NAT Revelio: Detecting NAT444 in the ISP

Andra Lutu, Marcelo Bagnulo,
Amogh Dhamdhere, kc claffy

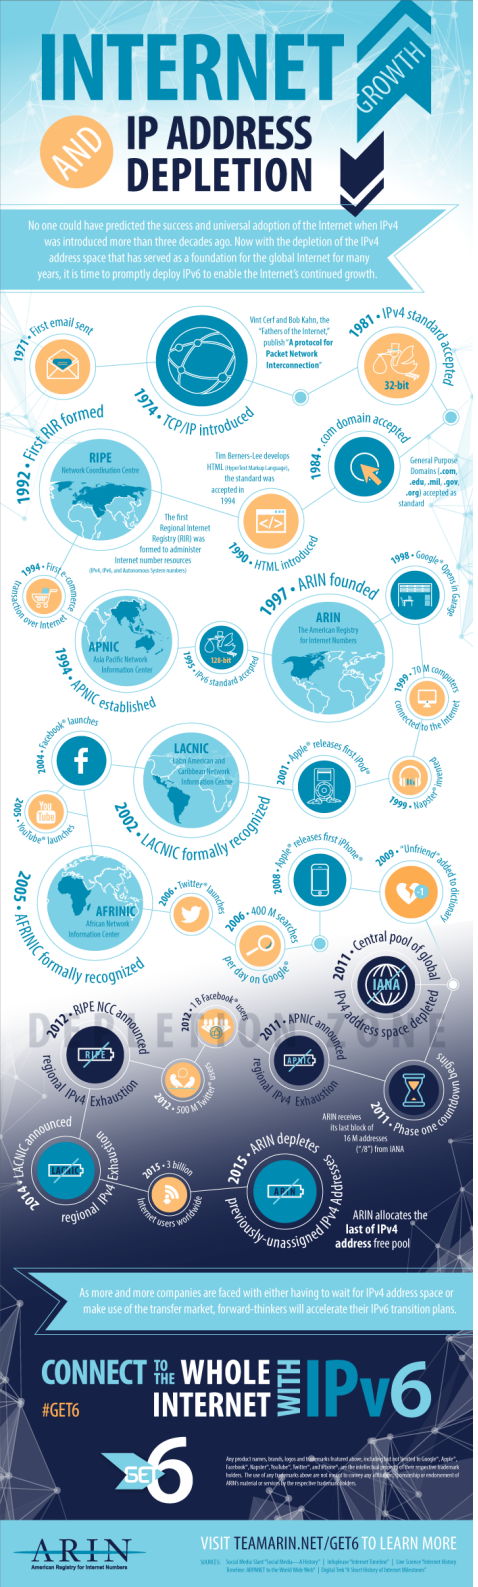

Network Address Translation (NAT)

- The success of the Internet led to the depletion of the IPv4 address space
- IPv6 - only viable solution, very slow adoption
- Network Address Translation – prolongs the life of IPv4, by enabling address sharing
- Criticism:
  - As broadband becomes prevalent, NAT devices turn into performance bottlenecks
  - NAT hinders certain applications (e.g., VoIP)
  - Breaks Internet end-to-end principle
  - Inhibits the conversion to IPv6 in the medium term

Traditional NAT (NAT44)

DSL Access Network mapped to the LMAP Reference Path

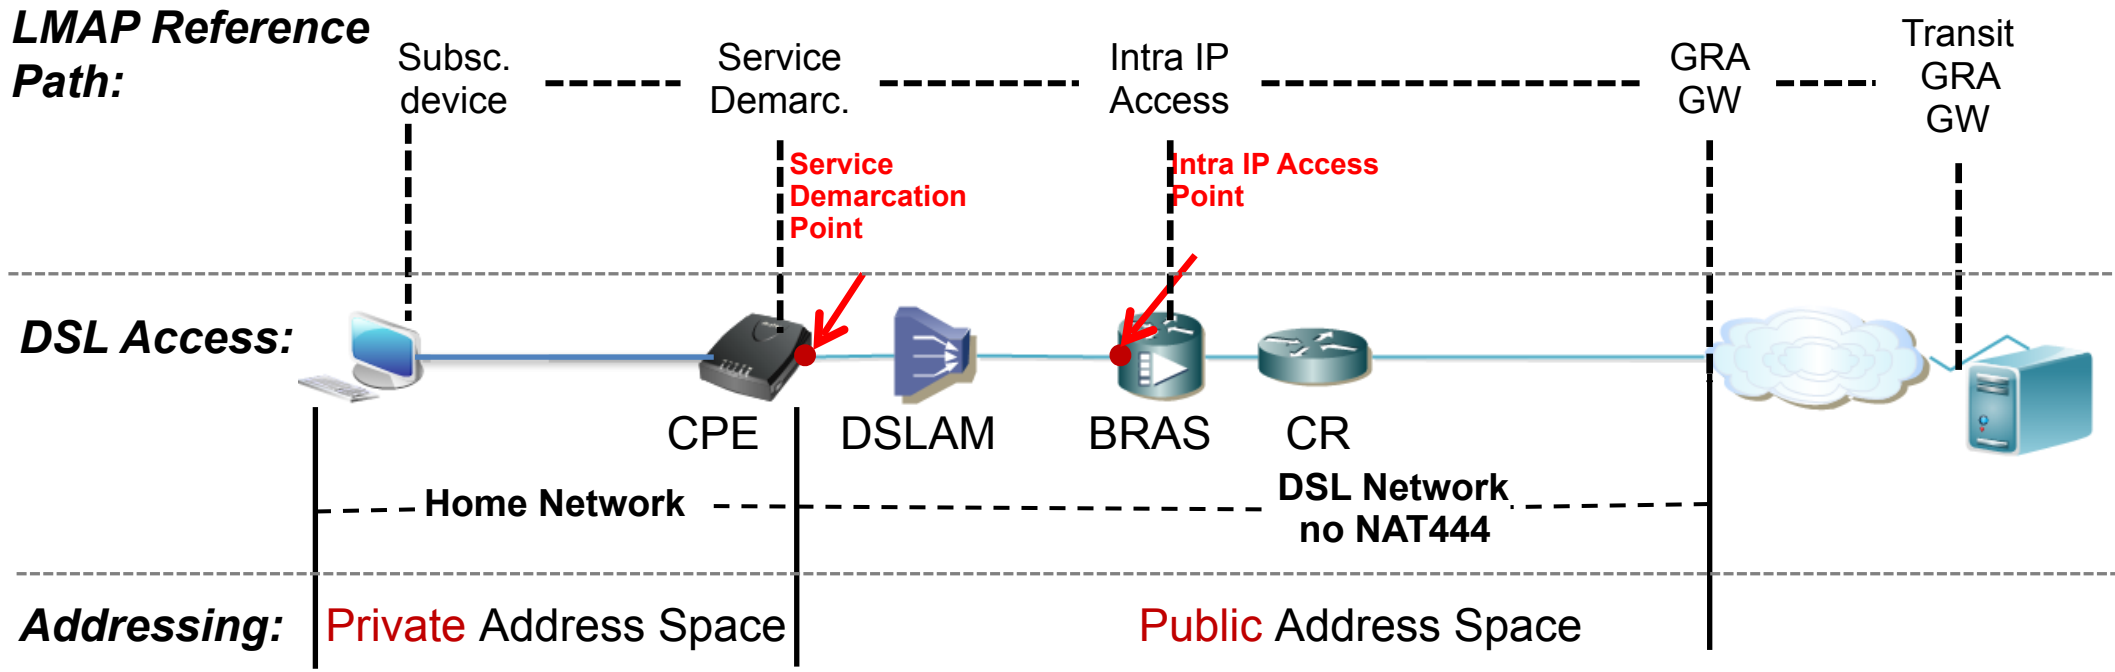

NAT444 / Carrier Grade NAT/ Large Scale NAT

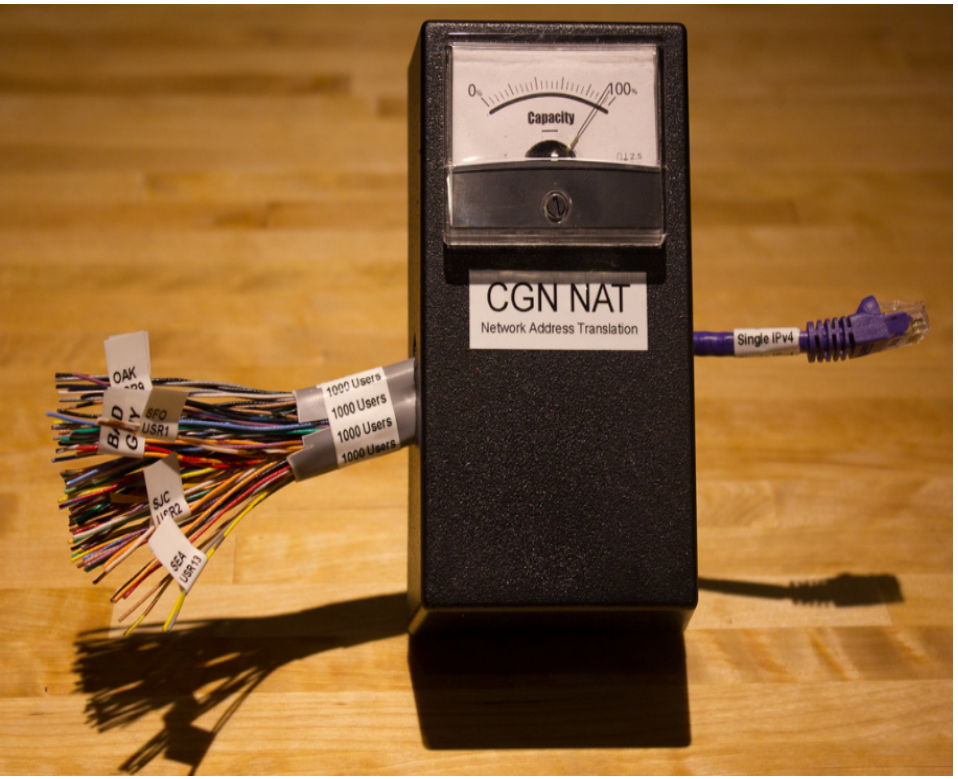

What it breaks (*RFC7021: Assessing the Impact of Carrier-Grade NAT on Network Applications*)

- On-line gaming
- Video streaming
- BitTorrent
- VPN & Encription
- VoIP
- … etc.

Large Scale NAT (NAT444)

