First Amended Report of AT&T Independent Measurement Expert: Reporting requirements and measurement methods

UC, San Diego's Center for Applied Internet Data Analysis
KC Claffy, Amogh Dhamdhere
in collaboration with independent contractors:
David Clark, Steven Bauer
Submit any feedback/concerns to att-ime-feedback@caida.org

June 28, 2016

Contents

1 Introduction

In compliance with the merger Order between AT&T and DirecTV, AT&T has engaged CAIDA to serve as the Independent Measurement Expert (IME) as described in the Memorandum Opinion and Order ("FCC Order"), Appendix B, V: Internet Interconnection Disclosure Requirements.[1] As described in this appendix, our task is to develop a proposed methodology for measurement and reporting of data, in the context of the following obligation imposed on AT&T:

> *2.b. Reporting Internet Interconnection Performance Metrics. Using a methodology developed and implemented as subpart 2.c., the Company must report, in accordance with the filing and service requirements set forth in Section VII.5. herein, on a schedule established by an Independent Measurement Expert, but no more than on a monthly basis, 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.*

...

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

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

8. *End points of measurements;*
9. *Placement of any devices; and*
10. *Frequency of disclosures.*

The structure of this document is as follows. Section 2 specifies the disclosure exemption. Section 3 described the reporting requirements on interconnecting parties, connectivity and traffic management details. Sections 4 and 5 catalog the data that AT&T must collect, and methods to collect it. We distinguish the process of collection from the reporting that results from the measurements. Section 6 specifies the nature of the reporting.

In a separate filing, titled "Report of AT&T Independent Measurement Expert: Background and supporting arguments for our analysis" we provide our rationale for our data sharing requirements and specific measurement methodologies.

2 Disclosure exemption

The FCC has required AT&T to report on all interconnections with peers and with on-net only customers (similar to what has been called "paid peering" in the past). The set of interconnecting parties is addressed in the merger order, Appendix B, Section II, as follows:

> *"Internet Interconnection Points" means the facilities over which traffic is exchanged between the Company's network that carries Broadband Internet Access Service traffic and (1) peering networks or (2) customers that purchase on-net only services to deliver traffic to and from the Company's end users over the company's network.*

AT&T is required to report on all such interconnection points "unless the volume of traffic exchanged with the interconnecting party is less than a *de minimis* threshold, as specified by the Independent Measurement Expert". AT&T and the FCC have identified a proposed minimum volume criterion to trigger reporting for their interconnecting parties, which is that the party connect with at least 80 Gbps of capacity. We have used data from AT&T and CAIDA's AS Rank system to infer the list of AT&T peers, determined the difference between this inferred list and AT&T's list of peers and on-net customers with at least 80 Gbps capacity, and concluded that this proposed threshold is reasonable.

1. AT&T shall collect data and report on any interconnecting party with at least 80 Gbps of interconnection capacity in sum across all interconnection points at the beginning of each report period. In computing the total capacity, AT&T may exclude from the computation any individual LAG with capacity of less than 10 Gbps.

2. With its first monthly report, AT&T shall provide to the FCC a list of all interconnecting parties with interconnection capacity of at least 80 Gbps in sum across all interconnection points. In each subsequent monthly report, AT&T shall revise this list to include any interconnecting party whose capacity exceeds this threshold, and delete interconnecting parties whose capacity falls below this threshold. Measurements and reporting on a new interconnecting party will commence within 90 days of that party exceeding the threshold.

3 Reporting of interconnecting parties, connectivity, and traffic management details

3. For each qualifying interconnecting party (as described in Section 2, and hereafter referred to as *interconnecting party*), AT&T shall include in its monthly report to the FCC the location and capacity of all 10 Gbps or greater links (or LAGs) connecting AT&T to that party at border routers. Link Aggregation Groups, or *LAGs*, are multiple parallel circuits connecting AT&T to an interconnection party. In cases where links are combined into LAGs, AT&T shall report the required data for the combined LAG. In cases where links are not combined into LAGs, AT&T shall report the required data for such individual links. We use the term *LAG* to capture both cases.

