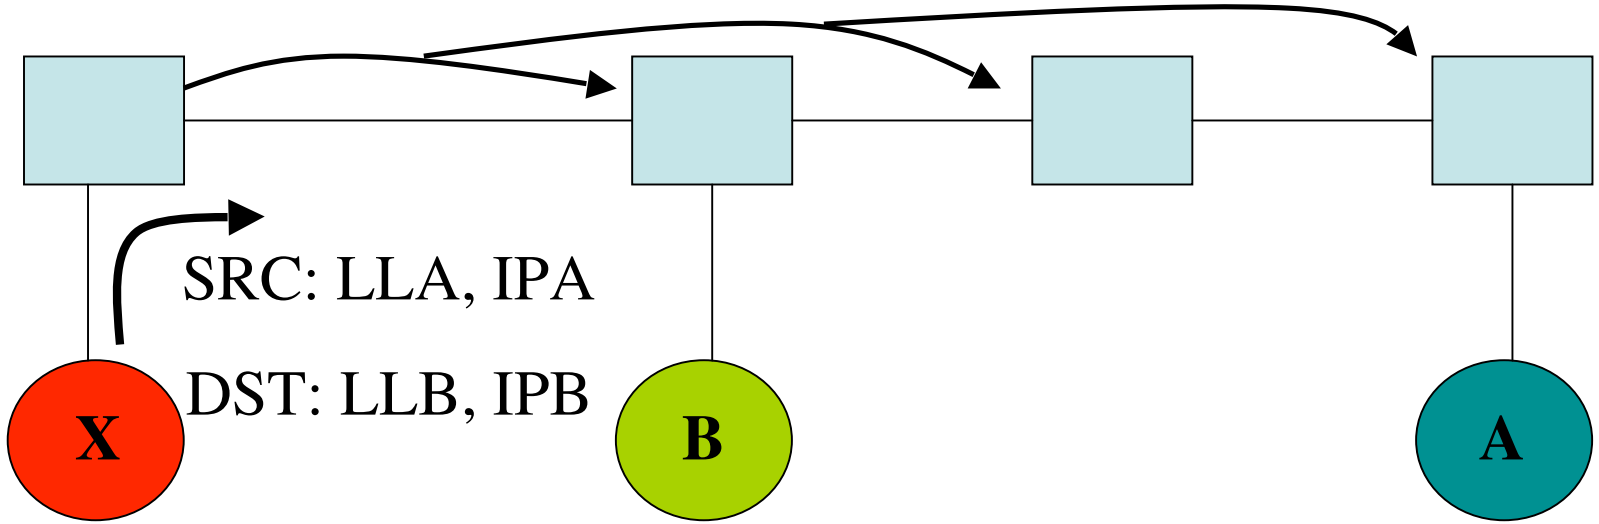

Attack RB.1

- On-campus attacker X wants to establish a new communication with a node B pretending to be node A