DSL Access Network *with NAT444 deployment*

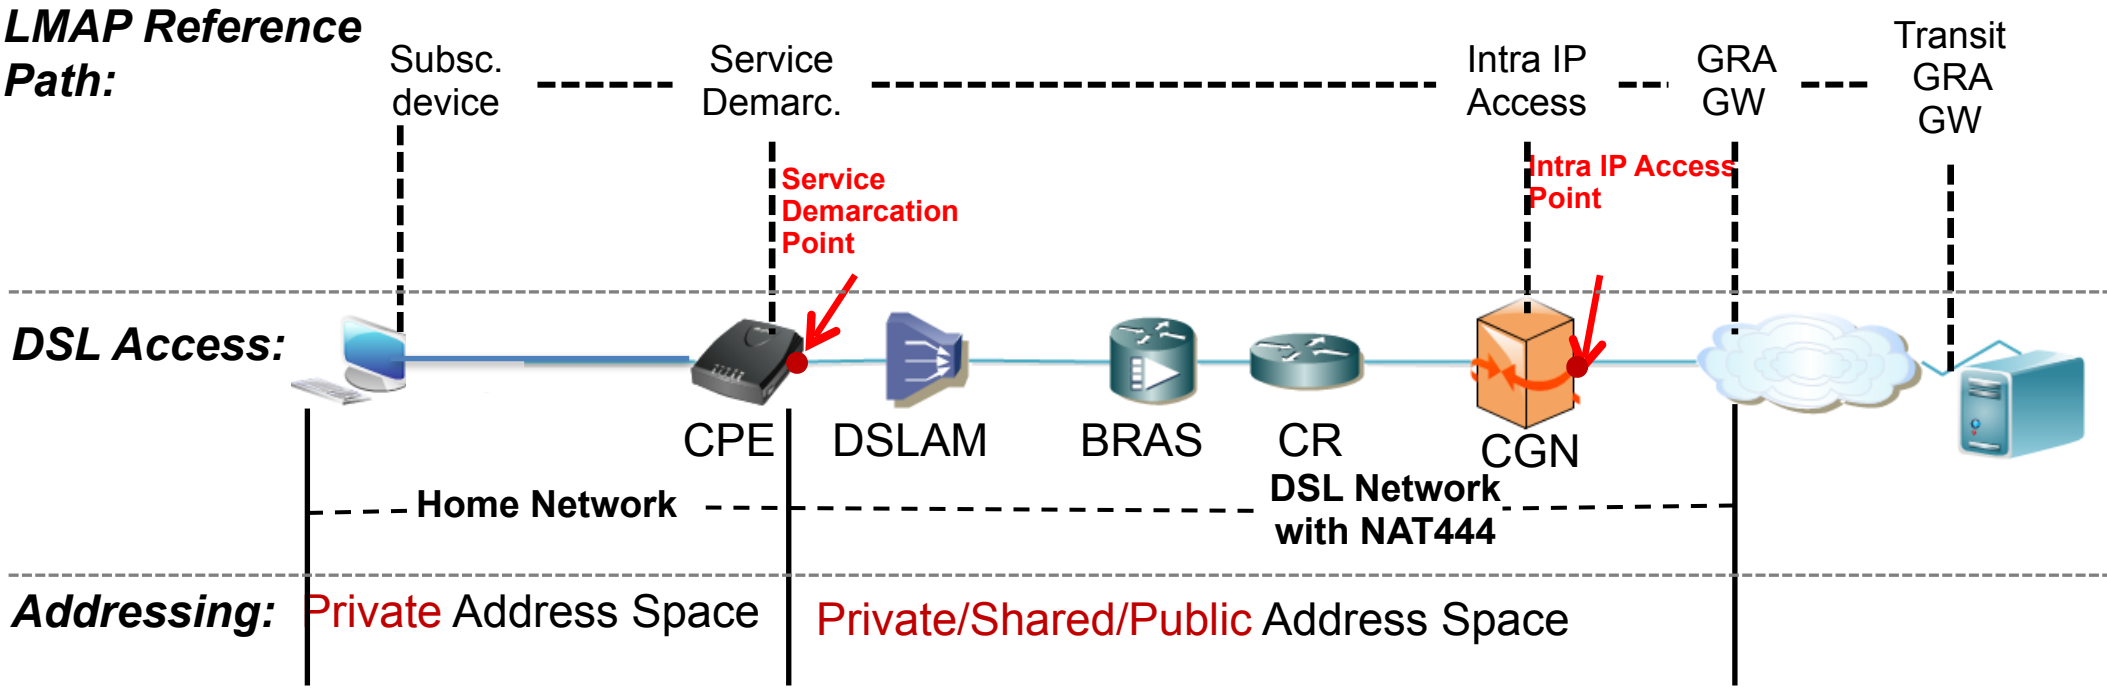

NAT Revelio

for Measuring Broadband America (MBA)

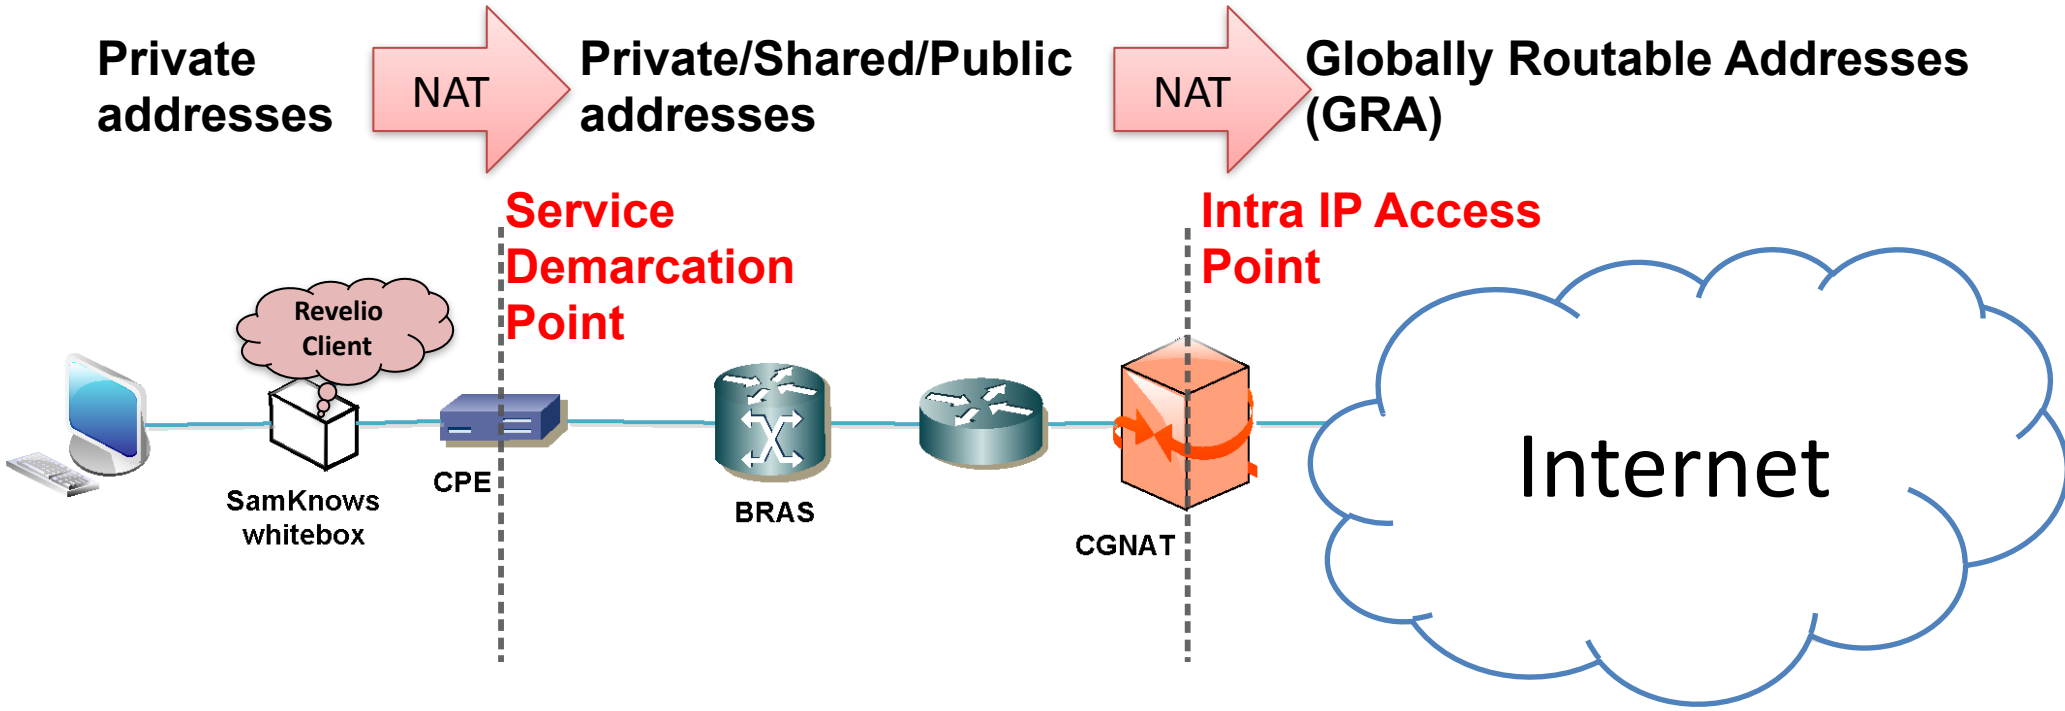

- Detect the usage of *private/shared address space* beyond the Service Demarcation device (CPE), in the ISP access network
- Detect the location *(home network or ISP access network)* device doing the translation to the GRA of the subscriber

NAT Revelio: Design Challenges

- Avoid NAT444 false positives:
  - Diverse home network configurations, e.g. In-home cascaded NAT, with probe ***NOT*** connected directly to the CPE (that is the Service Demarcation device)
  - Diverse ISP configurations and deployments e.g. use of private IP addresses internally even if they don't do NAT444

Incorrect Mapping with the LMAP Reference Path:

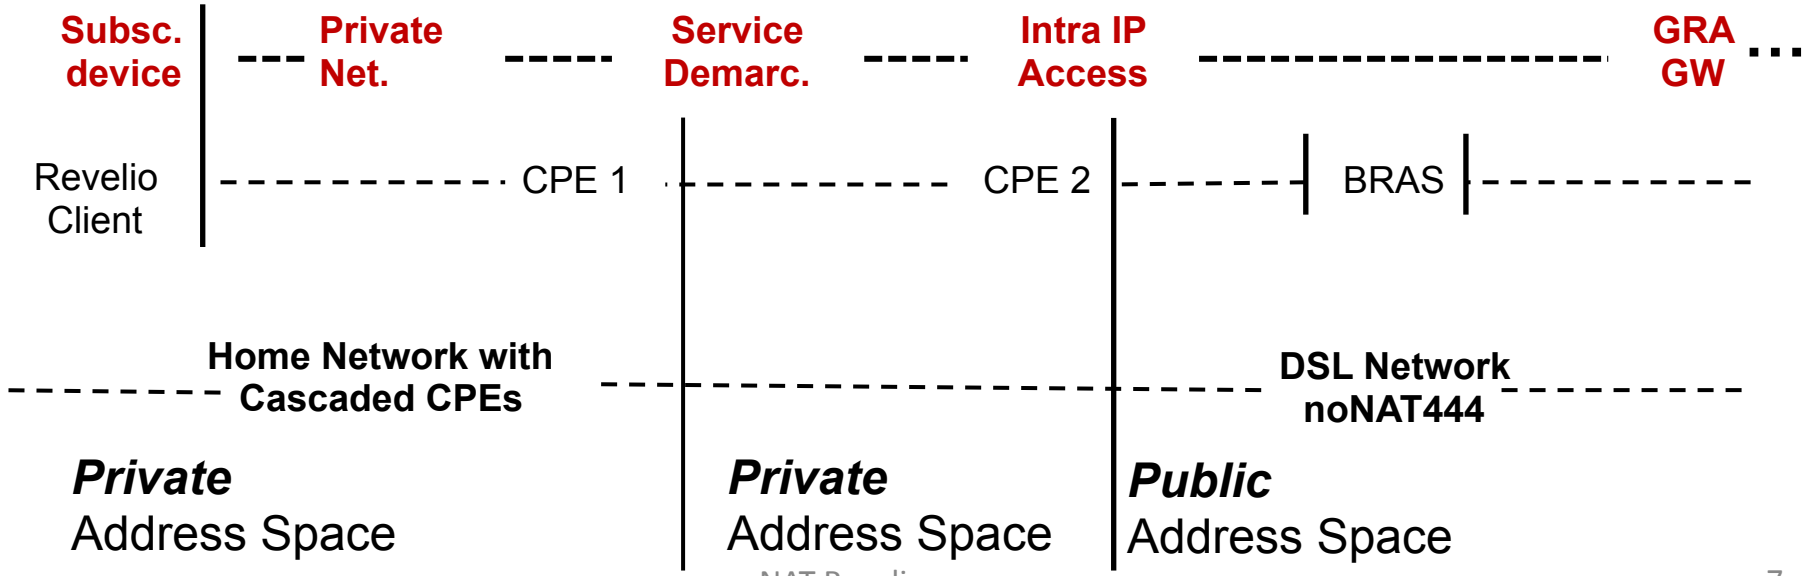

NAT Revelio: Design Challenges

- Need to detect the ***access link,*** to further delimit the ***access network*** and the ***home network***
- Allows to eliminate some false positives

