Internet performance measurement: Where do we stand and the road ahead

Methodological Approaches

Monday, April 22, 2024

Bill Woodcock

Packet Clearing House

My Goal With These Slides

…and in this workshop generally, is to assess the degree of consensus that exists among experts on the specific subject matter of an implementable norm or best-practice for methodologically-defensible, globally-comparable, Internet performance measurements.

Please be liberal in signaling agreement or disagreement with the principles I'm listing here. We have the rest of the workshop to discuss them in greater detail.

Terminology

(Just what I'm using in these slides, I'm not trying to pick any fights here.)

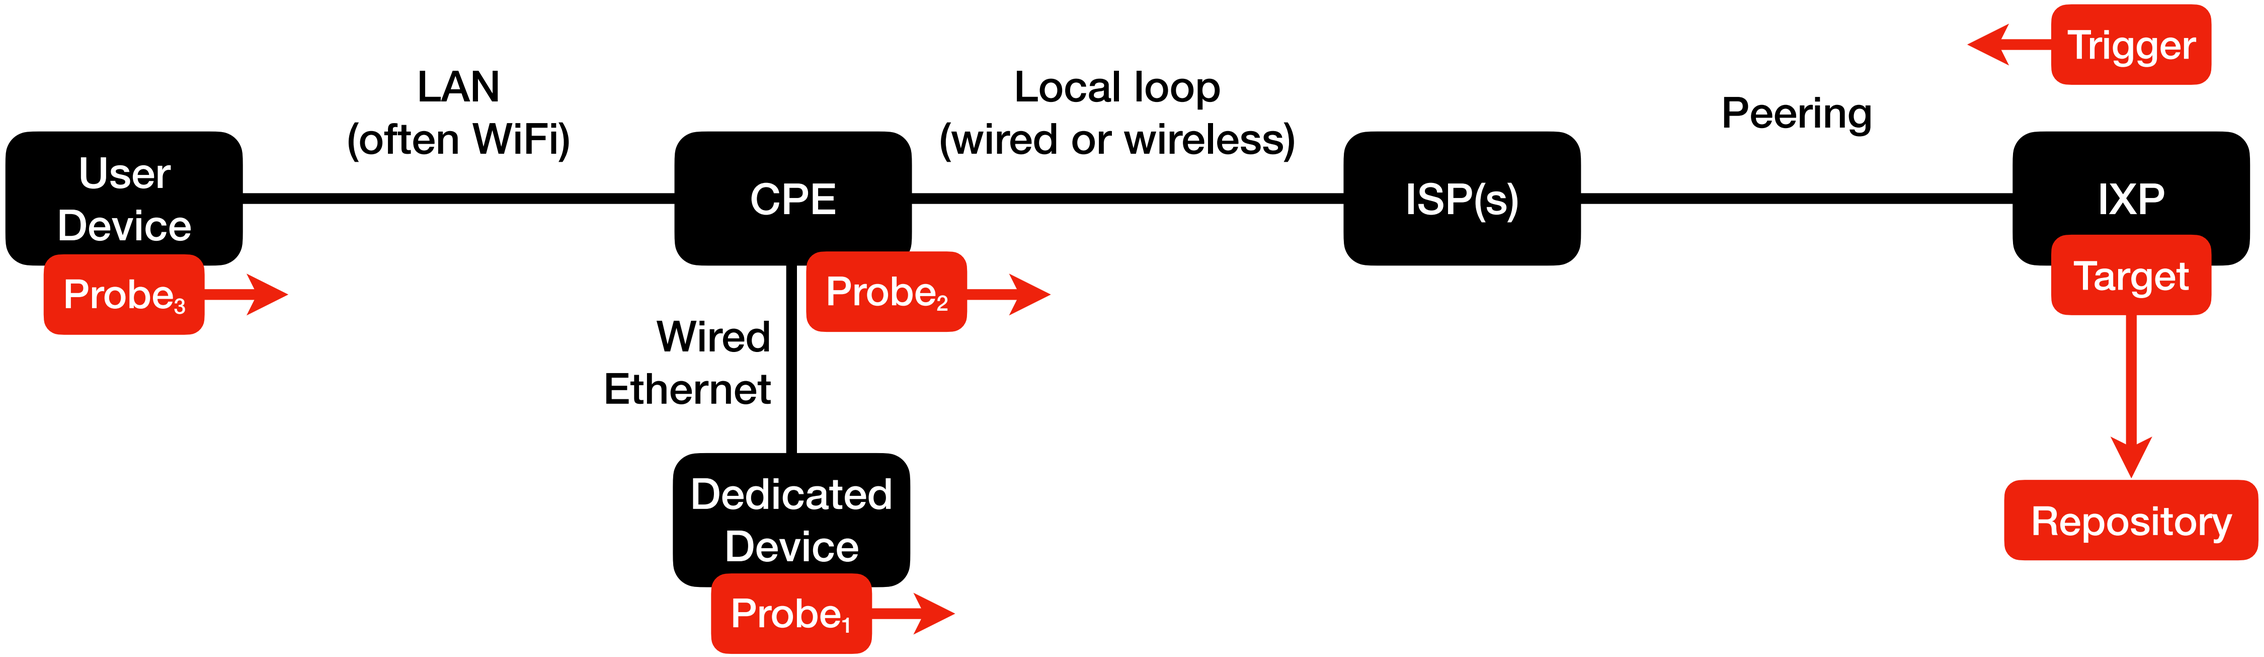