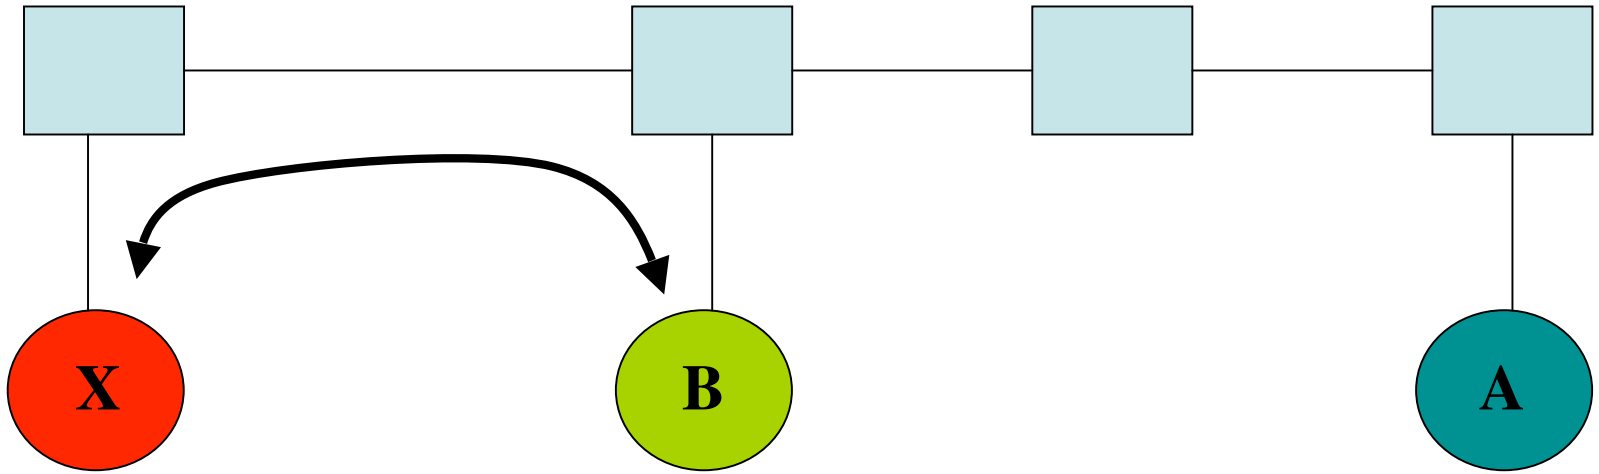

Attack RB.1

- The attack is effective if:
  - No other info about A is available or,
  - Dst(X,B) < Dst(A,B)

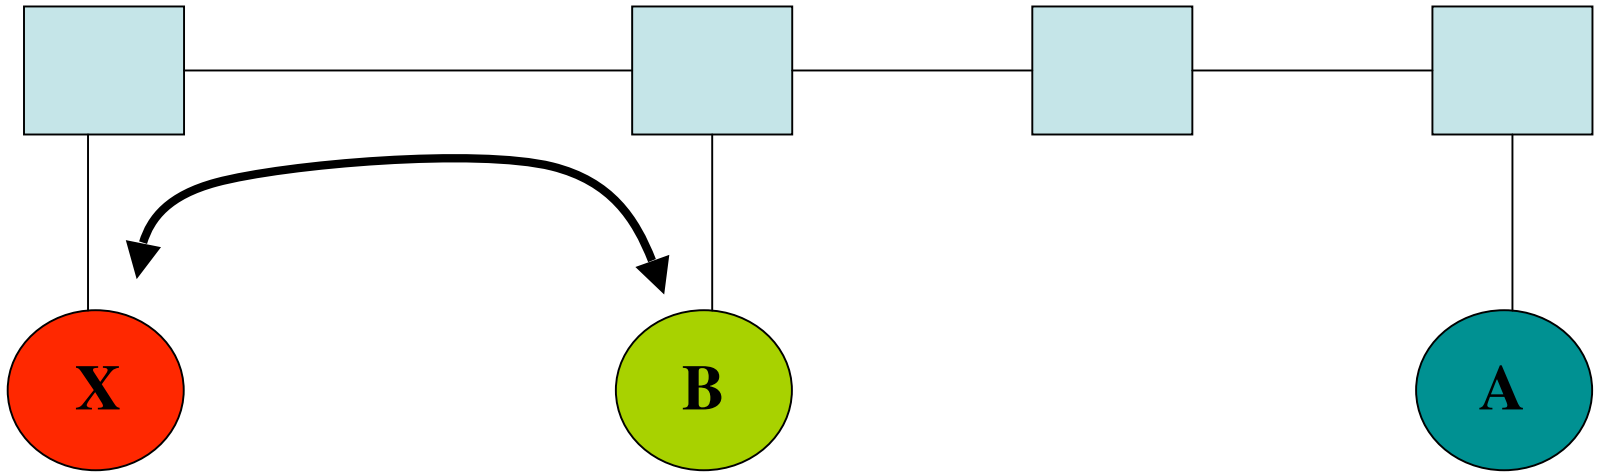

Attack RB.1

- The attack is effective if:
  - No other info about A is available or,
  - Dst(X,B) < Dst(A,B)