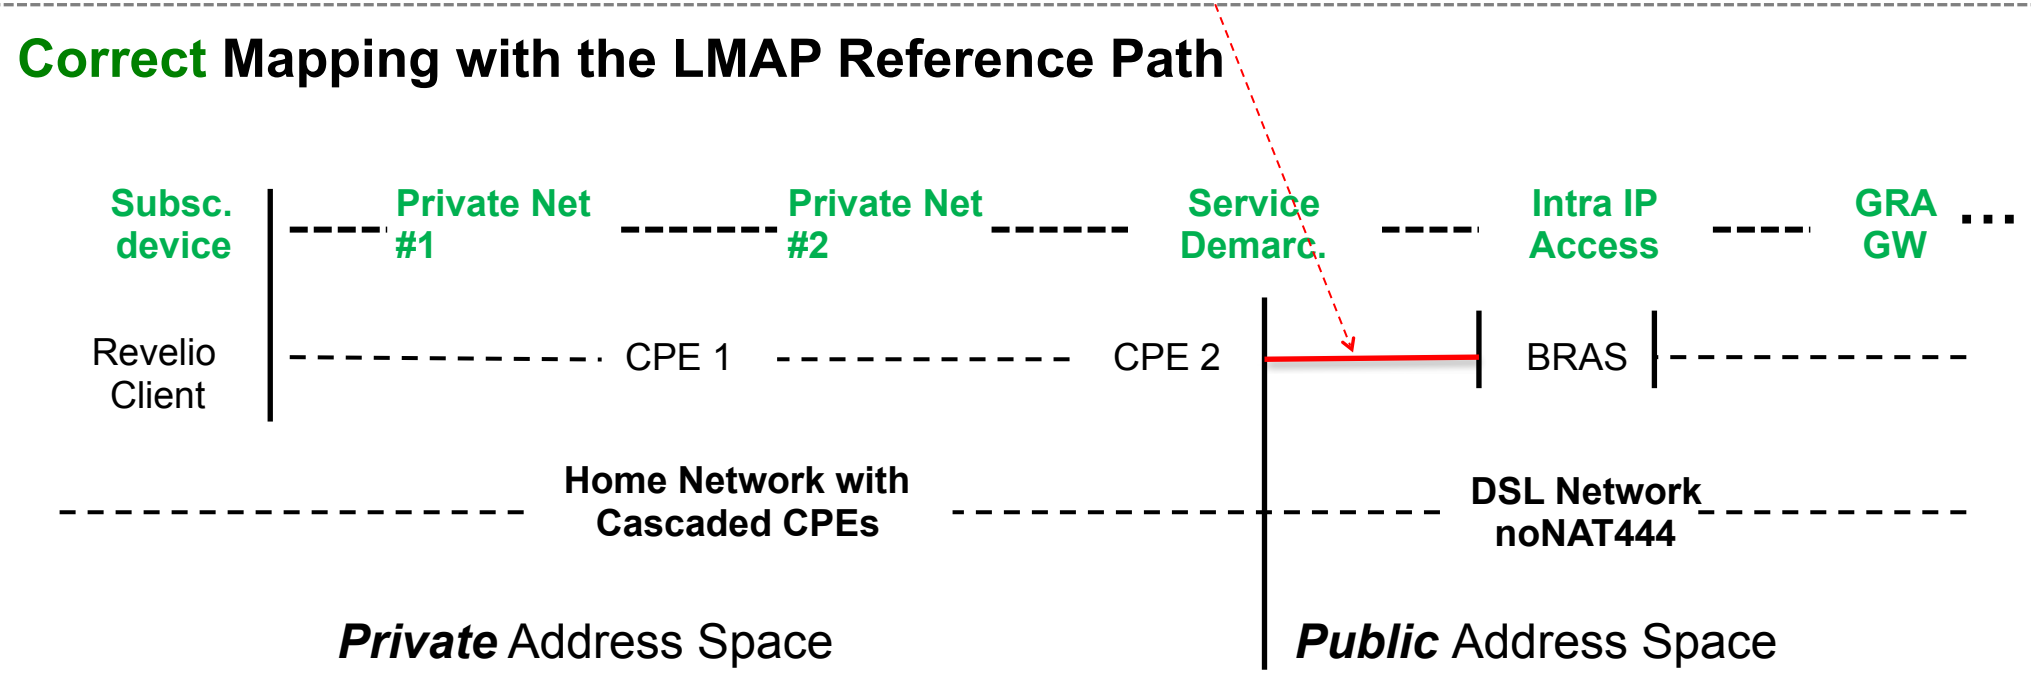

NAT Revelio

- The NAT Revelio test suite includes 2 phases:
  - ***Environmental Characterization***
    - *Understand the environment hosting the device running the Revelio Client*
  - ***NAT444 Discovery***
    - *Detection of signals that the ISP might deploy a NAT444 solution in the ISP access network*

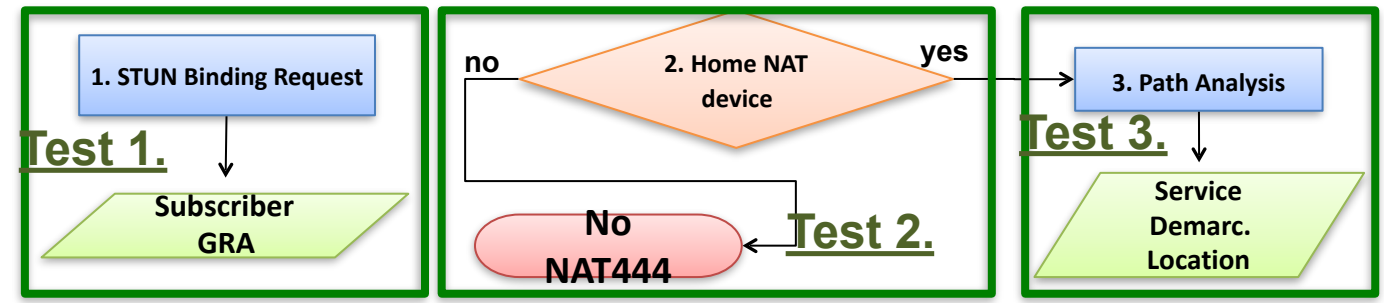

Phase 1) Environment Characterization

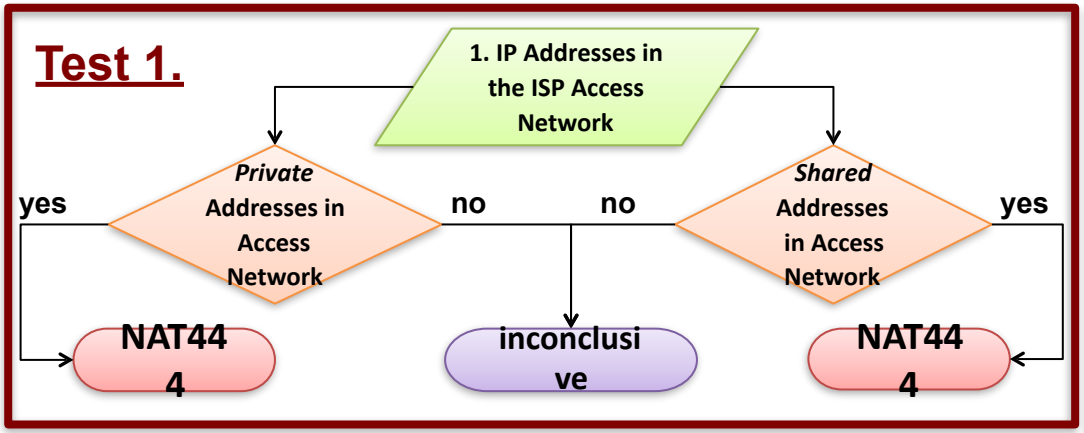

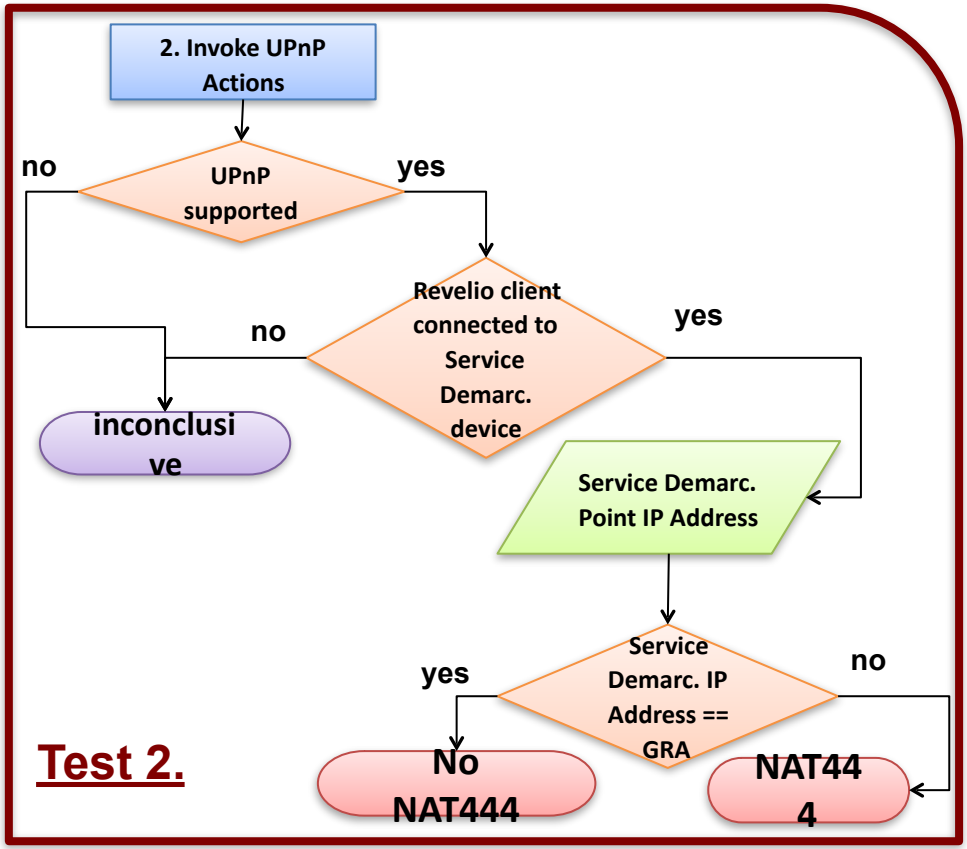

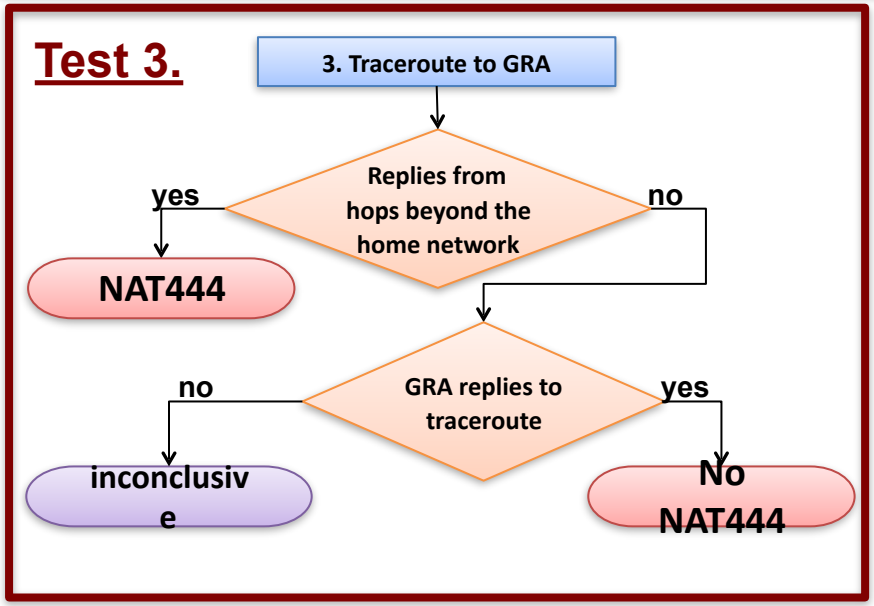

Phase 2) NAT444 Discovery

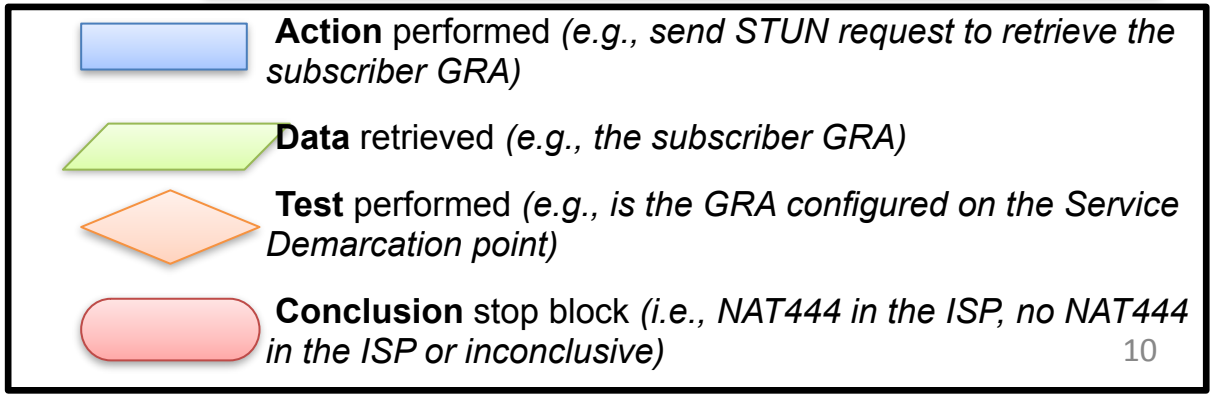

Environment Characterization

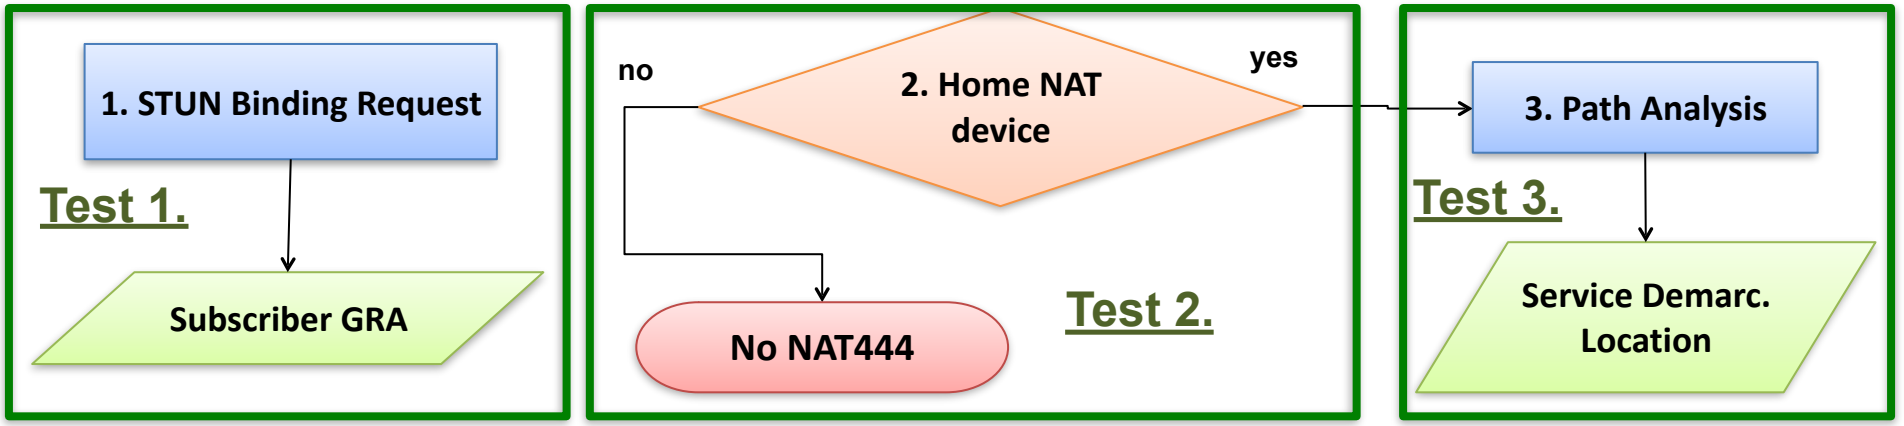

- This phase aims to determine:
  - **Test 1**: The GRA of the subscriber running the Revelio client
  - **Test 2:** Whether the subscriber is behind at least one level of NAT (i.e., the CPE performs the NAT function)
  - **Test 3:** Which is the position of the Revelio client related to the Service Demarc. Device (i.e., the position of the access link relative to the Revelio client)

