Metrics and Methods for IP Capacity

draft-ietf-ippm-capacity-metric-method-~~02~~ > 07

A. Morton, R. Geib, L. Ciavattone

Where we're headed: Next Steps

- DISCUSS any remaining issues today!
- Trigger any concluding reviews on new material with a IPPM WG 2-week review!!
- Reach ~~Consensus~~ Approval very soon!!!

Receiver Rate Measurement

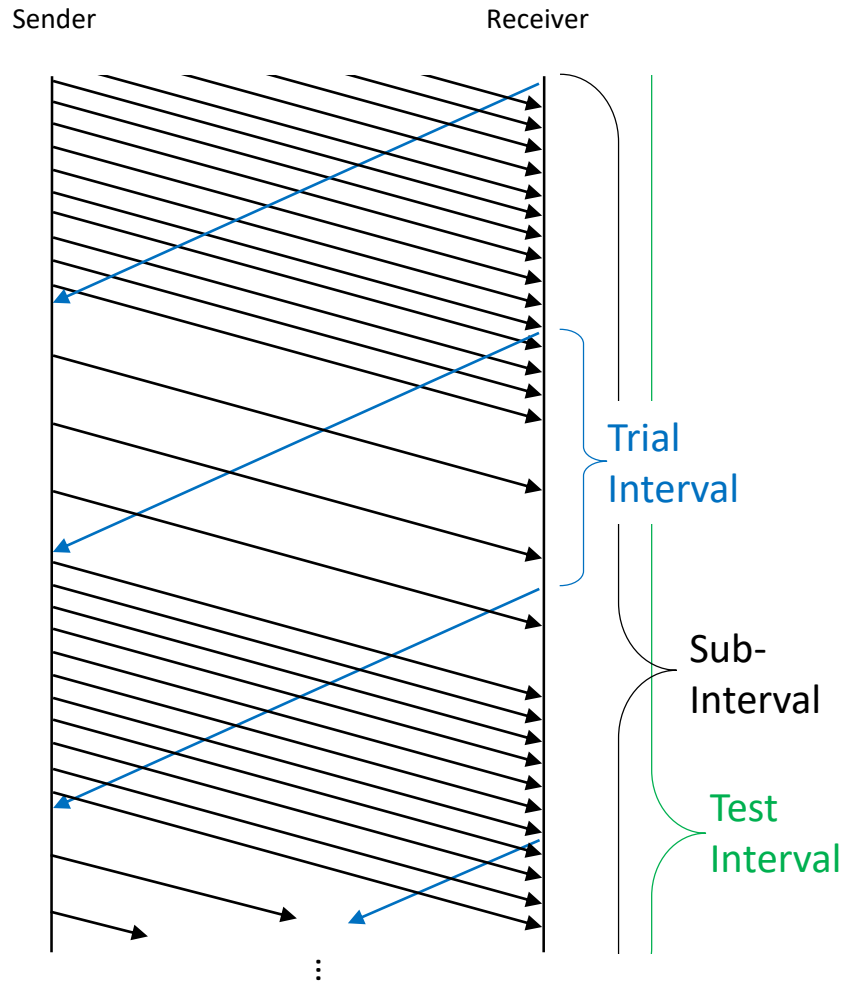

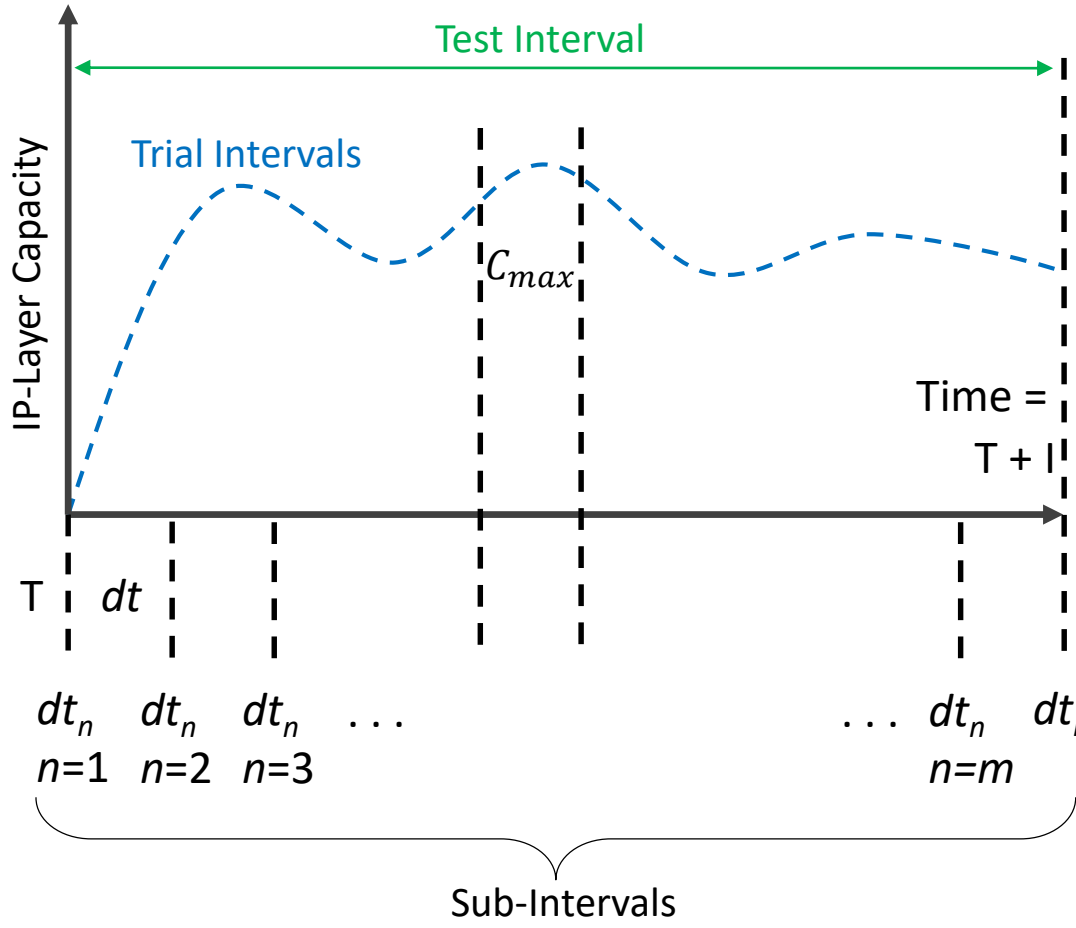

IPPM Draft Status

- Key Updates 02 -> 07: WGLC, Martin/AD, IESG:
  - New Applicability sub-section, for metric and for the method with its Load-Adjustment algorithm:
    - Reference to RFC7497 = “access measurement”
    - New Requirements for Load Adjustment/Search alg.
      - MUST only be used in the application of diagnostic and operations measurements as described in this memo
      - MUST only be used consistent with Sec 10, Security Consid.
  - Much more detail on Method & Load Rate Adj. alg.
    - Table of 18 Params: Defaults, Tested range, Expect. Safe range
    - New Params: Feedback Message Timer and Disconnect TO
    - Pseudocode and even more Parameters in Appendix
  - Metric Definition Requirements, 5.3 and 6.3:
    - the number of sub-intervals with duration dt **MUST** be set to a natural number m, …
  - Running Code Section – Release 7.1 March 5

Key Parameters (1)

- Don’t Keep sending when connectivity is lost!