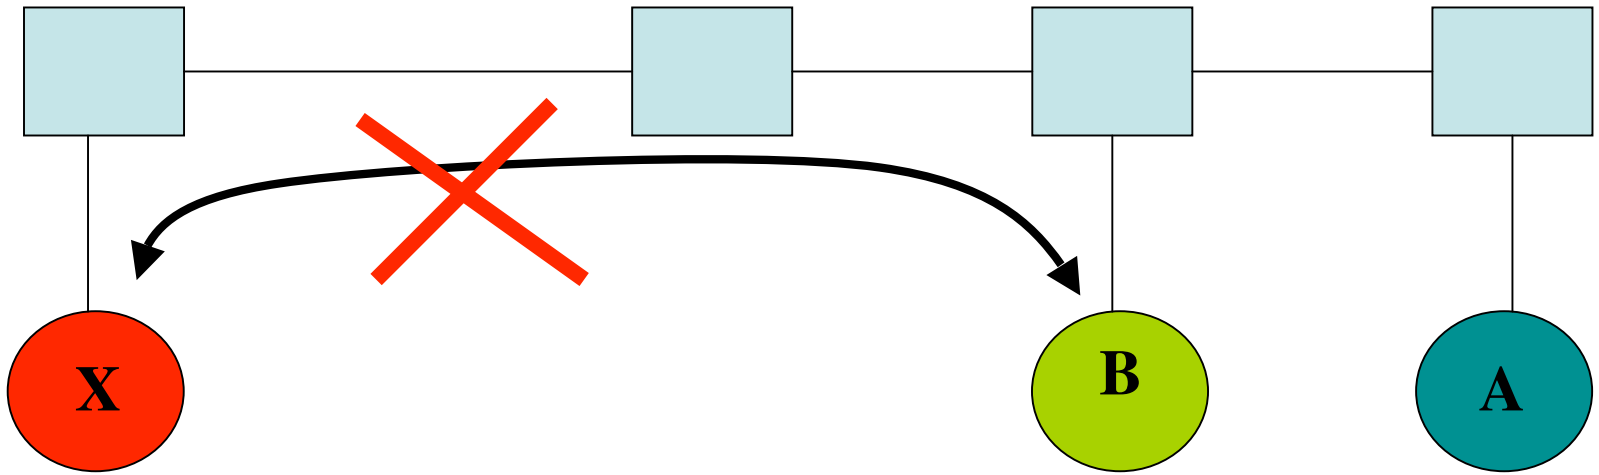

Attack RB.2

- On-campus attacker X wants to impersonate node A in any new communication established by node B.