Environment Characterization

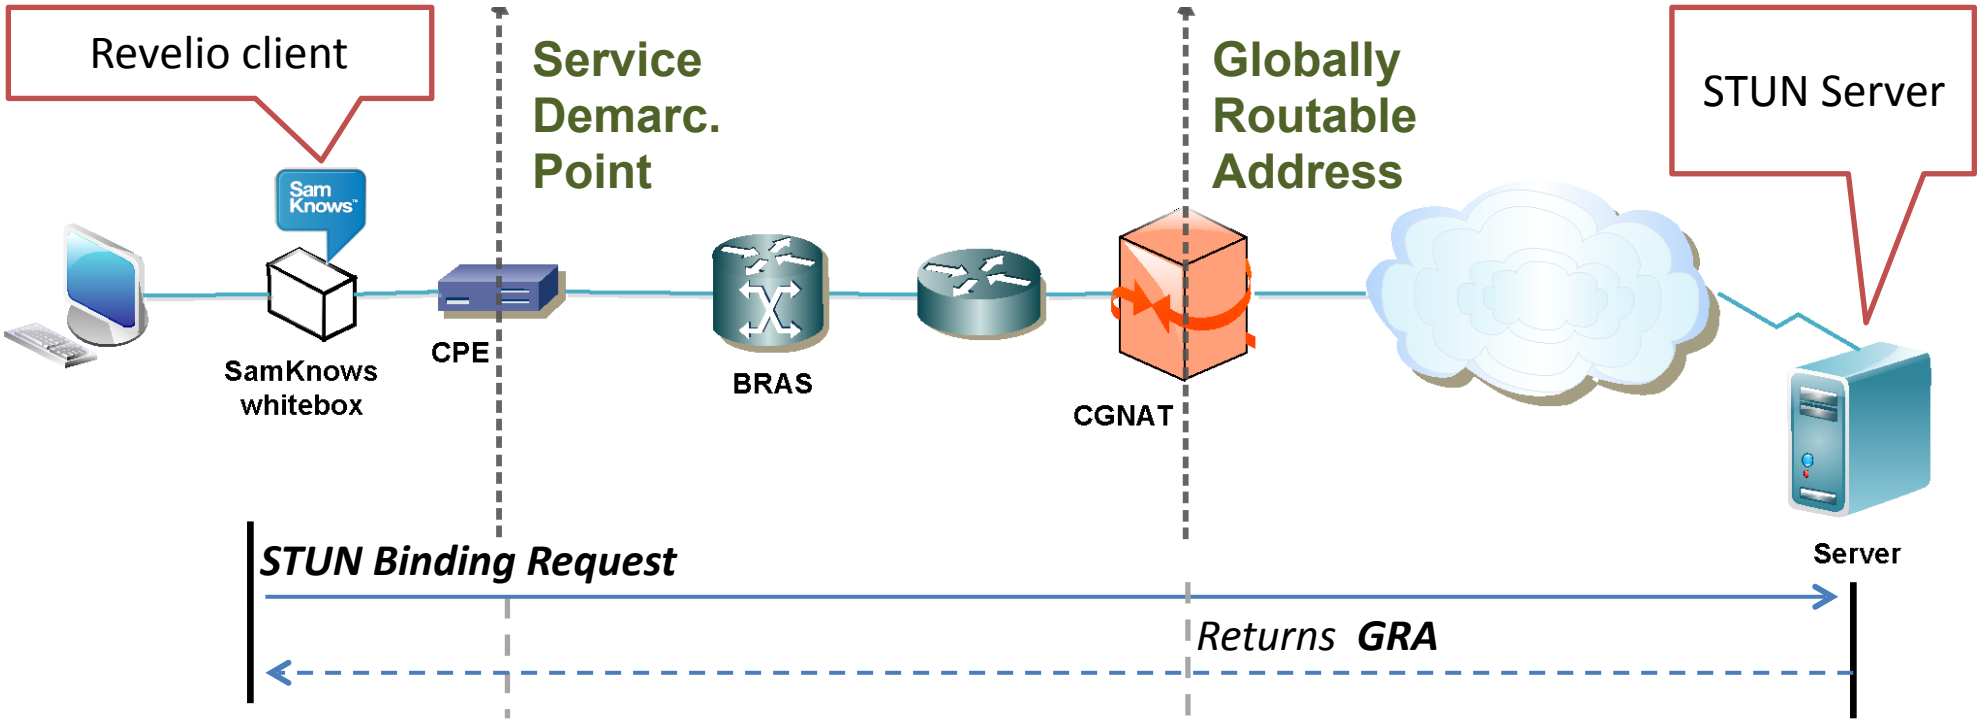

- **Test 1**: Subscriber GRA
- We send a STUN binding request to a public STUN server to retrieve the GRA of the subscriber
- We use this information in subsequent tests in the Revelio test-suite

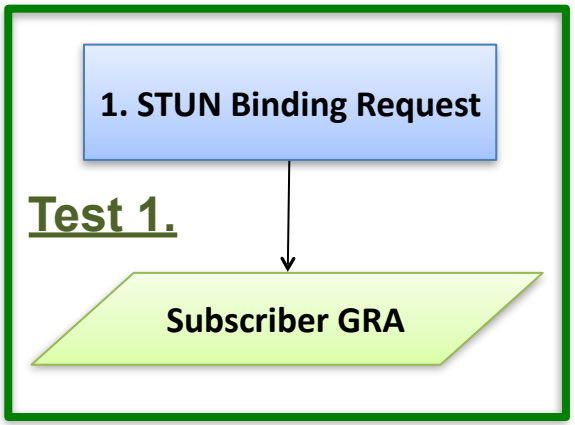

Environment Characterization

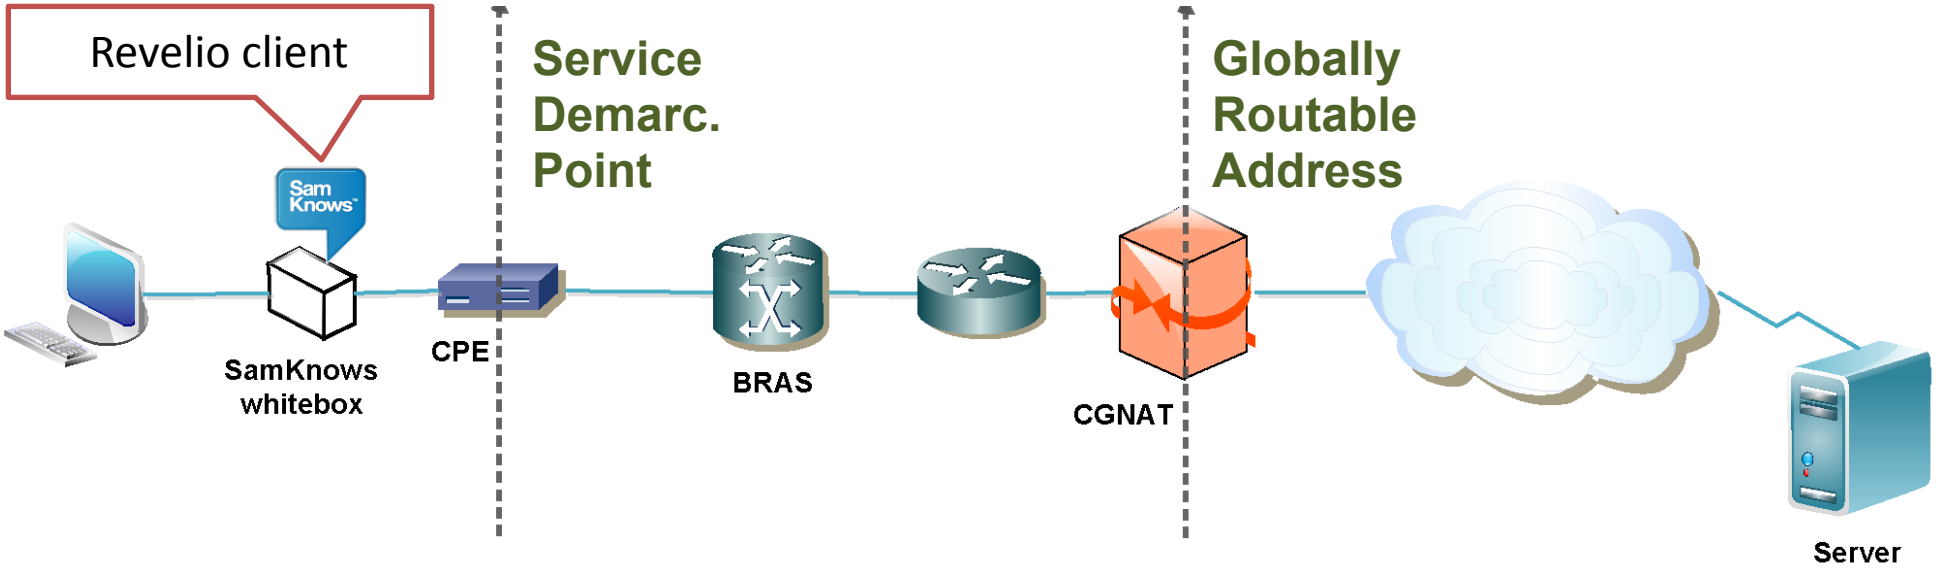

- **Test 2**: Home NAT device
- We compare the IP address of the device running the Revelio Client (e.g., SamKnows Whitebox) with the GRA we retrieve in the previous test
  - If the GRA is configured on the device running the Revelio Client (thus, no home NAT device), we conclude **no NAT444**
  - Otherwise, we continue testing