Parameter	Default	Tested Range or values	Expected Safe Range (not entirely tested, other values NOT RECOMMENDED)
FT, feedback time interval	50ms	20ms, 100ms	5ms <= FT <= 250ms Larger values may slow the rate increase and fail to find the max
Feedback message timeout (stop test)	L*FT, L=10 (500ms)	L=100 with FT=50ms (5sec)	0.5sec <= L*FT <= 30sec Upper limit for very unreliable test paths only
load packet timeout (stop test)	1sec	5sec	0.250sec - 30sec Upper limit for very unreliable test paths only

Knowledge of Protocol Helped Answer Safety Questions

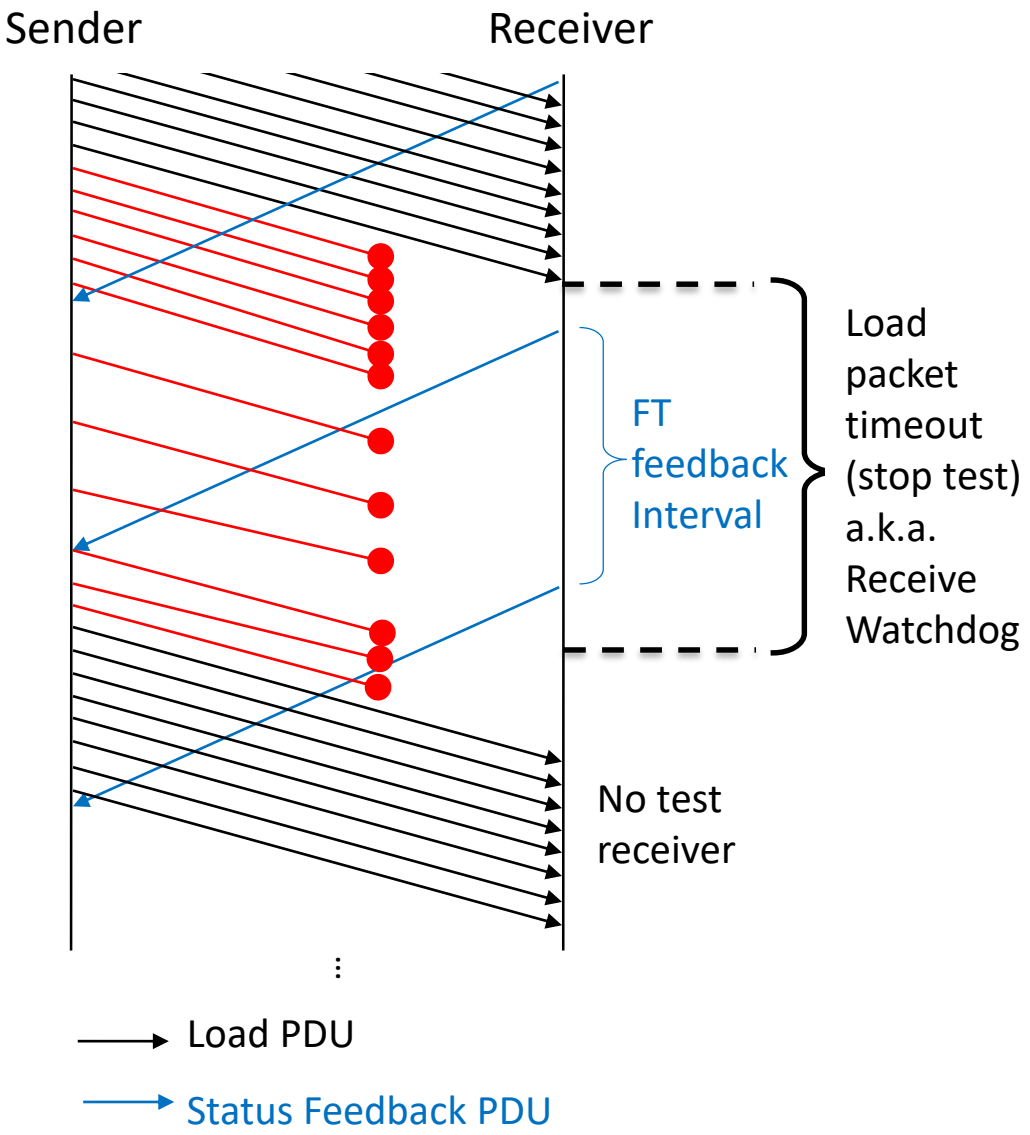

Operation: The load packet timeout SHALL be reset to the configured value each time a load packet received. If the timeout expires, the receiver SHALL be closed and no further feedback sent.

Knowledge of Protocol Helped Answer Safety Questions

Operation: The feedback message timeout SHALL be reset to the configured value each time a feedback message is received. If the timeout expires, the sender SHALL be closed and no further Load packets sent.

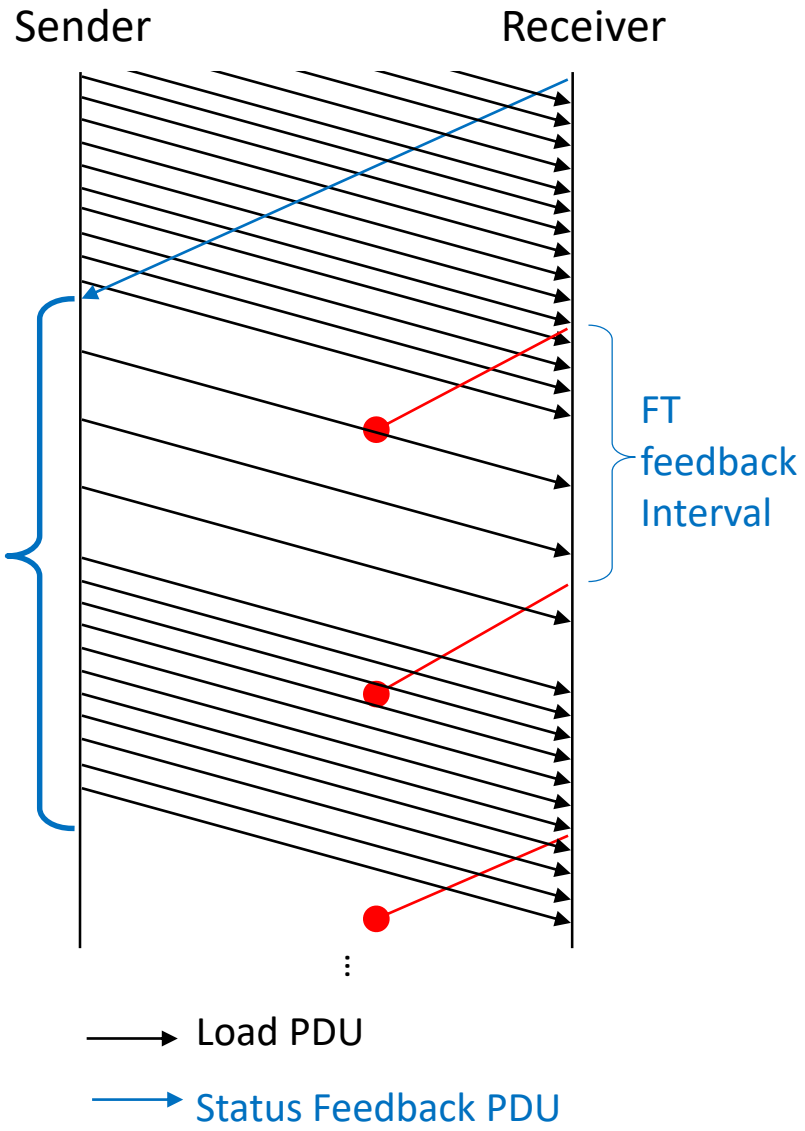

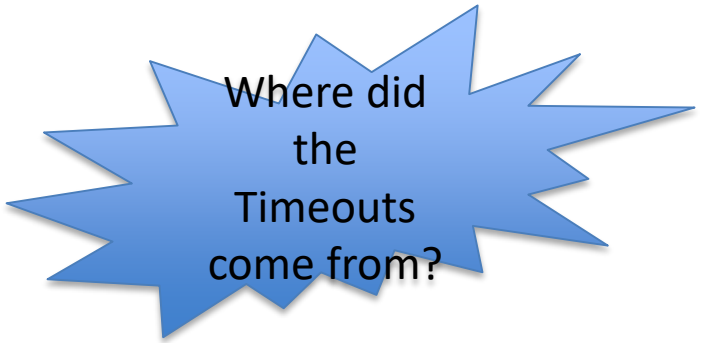

Brief Interlude on “Test Protocol for One-way IP Capacity Measurement”