4. For each interconnecting party, AT&T shall share the information above (item 3) with that party, so long as the interconnecting party wants to receive it and is willing to sign a confidentiality agreement prohibiting disclosure except to the FCC, and then only under the protective order issued in Docket 14-90. This sharing step is required as a mechanism for cross-checking the gathered data, although AT&T need not seek feedback from the interconnecting party regarding this information.

5. AT&T shall disclose to the FCC any AT&T use of tools for differentiated treatment of traffic across LAGs. Mechanisms for differentiated treatment of traffic, e.g., diffserv, across an LAG, can distort measurement results. For example, probing for packet loss and latency variation may fail to show overall link performance if probe packets end up in a preferred service.

6. AT&T shall ask its interconnecting parties to disclose in writing to AT&T whether they use such tools, with the overall goal of determining if the probing method used by AT&T is correctly measuring the character of the LAG. If AT&T is unable to obtain the interconnecting party's agreement to disclose their use of such tools to AT&T and the IME with suitable disclosure protection, AT&T shall advise the party that they may disclose this information directly to the FCC and to the IME. If the interconnecting party discloses this information, it will be covered by a confidentiality agreement prohibiting disclosure except to the FCC and the IME, and then only under the protective order issued in Docket 14-90.

7. AT&T shall provide a written list to the FCC and the IME of any of AT&T's interconnection LAGs that are carrying AT&T traffic other than Broadband Internet Access Service (BIAS) traffic. If any interconnection LAGs are carrying AT&T traffic other than BIAS during the period within which the reporting requirements apply, AT&T shall, for each such LAG, describe in writing to the FCC and the IME: (1) whether all traffic on the LAG is receiving identical treatment or how different services are treated differently; (2) how the LAG allocates capacity to the BIAS and other traffic (static or dynamic); and (3) how different classes of traffic may affect each other. For each such LAG, AT&T shall report utilization and packet loss separately for BIAS traffic and for the link as a whole. AT&T shall ensure that any measure based on probing will only measure loss and latency of the BIAS traffic. If there are technical barriers to compliance with this requirement, AT&T shall work with the IME to develop a suitable approach, which shall be subject to written approval by the IME.

4 Collection of existing data on utilization and loss

This section specifies utilization and loss data that AT&T will collect about each interconnection LAG (or link in the absence of a LAG), using information available from routers and related databases.

8. Using router interface counters and the SNMP protocol, or other protocol approved by the IME in writing, AT&T shall collect the number of bytes transmitted and received at each router in a LAG at a 5-minute granularity for the entire monthly reporting period, and compute 5-minute LAG utilization values from these numbers. For this purpose, we define *outbound utilization* as the number of bytes transmitted divided by the capacity of the LAG in the outbound direction, and the *inbound utilization* as the number of bytes received divided by the capacity of the LAG in the inbound direction. Any changes to reporting of capacity and byte counts must be synchronized (See Section 6).

9. For each interconnecting party, AT&T shall share these per-LAG transmitted/received byte counts with that party, so long as the interconnecting party wants to receive that information and is willing to sign a confidentiality agreement prohibiting disclosure except to the FCC, and then only under the protective order issued in Docket 14-90. This sharing step is required as a mechanism for cross-checking the gathered data, although AT&T need not seek feedback from the party regarding this information.

Definition of one-way packet loss. We define a *lost packet* as 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 that point. The *one-way loss rate* for some interval is the ratio of packets lost to the sum of packets lost and packets successfully delivered during that interval.[2] For the passive method described in the next item, we use the loss counters implemented as part of the router management system as an estimate of loss. We recognize that due to implementation issues, a router may not record all the events that lead to a packet being dropped.