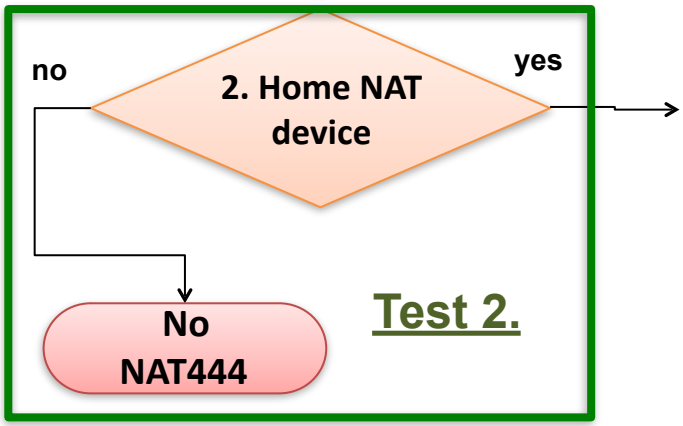

Environment Characterization

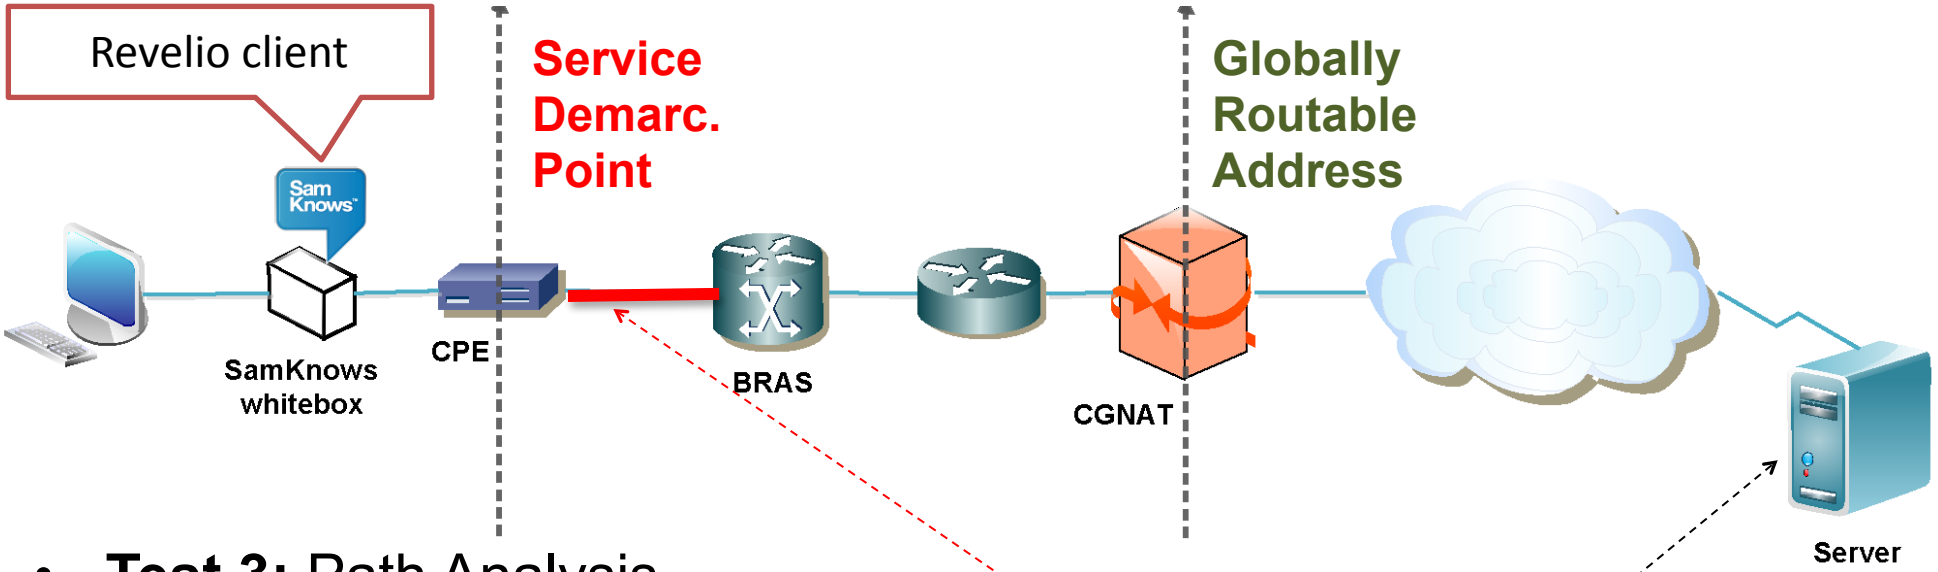

- **Test 3:** Path Analysis
- We determine the location of the ***access link (and Serv. Demarc. Point)*** relative to the Revelio client using repetitive traceroutes to an external target
- For this, we assume that the access link is the bottleneck (i.e., the link with the highest propagation delay)
- We measure latency per link and identify the one with an order of magnitude increase compared to the neighboring links

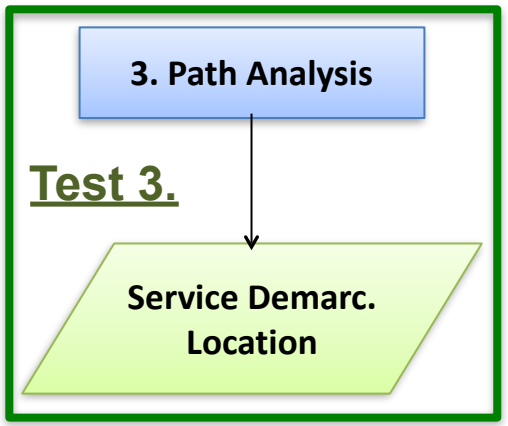

NAT444 Discovery

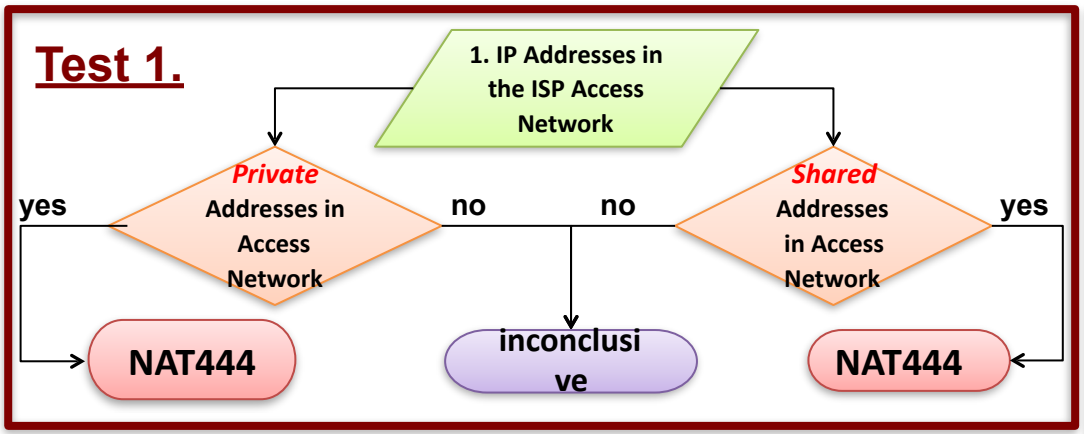

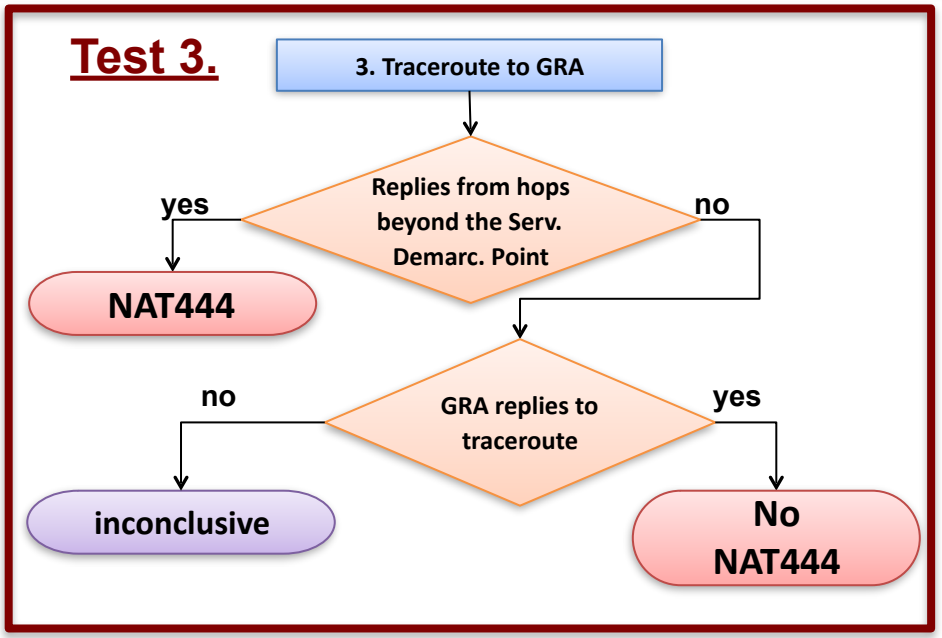

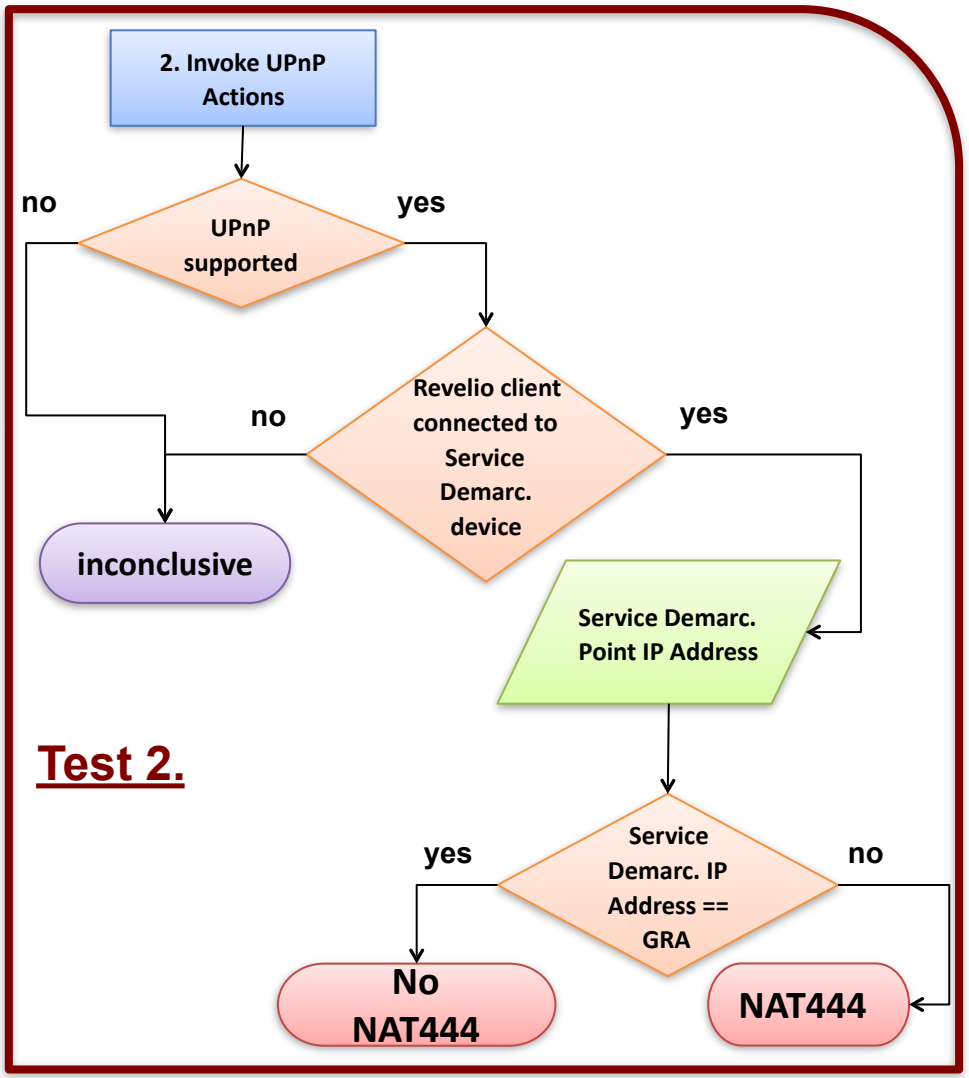

NAT444 Discovery

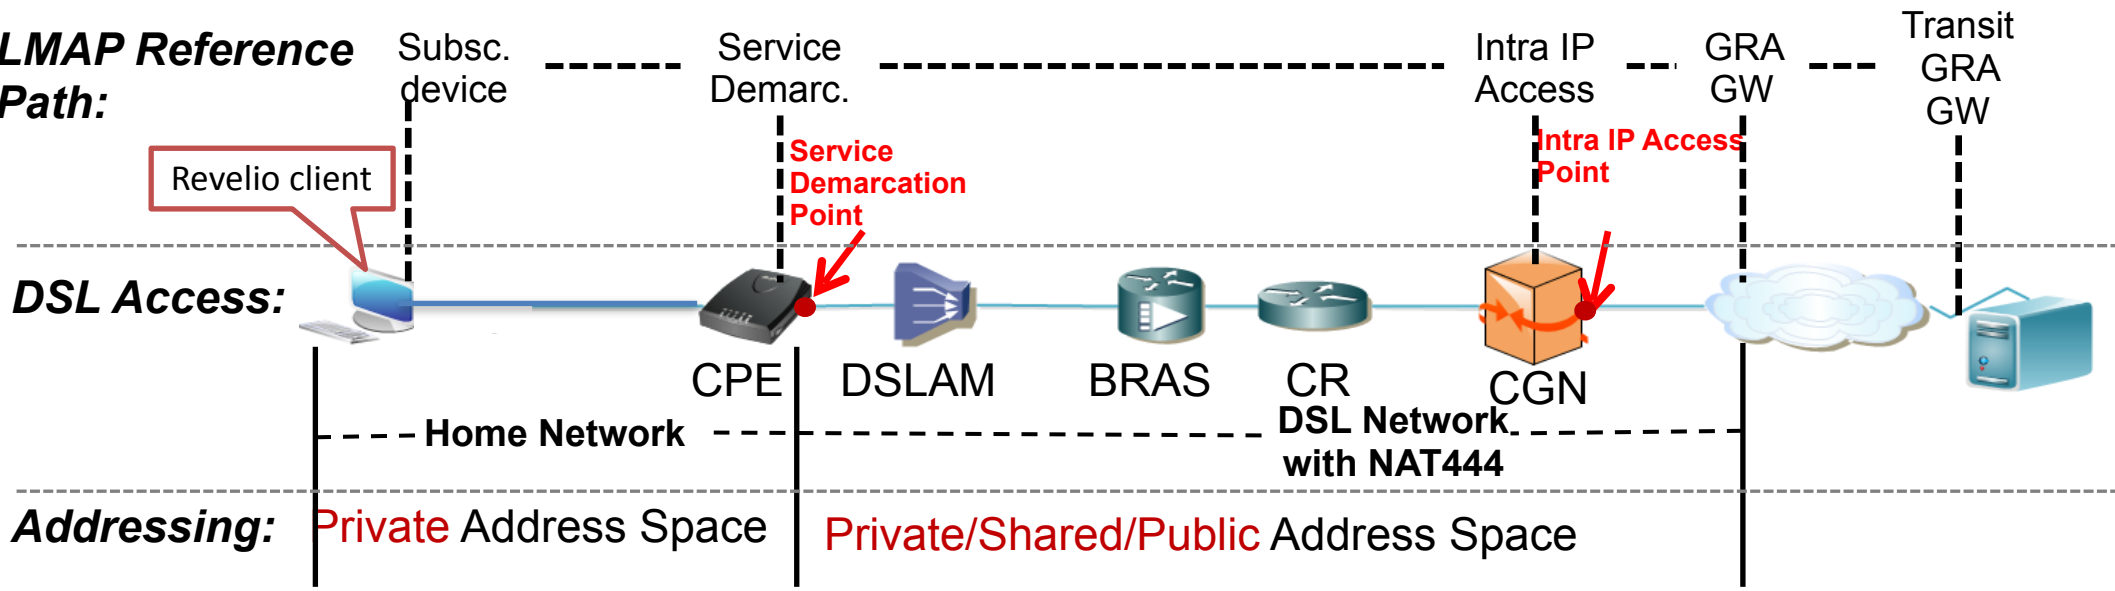

- **Test 1:** Detect private/shared IP Addresses in the ISP Access network
  - Traceroute to a external target (from Path Analysis in Environment Characterization)
  - Use the information about the location of the Service Demarc. Point
  - Detect private/shared IP addresses configured in the access network of the ISP

