Proxy ARP Detection in an IXP Peering LAN

Steven Bakker <Steven.Bakker@ams-ix.net>

Preparation

- For this exercise you will need:
  - A pencil and a piece of paper
  - An IPv4 subnet calculator
  - A set of dice (a d256 if you have it)
  - A host with an interface in your Peering LAN
  - The **`arping`** command

Why This?

On 11 August 2011, AMS-IX suffered a proxy ARP event that resulted in the tragic loss of many innocent BGP sessions.

Additional controls were put in place to prevent this from happening again.

These controls were recently enhanced.

Proxy ARP Refresher

- Wikipedia says (emphasis mine):
  - Proxy ARP is a technique by which **a proxy server on a given network answers the** Address Resolution Protocol (**ARP**) **queries for an IP address that is not on that network. The proxy** is aware of the location of the traffic's destination and **offers its own MAC address as the (ostensibly final) destination.** The traffic directed to the proxy address is then typically routed by the proxy to the intended destination via another interface or via a tunnel.
- *Source: https://en.wikipedia.org/wiki/Proxy_ARP*

What's the Problem?

- Proxy ARPers only send ARP *replies* for addresses that are not on the peering LAN.
- BGP routers only send ARP *requests* for addresses on the same network.
- Therefore, harmless on a peering LAN, right?
- Well...

What's the Problem?

- Proxy ARPers only send ARP *replies* for addresses that **they think** are not on the peering LAN.
- BGP routers only send ARP *requests* for addresses on the same network.
- Therefore, harmless on a peering LAN, right?
- Well...

The Problem with Proxy ARP

- Proxy ARPers only send ARP replies for addresses that **they think** are not on the peering LAN.
- Misconfigurations can lead to a confused router:
  - Wrong network prefix on the peering interface.
  - Wrong prefix length on the peering interface.

Causes of Misconfiguration

- Renumbering of peering LAN prefix:
  - IX has issued a new prefix and/or new prefix length, but customer has not reconfigured their interface yet.
- Quarantine (testing) IP address used in production:
  - IX moved port out of quarantine, but customer has not reconfigured their interface yet.

Deadly Combination

- **$bad_prefix + $proxy_arp =**
  $BIG_BLACK_HOLE

- Over time, ARP caches of peers will be poisoned.
- Traffic will be black-holed or diverted.
- BGP sessions *will* die.
- Can be slow or quick, depending on ARP cache timeout.
- Result is always painful.
- Very long recovery time.

☞ Early detection is key!

Proxy ARP Prevention/Mitigation

- Prevention:
  - Use inbound ARP filters on the customer ports.
  - Not available on all platforms.
- Mitigation:
  - Check often.
  - Use long ARP cache timeouts (≥ 4 hours).
    - Pro: Slows the spread of ARP cache poisoning.
    - Con: May slow the recovery of ARP cache poisoning.
  - Use automated tools to “unpoison” ARP caches.
    - The arpsponge’s **`inform`** command can help.

Proxy ARP Detection: Method #1

- During “port quarantine”:
  - Send ARP queries for a few “well known”, external IP addresses:
    - `www.google.com`
    - `www.cisco.com`
    - `www.microsoft.com`
    - `192.168.1.1`
- Anyone answering is obviously a proxy ARPer:
  - Do not move port to production until this is resolved.
- Problem:
  - No guarantee for what happens once a port is in production.

Proxy ARP Detection: Method #2

- Use Method #1 and add probing for ports in production.
- In the peering LAN, send ARP queries for a few “well known” IP addresses.
- Anyone answering is obviously a proxy ARPer:
  - Raise alarm and/or disable offending port(s).
- Problem:
  - Does not distinguish between harmless and fatal proxy ARP.
    - ☞ Simple annoyances result in **CODE RED** alerts.

