GNCC

MCS Test Process

MCS_X__BTJP_639

September 27, 2018

Introduction

This report is to promote discussion on MCS test model for GNCC implementation. The purpose of the test model is to address the core measurement requirements that were presented during the MCS onboarding course September 20-21, 2018.

The network measurements presented by GNCC for consideration during the MCS onboard training were as follows:

- Network Latency
- Network Jitter
- Network Bandwidth
- Network Packet loss
- Network Quality

The test definition is singularly the most important part of any network assessment. The validity and interpretation of all test results is solely dependent on the test settings and the test conduct. It is therefore important to ensure that the purpose and intent of each metric is well defined and accepted prior to constructing any test process. This must include a clear definition of the required filters for the pass, warning and failure value points.

The MCS solution provides the 5 metrics required as part of a much wider range of network assessment metrics. This document serves to provide a definition for each metric to aid discussion and agreement. This approach should allow GNCC to create the required test profiles to meet the provider and GNCC measurement objectives.

Network Latency

Latency is defined as a measure of delay (time) and therefore, by definition, latency is one of the most important measurements of connection stability and performance.

For a robust network assessment there are many different latency metrics to be considered. Each metric requires definition and acceptance.

TCP Latency (Data)

TCP Latency governs the user experience for all data related activities, the most common activity being browser applications. As latency is fundamental to performance, latency testing is an integral part of all MCS test types which includes the TCP bandwidth throughput and capacity test.

For data, TCP latency is a measurement that well defines a good network connection from bad network connection. Unlike bandwidth measurement, TCP latency defines the true performance of a connection because it is a measure of time as opposed to rate.

TCP is bidirectional and TCP latency is measured by the round-trip delay time (RTDT) in milliseconds for any two end points that establish a network connection. The minimum RTDT latency defines the equilibrium of a connection, therefore accuracy of the latency measurement is important. The RTDT result defines both the capacity limit and throughput of TCP data on a connection.

The RTDT latency metric in itself does not define an assessment of how a network connection will perform for the user experience, but it does define the limit to that performance. The more important metric derived from RTDT measurement is the variation of the RTDT result. It is the percentage change in RTDT that defines the latency jitter of a connection. High jitter percent (%) defines a bad connection.

To identify the impact of regulatory policy variations that artificially affect network traffic such as ramping and bursting, MCS integrates 3 independent trip measurement categories per test specification. These are:

1. Pre-socket binding, isolates policy changes, traps high latency TCP-ack spoofing common with satellite connections.
2. Post downstream, identifies provider downstream policy penalties
3. Post upstream test, identifies provider upstream policy penalties

MCS bandwidth test results report 4 metrics for RTDT latency:

- RTDT Min (ms)
- RTDT Max (ms)
- RTDT Avg (ms)
- RTDT Consistency (%, Jitter)

Additionally, MCS records the max delay latency of the stack which is spent waiting for data. This metric provides a vital indicator as to the quality and effectiveness of the connection being tested because it can be compared to the RTDT value. Large max delay values indicate packet integrity issues.

The formal declaration of the latency measure required by GNCC and what defines the latency pass/fail criteria for the latency being measured, requires confirmation.

UDP Latency (Jitter requirement?)

UDP latency governs the user experience for all streaming related activities, the most common application being video and voice traffic. Latency testing is an integral part of all MCS test processes which includes the UDP bandwidth utilization (available) and UDP quality (consistency). Both of these attributes are essential to UDP dependent applications.

UDP data protocol is specifically designed for packets that are dependent on the time of delivery. This means UDP data must operate to a well-defined flow pattern (codec) specific to the receiving application requirements (codec). UDP by definition is metered high priority data. Metered data defines the packet delay (time) of the packet-flow as it relates to each sequential packet. This is referred to as packet delay variation (PDV).

For UDP data, PDV is a measurement that well defines a good network connection from bad network connection. Unlike TCP, UDP is a unidirectional and loss tolerant. Meaning data only travels in one direction and if packets are lost the UDP protocol has no ability to recover (VPN encapsulation excluded).

PDV is naturally a well-defined jitter measure although the packet jitter is derived from five separate network attributes, namely packet delay (time), packet loss (time), packet out-of-order (time), packet-skip time and the end-to-end delay (time). All metrics are expressed in milliseconds.

There are many latency measures on networks, so it is important to clarify the jitter measure(s) required for unidirectional UDP packets.

There are two significant Jitter measures to consider for network assessments:

1. End-to-end delay (time)
2. Packet delay variation (PDV)

End-to-end delay only impacts the buffer size to accommodate for the data stream not the stream itself. For video the buffer has no perceivable impact except start time lag. However, voice is a bidirectional application and therefore the amount of buffer delay (common name jitter buffer) needed to accommodate the end-to-end latency introduces a lag time that is that sponsors over-talking.