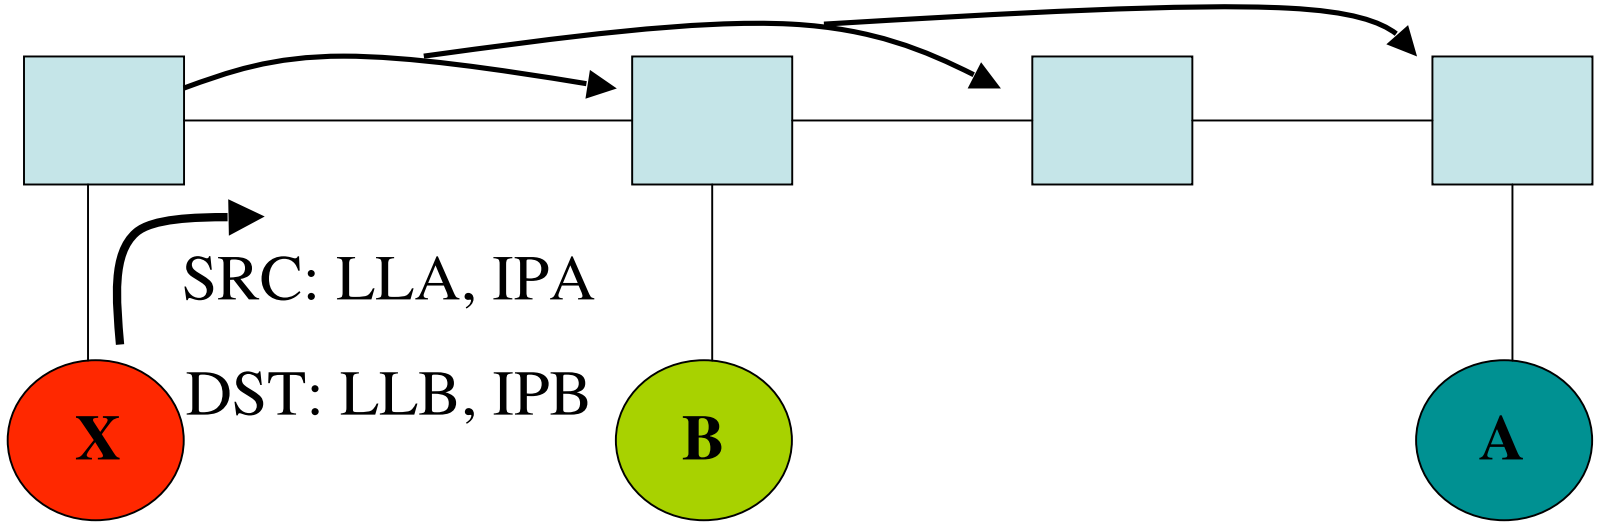

Attack RB.2

- On-campus attacker X wants to impersonate node A in any new communication established by node B.

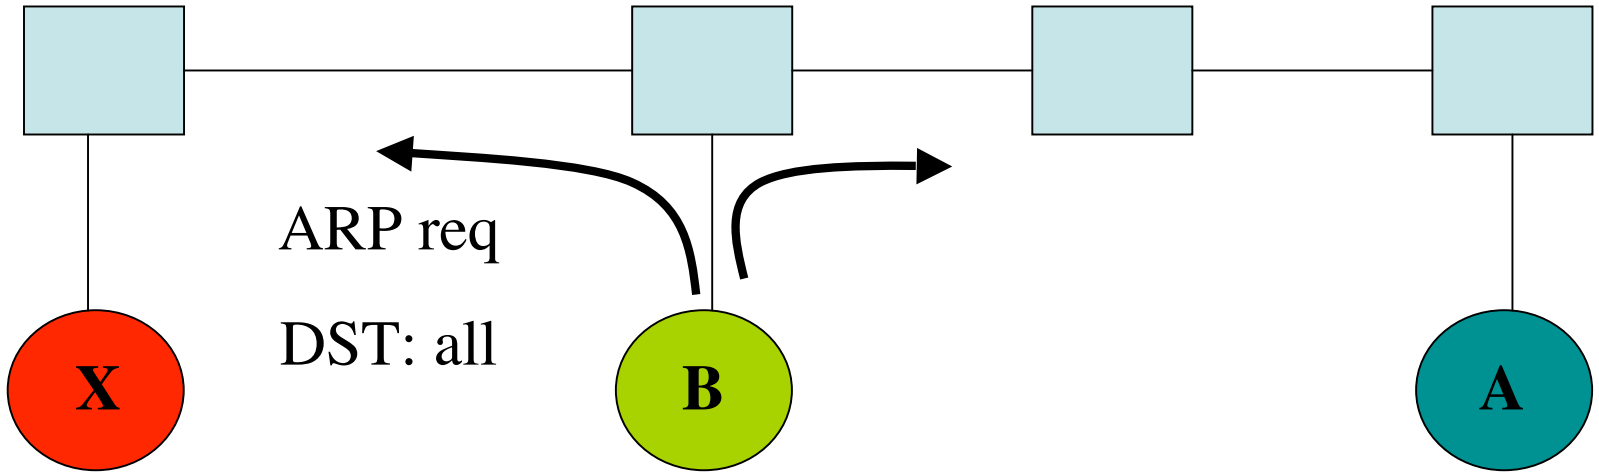

Attack RB.2

- On-campus attacker X wants to impersonate node A in any new communication established by node B.

