

Measuring Internet Interconnection Performance Metrics: an exercise to inform public policy

Center for Applied Internet Data Analysis

www.caida.org

2/9/2016



caida

Center for Applied Internet Data Analysis

An independent analysis and research group based at the University of California's San Diego Supercomputer Center



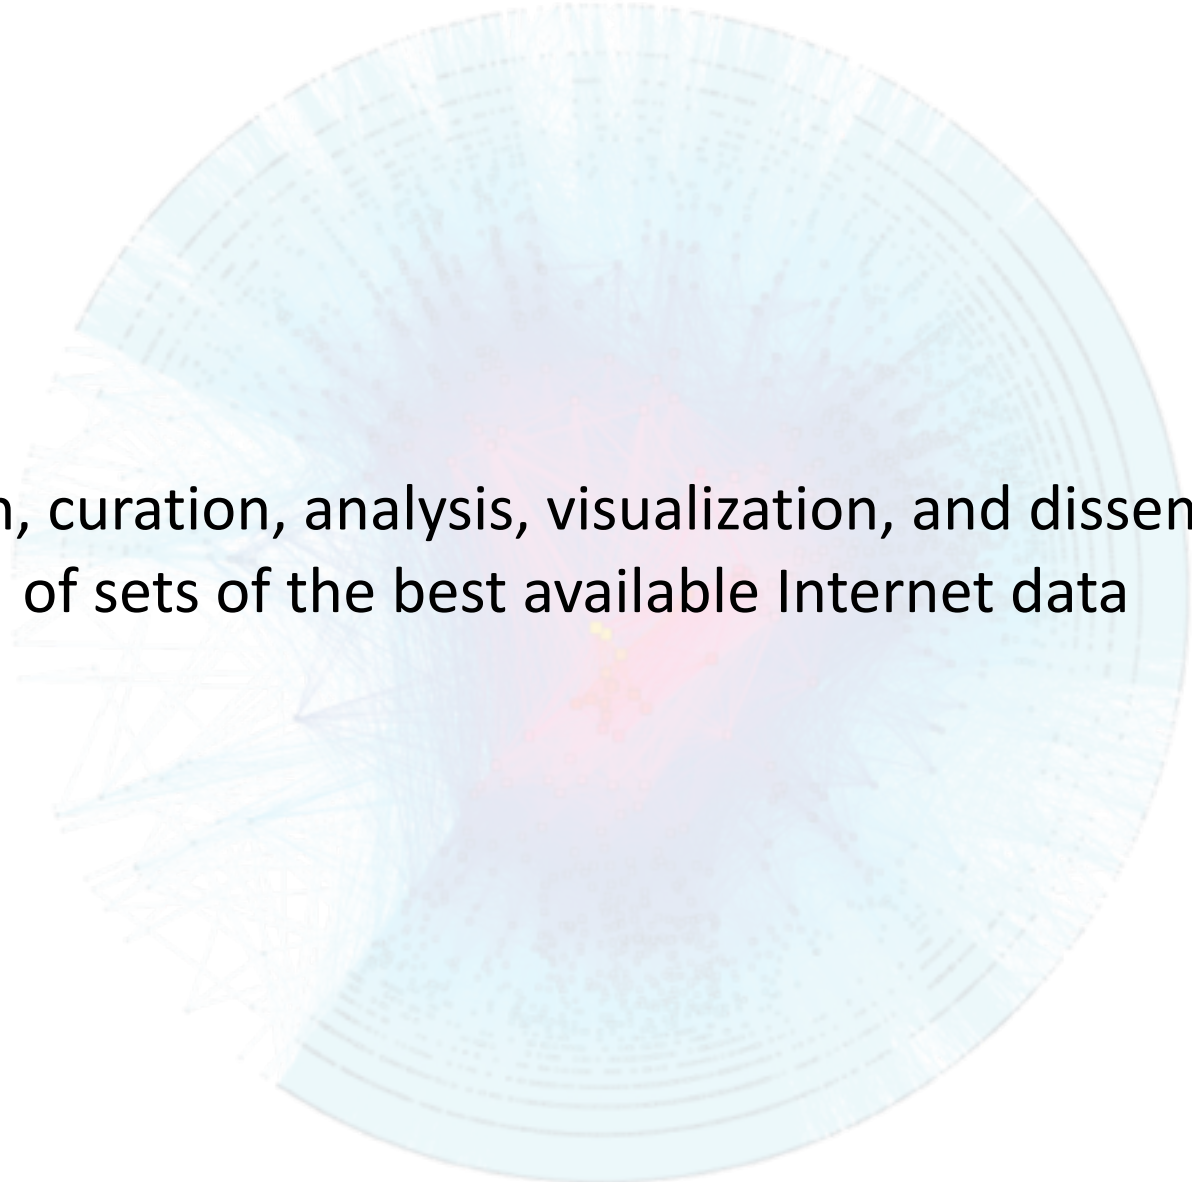
Providing macroscopic insight into the behavior of
Internet infrastructure worldwide