The Problem with Misaligned Incentives

Measurement companies typically sell both measurement targets and measurement data. ISPs buy targets in order to “look good” in the measurement data. Customers of ISPs which have bought targets are measured against a target which is artificially close, while customers of ISPs which have not bought a target are measured against targets which are artificially distant. Neither are representative of the general Internet transit the customer is buying.

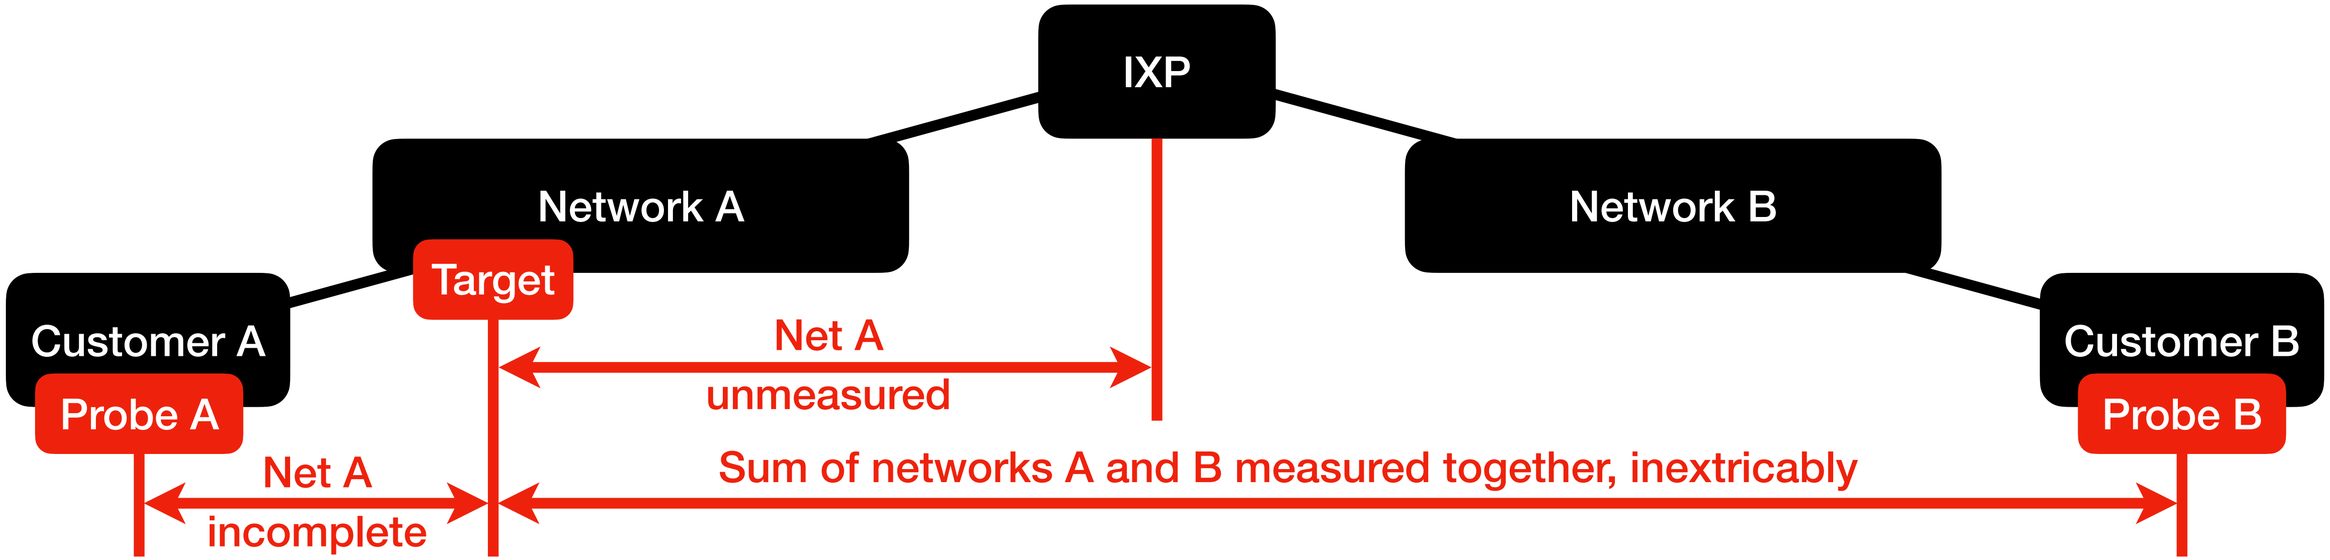

Place Targets only at the Sources of Bandwidth

BEREC, the FCC, and others have recognized that valid measurements can only be performed against targets which are located at IXPs, the sources of the bandwidth customers are consuming.