Proxy ARP Detection: Method #3

- Use Method #2 and add additional probing for ports in production.
- Send ARP queries for unused addresses on the peering LAN.
- Anyone answering is either a Proxy ARPer or an IP hijacker (or the `arpsponge`)
- Problem:
  - You may not have spare IP addresses on your peering LAN.
  - May fail to detect harmful proxy ARP.
    - Tested IP address may fit in offender's notion of the local prefix:
      - Correct prefix: 192.168.2.0/23
      - Proxy ARPer: 192.168.2.1/24
      - Test IP: 192.168.2.2

Detection Requirements

- Works for both quarantine and production.
- Does not require free or reserved IP addresses.
- Works for both wrong prefix and wrong prefix length.

Proxy ARP Detection: TL;DR

- Split peering prefix in two:
  - E.g.: 192.168.2.0/23 → 192.168.**2.0/24** + 192.168.**3.0/24**
- From each subnet, pick two random IP addresses.
- Send ARP queries for each of these four addresses.
- Any device answering for ≥ 2 of these IP addresses?

  ⟹ You found a harmful Proxy ARPer.

Proxy ARP Detection: “Proof”

- Offender’s interface prefix either falls:
  - Outside the peering prefix,
  - Somewhere in the lower half of the peering prefix,
  - Somewhere in the upper half of the peering prefix.
- Choice of random addresses guarantees that:
  - At least two IP addresses will be outside the offender’s prefix.
- Hence:
  - Offender will send proxy ARP replies for at least two addresses.

Case Studies

#1: Prefix OK

- Peering prefix: /21
- Offender’s prefix correct

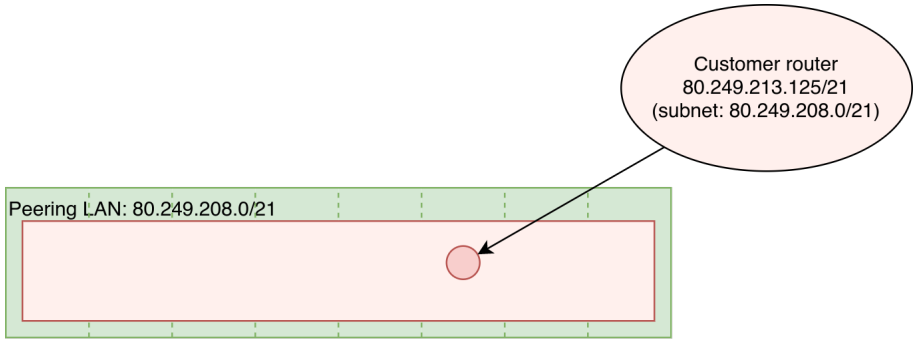

#1: Prefix OK

- Peering prefix: /21
- Offender's prefix correct
- All IP addresses are in offender's subnet
  - No proxy ARP replies
- "Harmless" proxy ARP