10. Using router interface counters and the SNMP protocol, or other protocol approved by the IME in writing, AT&T shall collect counts of packets sent and outgoing packets dropped during each 5-minute interval over the course of the month.

11. If packets are dropped at the incoming side of a congested LAG, i.e., *into* AT&T's network, AT&T's routers will not be able to record those losses. To account for losses in this direction, AT&T shall request in writing, with the IME copied on the correspondence, to obtain an agreement from each interconnecting party to share with AT&T, the IME and the FCC, the number of packets transmitted and packets lost in the direction to AT&T's LAGs from the party's router counters. If AT&T is unable to obtain the interconnecting party's agreement to share this loss data with AT&T, AT&T shall advise the party that they may disclose this information directly to the FCC or to the IME. Disclosure of this information to the FCC or IME will be subject to the protective order issued in Docket 14-90, or by a confidentiality agreement prohibiting further disclosure except to the FCC and the IME, respectively. AT&T is not responsible for the accuracy or completeness of this data, or the interconnecting party's failure to deliver this data to AT&T in a timely manner.

[2] A more technically detailed definition of one-way packet loss is available in RFC7680, A One-Way Loss Metric for IP Performance Metrics (IPPM), G. Almes, S. Kalidindi, M. Zekauskas, and A. Morton (Ed.) https://www.rfc-editor.org/rfc/rfc7680.txt.

At the commencement of reporting to the FCC, AT&T shall separately forward this data to the FCC. Upon identification of the initial set of cooperating interconnecting parties, AT&T and the IME will work with those parties to define a standard reporting format for this data. From the time a standard format is established, AT&T will have 60 days to integrate the interconnecting party's data in that format into its reports to the FCC. For those parties that choose to report their packet and loss counter data in this standard format, and provide the data to AT&T no later than 10 business days prior to each monthly reporting deadline specified in paragraph 21 below, AT&T shall integrate this data into the report that AT&T prepares for the FCC on the LAG or aggregate in question. AT&T shall not be responsible for validating the accuracy of any interconnecting party's data provided to AT&T and shall not be responsible for integrating any month's data that an interconnecting party provides to AT&T less than 10 business days prior to that month's reporting deadline.

5 Active measurements of loss and latency

In addition to the data capture described above, we define active probing measurement methods that AT&T shall implement to provide another measure of packet loss, as well as to provide data about latency and its variation. The general approach is to send a probe that results in a return packet. Failure to receive a return packet provides evidence of packet loss on the LAG. Our methodology uses observed failure to receive responses to probe packets as the basis to estimate the overall loss rate on the LAG. The timing of the return packet provides a data point about latency. Cooperation from the other side of the interconnection may allow one-way measurements that can discern direction of packet loss, and may be able to provide latency distributions in each direction across the interconnection LAG (See Section 6.4).

Definition of round-trip packet loss. In the context of probing as a measurement method, we define a *round-trip packet loss* as the event where the source sends a probe packet but does not receive a reflected packet from the destination within a reasonable time frame, either because the destination did not receive the packet, the destination did not reflect the packet, or the source failed to receive the reflected packet.[3] For active probing, the *round-trip loss rate* shall be the number of probe packets sent for which the sending source does not receive a reflected packet within a reasonable time frame, divided by the number of probe packets sent; this loss metric shall also include the number of probing packets sent, to enable computation of a confidence interval. Section 6 specifies three different intervals over which to report a loss rate: the peak periods of an epoch; the non-peak periods of an epoch; and overall for the reporting epoch.

Definition of latency. We define *round-trip latency* as the time required for a packet to travel from a source router to its destination (in this case, the interconnecting party's border router, or an adjacent probe target identified by the interconnecting party as per Item 16) and immediately back to the source router. Consistent with the terms of the CMaO, we define the probability distribution of individual latency samples to be equivalent to the probability density function (pdf), which describes the relative likelihood for a variable to take on a given value. Section 6.4 further describes the computing and reporting of the probability distribution.

