Question 1: DNS Private Nameserver

Suppose Alice, Bob, and Eve all share a wireless Internet hotspot that includes a private DNS server such that DNS poisoning attacks are prevented by using randomized query IDs, randomized ports, and a firewall that observes all DNS traffic and has a poisoning-detection mechanism.

Assume that the poisoning-detection mechanism works as follows: if an incoming packet has a pending question with an incorrect query ID, because an earlier one outgoing was not emitted with the same query ID, then it ignores any DNS response until a few seconds pass without any incoming DNS response. It then re-issues the original DNS query and tries again.

- How can Eve perform DNS poisoning on Alice despite the firewall safeguard?
- How can Eve exploit the firewall safeguard to perform a denial of service attack?
- How can Eve exploit the firewall safeguard to perform a denial of service attack on only Bob's traffic?
- How can Mallory, who is outside of the network and cannot see any traffic, exploit the firewall safe guard to deny Alice, Bob, and Eve access to www.cbc.ca?

Question 2: DNS and TLS

Suppose Eve uses DNS poisoning to make others believe that bank.com is on Eve's IP. Suppose that Alice only have trusted root certificates in her trusted store, and Eve does not have the private key corresponding bank.com's public key. Suppose that Alice has already visited bank.com securely and bank.com uses HSTS. How can Eve can masquerade as bank.com and not set off any alarms on Alice's computer? Explain the steps in the process. How might this masquerading end up being detected?