draft-morton-ippm-capacity-metric-protocol-00

L. Ciavattone, A. Morton

Protocol: Setup and Activate Test

draft-morton-ippm-capacity-metric-protocol-00

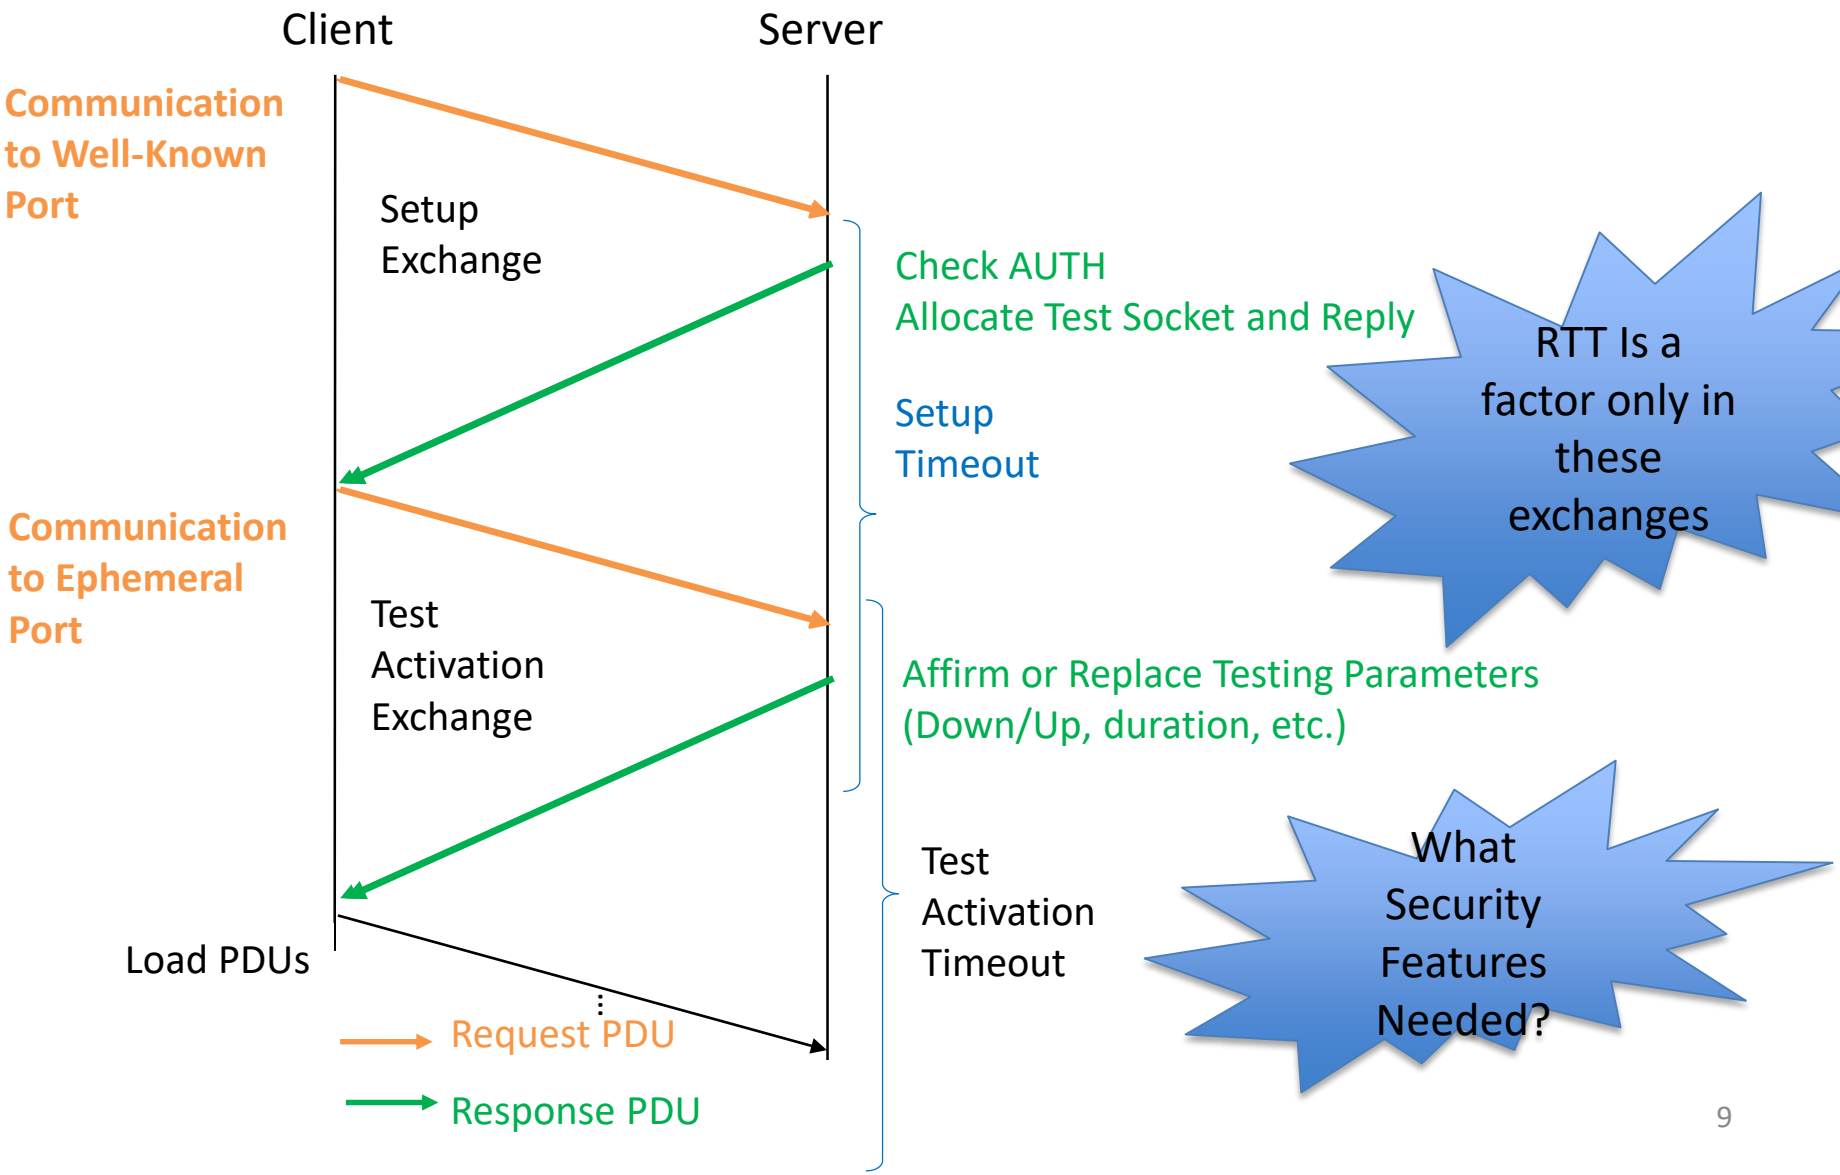

Key Parameters (2)

- Step Sizes in the Table of Sending Rates

Parameter	Default	Tested Range or values	Expected Safe Range (not entirely tested, other values NOT RECOMMENDED)
table index 0	0.5Mbps	0.5Mbps	when testing <=10Gbps
table index 1	1Mbps	1Mbps	when testing <=10Gbps
table index (step) size	1Mbps	1Mbps - 1Gbps	same as tested
table index (step) size, rate>1Gbps	100Mbps	1Gbps - 10Gbps	same as tested
table index (step) size, rate>10Gbps	1Gbps	untested	>10Gbps

Key Parameters (3)

- Load-Rate Alg: Seq. Errors, Delay Range Thresh

Parameter	Default	Tested Range or values	Expected Safe Range (not entirely tested, other values NOT RECOMMENDED)
ss, UDP payload size, bytes	none	<=1222	Recommend max at largest value that avoids fragmentation

>
> So isn't there a mismatch between the metric and the load algorithm values
> here? With the rate definition in Section 8.1 being defined as based on
> "ss" that UDP payload bytes, rather than IP packet sizes that are used?
[acm]
Not really, UDP is mandatory in the metric definition.
>
> I understand that one want to ensure that one measure using a size that
> actually works in the path. However, I think one should be warned that one
> might run into packet rate limitations rather than byte limits if one
> would use too small.
[acm]
Ok
"Use of too-small payload size might result in unexpected sender limitations."

Key Parameters (4)

- Load-Rate Alg: Seq. Errors, Delay Range Thresh

Parameter	Default	Tested Range or values	Expected Safe Range (not entirely tested, other values NOT RECOMMENDED)
low delay range threshold	30ms	5ms, 30ms	same as tested
high delay range threshold	90ms	10ms, 90ms	same as tested
sequence error threshold	0	0, 100	same as tested
consecutive errored status	2	2	Use values >1 to avoid misinterpreting transient loss

Standards High-Level Status: IP-Layer Capacity Metric and Meas.