[3] A more technically detailed definition of metrics for round-trip packet loss is available in RFC 6673 at https://tools.ietf.org/html/rfc6673.

5.1 Probing methods

We define a primary method that relies on cooperation of the interconnecting party, and a fall-back scheme if the interconnecting party declines to cooperate. The preferred method uses a measurement protocol, e.g., TWAMP or IPSLA. The fall-back method uses a less accurate and less precise protocol (ICMP).

12. Within 30 days after the methodology is approved by the FCC and prior to initiating measurement of an interconnecting party's LAGs, AT&T shall seek cooperation from each interconnecting party, in writing and copying the IME on all correspondence, to determine if that party is interested in supporting TWAMP/IPSLA measurements by enabling the far-side router, or other adjacent machine as identified by the interconnecting party, as a responder for TWAMP/IPSLA probes. AT&T shall enable and participate in a dialogue with the interconnecting party and the IME regarding any concerns the party has regarding this cooperation. If the party agrees to support TWAMP/IPSLA, this protocol will be the first probe method for active measurement of loss and latency. Once an interconnecting party agrees to participate in TWAMP/ISLA, AT&T will work with the interconnecting party to perform the testing and certification necessary to implement the method across the interconnecting link. AT&T shall have 120 days from the date the interconnecting party enables the TWAMP/IPSLA responder to begin reporting data as described below. This process shall not delay the reporting of any other measurements.

13. AT&T shall use its interconnecting router as a TWAMP/IPSLA and ICMP client executing measurements. Alternatively, upon written approval of the IME, AT&T may measure using a probe server placed suitably near each location where they interconnect with an interconnecting party. If AT&T does not use its interconnecting router for the active measurements, it must report and get written approval from the IME for the following: any devices used for measurement; endpoints of measurement, including physical and topological placement of any devices; and protocol used for measurements from the probe server.

14. With each interconnecting party that agrees to enable TWAMP/IPSLA, and for so long as such enablement continues, AT&T shall share with that party the results of probing that party using this method, so long as the interconnecting party wants to receive that information and is willing to sign a confidentiality agreement prohibiting disclosure except to the FCC, and then only under the protective order issued in Docket 14-90. This sharing step will allow the interconnecting party to optionally verify or cross-check the observed data. If the interconnecting party indicates to AT&T at the time that it enables TWAMP/IPSLA for AT&T that it wants AT&T to enable a responding TWAMP/IPSLA service on AT&T's interconnection router, AT&T will do so if the interconnecting party is willing to sign a confidentiality agreement prohibiting disclosure except to the FCC, and then only under the protective order issued in Docket 14-90.

15. In the case where TWAMP/IPSLA is infeasible, or if the interconnecting party does not enable TWAMP/IPSLA, AT&T shall use a fall-back ICMP-based probing method (e.g., echo request or TTL-limited probes) as the primary active method. This fall-back method constitutes the waiver process required by the FCC's Order.[4] In the event that an interconnecting

[4] *"For any performance metric contingent upon an interconnecting party's participation in the selected measurement methodology, a process for waiving the disclosure of that metric at points of interconnection where the interconnecting party declines to participate" https://apps.fcc.gov/edocs_public/attachmatch/FCC-15-94A1.pdf*

party later enables TWAMP/IPLSA, AT&T will begin reporting data for that interconnecting party consistent with the schedule described in Item 12 above.

16. AT&T shall ask the interconnecting party if it is willing to identify a machine (by IP address) as close as possible to the interconnecting router such that the IP address can serve as the target of an ICMP probe, in order that the probe packet pass through the router rather than terminate on the router. If the interconnecting party provides an IP address, AT&T shall use this IP address as the target of ICMP probes. Otherwise, AT&T will use the AT&T-facing interface on the interconnecting router as the target of ICMP probes. If, during the period that the IME is reviewing data as part of its validation of the method, a probe to a responder beyond the far side router shows variation in latency that indicates congestion, AT&T will also direct ICMP probes to the far side router and share this data with the IME, as a means to determine if the probes to the far side router show similar behavior. This additional probing will be temporary, as a part of the IME's validation of the method.