The RTDT Jitter referenced and reported in the MCS bandwidth test, defines the acceptable delay. RTDT values greater than 100ms will become material to voice conversation and RTDT values over 150ms should be considered a bad connection.

Trip jitter, the popular use name for RTDT, is fundamental to understanding a bad data network from a good data network.

PDV, like RTDT, is a measure of time, however it is a measure of delay variance for UDP packets travelling in a uni-direction, not a bi-direction. PDV is reported in milliseconds as the jitter metric in the MCS VOIP test. The MCS VoIP jitter result therefore defines a clear metric that identifies a good connection versus a bad connection for UDP metered data. The variation of time being the definitive assessment of networks ability to deliver time dependent packets consistently.

The MCS VoIP test supports most popular non-adaptive codecs such as G711, G728, G729 etc. In addition, the MCS VoIP test allows custom codecs to be defined for any codec-based application stream including video.

It is recommended that any test for UDP metered data processes for at least 10 seconds (default). Four Jitter measures are reported by the test:

- Downstream Jitter (ms) PDV
- Upstream jitter (ms) PDV
- Downstream Peak Jitter (ms) PDV
- Upstream Peak Jitter (ms) PDV

Peak Jitter is an important metric even though it is not a GNCC measurement requirement. Average Jitter of 1 or 2 milliseconds can hide Jitter peaks that can be an order of magnitude higher. If a number of peak events occur it can/will cause the jitter buffer to drain which dramatically affects audio or video quality. The MCS VoIP test also reports a Packet Discard metric (%) to isolate all packets where the jitter variation is high enough to cast a packet as too-late for use. The discard metric is reported based on the discard threshold time defined in the test. It is recommended that the discard threshold should be set to at least 50% of the jitter buffer size (expressed in milliseconds).

ICMP Delay Variation (Jitter)

ICMP delay variation is included here for completeness.

MCS provide a route assessment to measure the RTDT, using ICMP. The ICMP assessment process is designed to measure the latency of all the layer 3 hops that comprise a network connection between a client and server. The route assessment test provides 12 additional Jitter metrics as follows:

- Max, Min, Avg Trip Latency per layer 3 router upstream
- Max, Min, Avg Trip Latency per layer 3 router downstream
- Max, Min, Avg Trip Latency end-to-end router upstream
- Max, Min, Avg Trip Latency end-to-end downstream

The MCS Route analysis test identifies:

1. Source location (hop, IP, domain) of high latency threats
2. Source location (hop, IP, name) of high jitter threats
3. Out-of-context geolocation (hop) latency values
4. Dropped packets including firewall presence
5. Provider peering points
6. Ownership details if available
7. Route changes, both directions if available

Bidirectional latency assessment, minimum of 4 maximum of 7 per layer 3 hop recommended

Network Bandwidth

The connection between bandwidth and the user experience is not well defined, despite the prime recommendation for a bad experience is to increase bandwidth.

Bandwidth is a common metric that is seldom measured correctly. However, the user experience matters for a wide range of reasons, therefore the accurate measurement of bandwidth must be front and center for any assessment strategy. The enormous demand for more bandwidth is fueled directly from the belief that bandwidth is a speed and more speed will deliver a better user experience.

Unfortunately, bandwidth does not define a speed, only a rate and a capacity limit. It is the bandwidth rate/capacity relative to, the latency of a connection, available memory (sever/client), concurrency of use, regulatory policies and the transport medium that will ultimately dictate the user experience.

Measuring bandwidth correctly depends on the purpose for the measurement being conducted. Bandwidth assessment is not a simple test although it is seen as a simple test. There are 3 entirely different categories of bandwidth, these are:

- Bandwidth equilibrium
- Bandwidth capacity, usable
- Bandwidth capacity, available

MCS is able to measure bandwidth equilibrium (throughput) and bandwidth capacity, however, these measurements require different test methods. Knowing the purpose of the test is therefore fundamental to a valid and accurate test result.

Bandwidth Equilibrium (data throughput):
A throughput test is a bidirectional TCP data test where the payload assigned to the test is well defined. MCS conducts a bandwidth throughput test conformant to RFC6349. An RFC6349 throughput test delivers a true performance measurement because the test is assessed to the connections relative to the RTDT latency (time), i.e. the connections equilibrium is confirmed. Unlike a capacity test, a throughput test is not destructive but it can be depending on the bandwidth limit and the connection latency.

Bandwidth Capacity (usable):

A capacity bandwidth test (usable) is a bidirectional TCP test. MCS conducts a bandwidth capacity test conformant to RFC6349 as a multi-session test to assess bandwidth attainment to RTDT (latency) under a well-defined capacity load. TCP capacity defines a usable measure simply because TCP packet quality issues, that result from lack of bandwidth capacity, incur (heavy) time penalties over and above RTDT delays. Delay caused by quality issues are not easily separated from RTDT variations however it is important to note that these delays are not directly connected.