Improving the integrity of the field of Internet science



Improving Internet measurement technology

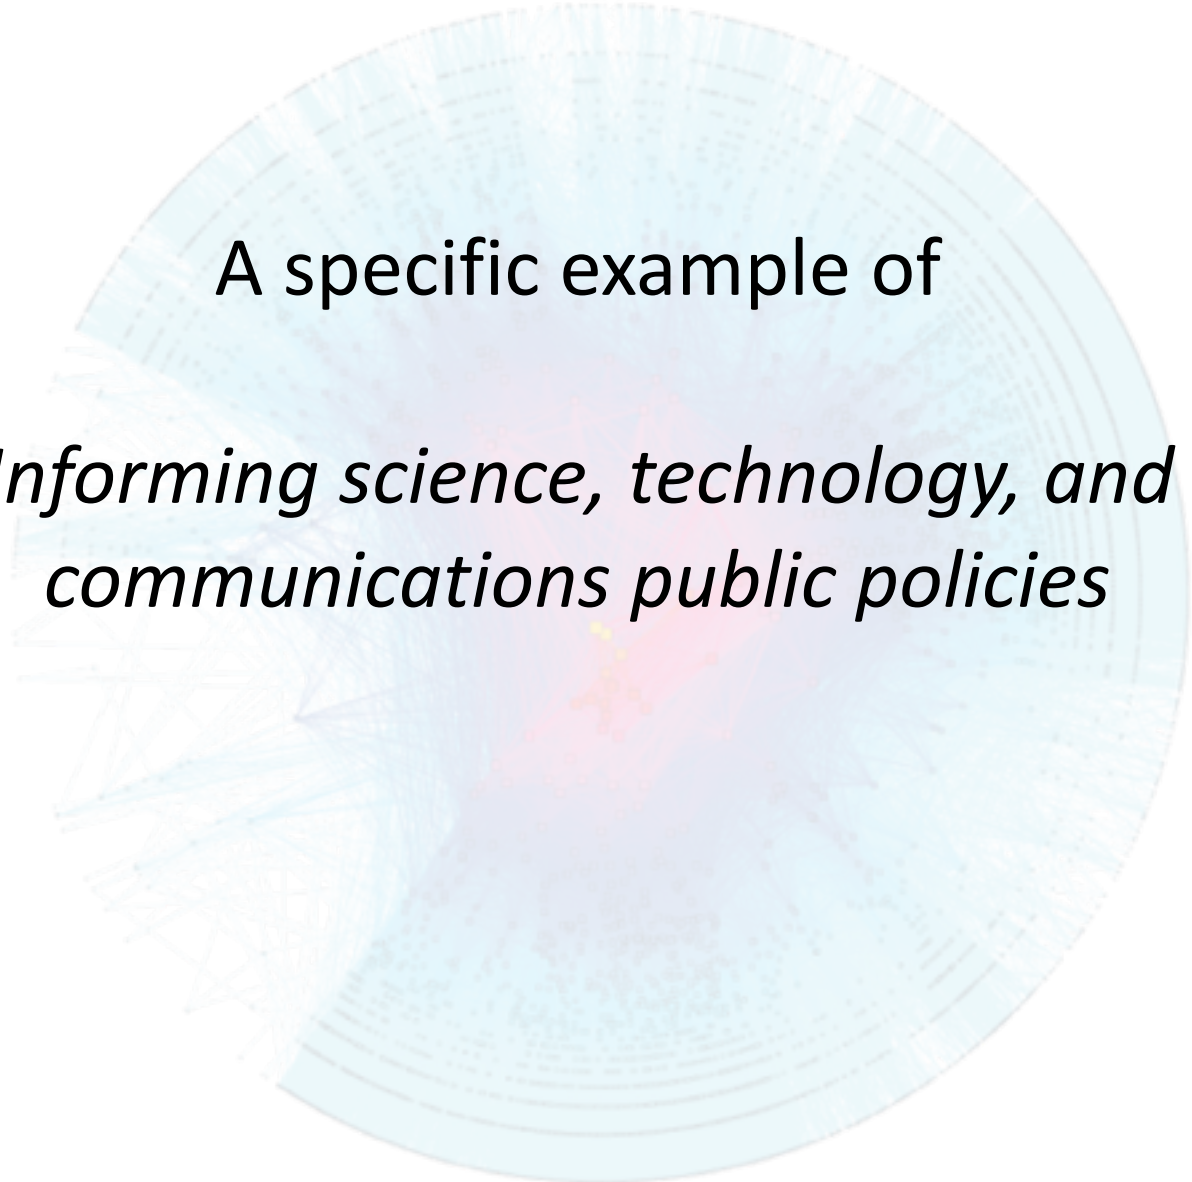


Collection, curation, analysis, visualization, and dissemination
of sets of the best available Internet data



Informing science, technology, and communications public policies

[more details on all of our activities at
<http://www.caida.org/home/about/annualreports/> .
but this is for another day.]



A specific example of
*Informing science, technology, and
communications public policies*



at&t



DIRECTV®



at&t



DIRECTV®

*The newly combined company – the largest pay TV provider in the United States and the world – will offer millions of people more choices for video entertainment on any screen from almost anywhere, any time. ..“We’re now a fundamentally different company with a diversified set of capabilities and businesses **that set us apart from the competition.**”*

“..that set us apart from the competition.”

that's what the competition is worried about

*Discussion. As stated in the 2015 Open Internet Order, “consumers bear the harm when they experience degraded access to the applications and services of their choosing due to a dispute between a large broadband provider and an interconnecting party.” Also, because OVD subscribers expect high-quality video, OVDs are vulnerable to degradation at the interconnection point with a broadband Internet access service provider’s last mile network. Thus, as stated in the 2015 Open Internet Order, we find that “**broadband Internet access providers have the ability to use terms of interconnection to disadvantage edge providers and that consumers’ ability to respond to unjust or unreasonable broadband provider practices are limited by switching costs.**” We appreciate commenters’ concerns in this area.*

Memorandum Opinion and Order, In the Matter of ATT/DirecTV

http://transition.fcc.gov/Daily_Releases/Daily_Business/2015/db0728/FCC-15-94A1.pdf

T/TDV *merger conditions*

AT&T has agreed to the following conditions for the next 4 years:

1. Fiber Internet access to 12.5M customers;
2. 1 Gbps to any E-rate school or library in AT&T's fiber footprint;
3. Discounted broadband service to low-income households;
4. Broadband Internet services will not favor its own online video programming services. (discounted bundling ok);
5. Submit to the FCC new interconnection agreements;
6. ***Develop, in conjunction with an independent expert, a methodology for measuring performance of its Internet traffic exchange and regularly report these metrics to the FCC.***

history in the making..