17. To detect the situation where the return path is not through the LAG being measured, AT&T shall probe once per hour the path to the far-side router with the IP Record Route option, if AT&T can operationally support that option. If AT&T cannot operationally support the IP Record Route option, AT&T will request the interconnecting party to provide a return traceroute to confirm no assymetrical routing exists along the measurement path. If AT&T detects an asymmetric route while using the approved method, and AT&T cannot successfully probe the LAG, AT&T shall disclose this condition to the FCC and annotate the probing data in its reports with this information for so long as the routing asymmetry persists.

18. In cases where the interconnecting party cooperates in the deployment of TWAMP/IPSLA, AT&T shall execute both methods above (TWAMP/IPSLA and ICMP-based), in order to cross-check the data and inform confidence levels in the second method. If the IME finds the measurements from the second method to be evidently inaccurate, the IME may, in writing, permit AT&T to suspend probing of this connection with this method.

19. If AT&T is unable to fulfill any requirement described in this Report (e.g., as a result of any significant operational issue on AT&Ts network or an interconnecting party's network, or lack of sufficient cooperation from an interconnecting party), AT&T shall (i) notify the FCC's Office of General Counsel and Chief Technology Officer, and the IME, within seven days of any suspension of measurement or determination that AT&T will not meet a deadline for initiation of measurement; (ii) propose to these parties a solution within 30 days that will allow AT&T to comply with its reporting obligations under this Report. This proposed solution must include timelines for complying with its reporting obligations, and must be approved in writing by the IME.

5.2 Probing rate

20. For both the preferred and fall-back schemes, AT&T shall use a probing rate with a mean of 1 second, with a Poisson distribution of sample intervals to improve the sampling method. Poisson sampling is described in RFC-7680.[5] If AT&T is not able to support Poisson sampling on their routers, they may use periodic (1 per second) sampling. AT&T shall include in its monthly report to the FCC the sampling interval they are using.

[5]RFC 7680, "*A One-way Packet Loss Metric for IP Performance Metrics (IPPM)*", G. Almes, S. Kalidindi, and M. Zekauskas, and A. Morton, January 2016. http://tools.ietf.org/html/rfc2680.

6 Reporting requirements

21. Unless otherwise specified in this report, AT&T shall produce monthly reports of the measurements required in this methodology, in the format described in Section 6.5. AT&T shall prepare and deliver such reports within 30 days after the end of relevant reporting period. As described in Section 3, AT&T will report, for each interconnecting party, the location and capacity of all LAGs for such parties. Changes to capacity, or the number and location of LAGs for an interconnecting party must be reported no later than the monthly reporting period following the period within which the change occurred.

6.1 Reporting packet and byte counts and utilization

22. For each LAG, AT&T shall obtain the inbound and outbound packet and byte counts as described in Section 4, and compute the inbound and outbound utilization (as a ratio of byte counts over LAG capacity) for each 5-minute interval in the month. AT&T shall retain the raw data for packet and byte counts and utilization in each 5-minute interval, for one year after the expiration of the reporting requirements. Any changes to reporting of capacity and these metrics must be synchronized, e.g., if AT&T adds a new link to a LAG or adds a new LAG, AT&T must update all metrics at the same time.

23. For (1) each LAG, (2) all LAGs with an interconnecting party within a city, and (3) all LAGs with an interconnecting party throughout the United States, AT&T shall report the peak value of the 5-minute utilization values, the 95th percentile utilization value, and average utilization over a month, and time-series utilization plots (e.g., MRTG) of 5-minute granularity utilization data, inbound and outbound, over a month.

