Model Based Metrics

IPPM Working Group
IETF 93, July 2015

Matt Mathis <mattmathis@google.com>
Al Morton <acmorton@att.com>

Outline

- Document status
  - Heavy revision after 1st WGLC
- Why Model Based Metrics are important
- Next steps

Document status

- WGLC Reviews:
  - Ruediger Geib
  - Mirja Kühlewind
- Clearly major problems
  - The big picture was misunderstood
  - Lots of feedback about inconsistent and non standard terminology

Document changes: -04 to -05 to -06

- Interim -05 draft submitted June 13th, -06 at Draft Cutoff
- Document restructuring
  - Split the introduction
  - New introduction
    - High Level view in 4 paragraphs
  - New Overview
    - One paragraph per concept preview of the entire document
    - New "system" diagram
  - Two minor subsections were reordered
- Major terminology overhaul
  - Aligned with other IPPM documents
  - Better self consistency
- See the document change log for more details
  - In the intro

High level view

- MBM is a framework
  - Maps predetermined transport (TCP) performance targets
  - Into a Targeted Diagnostics Suite of IP tests
- The Targeted Diagnostic Suite (TDS)
  - Pass fail/tests of IP performance
  - (Independent) tests of multiple packet delivery properties
    - Sufficient IP capacity (data rate)
    - Sufficient queue space to smooth and deliver bursts
    - Sufficiently low background packet loss ratio
    - etc
  - Failing **any** IP test means that some users **will fail** to attain the target
- This solves problems caused by TCP "equilibrium behavior"
  - Every detail affects every measured parameter
    - Even things that are explicitly out of scope, such as MP location
  - This is the unsolved problem in BTC Framework [RFC 3148]