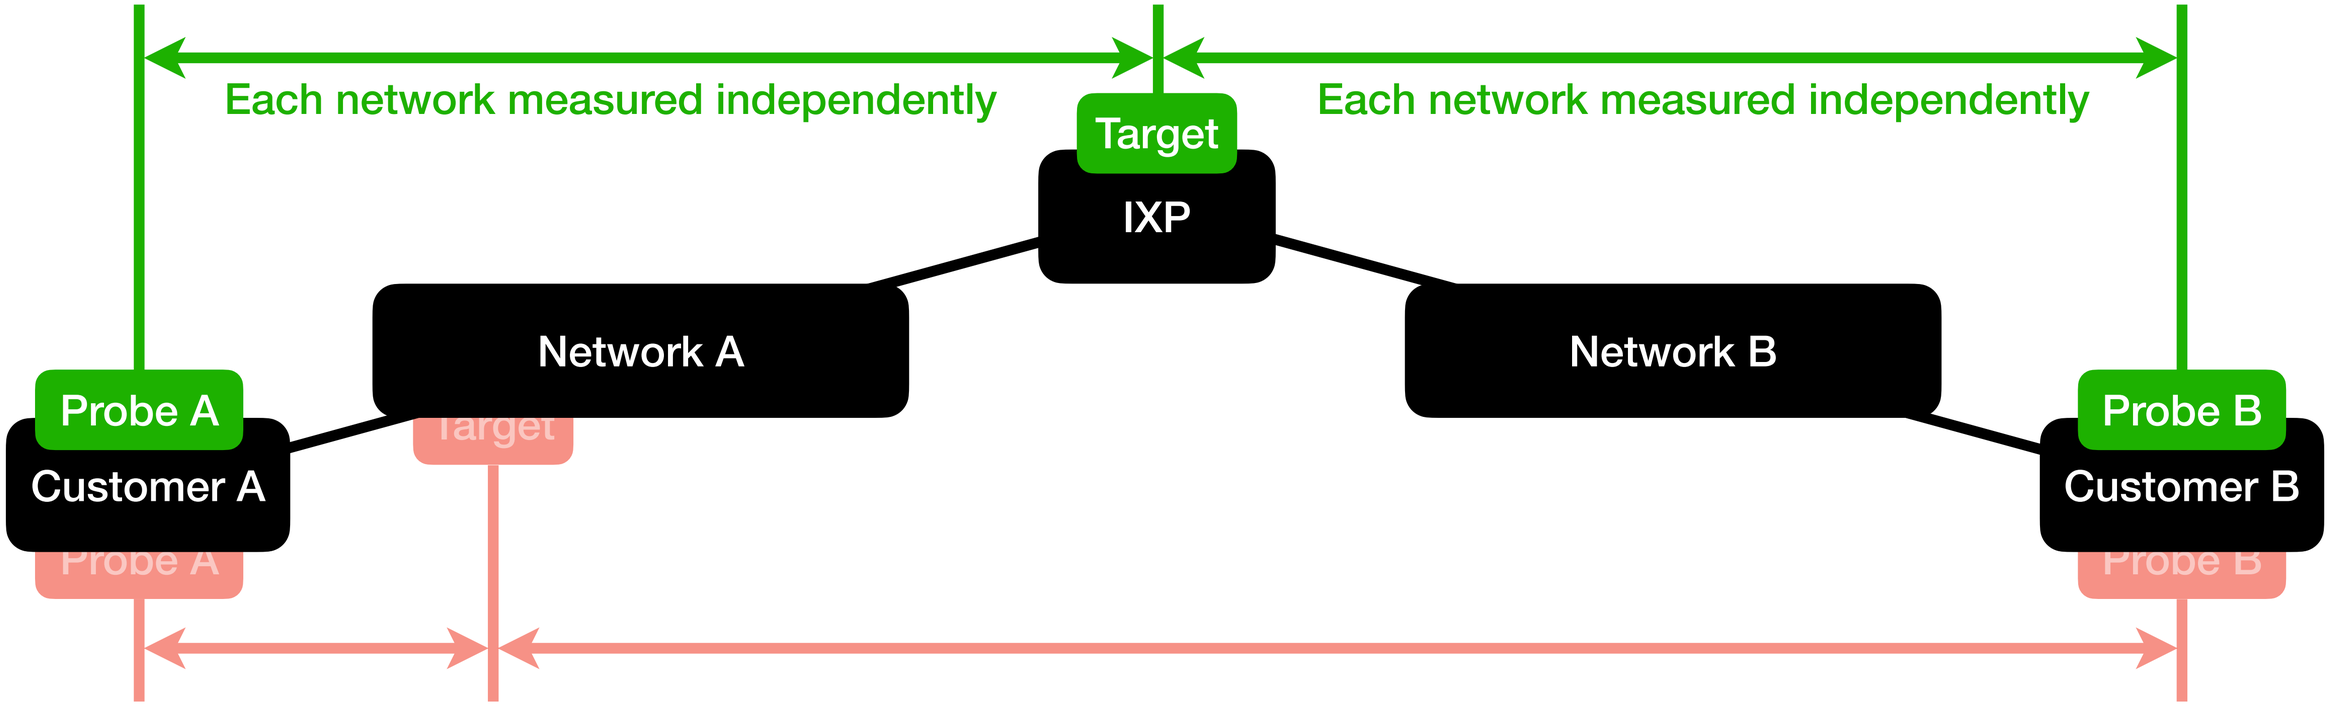

Place Targets only at the Sources of Bandwidth

BEREC: “Where measurements are performed against a test server, this server should be located outside the IAS network. Typically, this can be achieved by locating the measurement server at, or close to the national internet exchange point (IXP). Depending on the specific national situation, measurement servers may be located at more than one IXP location.

The hardware running the measurement server(s) should be connected as close to the IXP switch as possible. This means that the number of hops between the main IXP switch and the test server should be kept at a minimum. This is applicable whether the implementation runs inside the network of a hosting provider or directly on hardware under control of the NRA itself.”

Place Targets only at the Sources of Bandwidth