- ITU-T Study Group 12 - **Approved**
  - Question 17 on Packet Network Performance the Metric and Method of Measurement to Rec. **Y.1540 - 2019 (Annexes A and B)**
    - Considerable background (test results; research) in Appendices X thru XIII
    - **Approved New Supplement 60 on Interpreting IP-Layer Capacity Results**
- ETSI TC Speech and Multimedia Transmission Quality (STQ)
  - **Approved** the Metric in **TS 103 222 Part 2** on High Speed Internet KPIs
  - Reference to Rec Y.1540 for all other material
- Broadband Forum (BBF) – **Fully Approved**: **TR-471**
  - Standardize the Metric and Methods with details on Measurement Points and Info Model: control & reporting. (**Issue 2** updates in 2021)
- IETF IP Performance Measurements (IPPM) Working Group
  - **Internet Draft** Adopted by WG, Now seeking IESG Approval
- **ETSI STQ MOBILE – Approved TR 103 702,** 5G Performance Measurements and QoS
- **OpenBroadband udpst project: GitHub Mirror 3rd Rel. Mar 5**

Next Steps

- Reach Approval very soon, so we can take-up protocol support draft(s)!
- Trigger any concluding IPPM WG comments with 2-week review.
- Our team's approach ==>>

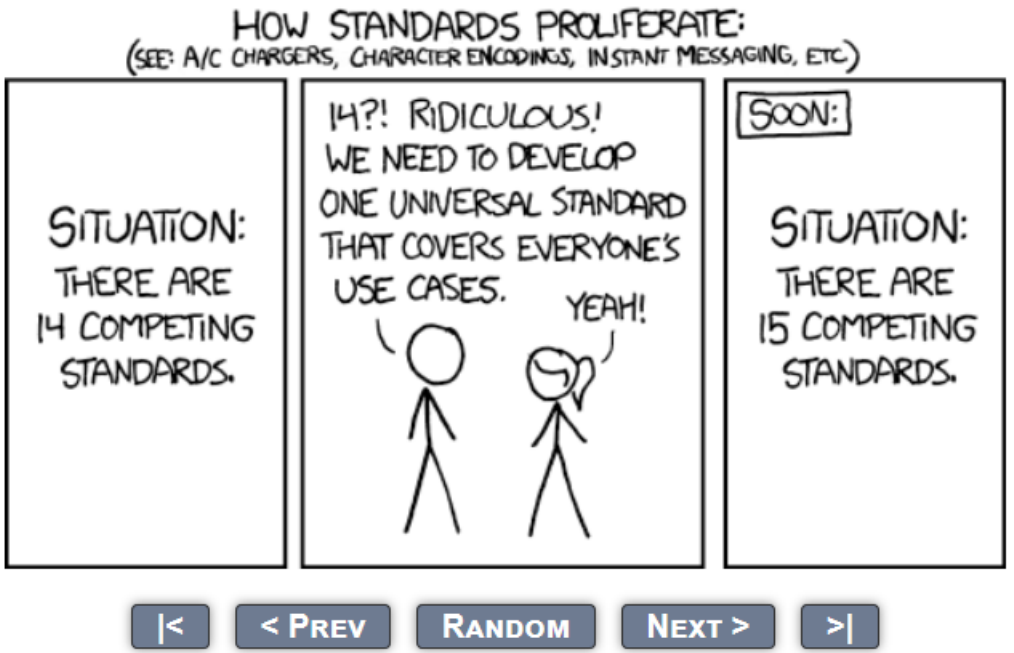

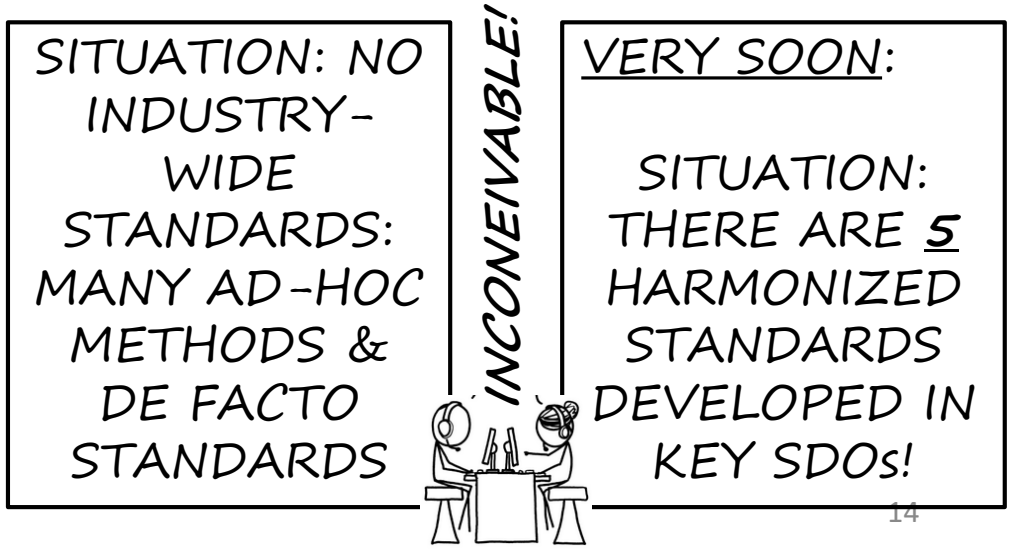

BACKUP

REFERENCES

- Hackfest 106 Slides: Test Results
- Hackfest 105 Slides: Test Results
- Liaisons from ITU-T SG 12 and ETSI TC STQ – see email for links, or
- https://datatracker.ietf.org/liaison/1645/
- https://datatracker.ietf.org/liaison/1643/
- https://datatracker.ietf.org/liaison/1634/
- https://datatracker.ietf.org/liaison/1632/
- More Test results in the Liaison attachments

Knowledge of Protocol Helped

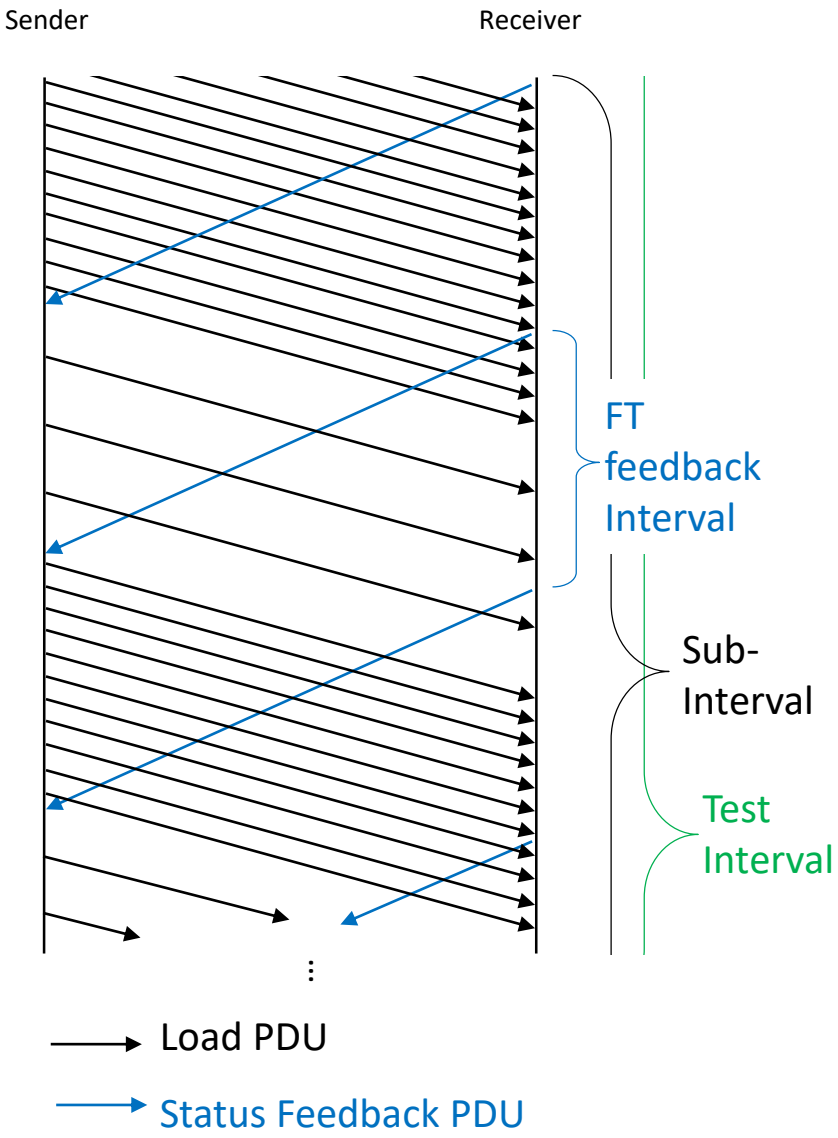

Key Parameters

- Other

Parameters

Parameter	Default	Tested Range or values	Expected Safe Range (not entirely tested, other values NOT RECOMMENDED)