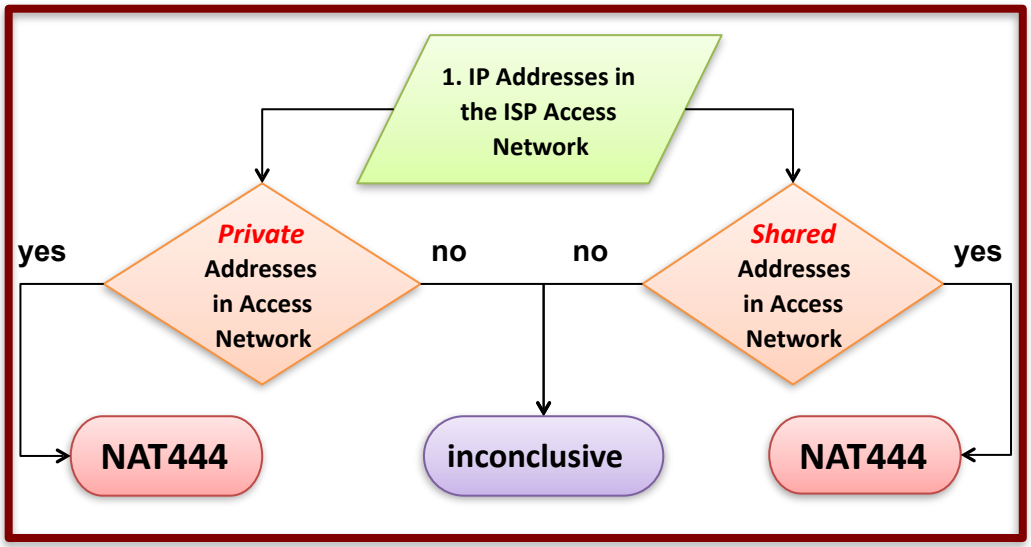

NAT444 Discovery

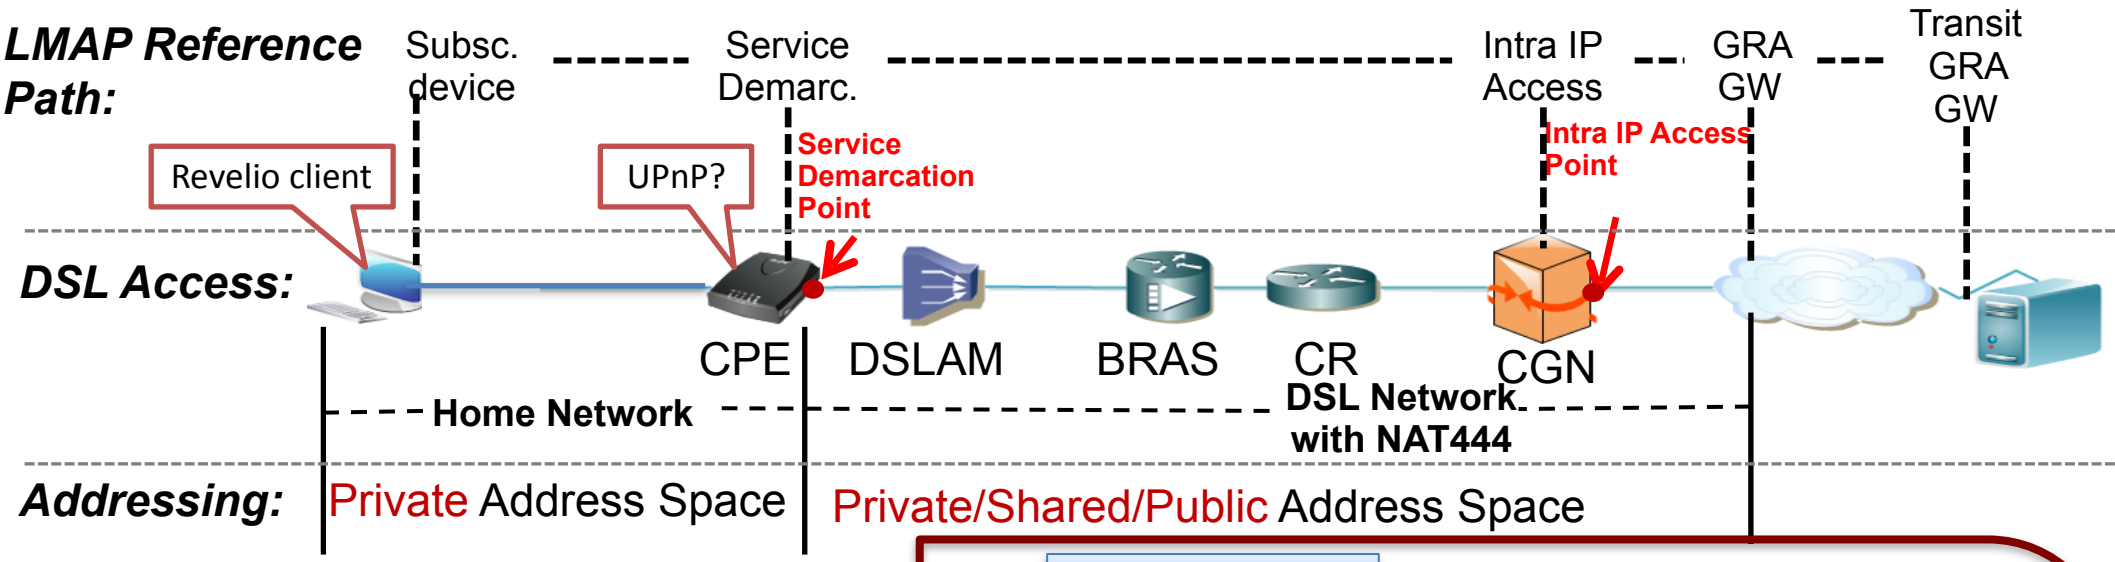

- **Test 2:** Invoke UPnP actions
  - If the Service Demarc. device (CPE) supports UPnP, then the Revelio Client acts as a *UPnP control point* and sends a request to learn the IP Address at the Service Demarc. Point

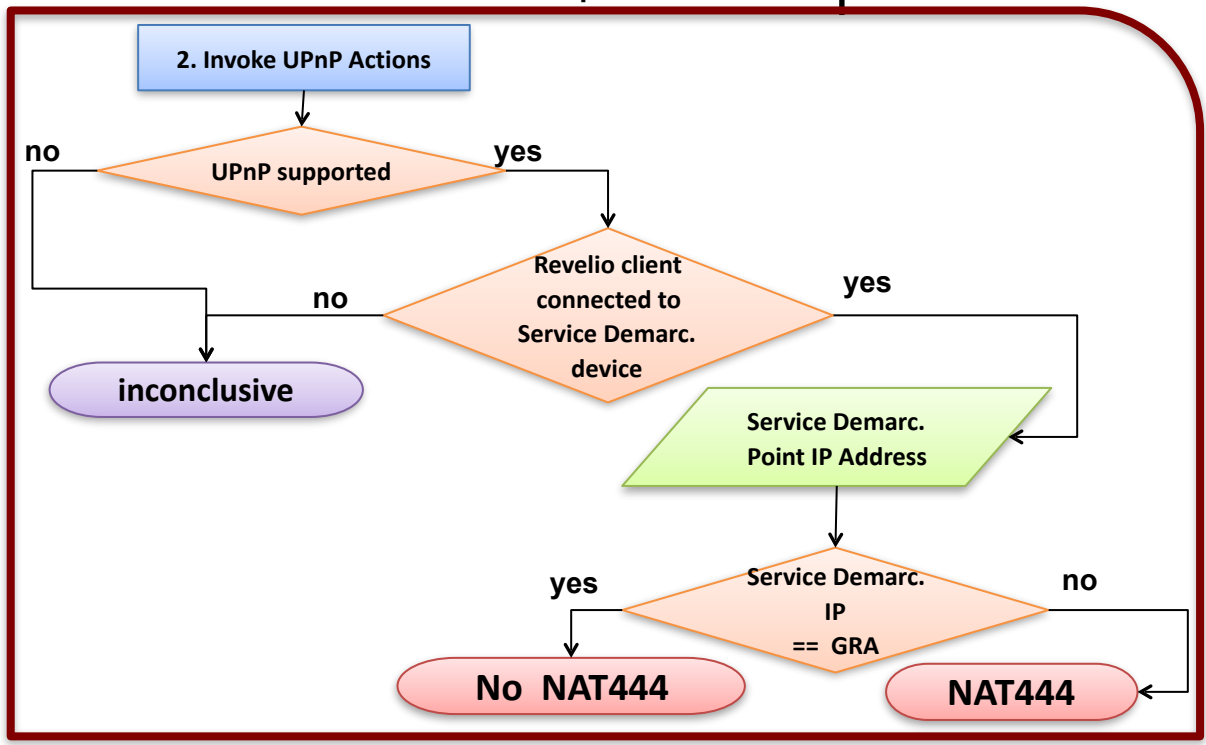

NAT444 Discovery

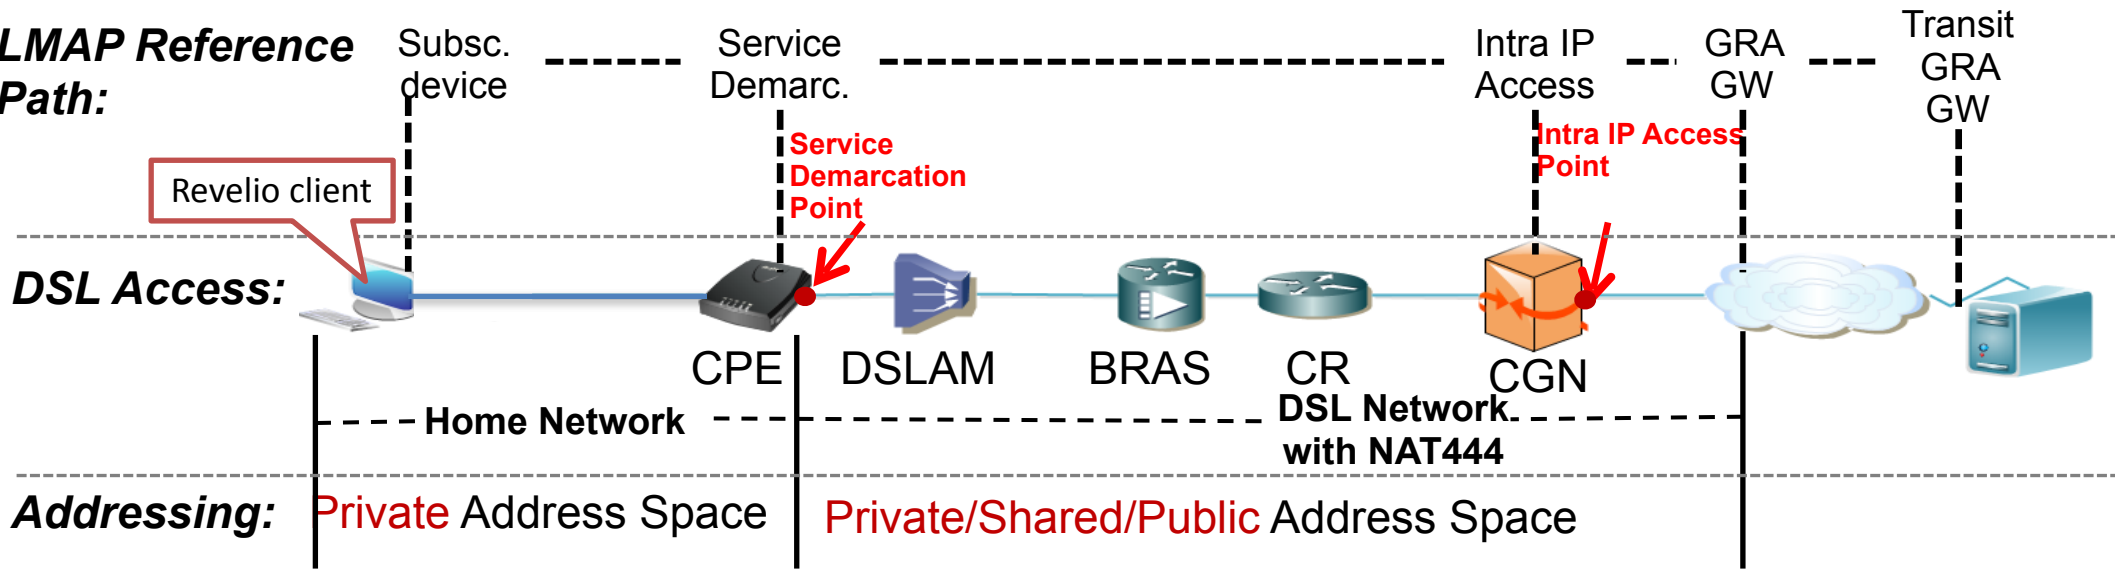

- **Test 3:** Traceroute to the GRA
  - Allows us to count the number of hops between the Revelio client and the device assigning the GRA
  - If this is larger than the distance between the Revelio client and the Service Demarc. Point (which we know from the Environment Characterization) => there is a NAT444 device in the ISP (e.g., CGNAT)