measurement conditions

....**the Company must report**, in accordance with the filing and service requirements set forth in Section VII.5... **the following performance characteristics of traffic exchanged at Internet Interconnection Points located within the United States**, unless the volume of traffic exchanged with the interconnecting party is less than a *de minimis* threshold, as specified by the Independent Measurement Expert:

- (i) **The probability distribution of latency** between the border router of the interconnecting network and the Company's border router ("Latency"), as defined by the Independent Measurement Expert;
- (ii) **The percentage of packets dropped** at or between the border router of the interconnecting network and the Company's border router ("Packet Loss"), as defined by the Independent Measurement Expert; and
- (iii) **The percent usage of each Internet Interconnection Point** ("Utilization"), as defined by the Independent Measurement Expert.

measurement conditions

Within sixty (60) days of engaging the Independent Measurement Expert, the Company, in consultation with the Independent Measurement Expert, will submit for approval by the Commission's Office of General Counsel, in consultation with the Wireline Competition Bureau and the Chief Technologist, a report describing the Independent Measurement Expert's proposed methodology for the measurement of the performance metrics described herein. Such report shall also be submitted to the Independent Compliance Officer. The proposed methodology should, at a minimum, address the following criteria:

[https://apps.fcc.gov/edocs_public/attachmatch/FCC-15-94A1.pdf]

measurement criteria


1. Identification of Internet Interconnection Points, including the identity of the interconnecting parties and **location and capacity of each interconnection point**;
2. Identification of a disclosure exemption threshold for a de minimis volume of traffic exchanged between the Company and interconnecting parties;
3. A definition of **“Latency,”** which shall include the probability distribution;
4. A definition of **“Packet Loss”**;
5. Time of measurements, which shall, at a minimum, include an identified window within **peak usage periods**;
6. For any performance metric contingent upon an interconnecting party’s participation in the selected measurement methodology, a **process for waiving** disclosure of that metric at points of interconnection where the interconnecting party declines to participate;
7. Frequency and duration of measurements;
8. Any devices used for measurement;
9. End points of measurements;
10. Placement of any devices;
11. Frequency of disclosures.

Independent Measurement Expert (IME) Identified 1 October 2015

The Center for Applied Internet Data Analysis (CAIDA), a research group at University of California at San Diego, has been identified as the Independent Measurement Expert, “by an agreement between AT&T and the Commission’s Office of General Counsel, which has approved the selection, as required by Section V(2)(c)(i) of Appendix B of the Order.” (Contract in place 1 November 2015.)

[http://transition.fcc.gov/Daily_Releases/Daily_Business/2015/db1001/DA-15-1116A1.pdf]

Within sixty (60) days of engaging the Independent Measurement Expert, the Company, in consultation with the Independent Measurement Expert, will submit for approval by the Commission's Office of General Counsel, in consultation with the Wireline Competition Bureau and the Chief Technologist, a report describing the Independent Measurement Expert's proposed methodology for the measurement of the performance metrics described herein.



publicly available in FCC ECFS (**BUT DRAFT NOT APPROVED BY THE FCC, WILL CHANGE**)
<http://apps.fcc.gov/ecfs/document/view;ECFSSESSION=8c3fWJkN7DyfNyhh24njsbLqnhLtQjnHpqpGLTM4DhG3PbqfKQs!1951721665!-1566059965?id=60001396715>

what did we suggest?

publicly available in FCC ECFS (**BUT DRAFT NOT APPROVED BY THE FCC, WILL CHANGE**)
[http://apps.fcc.gov/ecfs/document/view;ECFSSESSION=8c3fWJkN7DyfNyhh24njsbLqnhLtQjnHpqpGLTM4DhG3PbqfKQs!
1951721665!-1566059965?id=60001396715](http://apps.fcc.gov/ecfs/document/view;ECFSSESSION=8c3fWJkN7DyfNyhh24njsbLqnhLtQjnHpqpGLTM4DhG3PbqfKQs!1951721665!-1566059965?id=60001396715)

first, why did we take this on?