Receiver Rate Measurement

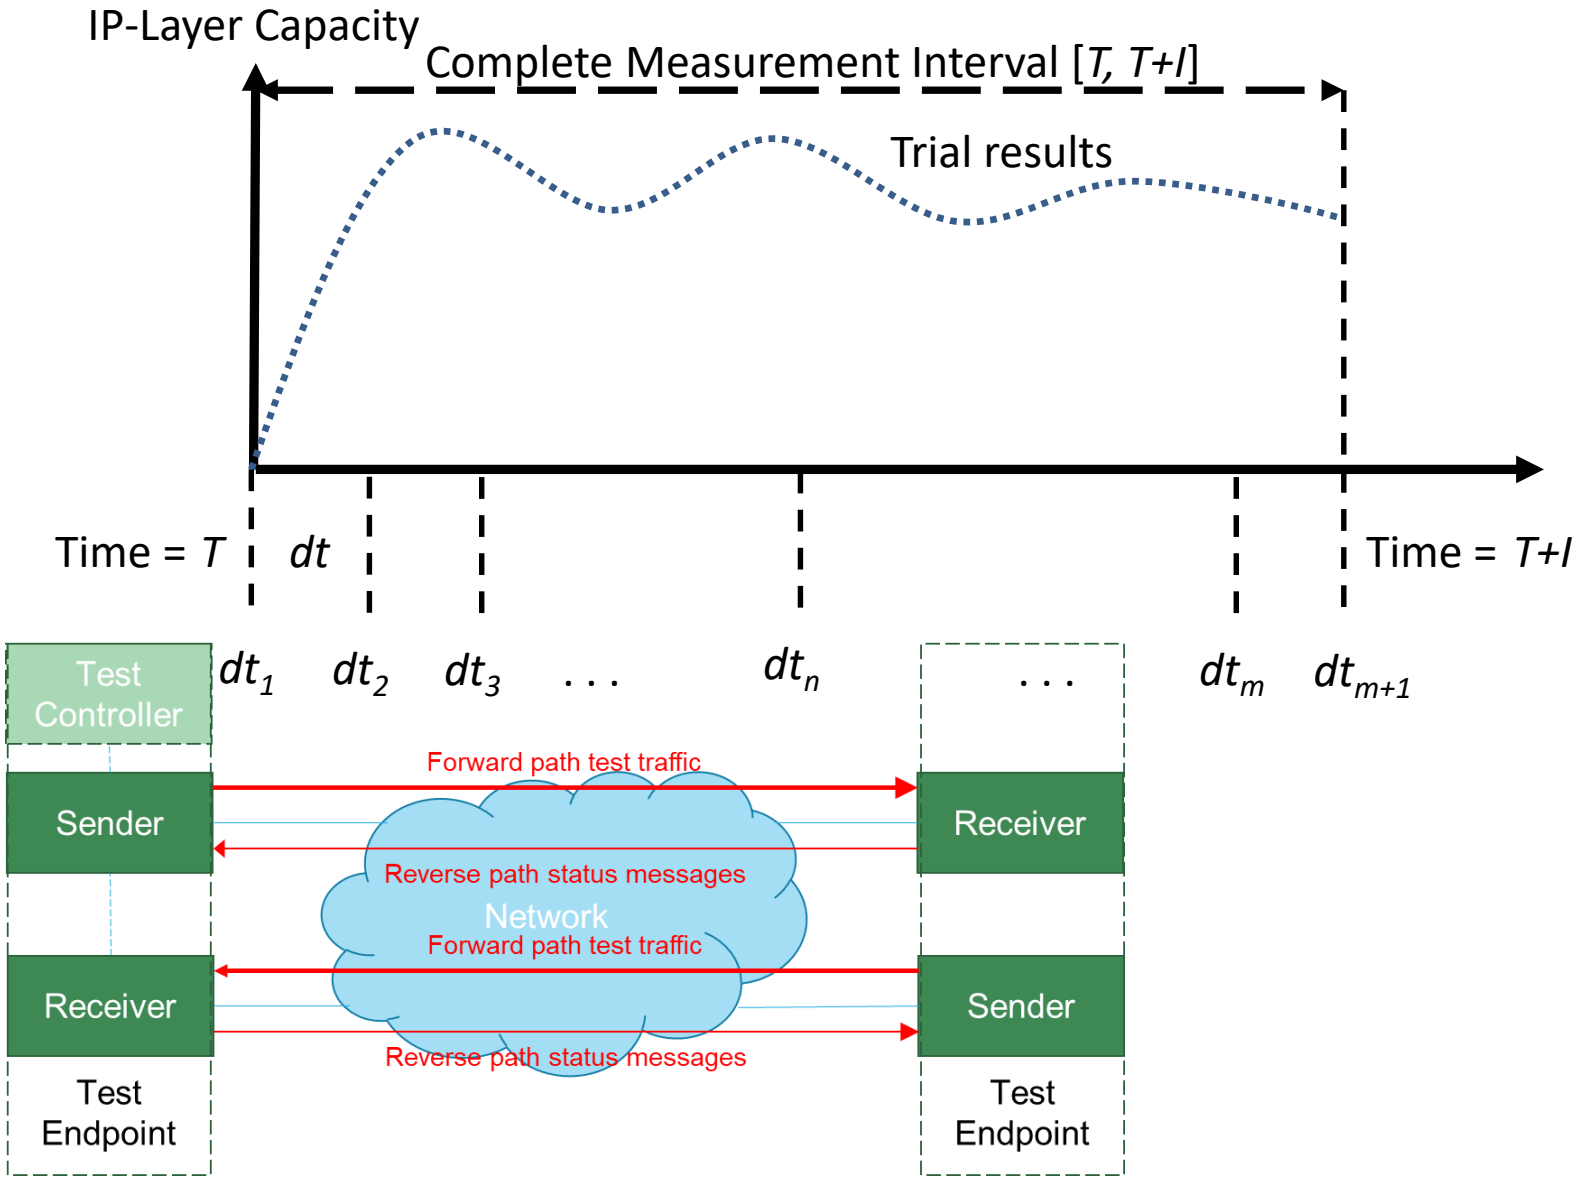

Define the Metric

- Maximum IP-Layer Capacity (includes headers + UDP payload)
- Follows IPPM Singleton, Sample, Statistic framework
- Def. in Words and an Equation (with variables explained)

$$\text{Maximum\_C}(T,I,\mathbf{PM}) = \frac{\max_{[T,T+I]} \left(n0[\mathbf{dtn},\mathbf{dtn+1}] \right)}{dt}$$

where:

T < ------ Measurement Interval ------> T+I

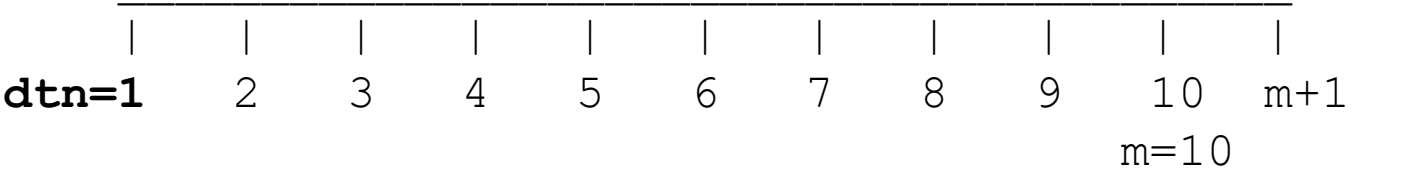

IPPM Draft Status

- Additional comments and reviews have resulted in a very complete draft.
  - Review in ETSI STQ-MOBILE, BBF, FCC WG on Gbps
  - Four New Members of ITU-T SG12 (testing co's)
- Key Updates 02:Measurement Considerations
  - access policies may limit the IP-Layer Capacity depending on the Type-P of packets
  - New References for Load Adjustment/Search
  - New Running Code Section – Release 1.0 ASAP!
- S 9: Info Model Config&Reporting: BBF TR-471

8.3 Meas. Considerations (new)

Conditions which might be encountered during measurement, **where packet losses may occur independently** from send rate:

1. Congestion of an interconnection or backbone interface may appear as packet losses distributed over time in the test stream, due to much higher rate interfaces in the backbone.
2. Packet loss due to use of Random Early Detection (RED) or other active queue management.
3. There may be only small delay variation independent of sending rate under these conditions, too. THIS IS A "TELL"
4. Persistent competing traffic on measurement paths that include shared media may cause random packet losses in the test stream.