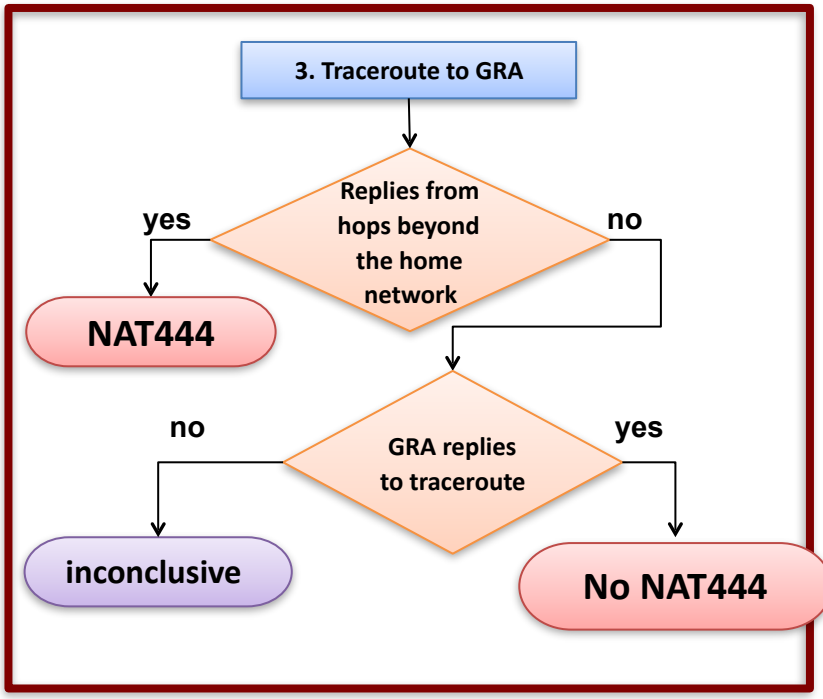

NAT444 Discovery - confidence levels

- ***Tests 2 (Invoke UPnP actions)*** has ***high confidence***
- Does not rely on inferred location of the access link
- Does not rely on the assumption that private addresses in the access network implies NAT444 in the ISP

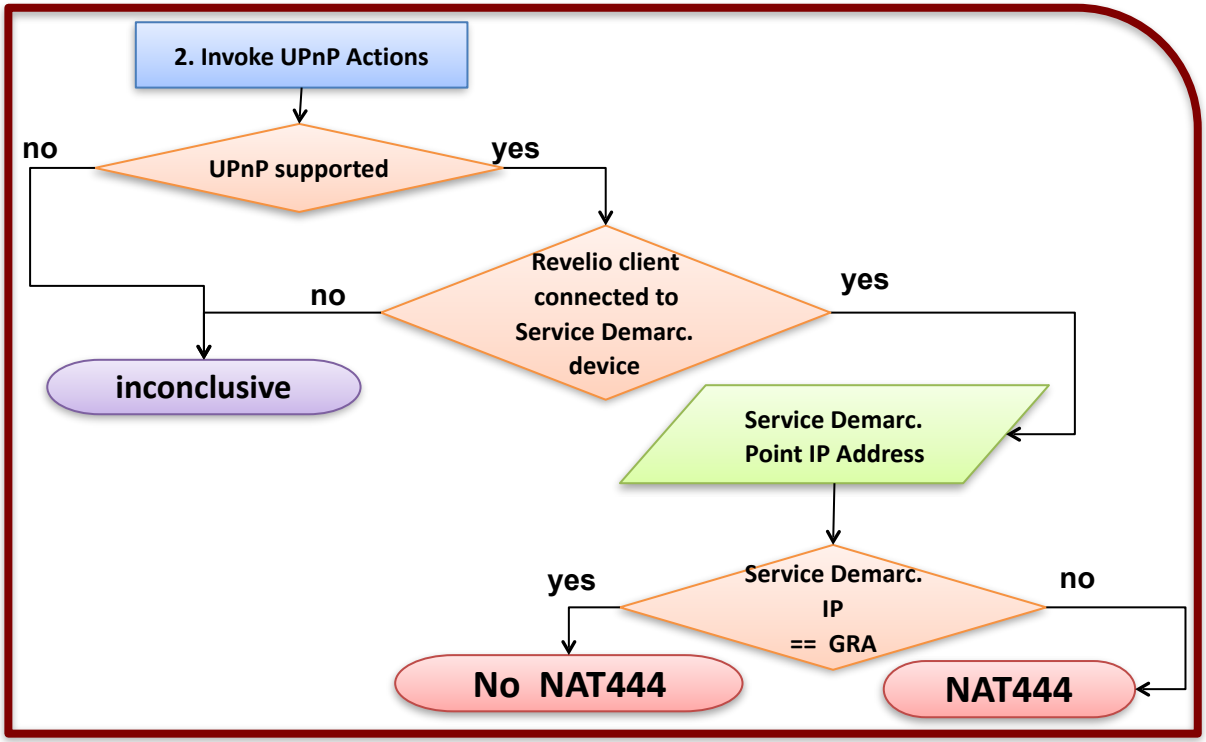

NAT444 Discovery - confidence levels

- ***Tests 1 (Private IPs in the ISP)*** and ***Test 3 (Traceroute to GRA)*** have ***low confidence***
- ***Test 1 (Private IPs in the ISP)*** relies on accurate location of the access link and the assumption that private IPs in the ISP implies NAT444
  - not the case for ***Shared IPs*** in the ISP (this is specific for NAT444)
- ***Test 3 (Traceroute to GRA)*** relies on accurate location of the access link (which might fail for fast media)

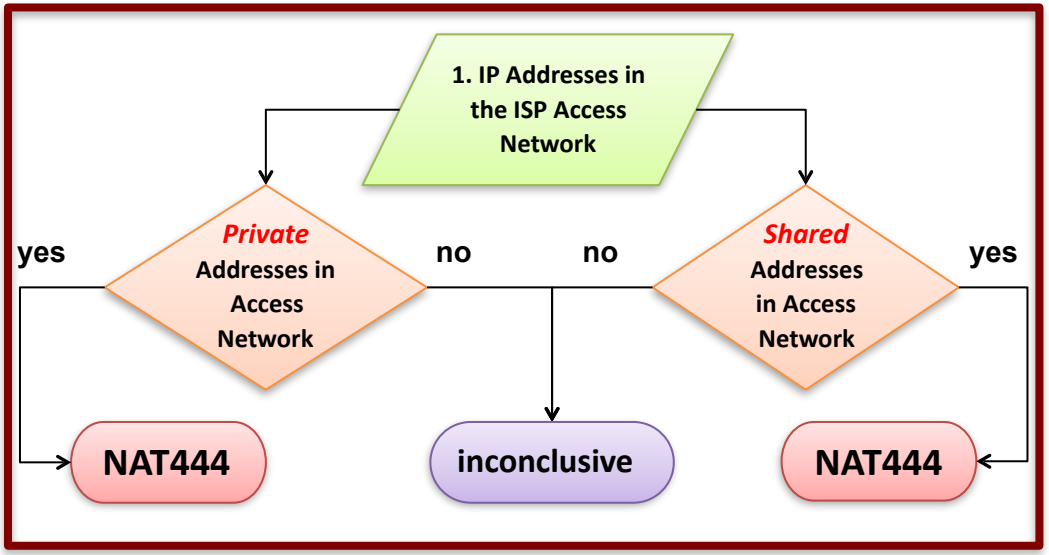

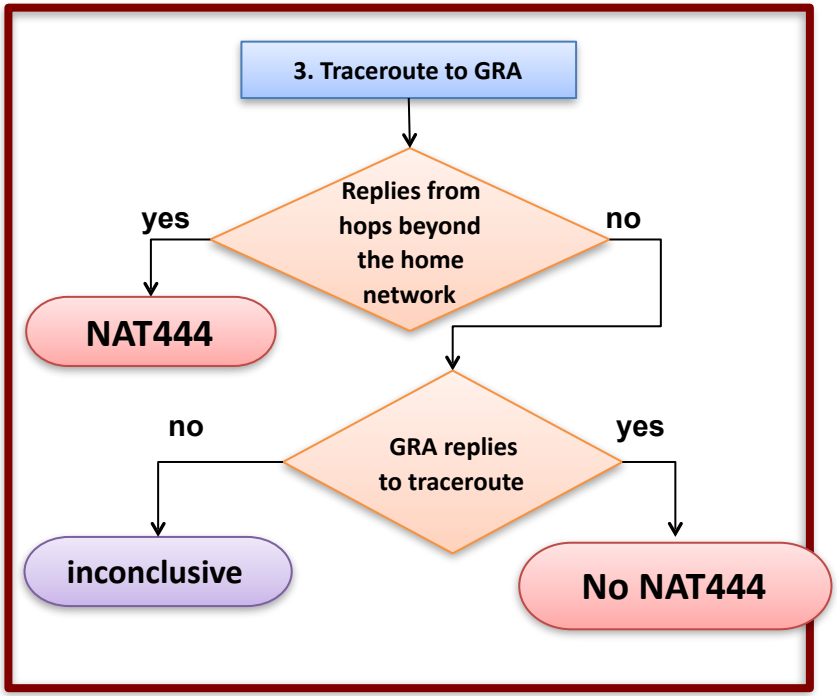

NAT Revelio Validation

- Tested Revelio in ***controlled environment***
  - 6 subscribers of a large UK ISP, 2 involved in a trial deployment of a NAT444 solution in the ISP
  - 24 residential DSL from Italian ISP that does not deploy NAT444 solutions
- 6 UK subscribers: 2 behind a NAT444 device
  - All tests in the Revelio Discovery Phase correctly identified the lines behind NAT444
- Even if the Italian ISP uses private IP addresses in its configuration, Revelio successfully identified that there is no NAT444 solution in the ISP
  - We use multiple tests and when results are conflicting, we prioritize the negative result for NAT444 presence

Large-Scale Measurement Campaigns

- ***SamKnows deployment***
  - **June2014:** 2,000 devices in 26 ISPs in UK
    - 10 lines tested positive for NAT444 — 5 providers
  - **October 2015:** 1,500 devices in 26ISPs in UK
    - 3 ISPs (out of the previous 5 ISP) detected in this second phase
- ***BISmark deployment***
  - **February 2015:** 37 devices in 24 ISPs over 13 countries
  - Revelio identified NAT444 in 3 ISPs (Vodafone Italia, Embratel and Comcast)
    - low confidence for Embratel and Comcast (we only found private IPs after the Service Demarc. Point)
    - high confidence for Vodafone Italia — two test in Revelio Discovery gave positive results

Conclusions

- We propose ***NAT Revelio*** to detect the presence of NAT444 solutions in the ISP
  - **OPEN CODE: *https://github.com/alutu/revelio***
- Some of the tests might fail at times, depending on the network we test (e.g., traceroute blocked, CPE does not support UPnP)
- In case of conflicting results from different tests, we prioritize the negative result for NAT444 in the ISP
- We do not *categorically* detect a NAT444 solution — we determine the likelihood that there is one in the ISP, relying on the results from multiple tests

FAQ

How does Revelio perform with different access technologies?

- The Path Analysis (Env. char. - test 3) might fail for very fast media (cable/FTTx)
  - thus, the location of the access link might be inaccurate
- Revelio is a test-suite — other test might still work well to detect the NAT444 in the ISP, even without accurate location of the access link (UPnP test)

FAQ

Why do you need different confidence levels?

- They show the reliability of the test
- Some tests use *information* that might be inaccurate in some conditions (i.e., location of the access link) or *assumptions* that are not correct at all times (i.e., private IPs in the ISP not always mean that there is a NAT444 in the ISP)
- The strength of *Revelio* lies in putting together the results from all the tests

FAQ

Would Revelio work for mobile broadband (MBB) networks?