Context

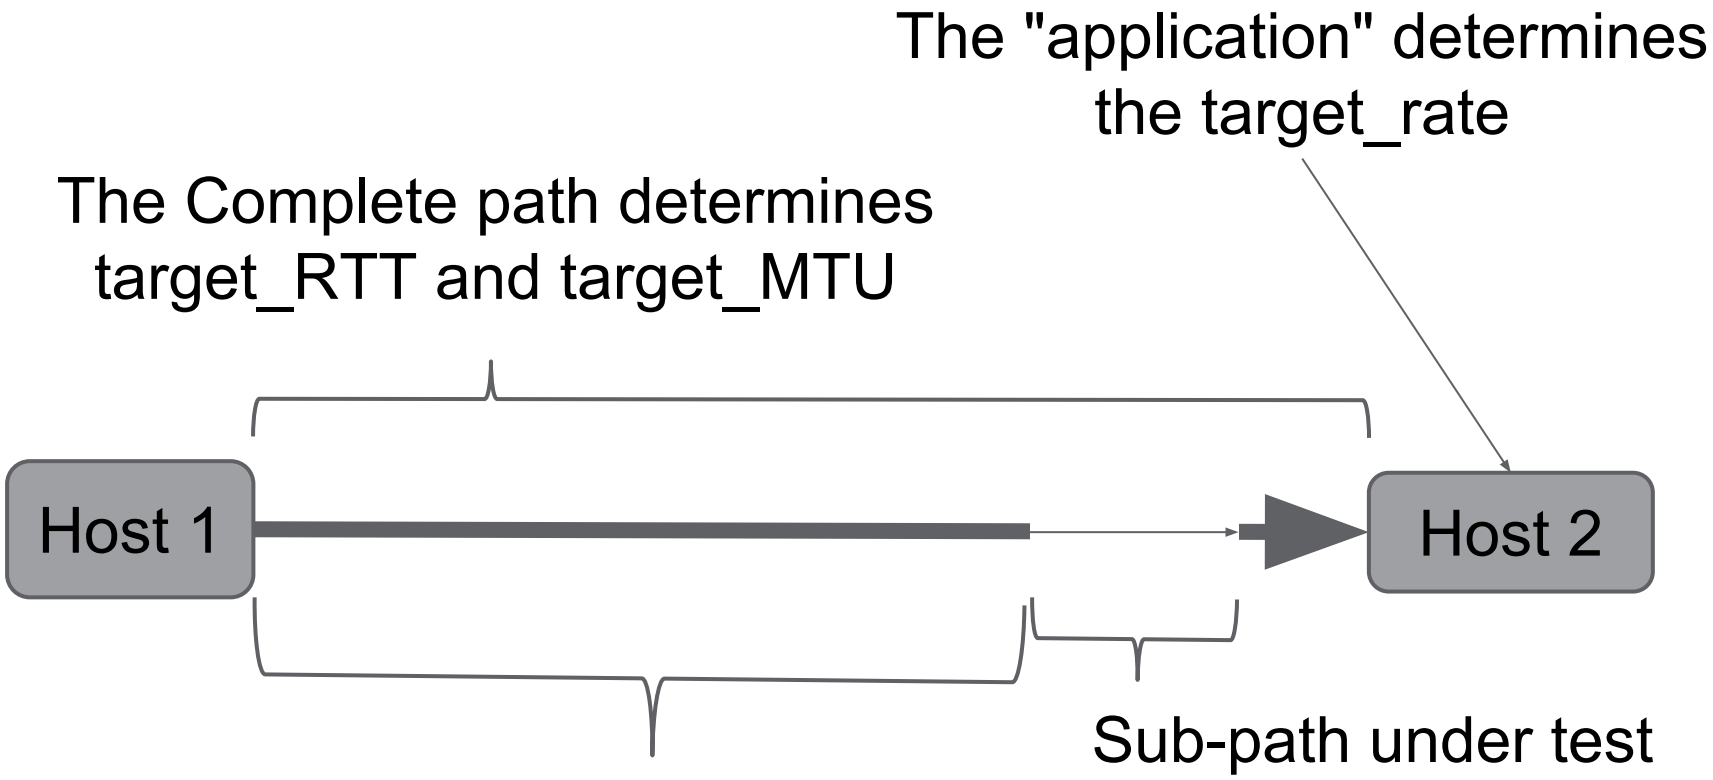

Each sub-path must pass all IP diagnostic tests of a Target Diagnostic Suite (TDS).