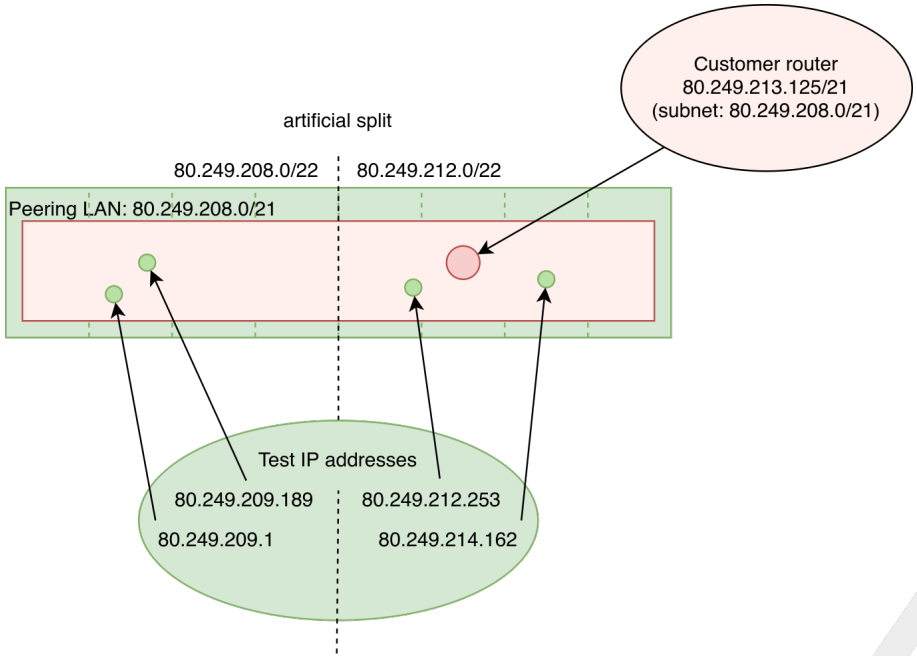

#2: Outside Peering LAN

- Peering prefix: /21
- Offender’s prefix outside peering LAN

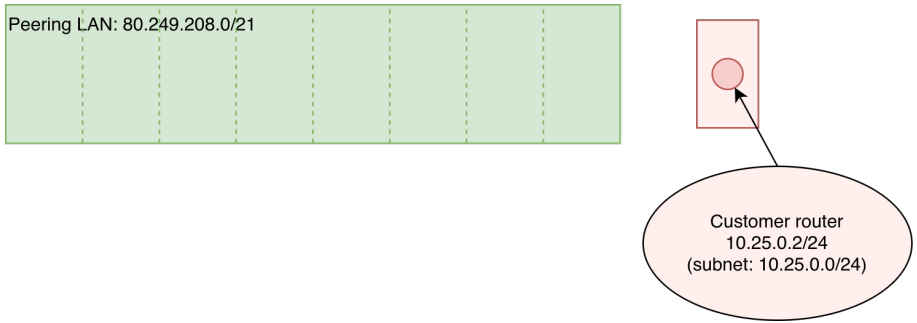

#2: Outside Peering LAN

- Peering prefix: /21
- Offender's prefix outside peering LAN
- Proxy ARP for all four addresses
- Happened at AMS-IX on 11 August 2011