FCC: “Carriers must test speed and latency from the premises of subscribers to a remote test server located at an FCC-designated Internet exchange point (IXP). ISPs also serve communities that may be geographically distant from the IXP locations used.”

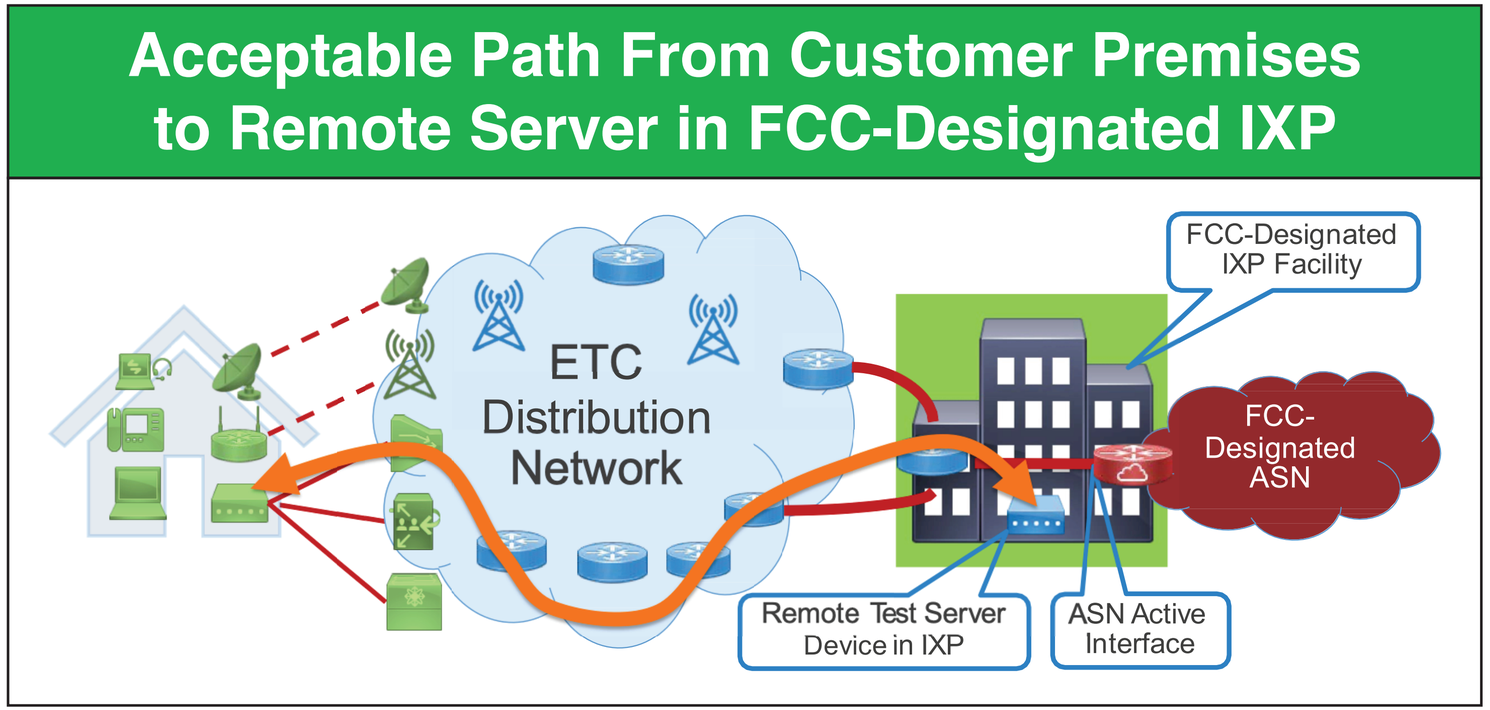

Source: https://www.usac.org/wp-content/uploads/high-cost/documents/Tools/PMM-Test-Paths-and-Remote-Server-Locations-1.pdf

Place Targets only at the Sources of Bandwidth

The existing norm can be summarized:

- Measurements should be conducted between customer premises and the multiple IXPs from which the customer is deriving their bandwidth.
- These IXPs may in many cases be geographically distant from the customer.
- The equipment performing the measurement should be connected directly to the IXP switch fabric, not on a separate network.
- The equipment performing the measurement should be operated by a national telecommunications regulatory agency (TRA) or other independent and impartial party of similar competence.

