Huawei
HTR 4K Service Acceleration
TCP Throughput Optimization, IPv6 Support, Availability

Introduction

In the recent decades we saw rise of new networking technologies to deliver new services to the customers, but we can also see a constant change in abilities, requirements, and how those services are applied and used. The underlying technology constantly has to adapt to these changing conditions. Underlying principles that were developed long time ago may suddenly encounter trouble adapting to these new conditions. New solutions need to be found to support this process. Such is the TCP acceleration technology RACE implemented by Huawei which we present in this review.

As a good example of such evolution in the technology, we can consider the video delivery services. It's not long ago, that the IPTV and VoD delivery became one of the standard services offered to the broadband users. The IPTV delivery networks were tightly integrated into provider's infrastructure in order to ensure a high quality of the service. It commonly uses MPEG Transport Stream encapsulation over IP/UDP, and requires a suitable infrastructure. The IPTV is delivered with strict QoS requirement, and uses sufficiently reserved bandwidth. The source is located within the provider's network and the end-to-end forwarding latency is low.

In just recent years, we saw a major evolutionary shift in this area. More and more people are moving away from the traditional IPTV services to the novel third-party video programming. We saw emergence of the video platforms like Youtube or Netflix, which can be generally described as OTT (Over-the-top) video services.

The major difference is that video content is delivered over the common Internet and has to share the bandwidth and other network conditions with the common Internet traffic. Strict QoS is no longer enforceable, the latencies grow, and the bandwidth is not reserved. The service is confronted with the congestion situations in the network.

The subscriber's home network have been changed as well. More users are switching to and wireless access. Congestion and interference may quickly affect video traffic which is very sensitive to loss.

Test Highlights

- ➔ **Improved throughput performance for download or 4K video streams**
- ➔ **Performance suitable for 4K streaming video delivery even under heavy impairment**
- ➔ **TCP performance for IPv6 impaired network showing 0.01% loss and 50 ms delay: Throughput increased from 38 Mbit/s without HTR to 99.5 Mbit/s with HTR**
- ➔ **Slot redundancy**
- ➔ **Transparent bypass in case of a acceleration module failure**

As a result, these new emerging video delivery services have to switch to the reliable TCP Transport. Still, new challenges await.

TCP Congestion Control

The TCP is a transport protocol now several decades old, but its development never stopped. In the past decades, many efforts were made to implement more advanced congestion avoidance algorithms and to adapt it to the ever-changing conditions of the Internet and new emerging services.

The TCP Reno for a long time was a de facto standard congestion avoidance method in most operating systems and devices. Recent analysis however shows major deficiencies that would effectively prevent its use for the transport of today's services over the Internet, as it fails to utilize the available bandwidth efficiently.

The cause for this is the specific mechanism used by TCP Reno called Additive Increase Multiplicative Decrease (AIMD). The CWIN size and so the throughput are incremented slowly, yet in a case of congestion, the decrement is rapid. The goal is to quickly eliminate congestion situations and not allow them to build up due to additional retransmission traffic. The result is a typical zig-zag variation of the throughput.

TCP Reno assumes packet loss as the indication of congestion, but does not differentiate it from the random packet loss. Simulations show that on a link with realistic delay and loss, as little as 1/4 of the capacity would be used. A seemingly well dimensioned 100-Mbit/s link may turn up unsuitable for assured 4K video delivery..

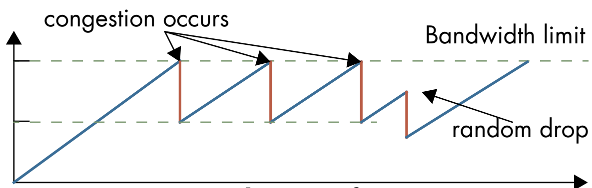

Figure 1: Behavior of TCP Reno

More advanced algorithms like H-TCP, BIC or CUBIC and such as CTCP that use latency analysis for congestion detection brought improvements, producing a more balanced adjustment and quicker recovery after congestions. Still, random packet loss negatively affects them, causing unnecessary throughput reductions.

The congestion avoidance was absolutely necessary in the networks of the past, but today the challenges lie elsewhere. The transport networks usually have sufficient bandwidth capacity, but we are more likely to observe random losses occurring on users' networks or in the access area.

The use of the Explicit Congestion Notification (ECN) also does not provide a suitable solution. Aside from the need to support it on the network elements and the endpoints, the ECN would only provide notifications about congestion, but not the random packet loss.

A more intelligent solutions to the TCP bandwidth control are necessary.

Huawei's RACE TCP Acceleration

Huawei's RACE is a novel TCP acceleration technology implemented for its line of NE40E Universal Service Routers. This technology is aimed to mitigate the negative effects of random packet loss in the access network on the TCP traffic and improve its performance. When the TCP traffic is redirected to a RACE-enabled HTR, it performs a detailed analysis of a large number of TCP connections in parallel and derives the best way to optimize the traffic flow.