The Mode Based Metrics framework

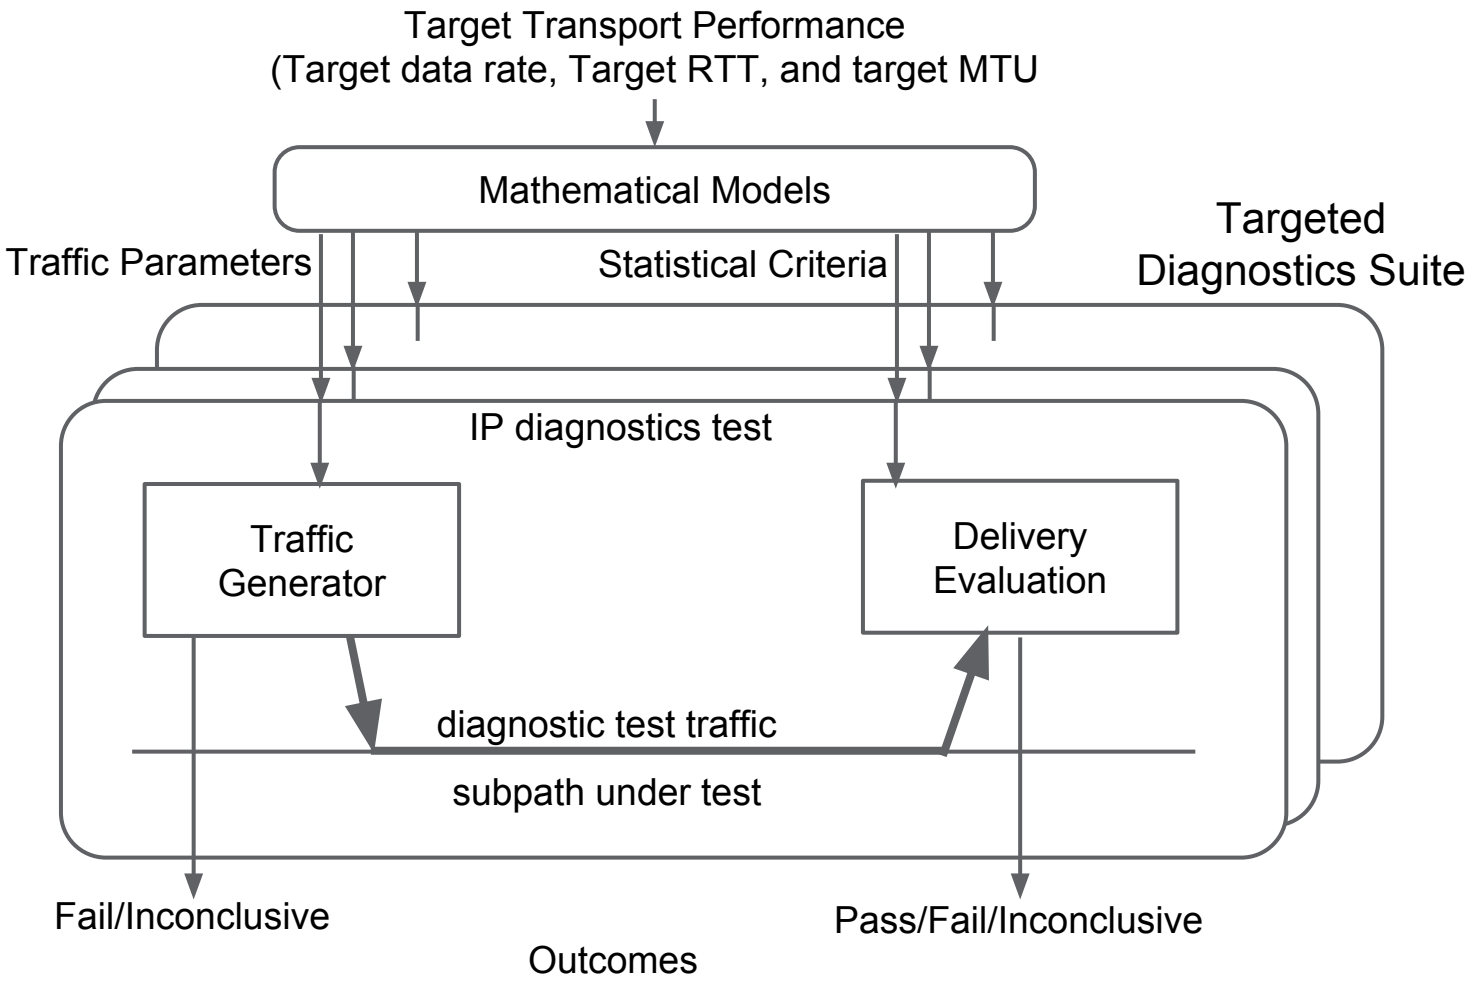

Elements of the Framework

- Target Transport Performance - what the user or application wants
  - Target data rate over the complete path
  - Target RTT and Target MTU are just as important
    - They determine how hard TCP and the network have to work
- Mathematical Models are used to calculate:
  - Traffic parameters (rates, burst sizes, etc)
  - Statistical criteria (bounds on packet loss ratio)
- Targeted Diagnostic Suite consisting of multiple:
  - IP diagnostic tests
    - Each measure one (or few) IP properties
    - Many based on existing IPPM metrics
    - With the addition of traffic controls and delivery evaluation

Building the individual IP diagnostic tests

- Traffic generation mimics TCP over a long path (bursts etc)
  - A longer Target RTT implies larger bursts
  - Subpath properties are prevented from affecting traffic patterns
  - May be built on top of existing IPPM metrics and tools
- Estimate and verify packet loss ratio
  - A longer target RTT requires a lower (better) packet loss ratio
  - Use Sequential Probability Ratio Test (SPRT)
    - Count delivered and lost packets
    - Stop when either hypothesis is confirmed or at a maximum count
- Outcomes
  - Pass or Fail
  - Inconclusive
    - Traffic generation was not accurate
    - Neither result is statistically confirmed
    - Something else interfered with the test

IP Properties Required to deliver Target TCP performance

- The IP capacity is above the Target Data Rate by sufficient margin
  - Capacity for all TCP/IP overhead, including rate hunting
- The observed packet loss ratio is low enough
  - Background losses caused by other cross traffic
- Sufficient buffering to absorb slowstart bursts
  - Full target_window_size at twice the bottleneck rate
- Sufficient buffering to absorb sender interface rate bursts
  - Partial target_window_size at full server interface rate
- Onset of packet loss has to be appropriate (Engineering)
  - This implies something AQM like
- Bound on how the data and ACKs interact (Engineering)
  - Channel arbitration must honor protocol self clocks

TCP Performance Guarantees

- If any subpath (link, device or interface, etc) fails any IP diagnostic test in a TDS, then some users will not be able to attain the target performance through that subpath.
- Implied goal: no failing tests for any subpaths
- There is the potential for corner cases (false results)
  - Validation procedure to help refine the TDS
  - The metrics are naturally slightly conservative
    - A fully passing subpath is likely to do better than the Target for some users

A few words about TCP & Standard Congestion control

- Keep pushing faster, until the network drops packets
- TCP and the network find a balance between
  - rate or window (determined by TCP)
  - loss and queuing delay (determined by the network)
- This is classic example of equilibrium behavior
  - It has loops in its dependency graph
  - Some action are non-linear
  - Therefor all parameters have non-linear sensitivity to everything
- Exported Parameters (measurements) have no predictive value
  - Thwarts "A-Frame" in RFC 2330
  - And Bulk Transport Capacity RFC 3148

Next steps

- WGLC, take 2