    To aggregate byte counts for a single interconnecting party by city, separately for inbound and outbound, AT&T shall add the byte counts recorded on each LAG connecting with the interconnecting party in that city. To compute utilization for the interconnecting party in a city, for each direction separately, AT&T shall divide the aggregate byte counts by the total capacity of all LAGs connecting to that interconnecting party in that city. AT&T shall follow the same procedure across all LAGs connecting to that interconnecting party in all cities to compute the aggregate byte counts and utilization for the interconnecting party as a whole.

24. For each interconnecting party, AT&T shall share per-LAG packet and byte counts and the time-series utilization graphs with that party, so long as the party wants to receive that information and is willing to sign a confidentiality agreement prohibiting disclosure except to the FCC, and then only under the protective order issued in Docket 14-90.

6.2 Peak usage period reporting

25. AT&T shall use the data for each LAG to define the peak usage periods for that LAG. The peak usage periods for any LAG are those times at which the 5-minute sampled utilization is above 80%. AT&T shall report packet loss and latency measurements separately for every LAG for those periods we define as peak (Section 6.3 and 6.4). For each LAG, AT&T shall also report the number of 5-minute samples for which the LAG utilization exceeds 80%, in order to gauge statistical significance of the reported value.

6.3 Reporting packet loss

AT&T shall separately report on packet loss derived from the various measurement methods defined in Section 4 and Section 5:

26. For the loss measurements derived from counters in routers (Section 4), AT&T shall report packet loss count and loss rate in 3 different ways for each LAG: for peak usage periods (as defined in Section 6.2), off-peak usage periods, and overall for the monthly reporting period. If the interconnecting partner agrees to report packet counts and packet loss counts from their router to AT&T using the standard format and procedures described in item 11 above, then AT&T shall report these counts as well as the loss rate derived from these counts, for each of these 3 periods. AT&T shall not be responsible for validating the accuracy of any interconnecting party's data provided to AT&T and shall not be responsible for integrating any months data, or calculating loss rates derived from such data, that an interconnecting party provides to AT&T less than 10 business days prior to that month's reporting deadline.

27. In addition to reporting of router counters, AT&T shall separately report the loss rate as captured by the active probing method(s) as a cross-check. If router counter data from the interconnecting party is not available, AT&T shall report only the loss rate(s) based on the probing method(s).

    For active probing methods (both TWAMP/IPSLA if available and ICMP), AT&T shall compute the loss rate as the ratio of probe responses not received to probes sent during an interval, separately computing and reporting this metric for peak and off-peak times and overall for the LAG, as described for the reporting of the counter-based measurement. For both probing methods AT&T shall also report the number of sample probes used to compute this loss rate, to enable computation of confidence intervals. If using TWAMP/IPSLA, and the available implementation has the capability, AT&T shall report inbound and outbound loss rate separately. For the ICMP-based method, AT&T shall report the round-trip loss rate.

28. For each LAG, AT&T shall compute the loss rates via both the passive and active methods described above, for each 5-minute interval, and provide a plot of these computed samples in a way that allows comparison with the utilization plots required in Section 6.1. If loss data is available in both directions, AT&T shall plot two separate time series corresponding to loss rate for the inbound and outbound directions.

6.4 Reporting distribution of latency

Section 5 defines *round-trip latency* as the time for a packet to travel from a source router to its destination (in this case, the interconnecting party's border router or otherwise specified device) and immediately back to the source router. This definition relates to a given packet. The FCC order requires AT&T to report a probability distribution function of these measured latencies.

29. For each LAG, AT&T shall report the probability distribution of latency, using the latencies recorded from the probing with IPSLA/TWAMP (if available) and ICMP separately during peak and off-peak periods as defined. (a) If using TWAMP, and it provides this capability, AT&T shall compute two distributions corresponding to latency in the two directions. (b) If the two directions cannot be separated, AT&T shall compute one distribution corresponding to the round-trip latency. In both cases AT&T shall report the number of samples used to compute the probability distribution function.