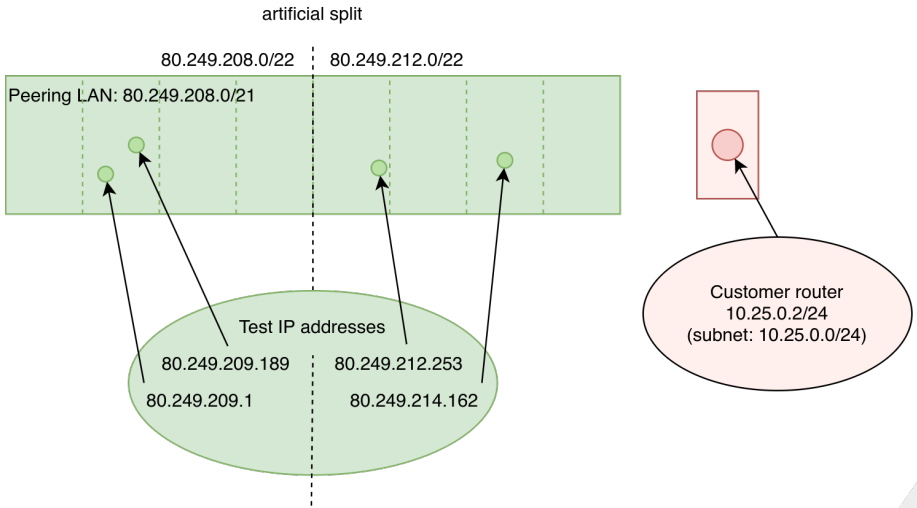

#3: Prefix Length Off By 1

- Peering prefix: /21
- Offender's prefix: /22

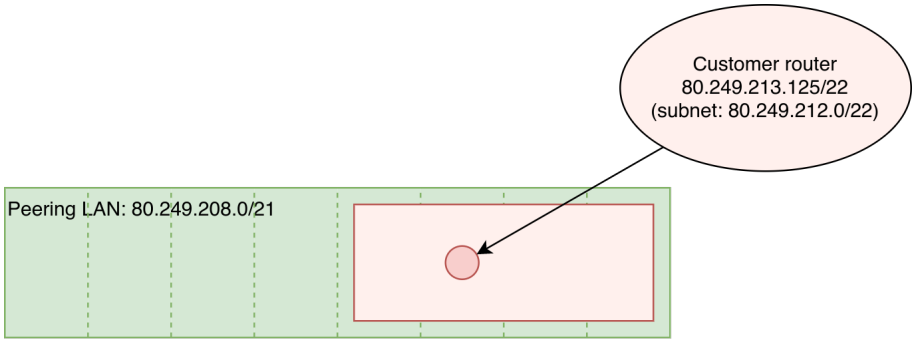

#3: Prefix Length Off By 1

- Peering prefix: /21
- Offender's prefix: /22
- Proxy ARP for two addresses

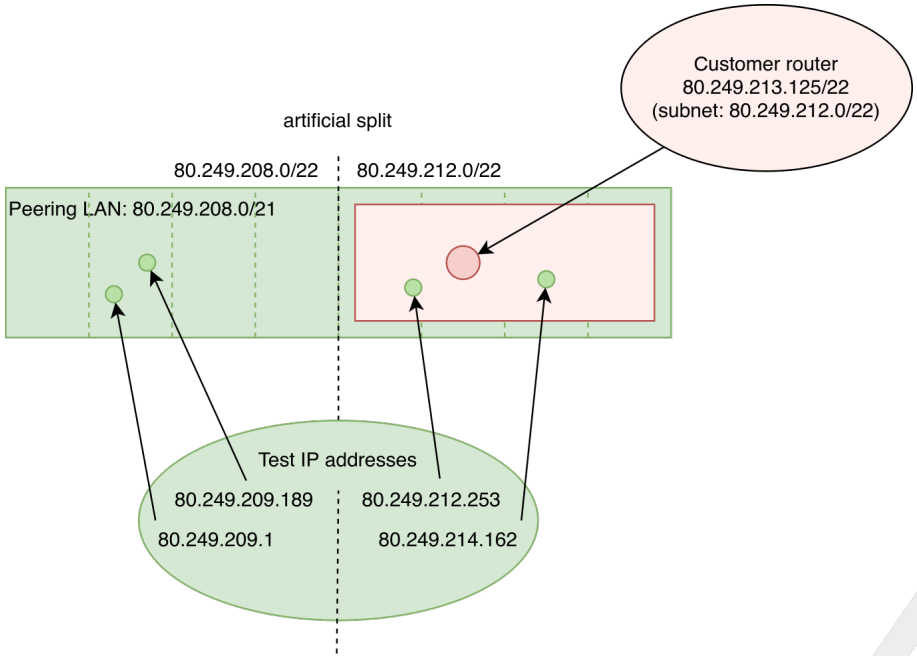

#4: Prefix Length Way Off

- Peering prefix: /21
- Offender’s prefix: /24
- Proxy ARP for:

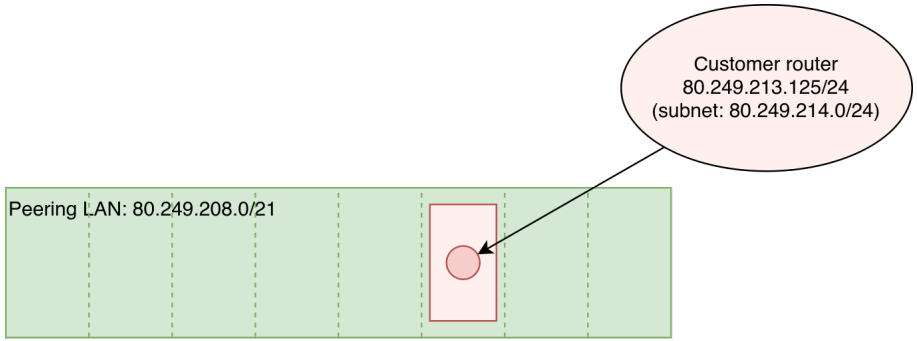

#4: Prefix Length Way Off

- Peering prefix: /21
- Offender's prefix: /24
- Proxy ARP for:
  - 4

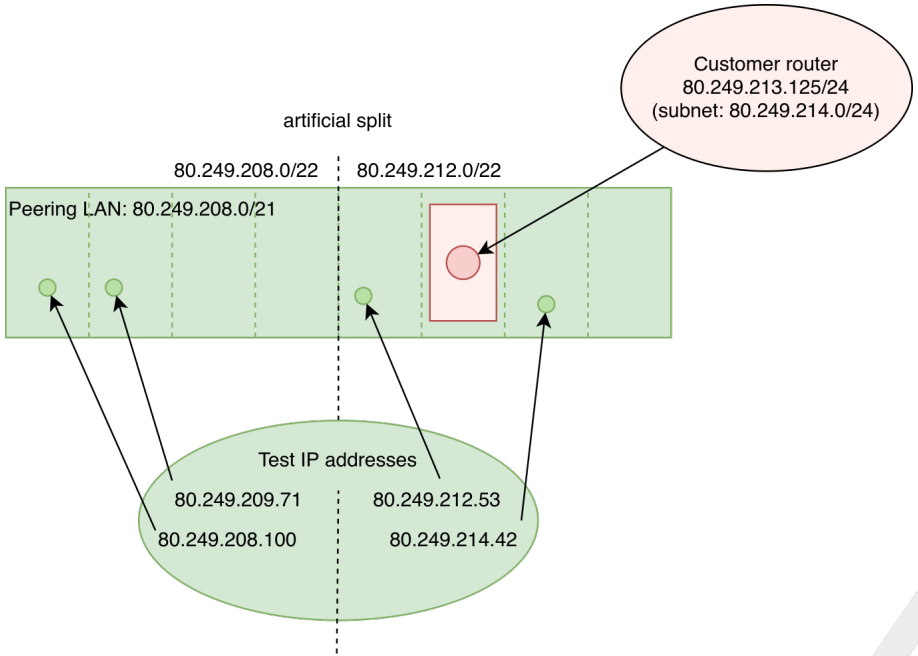

#4: Prefix Length Way Off

- Peering prefix: /21
- Offender's prefix: /24
- Proxy ARP for:
  - 4, 3

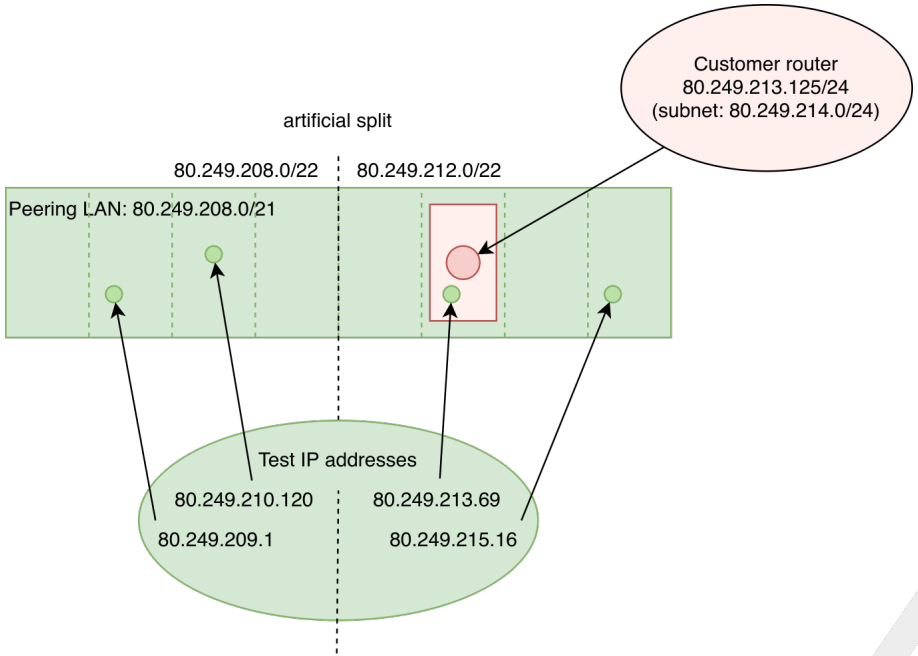

#4: Prefix Length Way Off

- Peering prefix: /21
- Offender's prefix: /24
- Proxy ARP for:
  - 4, 3, or 2 addresses

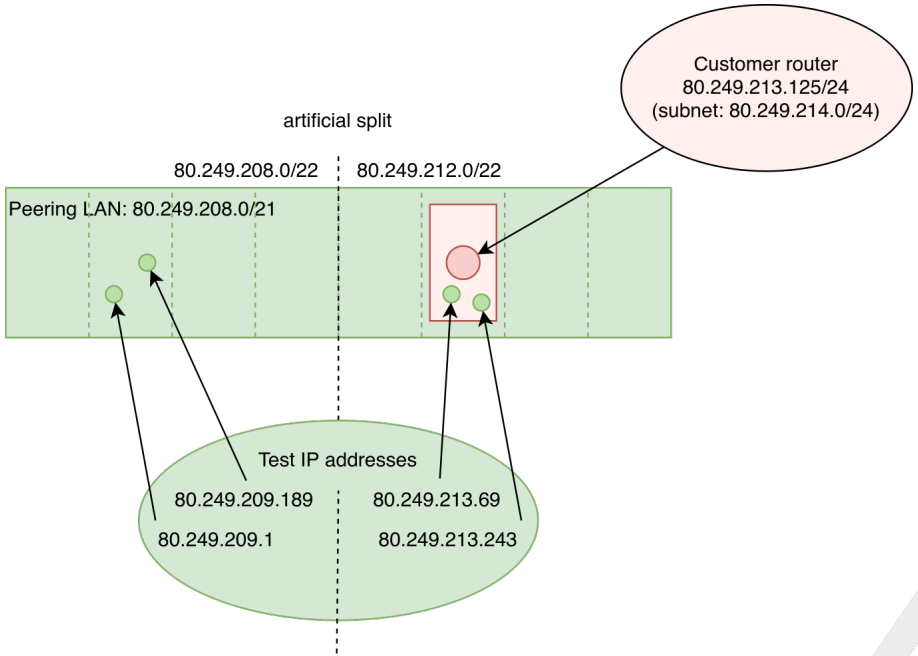

Proxy ARP Detection at AMS-IX

amsix

Proxy ARP Detection at AMS-IX

- Quarantined ports cannot go to production if any proxy ARP is detected.
- Every 3 minutes:
  - Check for “harmless” proxy ARP in peering LAN:
    - Open ticket, send nag mails.
  - Check for harmful proxy ARP:
    - Disable port, open ticket, raise alarm.
- Only found “harmless” proxy ARPers since the 2011 incident.
  - Most during port quarantine, a few in production.

Conclusion

- Proxy ARP on a peering LAN can potentially be fatal.
- In absence of adequate ARP filtering, early detection is crucial.
- Checks during port quarantine help a lot.
- Frequent checks in production LANs remain necessary.

- Questions?
  - Steven Bakker <Steven.Bakker@ams-ix.net>