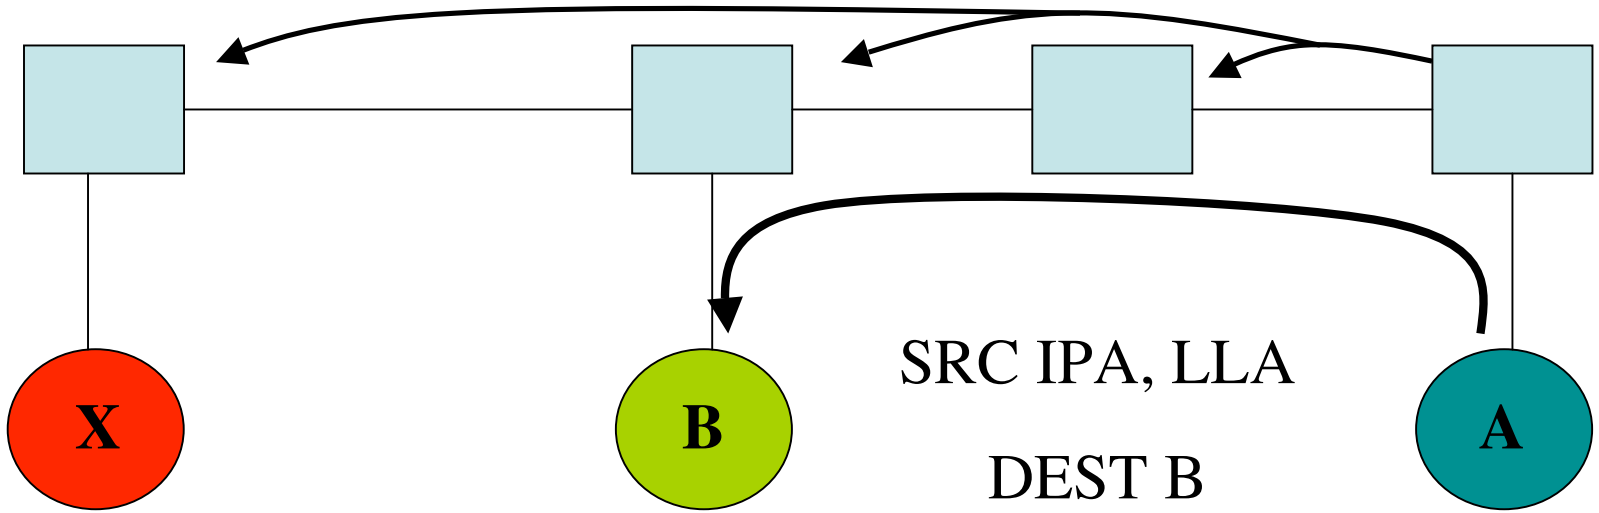

Attack RB.2

- The attack is effective if:
  - Dst(X,B) < Dst(A,B)
- Flooding optimization: may imply that the attack affects the whole campus, since A would not receive ARP requests

Attack RB.3

- The attacker wants to hijack an ongoing communication
- Same procedure
- The attack is effective if:
  - Dst(X,B) < Dst(A,B)

Attack RB.4

- Off-campus attacker X sends packets with a spoofed IP source address.
- Assumes that inter-rbridge forwarding is done based on IP addresses (not clear if true)
- Can cause packets to be directed to the ingress router
- No problem if IP addresses are not used for forwarding, or ingress filtering is in place

Overview

- identify possible attacks to current bridges.
- threats related to the End-node Location Discovery Mechanism of rbridges.
- **threats related to the Link- State Protocol**
- security aspects that limit the usage of the rbridges beyond the scope of application of current bridges.

Threats related to the Link-State Protocol

- Possibility to induce the rbridges to believe any topology
- Potential to extend the attacks to those nodes that are far away
- More analysis of specific routing protocol and its application to the rbridge is needed
- Not clear how worse is this w.r.t. bridged case where X sending periodic packets to random destinations
- In addition, possible attacks to the spanning tree similar to those to bridges
- Need to explore the need of configuring a password

Comparison with bridges

- Bridges: last one wins
- Rbridges: closer one wins, may be extended attacking the link state protocol
- Different characteristics, not obvious that one is better or worse

Overview

- identify possible attacks to current bridges.
- threats related to the End-node Location Discovery Mechanism of rbridges.
- threats related to the Link- State Protocol
- **security aspects that limit the usage of the rbridges beyond the scope of application of current bridges.**

Going beyond bridges

- Broadcast storms: All the campus is a single broadcast domain. Gabriel Motenegro
- Larger (campus-wide?) subnets means that spoofing inside a subnet is also easier, and ingress filtering granularity ("in-prefixspoofing") is more coarse, leading to more difficult user tracking. (Pekka Savola)
- Larger subnets do not mean good for firewalling between segments.(Pekka Savola)