30. AT&T shall produce a probability distribution report that includes a plot of the resulting probability distribution, as well as the following statistics of the distribution: mean, standard deviation, 10th, 25th, 50th, 75th, 90th, 95th, 99th percentiles, and the maximum.

    If using TWAMP and the two directions can be distinguished:

    (i) Plot of probability distribution of outbound latency during a) peak times and b) off-peak times;

    (ii) Plot of probability distribution of inbound latency during a) peak times and b) off-peak times;

    (iii) Mean, standard deviation, 10th, 25th, 50th, 75th, 90th, 95th, 99th percentiles and maximum outbound latency during a) peak times and b) off-peak times; and

    (iv) Mean, standard deviation, 10th, 25th, 50th, 75th, 90th, 95th, 99th percentiles and maximum inbound latency during a) peak times and b) off-peak times.

    Otherwise, and for ICMP probing:

    (i) Plot of distribution of round-trip latency during a) peak times and b) off-peak times.

    (ii) mean, standard deviation, 10th, 25th, 50th, 75th, 90th, 95th, 99th percentiles and maximum round-trip latency during a) peak times and b) off-peak times.

    AT&T has proposed a specific approach to complying with the FCC Order's requirement to produce a probability distribution, which is set forth in Appendix A. The IME has approved AT&Ts proposed probability computation approach. If at any time during the reporting period this method becomes infeasible, AT&T will propose an alternative approach, which must be approved by the IME in writing.

6.5 Data Reporting Format and Procedures

31. AT&T shall submit monthly to the FCC, in formats to be approved in writing by the IME:

    (i) a PDF of the graphs and metrics required by Section 6;

    (ii) an Excel workbook of the calculated values;

    (iii) digital version of all reported data in a structured format (e.g. HDF5, RRD, JSON, XML, CSV),

    (iv) a digital version of the raw data;

    AT&T shall retain a digital version of all raw data for one year after the end of the reporting period. If AT&T discovers or is notified of an error in data or calculations, AT&T shall report the errors to the FCC, and submit corrected reports to the FCC.

32. For any data that AT&T is required to share with the interconnecting party, the format must be the same as that presented to the FCC.

Appendix A

FCC Interconnect Measurements

Computing Aggregate PDFs and Percentiles

Version 1.0

June 2, 2016

Authors:

Gomathi Ramachandran
Len Ciavattone

Table of Contents

Tables

1 Reporting

This report details the data collected in each 5-minute sample and the summary computations computed from it. This report does not cover graphical presentation of this data. Section

1.1 Report Format

The fields in each row of 5-minute data are shown below.

Field Number	Field for report…
1	Routername
2	IP Address of Interface (source)
3	Interface Name (e.g. GE1/0/0)
4	Time of first poll
5	Seconds between polls
6	Link status at first poll
7	Link Status at second poll
8	Reported Speed at first Poll
9	Reported Speed at second Poll
10	Common Name (Link Description)
10a	Target IP address for active measurement
10b	ASN
10c	Cust/peer
11	In Discards
12	In Errors
13	Out Discards
14	Out Errors
15	In Octets
16	In Unicast Packets
17	Out Octets
18	Out Unicast Packets
19	In Utilization
20	Out utilization
21	Packets sent (active measurement)
22	Lost packets
23	Loss percentage
24	Mean Latency
25	Standard Deviation
26	Current Hourly Max
27	Current Hourly Min
28	10th %ile
29	25th %ile
30	50th %ile
31	75th %ile
32	90th %ile
33	95th %ile