It is possible to mitigate these conditions… but try locating measurement points as close as possible, first!

8.3 Meas. Considerations (new)

where packet losses occur independently from send rate:

Mitigate using parameters of search alg. described in Section 8.1 (tuning specific parameters, **more flexibility than typical CCA**).

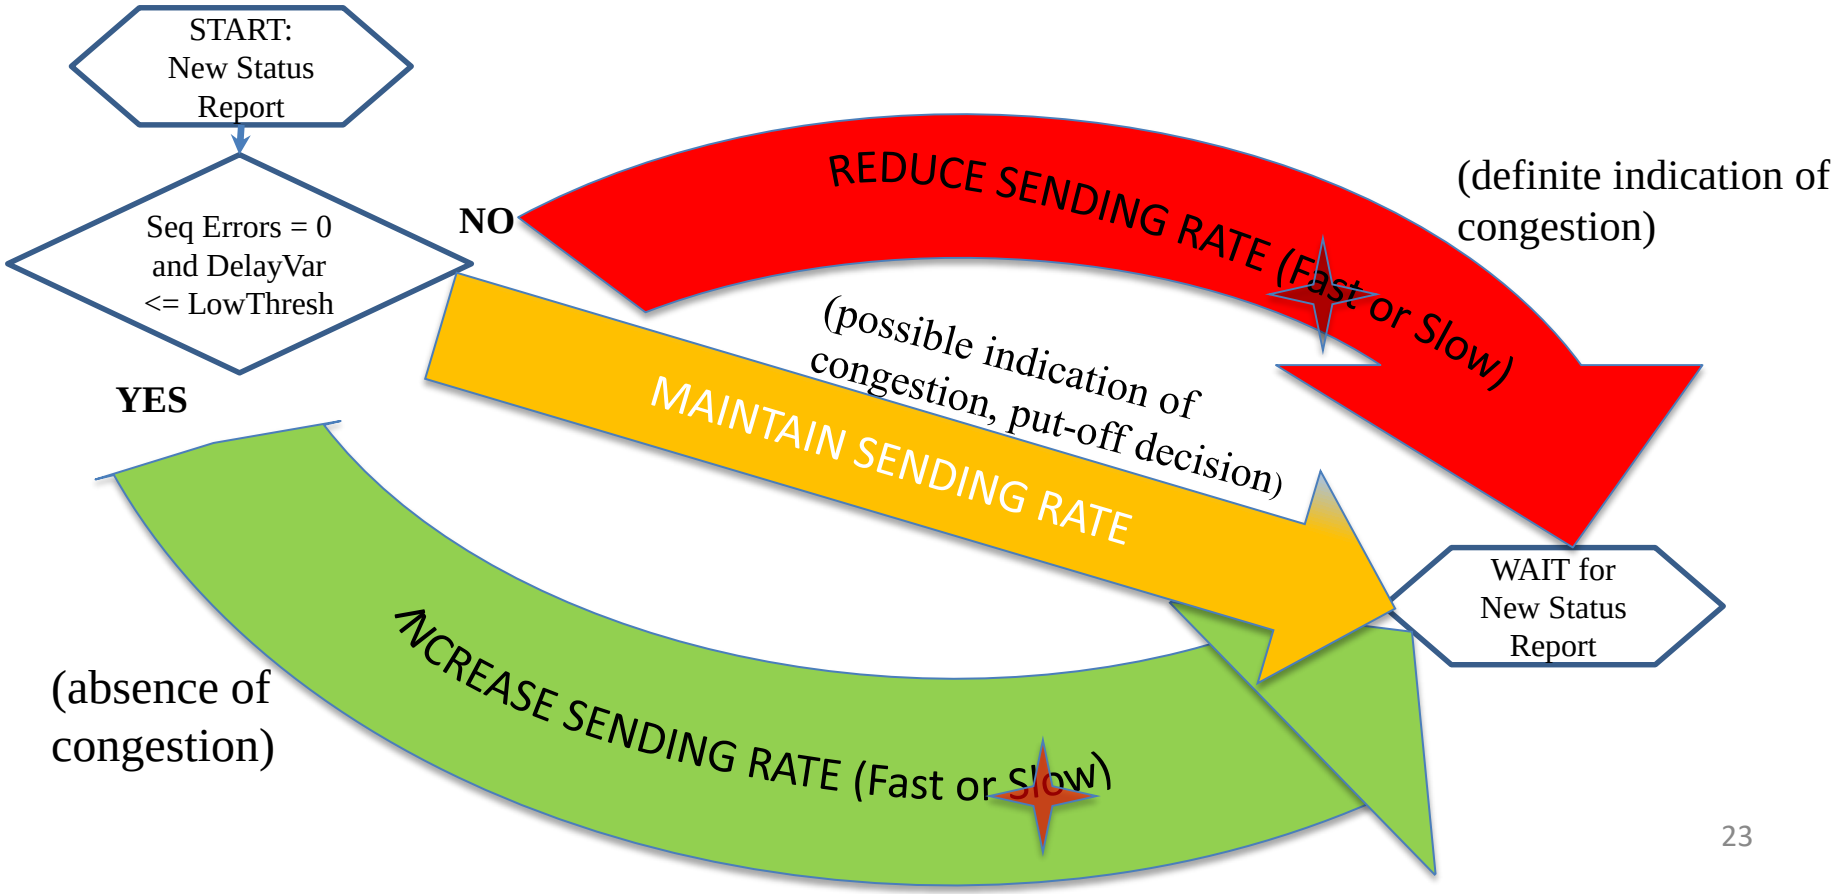

Results Reporting Considerations

- "Turbo-mode" concept (Matt Mathis' testing)
- Report separate results for repeatable modes

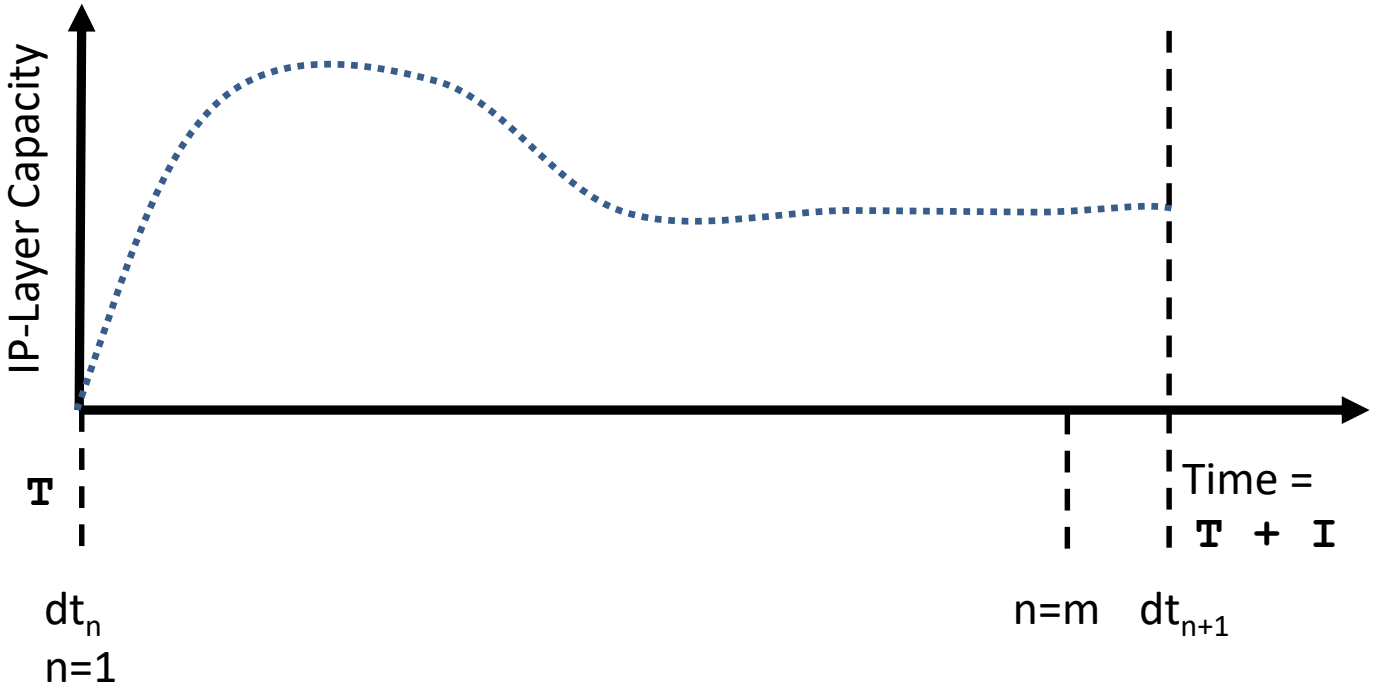

- Other modes may be encountered (repeatable?)
- Radio constellations, Cellular modes, weather

9. Reporting Format Elements (Others?)

The Singleton IP-Layer Capacity results SHOULD be accompanied by the context under which they were measured.