The HTR acts as a TCP proxy, dynamically modifying the relevant TCP parameters and so controlling the traffic flow for each individual connection. Unlike other TCP congestion avoidance algorithms, that gradually try to increase the window size and the traffic rate, RACE tries to adjust the flow to the ideal calculated throughput.

RACE utilizes a mechanism for distinguishing congestion from random packet losses by monitoring the round-trip delays and other parameters. Congestions can be recognized by an delays increasing above normal, an indication of overfilled queues in the network nodes, while random packet drops won't typically cause such effect.

Huawei developed the Versatile Service Board (VSUI-160-E) to efficiently implement the RACE technology in their NE40E router platform which was the focus of our HTS test.

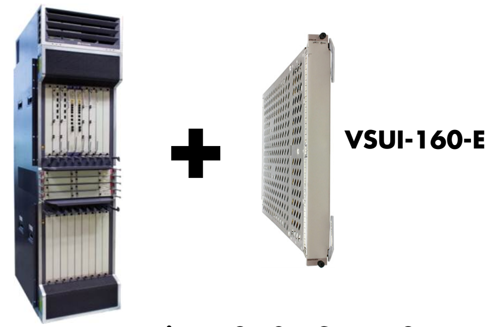

Figure 1: Huawei NE40E & VSUI-160-E

Test Equipment

In the core network, we provisioned a nginx-based FTP- and a video streaming server (Apple HTTP Live Streaming) equipped with a set of test files. These servers were accessed in our test to simulate large downloads and over-the-top video services used by the broadband subscribers. As video and FTP server platform Huawei decided to use a SUSE Enterprise 11 SP1 server which has been published in June 2010 as this can still be found in many of the content provider networks these days.

An Ixia ANUE GEM impairment generator included in the link between BNG and the HTR and operating transparently, provided simulation of different loss and delay conditions occurring in the access network.

A Windows workstation served as a client equipment for FTP downloads and a 4K TV and a set-top box for the playback of the 4K content.

For the video streaming tests, we used an IneoQuest IQD video quality analyzer, which was connected at a mirroring port on the BNG for video traffic monitoring.

Test Setup and Methodology

As our test setup, we simulated a broadband access scenario with GPON. By inserting an impairment generator into the link between the BRAS and the core network, we simulated packet loss, delay and delay variation, that would occur in the broadband access infrastructure, or in subscriber's network.

The HTR was situated in the core network. Depending on scenario, TCP traffic was either redirected to the TCP acceleration module or forwarded directly to the subscriber.

We simulated file downloads and video streaming by the subscriber, accessing external servers and measuring the average throughput. We conducted a series of tests with different level of impairment, and with or without HTR performing optimization of the TCP traffic.

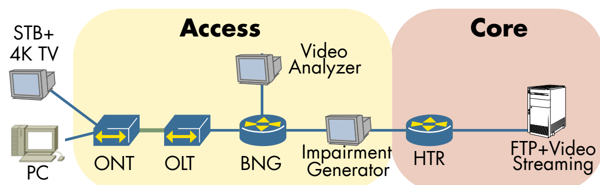

Figure 2: Huawei HTR Test Setup

The focus of the test was to test HTR's positive effects on throughput, not its reaction under critical situations or maximum load.

Test Results - File Download

In order to measure the effects of the impairment on the TCP traffic, and how RACE acceleration technology can mitigate them, we ran a series of comparison tests, using identical test setup, but different impairment settings.

The throughput was measured by downloading a 1.2 gigabyte large file from the FTP server to the client.

A baseline test conducted without any impairment showed nearly ideal throughput performance of 98.7 Mbit/s in both cases and confirmed that the test bed itself had no issues that could invalidate our results.

The test conducted with the plain TCP show that even a slight loss and delay immediately severely affect the performance. Even in the simplest case of 0.01% loss and 20ms Round Trip Time (RTT) we observed a reduction of throughput to just one third of the theoretically possible.

On the contrary, the measurements performed with the activated HTR showed a throughput performance above 95% of the baseline test except for the cases with the most severe delay (200ms), where we still achieved more than 65% of the original performance.

The diagram and the table below summarize the achieved throughput results.