34	99th %ile
35	Count for 0<5ms
36	Count for 5<10ms
37	Count for 10<15ms
38	Count for 15<20ms
39	Count for 20<25ms
40	Count for 25<30ms
41	Count for 30<35ms
42	Count for 35<40ms
43	Count for 40<45ms
44	Count for 45<50ms
45	Count for 50<55ms
46	Count for 55<60ms
47	Count for 60<65ms
48	Count for 65<70ms
49	Count for 70<75ms
50	Count for 75<80ms
51	Count for 80<85ms
52	Count for 85<90ms
53	Count for 90<95ms
54	Count for 95ms and greater
55	Sum of RTT for 0<5ms
56	Sum of RTT for 5<10ms
57	Sum of RTT for 10<15ms
58	Sum of RTT for 15<20ms
59	Sum of RTT for 20<25ms
60	Sum of RTT for 25<30ms
61	Sum of RTT for 30<35ms
62	Sum of RTT for 35<40ms
63	Sum of RTT for 40<45ms
64	Sum of RTT for 45<50ms
65	Sum of RTT for 50<55ms
66	Sum of RTT for 55<60ms
67	Sum of RTT for 60<65ms
68	Sum of RTT for 65<70ms
69	Sum of RTT for 70<75ms
70	Sum of RTT for 75<80ms
71	Sum of RTT for 80<85ms
72	Sum of RTT for 85<90ms
73	Sum of RTT for 90<95ms
74	Sum of RTT for 95ms and greater

Table 1-1 Sample Report Field Format

1.2 Computing Monthly Peak and Off-Peak PDF

Let the peak period be designated as all time periods where utilization is larger than the designated threshold of X% in either the incoming or outgoing direction.

Let r denote any individual record, and let R_{peak} denote the set of records corresponding to the peak period. Let c_i^r denote the count of RTT samples in bin i for a record r and s_i^r denote the sum of the RTT samples in bin i for record r. Bin i corresponds to one of 20 intervals between 0 and 95 milliseconds and 95 milliseconds to the maximum value e.g. $\{(0,5], (5,10], \dots (90,95], (95, max)\}$

1. Compute the total count in each i over all records in R_{peak}:

$$c_i^{agg} = \sum_{r \in R_{peak}} c_i^r$$

2. Similarly compute the total of the RTT sums in each bin i over all records in R_{peak}:

$$s_i^{agg} = \sum_{r \in R_{peak}} s_i^r$$

3. Compute the total number of RTT samples over all bins of the aggregate as:

$$c_{total} = \sum_{i=1}^{n} c_i^{agg}$$

The values of c_i^{agg}, s_i^{agg}, and c_{total} can be used to approximate a probability mass function as follows:

4. For each bin i obtain the average RTT in that bin as:

$$a_i^{agg} = \frac{s_i^{agg}}{c_i^{agg}}$$

5. The probability mass at a_i^{agg} is then:

$$PMF(a_i^{agg}) = \frac{c_i^{agg}}{c_{total}}$$

1.3 Computing Percentiles of the Combined PDF

To compute percentiles of the combined probability mass function:

1. For each fractional percentile $f \in \{0.1, 0.25, 0.50, 0.75, 0.90, 0.95, 0.99\}$ compute the sample which corresponds to that percentile. (Round to the nearest integer.)

$$c_f = [c_{total} \times f]$$

2. Find the bin f that corresponds to count c_f. This will be the first bin where the cumulative sum of counts in f and all bins with lower exceeds c_f:

$$\sum_{j=1}^{f-1} c_j^{agg} < c_f \leq \sum_{j=1}^{f} c_j^{agg}$$

3. Report the percentile as the average RTT a_f^{agg} as computed previously.

1.4 Data Collection Errors

Testing has found cases where data collection errors have occurred where count data is properly incremented but the corresponding RTT data is not accumulated. To compensate for this error instead of computing a_i^{agg}, the average of samples for the given interval (e.g. 15ms - 20 ms) the mid point (e.g. 17.5 ms) of the interval can be used for a_i^{agg}. In the last bin (95msec to max), where there is no natural mid point to select, the maximum value in the hour will be employed.