Targets at the Sources of Bandwidth

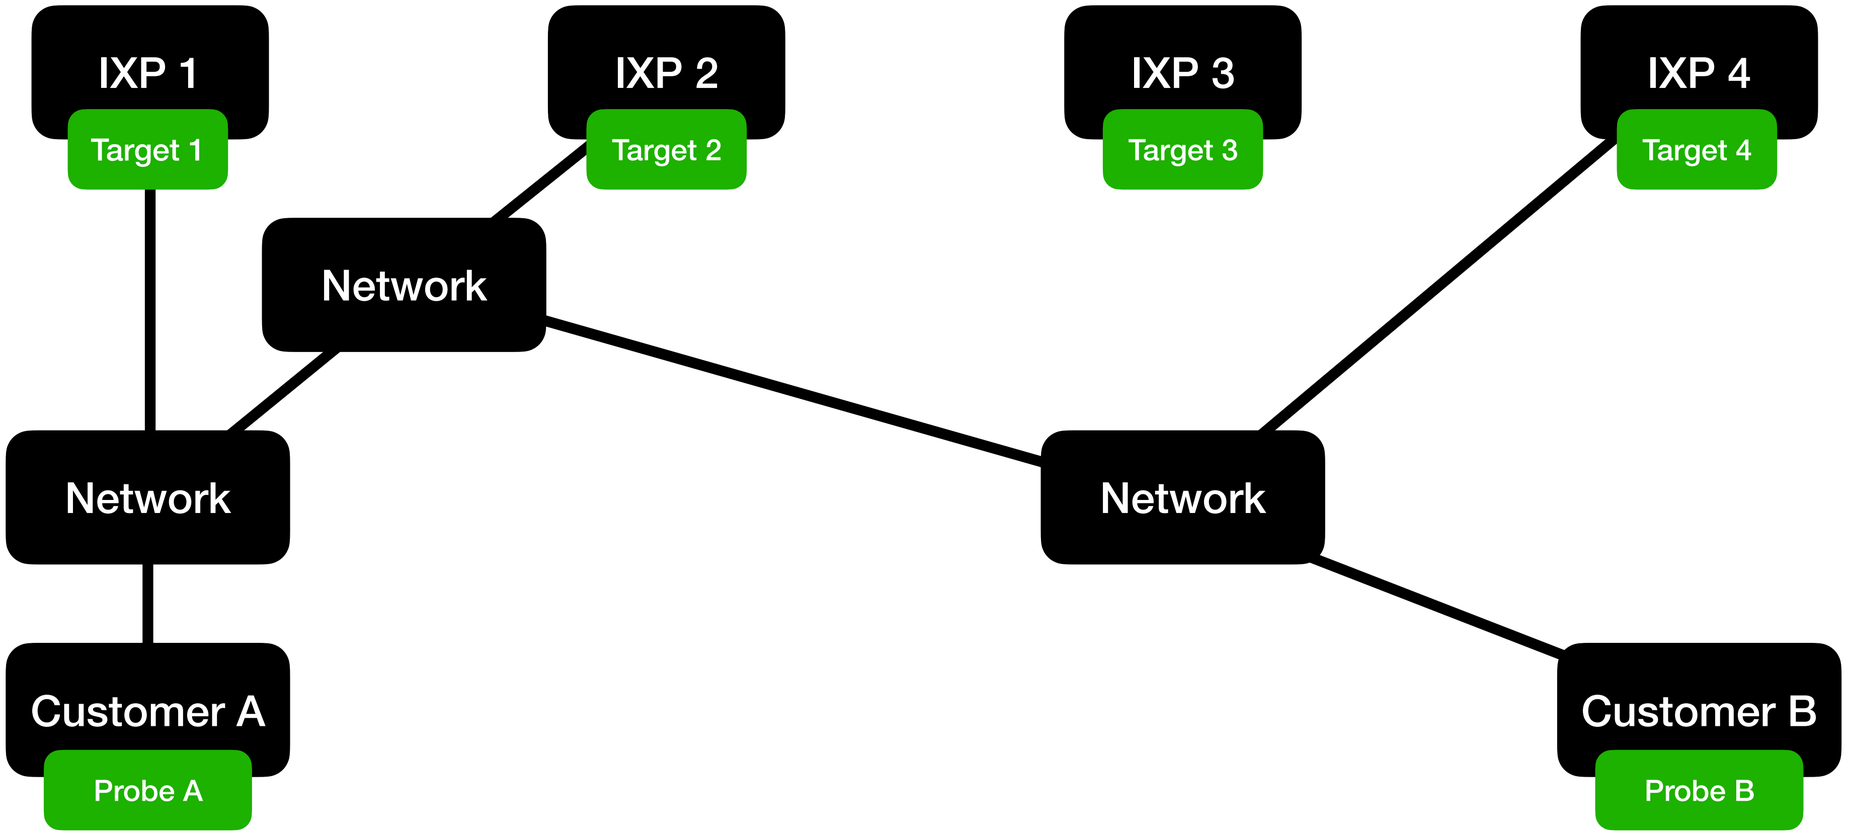

Targets at the Sources of Bandwidth

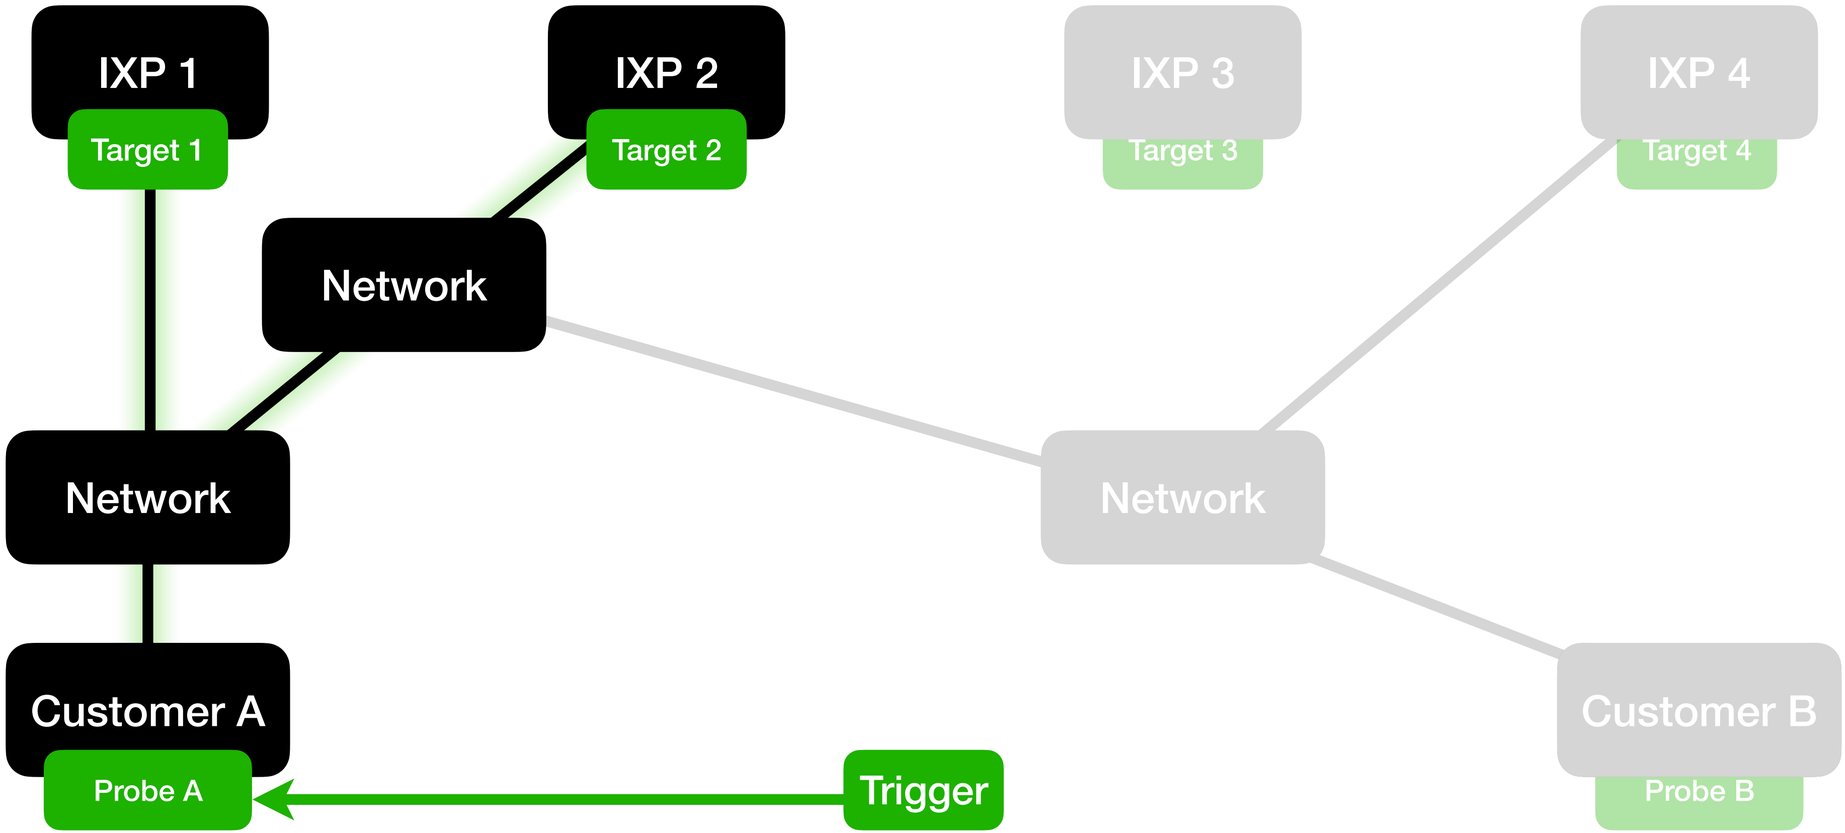

Targets at the Sources of Bandwidth

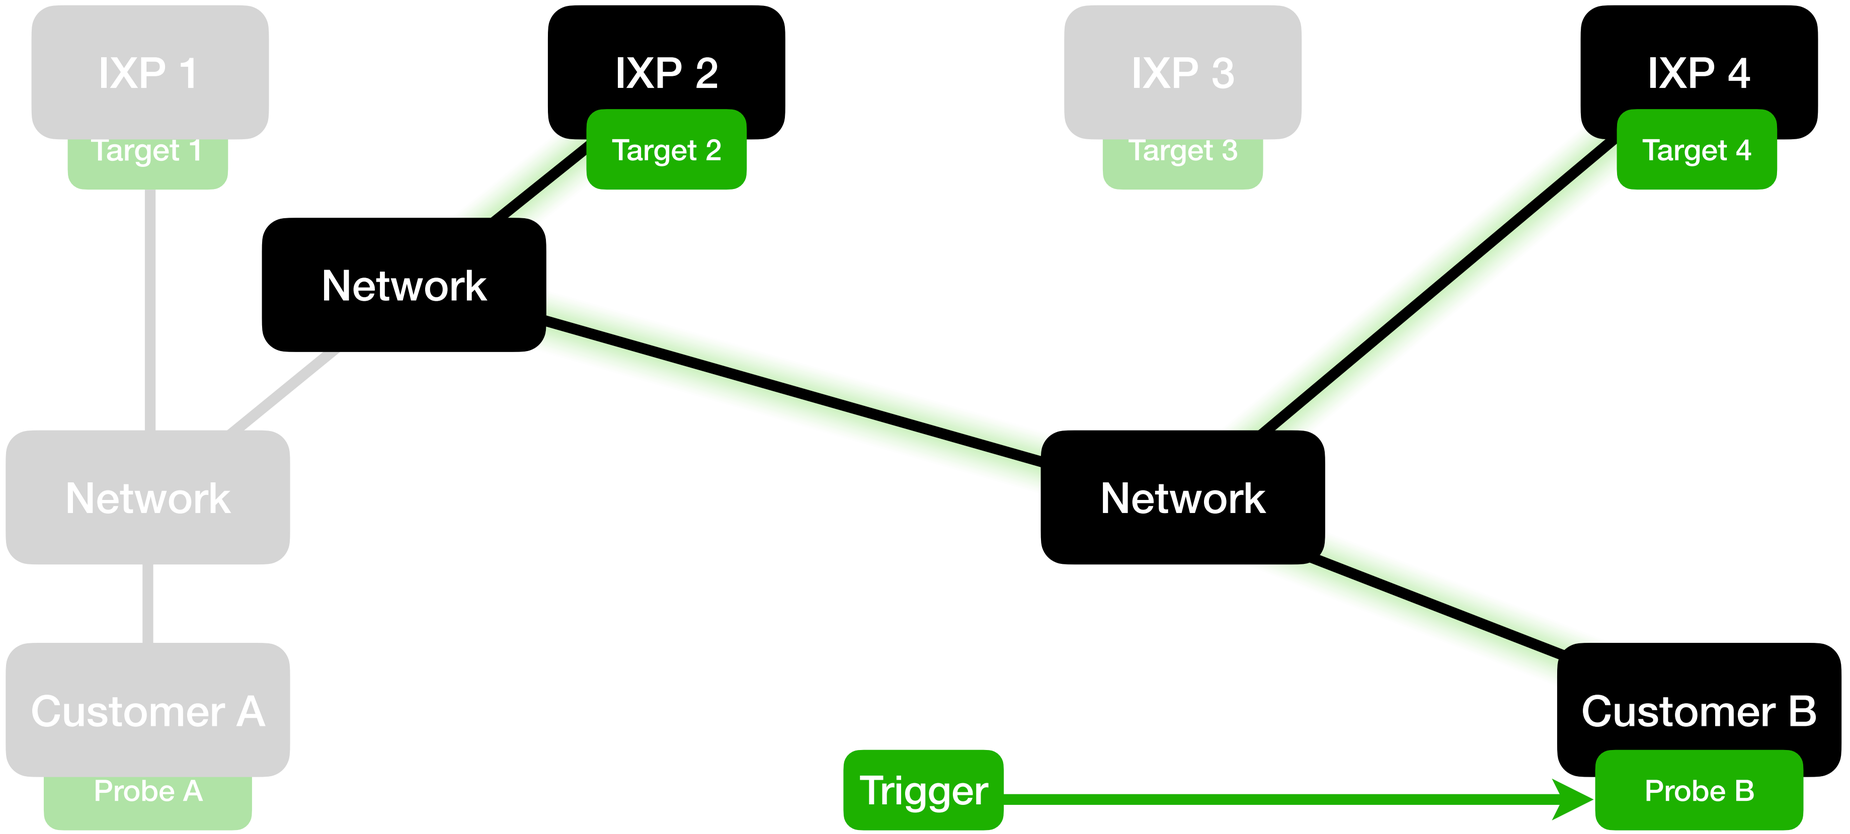

Non-Operational Principles

To maximize value, test results should be published and aggregated in a public repository (see RIPE Atlas as an example of success).

To minimize risk of gaming or falsification, measurements should be performed by a device under the control of a national telecommunications regulator or independent and impartial organization of similar competence. The device should sign results before submitting them to the repository.

> The implication is that measurements should be collected from the target rather than the probe, unless the probe is dedicated hardware/software with a TPM, running signed code and signing results.

It should be possible for anyone to trigger a measurement, but the mechanism should be particularly streamlined for the use-case of TRAs.

Measurements should be non-destructive (“first do no harm”): the act of measurement should not significantly degrade the properties being measured. i.e. no more “dump a big file on the network and see what happens.” We’ve known better since pathchar (1997).