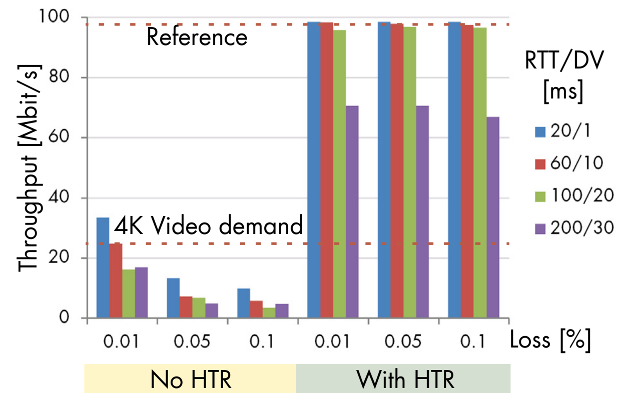

Figure 3: TCP Throughput under Impairment

Loss [%]	RTT [ms]	Delay Variation [ms]	Throughput [Mbit/s]	
			No HTR	HTR
Reference (no impairment)			98.7	98.7
0.01	20	1	33.4	98.4
	60	10	24.8	98.3
	100	20	16.1	95.7
	200	30	16.9	70.6
0.05	20	1	13.3	98.5
	60	10	7.2	97.8
	100	20	6.8	96.8
	200	30	4.8	70.6
0.1	20	1	9.8	98.4
	60	10	5.7	97.4
	100	20	3.5	96.5
	200	30	4.7	66.8

Test Results - Video Streaming

We conducted the next series of tests using real 4K video traffic (HTTP Live Streaming) requested by a 4K television set and set-top-box and streamed 4K content from the HLS server located in the core network.

The video experience quality was asserted visually on the TV, but also analyzed by the IneoQuest IQD. The analyzer measured certain streaming-related KPIs, such as buffering issues.

The measurements performed without impairment confirmed good video quality, reaching VeriStream score of 4-5. After that, we introduced varying levels of impairments and evaluated the test video again, with or without HTR.

The diagrams presented below show a “KPI boundary” separating the impairment scenarios where the video quality was still acceptable from those where it was not.

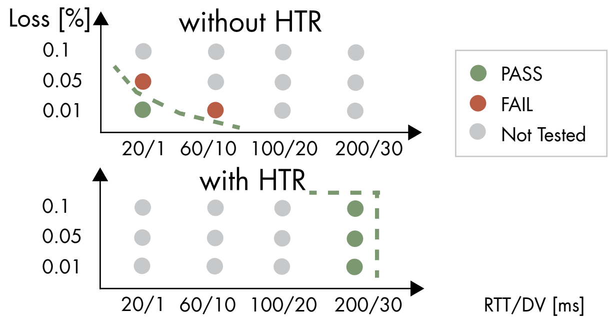

We based ther verdict on the minimum achieved VeriStream score that signifies the worst buffering event during the test. The test was evaluated as FAIL if a buffer underrun occurred in any of the video segments, corresponding to scores below 4.

IPv6 Support

We repeated a small subset of the test cases using file downloads, this time using IPv6 traffic. We achieved very similar performance values: 37.9 Mbit/s throughout without HTR and 99.5 Mbit/s throughput with HTR respectively, when an impairment profile with 0.01% loss and 50 ms delay was configured. This test confirms that Huawei’s solution is prepared for upcoming IPv6 deployments.

Availability

Finally, we also examined the availability-related features. Regarding HTR’s role, where it has to forward, analyze and manipulate TCP, i.e. most of the Internet traffic, an extremely important aspect.

As the first test, we examined the behavior of the HTR running 2 TCP acceleration modules in case of one module’s failure.

The test was performed by powering off one of the modules while video traffic is transmitted through the network. In addition, we generated background HTTP traffic between the BNG and the core network using Huawei’s XStorm traffic generator in order to load the HTR with up to 10 Gbit/s.

We observed that the TCP sessions registered on the affected card have failed, and were re-established through the streaming process. Since the stream was buffered on the STB, no defects could be observed visually, and since the new TCP sessions were maintained by the second module, good TCP performance was maintained as well.

As the next test, we tested the emergency bypass feature of the HTR in the case of a complete failure of the TCP acceleration modules. After we disabled the remaining module, the HTR switched to bypass mode, forwarding all traffic as is. As expected, the performance of TCP was degraded due to impairment, but the connectivity was still maintained.

Summary

Through the series of tests with data download and streaming 4K video, we could verify the ability of Huawei’s RACE TCP acceleration technology to maintain good TCP performance in a network significantly impaired by random packet loss and high delay and delay variation. The tests have shown that a HTR-enhanced network would allow a stable 4K video streaming experience to the subscribers even in such unfavorable conditions.

About EANTC

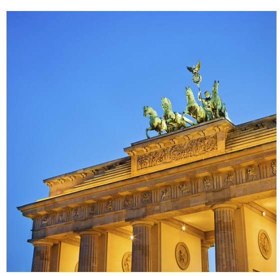

The European Advanced Networking Test Center (EANTC) offers vendor-neutral network test services for manufacturers, service providers and enterprise customers. Primary business areas include interoperability, conformance and performance testing for IP, MPLS, Mobile Backhaul, VoIP, Carrier Ethernet, Triple Play, and IP applications.

EANTC AG
Salzufer 14, 10587 Berlin, Germany
info@eantc.com, http://www.eantc.com/

[version vF3] [date 20151018], [TS]