- Goal: rigorously demonstrate that the performance of AT&T's points of interconnection with partners are not a source of performance impairment for consumers.
- Also a bigger picture here: Open Internet Order (if it stands) explicitly asserts the authority to regulate interconnection....maybe later.
 - Already requires loss measurements on access links
- Opportunity to help FCC get the best available data and understanding of the richness and complexity of interconnection.

second, **not** typical academic research

- **We** don't measure here
- No iteration before method published, thus some degree of 'ex ante' precision warranted
- Proposed comparison of different methods for measuring same parameter, to lend confidence in interpretation
- And although the IME reviews the first (few?) reports, the IME is not ideally placed to validate observations
- Best positioned party to validate: interconnecting partner.

➔ Thus, our method requires sharing some partner-specific data with each partner.

third, we recognize the sensitivity

Sharing of data implies release of data that may be considered proprietary to the firm.

We limited sharing to that which we believe is essential to validation, and not commercially significant, given what is already feasibly obtainable by the partner.

Primary objective of this exercise is that all parties feel the process has been balanced and fair.

We believe data sharing is a necessary component of this outcome, but the final decision lies with FCC.

Data sharing requirements

*“This condition will enable the monitoring of the combined entity’s future interconnection agreements terms to determine whether the combined entity is using such agreements to deny or impede access to its networks in ways that limit competition from third-party online video content providers. In addition, **this condition requires the combined entity to work with an independent measurement expert to report certain Internet interconnection performance metrics, and to the extent possible, make such metrics publicly available.**”*

Memorandum Opinion and Order, In the Matter of ATT/DirecTV

http://transition.fcc.gov/Daily_Releases/Daily_Business/2015/db0728/FCC-15-94A1.pdf

Data sharing requirements

“report certain.. interconnection performance metrics, and to the extent possible, make such metrics publicly available.”

We agree there is great value in releasing overall insights and summaries to the public, as well as releasing specific data about a given interconnection partner to that partner.

proposed method
(work in progress)

Identification of Interconnections

- Location, capacity and utilization of interconnection points
- Share above data with interconnection partner
 - Support cross-checking
- Report to FCC any traffic differentiation, or non-BIAS (public Internet) traffic on links
 - annotate data accordingly

de minimis threshold

- **Peers and on-net only customers**
- **N Gbps of capacity**
 - Selected threshold that provides useful data to FCC but not unreasonable reporting burden
 - Exact list covered by NDA, but partners will know.

performance metrics

- (i) Probability distribution of **latency**
- (ii) Percentage of packets **dropped**
- (iii) **Utilization**

will discuss in reverse order..