Should the IP addresses associated with results be truncated to /24 or BGP-advertised prefix?

Probe Types

Purpose-specific signed software running on purpose-specific hardware, with a TPM and a hardwired connection to the CPE (or mobile, for mapping mobile wireless)

Purpose-specific software running on the CPE

Purpose-specific software running on the user’s device with root permissions

Purpose-specific software running on the user’s device in user-space

In-browser test initiated by the user

In-browser test initiated by an ad bug

Result Weighting

If measurements are performed from the probe to the set of targets which exist at the IXPs at the tops of the transit cones the user lies within, the set of measurement results need to be weighted proportional to the degree that they consume bandwidth from each of those IXPs.

This requires summarized flow statistics collected either from the IXPs (as per the IIFQ project), or from within ISPs (as could be mandated by their responsible TRA), or by the end-user themselves. Each of these constitutes a flow “fingerprint” which may be private to its holder, and which can be applied to public measurement data.

Basic Quality Metrics

Loss (percentage relative to total observed)

Latency (minimum and median values)

Jitter (variation in latency)

Out-of-order delivery (percentage relative to total observed)

All of the above over IPv4 and IPv6 (perhaps ICMP, TCP, UDP, and QUIC independently)

Duplicate packets?

Loss period / Loss distance (variation or jitter in loss)

Path MTU?

Traceroute (RFC 5388 XML format)

Metadata

Number of packets
Start time
Duration
Version of standard and/or version of each measurement

RequestorID
TriggerID
TriggerSerial
ProbeID (unique, not IP, but perhaps not exposed in public dataset)
TargetID

Replay attack protection on the target side

Loss

As per RFC 4689, section 3.3.4.

Percentage of observed total

To two decimal places?

Latency

Lowest observed value

Median value

Expressed in whole-digit milliseconds?
Milliseconds to two decimal places?
Microseconds?

Jitter

Instantaneous jitter, per RFC 4689, section 3.4.4.8.

- Average over the period of measurement?
- Median over the period of measurement?
- Maximum over the period of measurement?

Peak-to-peak jitter, per RFC 4689, section 3.4.4.10, over the period of measurement

Expressed in whole-digit milliseconds?
Milliseconds to two decimal places?
Microseconds?

Loss Period / Loss Distance

The usability difference between the loss of every twelfth packet and the loss of 100% of packets for two hours per day is significant.

Per RFC 3357:

> “Loss period” captures the frequency and duration of loss events
>
> “Loss distance” captures the spacing between loss periods

Out-of-Order Delivery

As per RFC 4689, section 3.3.4.

Percentage of observed total

To two decimal places?

Additional Possible Measurements

What DNS recursive resolver is the probe using?

What delay is added by the recursive resolver?

What is the probe’s topologically nearest visible IXP?

Measurements Which May Require or Benefit from Purpose-Specific Probe Software

NAT detection

VPN detection

Probe-to-target traceroute (ground truth for reverse-path traceroute)

Port scanning (for blockage, redirection, (de)prioritization, or special handling)

Recursive resolver non-transparency

Geolocation (particularly for mobile wireless)

Outstanding Questions

Hash link to reverse-path traceroute in separate database?

What specific algorithm should be used to estimate path capacity?

pathchar (1997)?

Something more recent?

(Both will be discussed in the second session of this workshop.)