- In the current form — no
- Detecting CGN solutions in MBB networks is challenging — MBB providers already rely on a large NAT device in their core infrastructure
- It is not obvious how to make the difference between the CGN and this NAT device

pathchar to detect the access link

- Run UDP traceroute to a fixed target (router inside Level3 network with no rate limiting)
  - Used the well-known traceroute port range
  - 21 different packet sizes (from 120 to 1400 bytes)
  - One traceroute probe per TTL, max TTL of 30
- Run every hour, over 4 days => collected ***96 RTT samples per TTL*** and for each packet size

pathchar to detect the access link

- For each TTL:
  1) Minimum Filtering:
     - For each packet size, choose the minimum value of the RTT
       - Capture only the transmission delay and the propagation delay
     - RTT = packet_size/BW + LAT
  - 2) Line fitting
     - Using the 21 different points, fit a regression line for the RTT and determine the **slope [1/BW]** and the **intercept [LAT]**

pathchar to detect the access link

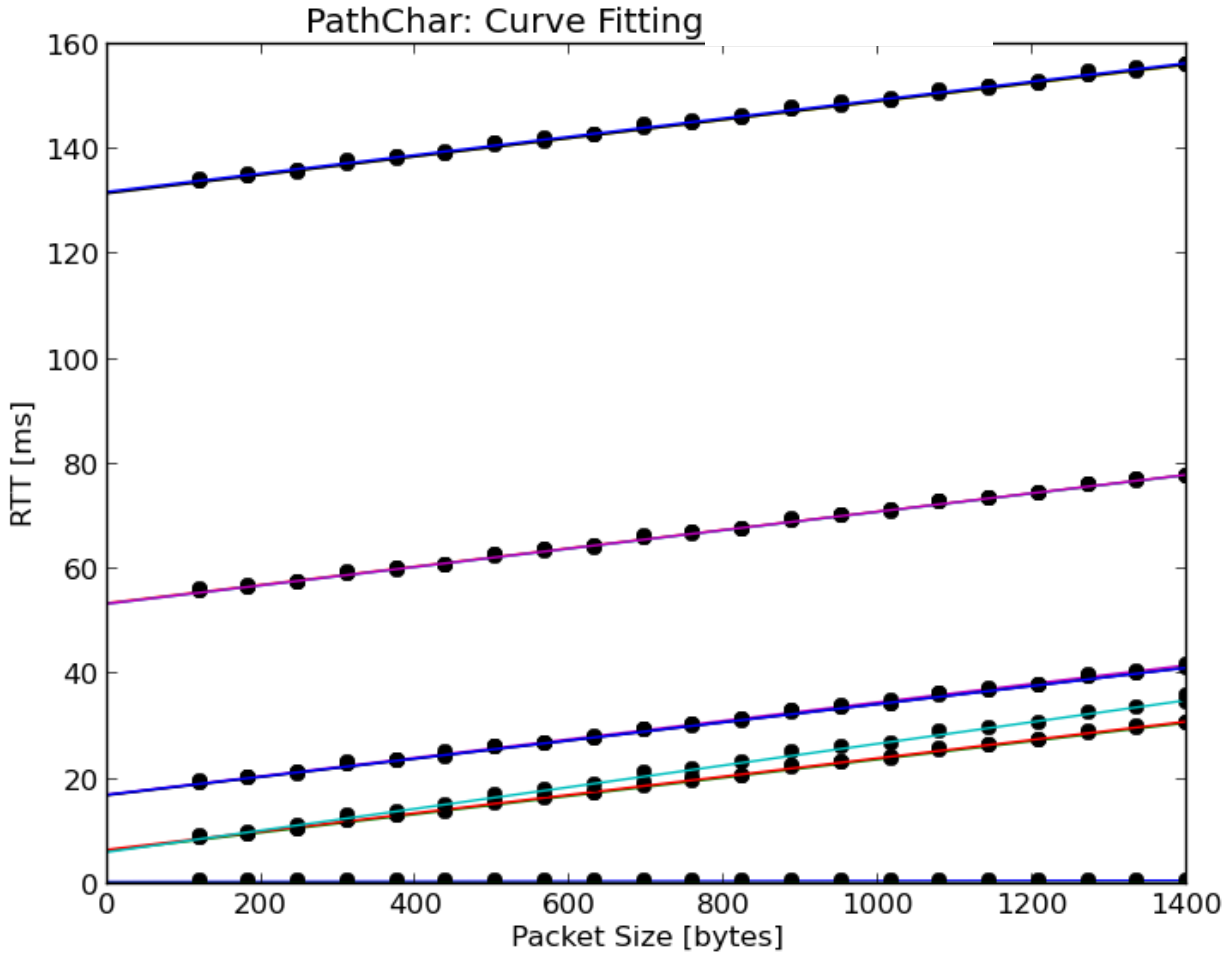

pathchar to detect the access link

3) Differencing

- Given the estimated cumulative parameters above, pathchar determines the per-link parameters (slope and intercept, i.e., 1/BW and LAT) by subtracting the consecutive fitted lines parameters

Invoke UPnP Actions

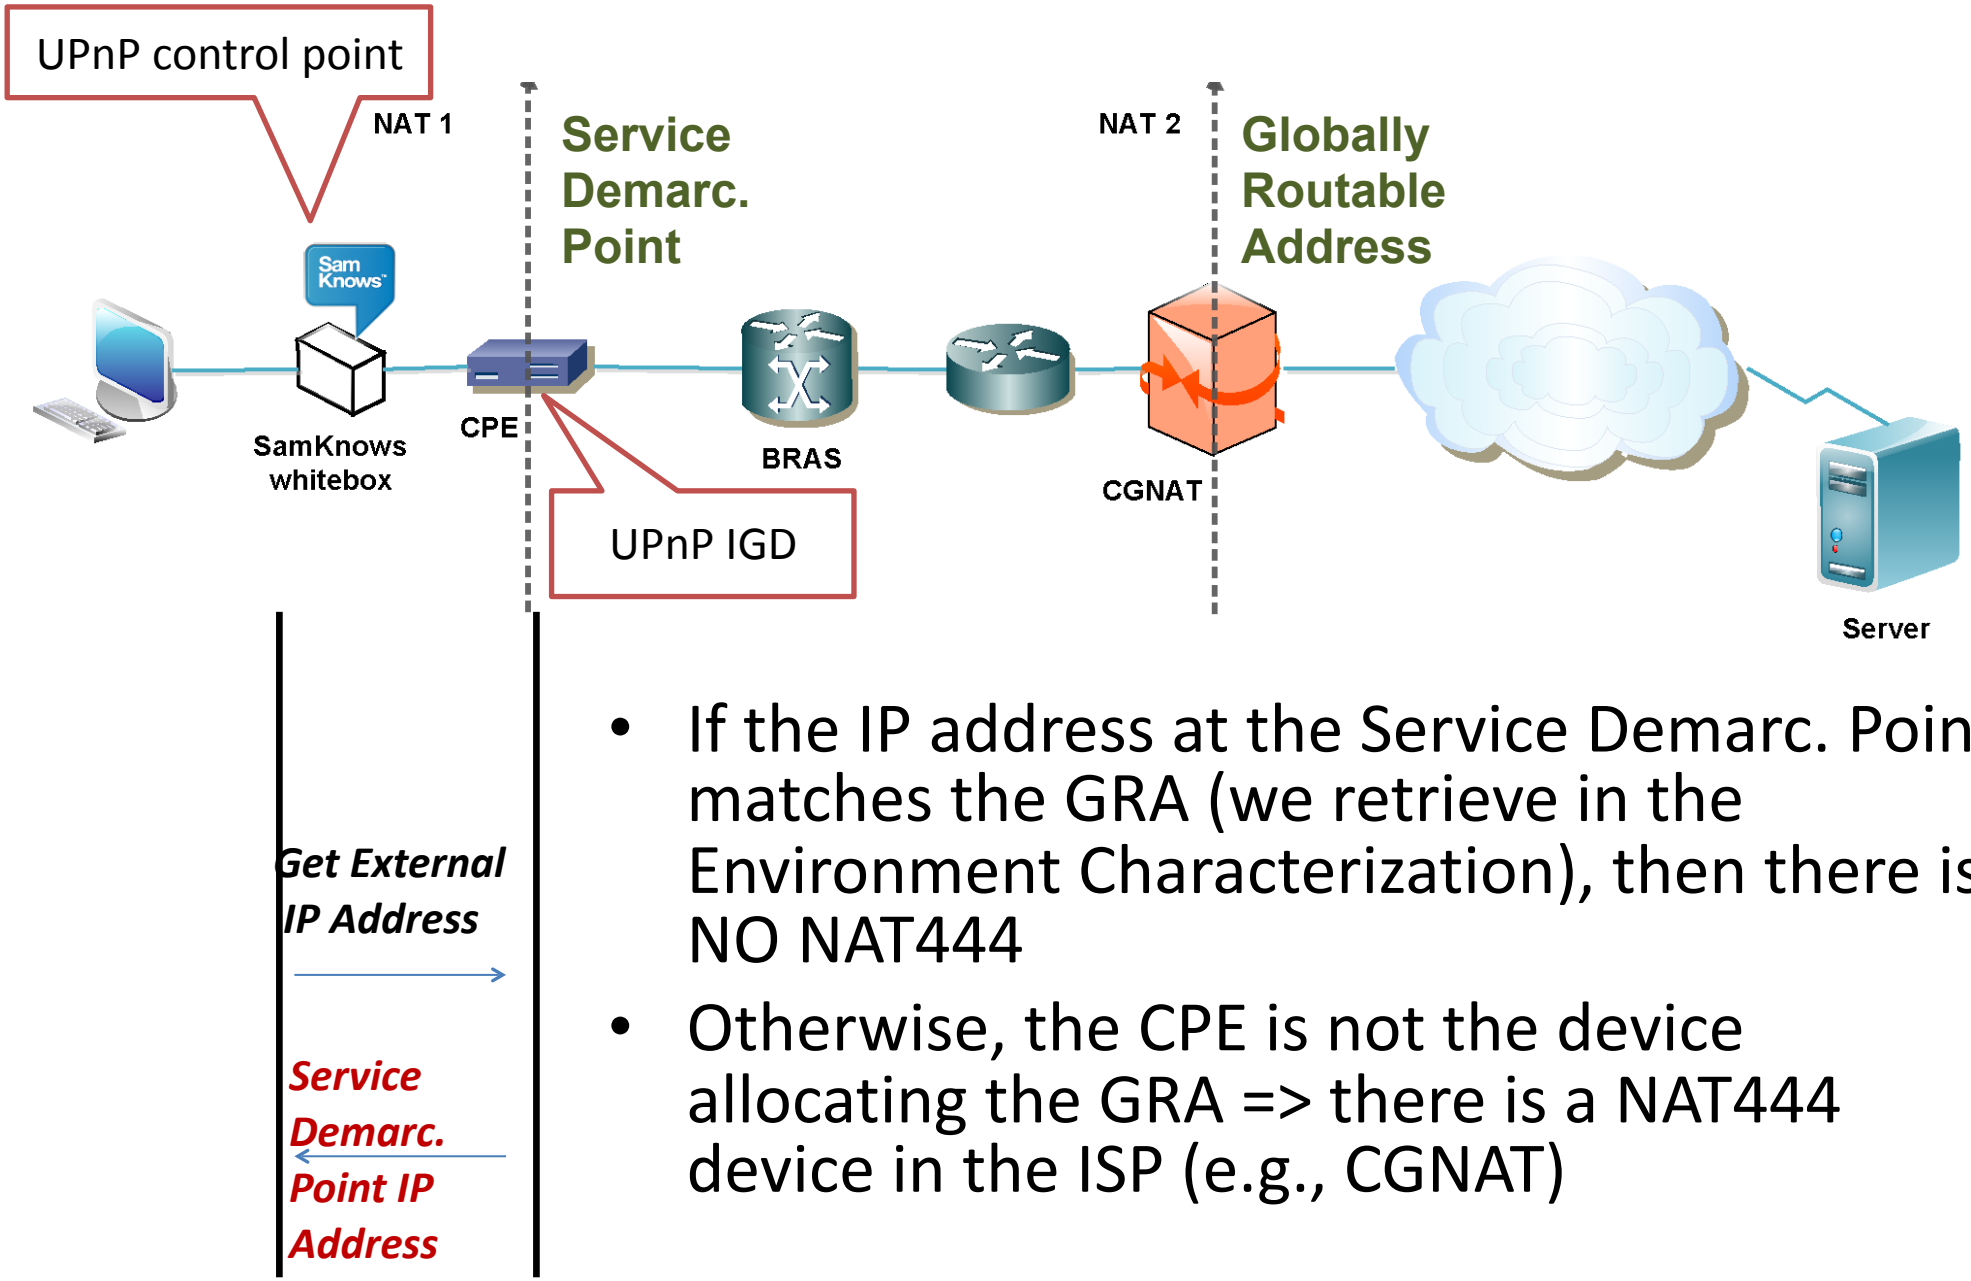

- If the IP address at the Service Demarc. Point matches the GRA (we retrieve in the Environment Characterization), then there is NO NAT444
- Otherwise, the CPE is not the device allocating the GRA => there is a NAT444 device in the ISP (e.g., CGNAT)