utilization metrics

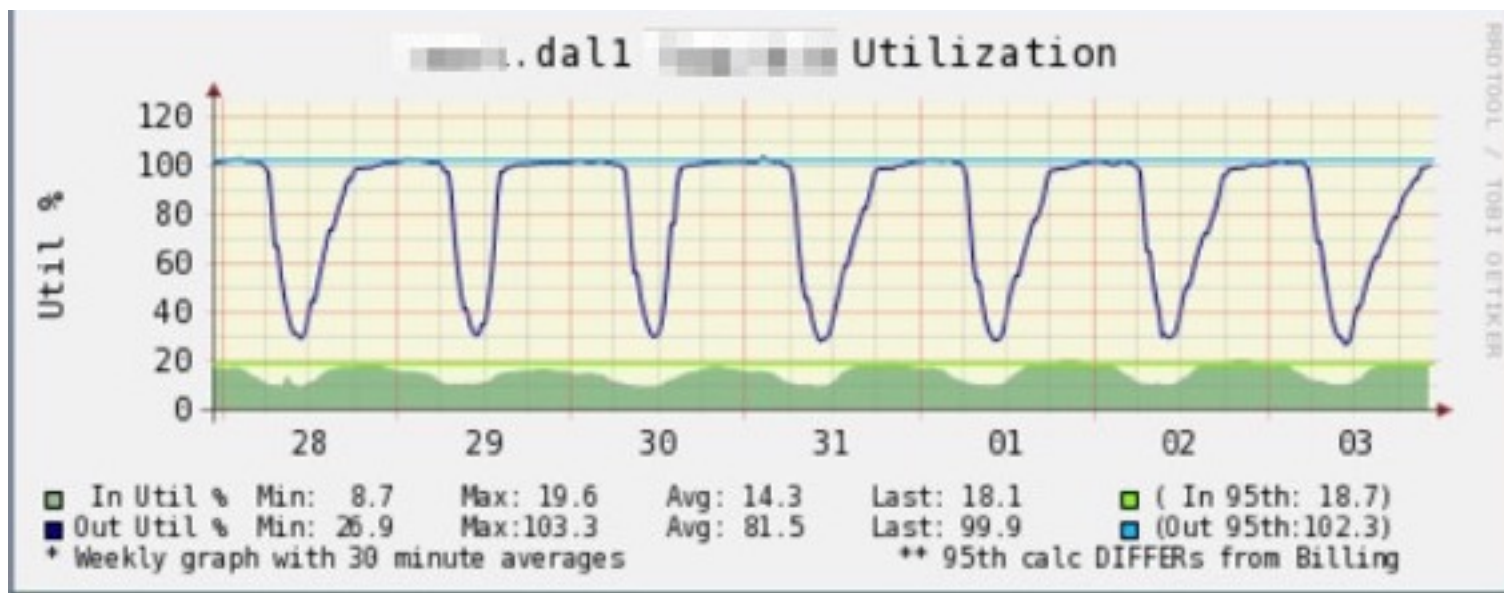
Number of bytes transmitted/outbound-capacity

Number of bytes received/inbound-capacity

(using router counters)


Why not just infer congestion from link utilization data?

- heavily loaded links often show a plateau (flat top) of utilization at or near the actual link capacity
- evidence (but not proof) of congestion
- could also be extremely effective traffic engineering



Why not just infer congestion from link utilization data?

- link utilization alone cannot reveal an accurate picture
- need *sufficiently accurate* metrics of loss and latency



<http://blog.level3.com/open-internet/observations-internet-middleman/>

what does sufficient accuracy require?
(part of the challenge: we are not sure.)

Loss metrics

- Definition: *lost packet*: one that arrives at one router of an interconnecting LAG, intended and valid for delivery to the router on the other side of that LAG, which does not successfully reach it.
- Definition: *loss rate* for some interval is the ratio of packets lost to the sum of packets lost and packets successfully delivered during that interval.

Packet loss: from router counters

- Report (to FCC only) loss on AT&T's outgoing interface
- **But** for many partners, bulk of traffic will flow into AT&T and losses will be seen on partner's router interface, not AT&T router
- If interconnection partners are willing to supply loss counter data for links, then it should be reported
- Caveats
 - Loss rate might be higher than counters indicate
 - Not all excess load is necessarily dropped at interconnection link

Passive vs. Active Loss Measurement

- “Passive” measurement: router counters
 - Counts only known losses
- Active measurement: send probe packets
 - Sampling challenges, at least
- Relative accuracy of these two approaches not clear
- We are requesting both
- In fact we are requesting two methods of active measurements: a superior method, that requires cooperation, and an inferior fallback method

Latency metrics

- Definition: *round-trip latency* is time required for a packet to travel from a source router to its destination and immediately back to the source router. (One-way latency if feasible.)
- Define probability distribution of latency samples as probability density function
 - relative likelihood for variable to take on a value

Loss and latency: active measurements

ICMP: Default method to probe Other People's Networks

Caveats

- Some routers may not respond to ping/TTL-limited measurements, leading to over-estimate of loss rate
- Some routers respond with highly variable delay
- Some routers treat probe packets differently
- Probing *past* router requires topology discovery
- Return route may be asymmetric

The IETF has developed something
much better than ICMP for this..