Bandwidth Capacity (available):

A capacity bandwidth test (available) is a unidirectional UDP metered data test. Metered data provides the ability to eliminate the impact of RTDT variation/delays. A such a UDP capacity defines a bandwidth available assessment because the UDP protocol is loss tolerant and any packet losses that occur during the test little or no impact on throughput. A bandwidth capacity test provides the ability to measure and assess available bandwidth regardless of quality issues that reduce the usable bandwidth.

Destructive assessments

A capacity test, by definition, is destructive and should not be frequently run. The destructive side effect increases the importance of the equilibrium throughput test for underwriting a healthy connection. It is the most effective test.

The prime purpose of the TCP capacity test is to understand how the bandwidth delay (latency) product changes the equilibrium as data demand climbs. This defines a connections efficiency. Establishing a threshold for acceptable delay is fundamental to a successful assessment strategy.

All three capacity tests documented above will deliver a bandwidth measure. Therefore, the GNCC purpose for the bandwidth test must be well defined. Discussion on this is recommended.

Network Packet Loss

The loss of packets naturally defines a poor network quality. The corollary to this defines that a good network does not drop packets. The assessment of packet loss is therefore important. With that said, networks discard packets for a wide range of reasons such as a regulatory policy, lack of buffer space and security. Loss of packets naturally has a detrimental effect on any application and with it, the user experience. The nature and magnitude of the resulting problem will depend on what packet, what protocol and the application need for the missing packets.

Packet loss comes in several forms and the measurement of packet loss required for good network assessment must be cognizant of the packet type as well as the loss.

The characteristics of packet loss vary by protocol and the impact of any loss affects each application differently based on protocol. The definition of packet loss needs to be well defined for GNCC assessment testing.

There are only three core protocol types that can encounter packet loss. UDP/RTP, ICMP and TCP.

UDP

UDP/RTP is a single direction, loss tolerant, protocol for packets that are time dependent. If a UDP packet is discarded for whatever reason, there is no recovery. The packet remains permanently lost (excluding VPN encapsulation attributes). As there is no recovery, UDP incurs negligible penalty for loss but the loss will impact other important performance metrics such as the PDV and packet order metrics.

ICMP

ICMP is a bidirectional diagnostic management protocol. ICMP packets can be intentionally discarded or priority disabled for reasons such as security, general policy or congestion interference. This makes the ICMP packet type useful for climate testing but unreliable for overall network assessment purposes.

TCP

TCP is a bidirectional data protocol. The TCP protocol packets are designed never to be lost including any VPN encapsulation. TCP packets are vitally important because they affect most data centric applications including all browser-based applications such as online banking, online reservations etc.

Measuring packet loss is important, however, the effect of loss on the different packet types is not the same. Significantly, TCP data packets incur the most severe penalties and therefore create the biggest threat to the user experience, whereas UDP is loss tolerant and has a far less destructive affect on the application. Unfortunately, TCP penalties occur only at layer 4 and are therefore not easily visible to any assessment measurement process.

Measuring loss of one specific protocol cannot and should not be applied to other protocols. Therefore, the assessment requirement for packet loss must be specific as to the packet loss metric(s) needed. UDP and TCP being a case in point. UDP is, time dependent, creates a small load on the network, is normally prioritized important and incurs negligible penalties from packets lost. Whereas, TCP is, not time dependent, creates significant load on the network, seldom prioritized and incurs significant time penalties from packets lost.

Network Quality

There is no clear definition of network quality that is understood for test measurement as it relates to the GNCC requirements. The definition of quality applies to many, if not all, measurements that are derived from network assessment testing. This includes, jitter, packet loss, bandwidth, latency etc.

That said, a network has an inherent definition of quality which is defined by a poor user experience. If the user experience is bad, the user does not care what the bandwidth test reports. In short, the events that affect the user experience are not easily measured by standard network testing. This does not diminish the importance of Jitter or bandwidth measures. On the contrary, the 5 measures targeted by GNCC are important. Namely:

- Latency
- Jitter
- Bandwidth
- Packet loss
- Quality

Not all packet delays on a network stem from congestion or over subscription. There are a wide range of issues that affect the user experience even when bandwidth looks healthy. Packet fragmentation, packet retransmissions, packets out of order, duplicate packets, packet timeouts etc. Lost packets (TCP not UDP), lost acknowledgements, CPU utilization etc. etc.

In addition to the metrics sought for the provider assessments, MCS supports the full measurement of TCP data quality including all the quality events listed above and more.

For a thorough network assessment, the inclusion of TCP quality metrics is an important and essential addition to the assessment processes as already defined. It is therefore recommended that GNCC carefully review and determine what comprises the quality requirement for the Visualware MyConnection Server project.