- timestamps

  (especially the time when the maximum was observed in dtn)
- source and destination (by IP or other meaningful ID)
- other inner parameters of the measurement (Section 4)
- outer parameters, such as “performed in motion" or other factors belonging to the context of the measurement
- result validity (indicating cases where the process was somehow interrupted or the attempt failed)
- a field where unusual circumstances could be documented
- a field for "ignore/mask out" purposes in further processing

Define the Method

- "PM" is short-hand for the performance constraints on the Load Rate Adjustment Alg.:

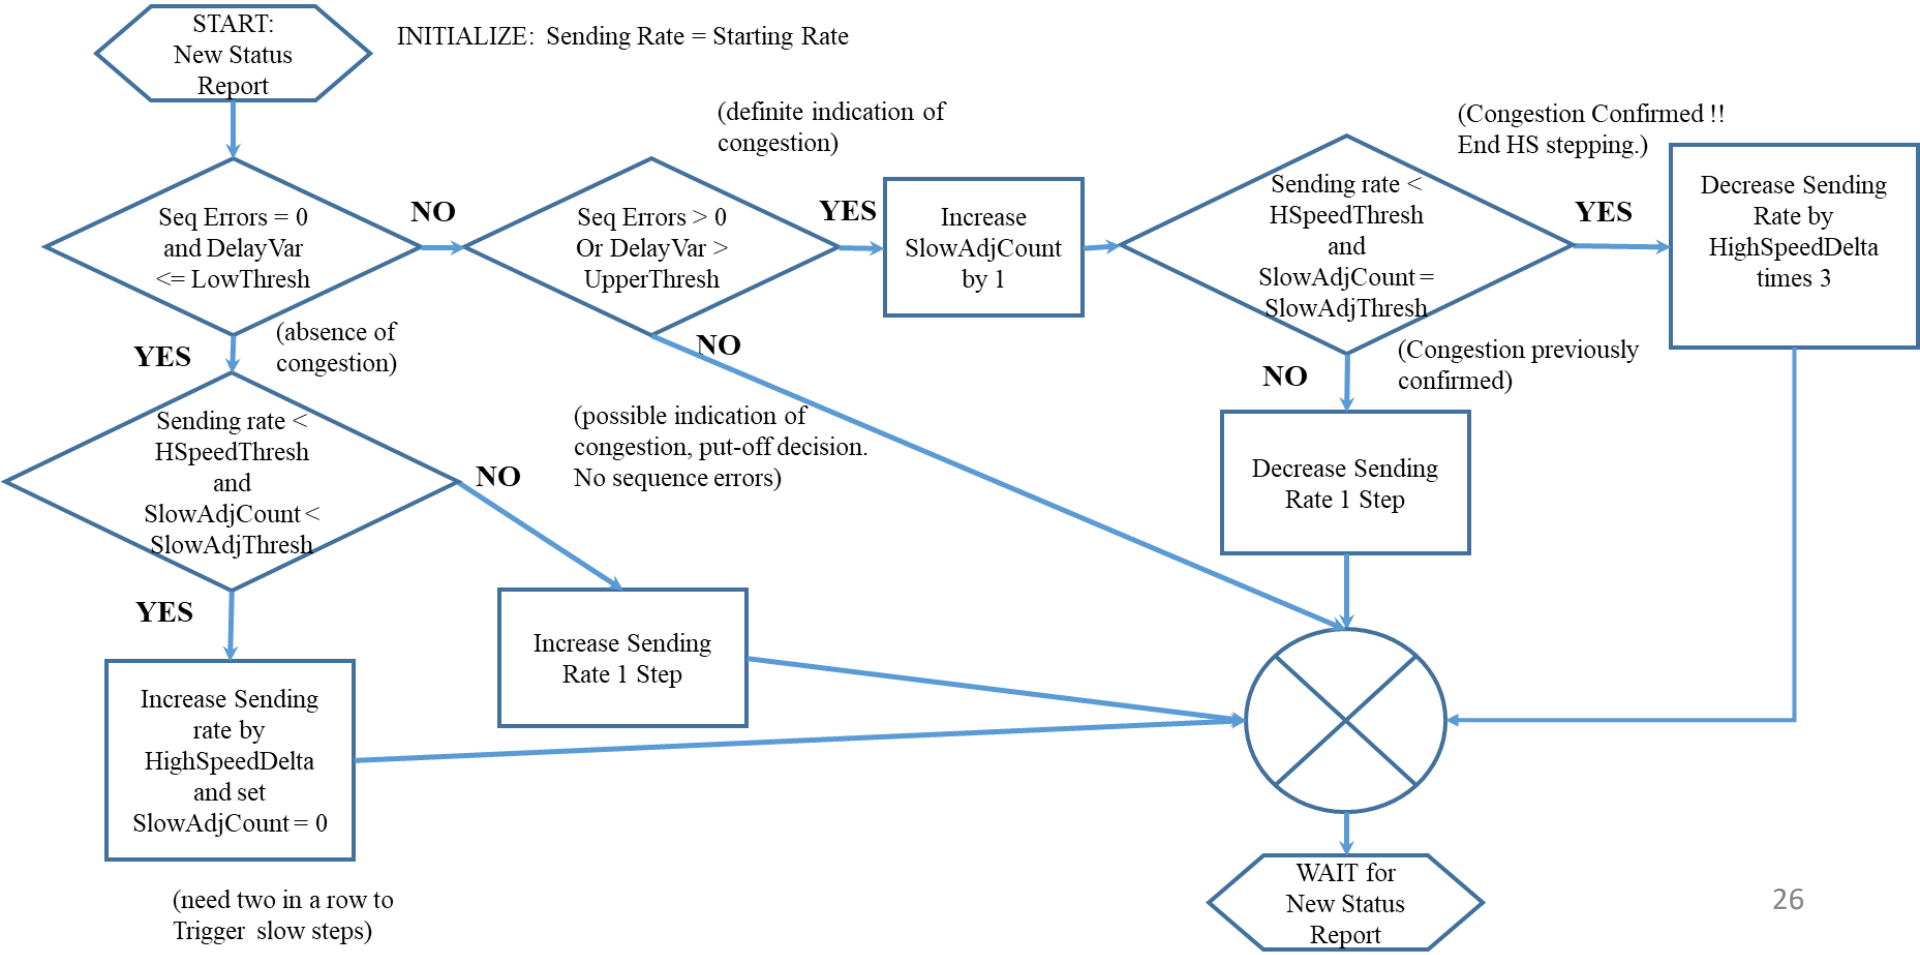

Recent Test Results

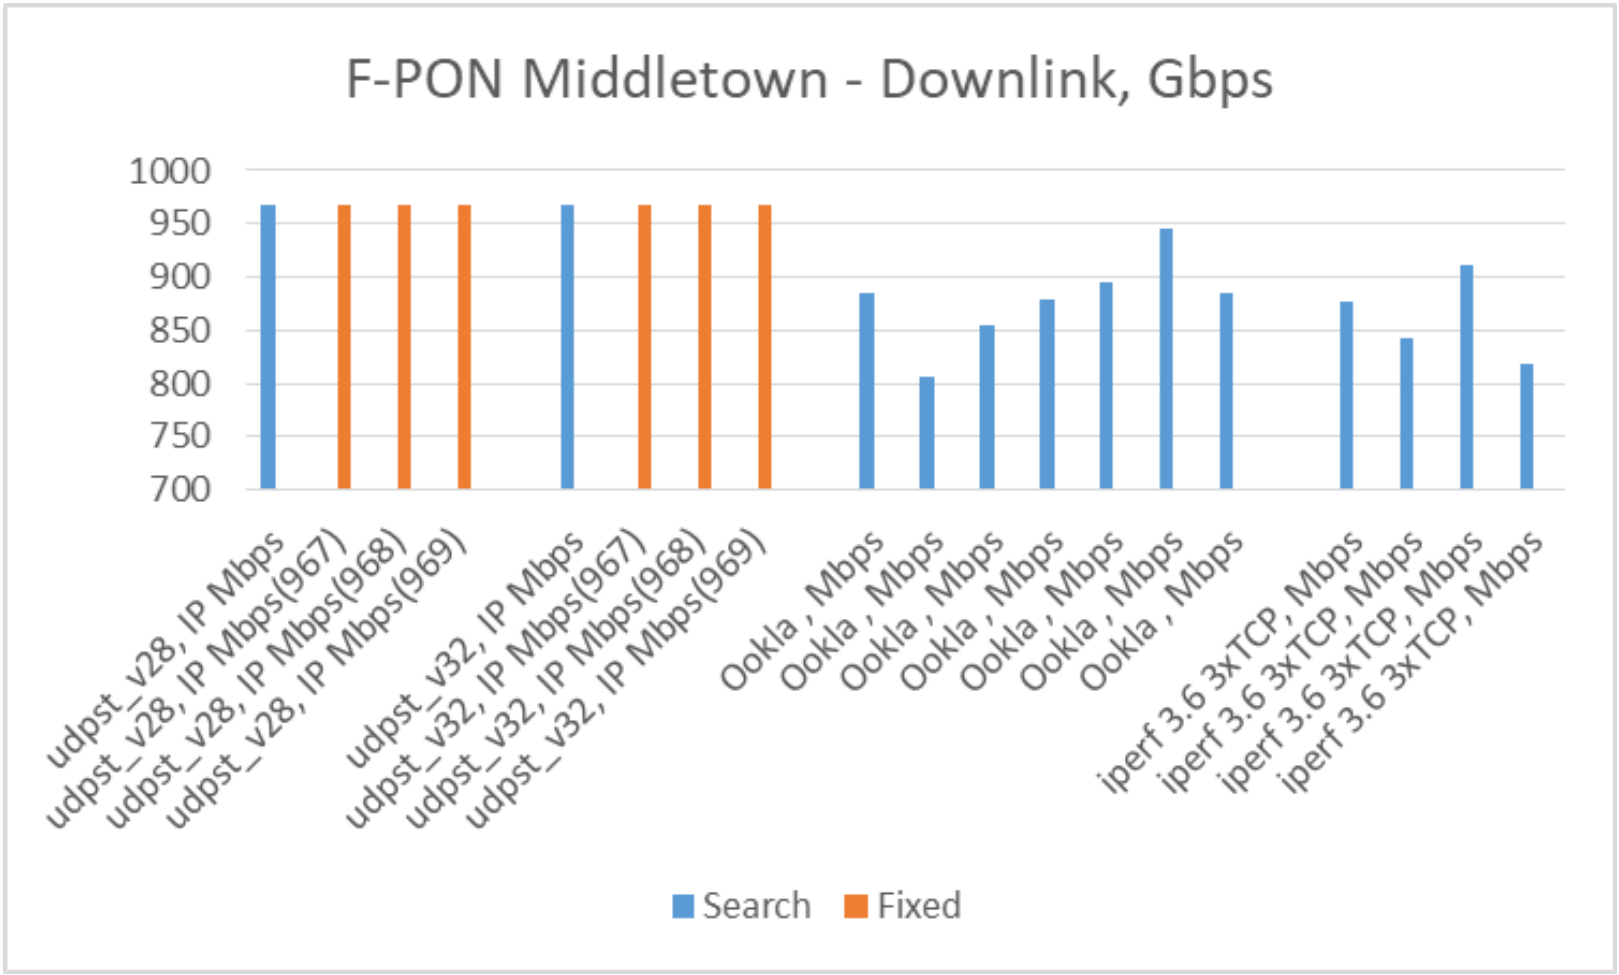

Udpst and Ookla Web Sockets Clients

Udpst and Ookla Web Sockets Servers

UDP-Speedtest

Middletown, NJ

IETF – Sept List Points raised (and addressed) -1

- @@@@ A clear take-away is that reporting must account for **bimodal** features, if/when measured.
- ++++ Covered in Section 6.6, reporting the Metric
- @@@@ Also, that wide-spread measurements will encounter wide-spread behaviors - testing should continue + expect some evolution.
- ++++ Covered in the Methods of Measurement Section
- @@@@ IMO, many of the above challenges fall on the measurement methodology: allow for traffic & time to initiate an on-demand access.
- @@@@ Also, results depend on the sending stream characteristics; we've known this for a long time, still need to keep it front of mind.
- ++++ both above covered in Methods of Measurement, Considerations.
- @@@@ Max IP-Layer Capacity and RFC 3148 BTC (goodput) **are different** metrics. Max IP-layer Capacity is like the theoretical goal for goodput.
- ++++ Section 1, Intro
- @@@@ This is a big one: when the path we measure is state-full based on many factors, the Parameter "Time of day" when a test starts is not enough info. We need to know the time from the beginning of a measured flow, and how the flow is constructed including how much traffic has already been sent on that flow, because state-change may be based on time or bytes sent or both. See RFC 7312.
- ++++ included in Measurement Considerations

IETF – Sept List Points raised (and addressed) -2

- @@@@ The **Singleton and Statistic** formulations of IPPM's framework RFC 2330 are still valuable in this context, possibly combined with results criteria ("stable" for X singletons, non-arbitrary threshold needed to define "stable").
- ++++ The Singleton, Sample and Statistic for IP Capacity are implemented.
- ---- "stable" needs more discussion, or may be resolved by Qualification below.
- @@@@ Measurements depend on the access network and the use case. Here, the use case is to assess the maximum capacity of the access network, with specific performance criteria used in the measurement.
- ++++ Covered in the Intro.
- @@@@ Goals made clearer in the next draft, if possible.