A Two-Way Active Measurement Protocol (TWAMP)

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

The One-way Active Measurement Protocol (OWAMP), specified in RFC 4656, provides a common protocol for measuring one-way metrics between network devices. OWAMP can be used bi-directionally to measure one-way metrics in both directions between two network elements. However, it does not accommodate round-trip or two-way measurements. This memo specifies a Two-Way Active Measurement Protocol (TWAMP), based on the OWAMP, that adds two-way or round-trip measurement capabilities. The TWAMP measurement architecture is usually comprised of two hosts with specific roles, and this allows for some protocol simplifications, making it an attractive alternative in some circumstances.

The IETF has developed something
much better than ICMP for this..

A Two-Way Active Measurement Protocol (TWAMP)

Status of This Memo

This document specifies an Internet standards track protocol for the

*but it requires cooperation from
both endpoints of the measurement!*

Abstract

The One-way Active Measurement Protocol (OWAMP), specified in RFC 4656, provides a common protocol for measuring one-way metrics between network devices. OWAMP can be used bi-directionally to measure one-way metrics in both directions between two network elements. However, it does not accommodate round-trip or two-way measurements. This memo specifies a Two-Way Active Measurement Protocol (TWAMP), based on the OWAMP, that adds two-way or round-trip measurement capabilities. The TWAMP measurement architecture is usually comprised of two hosts with specific roles, and this allows for some protocol simplifications, making it an attractive alternative in some circumstances.


Selected IETF RFCs and loss measurement

RFC	Title	Date
2680	A One-way Packet Loss Metric for IPPM	Sept 1999
3357	One-way Loss Pattern Sample Metrics	Aug 2002
3393	IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)	Nov 2002
3432	Network performance measurement with periodic streams	Nov 2002
3611	RTP Control Protocol Extended Reports	Nov 2003
6374	Packet Loss and Delay Measurement for MPLS Networks	Sept 2011
6534	Loss Episode Metrics for IP Performance Metrics (IPPM)	May 2012
6673	Round-Trip Packet Loss Metrics	Aug 2012
Internet Draft	Model Based Bulk Performance Metrics	Dec 2013 - Mar 2015

Active Measurements


Participating Partner

- IP SLA (Cisco)
- TWAMP



Non-participating Partner

- ICMP based



Execute and report on both methods initially

*Time, frequency, duration of measurements
(which must include peak usage periods)*

The peak usage periods for any LAG are those times at which the 5-minute sampled utilization is above 80%.

Report loss and latency measurements separately for every LAG for peak periods.

Frequency of probing: mean of 1 probe/sec,
Poisson if possible [RFC7680]

Utilization: 5-minute intervals across month.

Process for waiving disclosure of measurements that require cooperation

If AT&T does not obtain cooperation of the interconnection partner, they will submit a written explanation to the FCC.

The partner may also submit a written explanation.

If the FCC confirms that the outcome is reasonable, the IME will approve the fall-back scheme that does not require cooperation.

FCC still mulling this one over..

Measurement endpoints

If not measuring from the interconnection router itself, approval required for:

1. Any devices used for measurement;
2. End points of measurements;
3. Placement of any devices.

Limitations of these measurements

- A link without losses or queuing might still be bottleneck, e.g., rate adaptation occurs
- Only measuring interconnection links, not entire path
- Measurement inaccuracies possible with all methods, especially our fallback method

Monthly reporting requirements

- Interconnection locations & capacities for all partners
- Packet and byte counts, 5-minute utilization
 - peak, 95p, average, time-series plots
 - inbound and outbound
- Loss rate: 3 ways: router counters (+partners), TWAMP, ICMP.
 - plot for peak, off-peak, overall
 - allow comparisons with utilization plots
- prob distribution of latency: TWAMP and ICMP
 - peak, off-peak, and overall
 - two directions if possible

*graphs, spreadsheet, **digital** version of raw & summary data*

Three goals in telling you all this

(1) Let you know this is going on

(2) Get your feedback

Probably not the last transparency episode

(3) Convince you of the need for cooperation to make this episode useful for everyone

Thanks

att-ime-feedback@caida.org
(or nanog list)