- ++++ Covered in the Intro.
- @@@@ A qualification measurement for the search result is a subsequent measurement, sending at a **fixed 99.x %** of the Max IP-layer Capacity for I, or an indefinite period. The same Max Capacity Metric is applied, and the Qualification for the result is a sample without packet loss or a growing minimum delay trend in subsequent singletons (or each dt of the measurement interval, I). Samples exhibiting losses or increasing queue occupation require a repeated search and/or test at reduced fixed sender rate for qualification.
- Here, as with any Active Capacity test, the test duration must be kept short. 10 second tests for each direction of transmission are common today. In combination with a fast search method and user-network coordination, the concerns raised in [RFC 6815] are alleviated.
- +++ covered in the method of measurement section, subsection on Measurement Qualification and Verification

https://tools.ietf.org/html/draft-morton-ippm-capacity-metric-method-00

October List Discussion: Matt, Rüdiger, acm (1)

- Summary: Matt is saying (? Subject to confirmation)
  - @@@@@ RTT is a good singleton measurement interval (dt) to avoid "bursts & silence"
  - Use windowed Max of max_rate from BBR (but see our measurements)
- Rüdiger: "Len and acm meas. results show convergence to an LTE receiver bandwidth meas. with limited queuing and no drops."
  - Defaults of dt = 1 second, Δt = 10 sec
  - udpst tool sends feedback measurement at regular intervals = 50 ms
- acm thinking: sub-second rate meas. are more susceptible to the cases described by Matt, and by Joachim Fabini (time-slot service with full link capacity play-out of the queue: LTE, others).
- acm: But no assessment of loss with BBR, QUIC encrypt & aggregates

++++ We've added the defaults above with parameters when they appear, and more discussion in section 8.2

++++ Considerations for testing with parallel flows (sec 8).

++++ Default for the Sending rate measurement interval (sec 7, 0.05 s)

https://tools.ietf.org/html/draft-morton-ippm-capacity-metric-method-01

October List Discussion: Matt and Rüdiger

- It is fairly normal to see packets arrive in back to back packet trains, separated by periods of silence. Half- Duplex, Pkt Aggregation, …
- MM: simplistic meas. of LTE receive rates often see modes at 1Gb/s.
- BBRv2 uses rate measurement per RTT:
  - rtt_sample = delta(timestamp) # 1 RTT
  - rate_sample = delta(total data ACKed)/rtt_sample # one RTT's worth of data
- Effectively: Capacity(t, Δt, n, <no PM>) = n0[dtn-1,dtn]/(dt = RTTn)
- min_rtt and max_rate (used by BBR congestion control) are the windowed (?) max and min of rtt_sample and rate_sample above
- MM: I predict that max of BBR's max_rate will be a more robust and more accurate measure of the short duration maximum rate than anything you can do with UDP (except perhaps QUIC, BBR over UDP).