NetHealthCheck™* Packet Flow Performance Optimisation

Neil Davies and Peter Thompson

July 2012

PREDICTABLE NETWORK SOLUTIONS

Contents

List of Figures

*NetHealthCheck is a trademark of Predictable Network Solutions Limited.

1 Who Needs NetHealthCheck?

Any network operator that can prove that their network is optimised so as to deliver a predictable and optimal experience to all of their customers all of the time while operating at the lowest possible cost does not need NetHealthCheck!

What NetHealthCheck delivers is a measurement framework and the support to achieve the above goal, for operators who have either:

- self-built and continue to manage all the performance and system integration aspects of the network in-house;
- subcontracted out implementation and operations and now need a clear and concise way of assuring that their overall objectives are being met.

2 The NetHealthCheck Solution

NetHealthCheck is a service to uncover performance impairments and sub-optimal aspects of the way that a network is configured, managed and structured (recognising that optimality has different meanings for different Network Operators and even for the same Network Operator at different points in the market cycle). NetHealthCheck delivers either evidence that the network is "perfect" or one or both of:

- how to improve the network for little or no cost, including a quantification of the effect on subscriber perception of quality
- how to achieve cost reductions in operation or near- and long-term planning.

This is achieved by using a scientifically rigorous framework, backed up by powerful mathematical analysis that enables root cause isolation. This is combined with reproducible tests to assure that problems stay fixed, with the capability for ongoing monitoring to catch any recurrences.

NetHealthCheck uses a small set of probes to perform very precise measurements on typical end-to-end network paths, which are then analysed to extract information (as explained in Section 3 below) to identify performance impairments. These results are presented in a report that includes recommendations for reducing or eliminating the impairments. If required, further cycles of measurement and recommendation can be made, or alternatively the probes can be removed at the end of the engagement.

2.1 Observing Quality Attenuation Accumulation Across a Network

Figure 1 shows typical locations for the insertion of probes in a fixed and/or mobile network. Most network test solutions use many packets passing between two points; by contrast, NetHealthCheck measures the passage of relatively few packets past a sequence of observation points on typical network paths. Packets are passed all the way between reference probes R and/or S, connected at the operator/corporate interface, and edge probes A - E, and back. They are also observed at intermediate points such as I and J, and together these allow independent observations of performance impairments in both upstream and downstream directions and in different sections of the network.

When issues are uncovered, additional probes may need to be inserted (or original ones moved) to isolate the cause precisely.

Performing measurements this way rather than simply exchanging packets between pairs of points allows more subtle behaviour to be observed and diagnosed. The focus is on the contribution of each section of the network to the end-to-end performance, rather than obtaining (mostly average) measures of individual network elements.

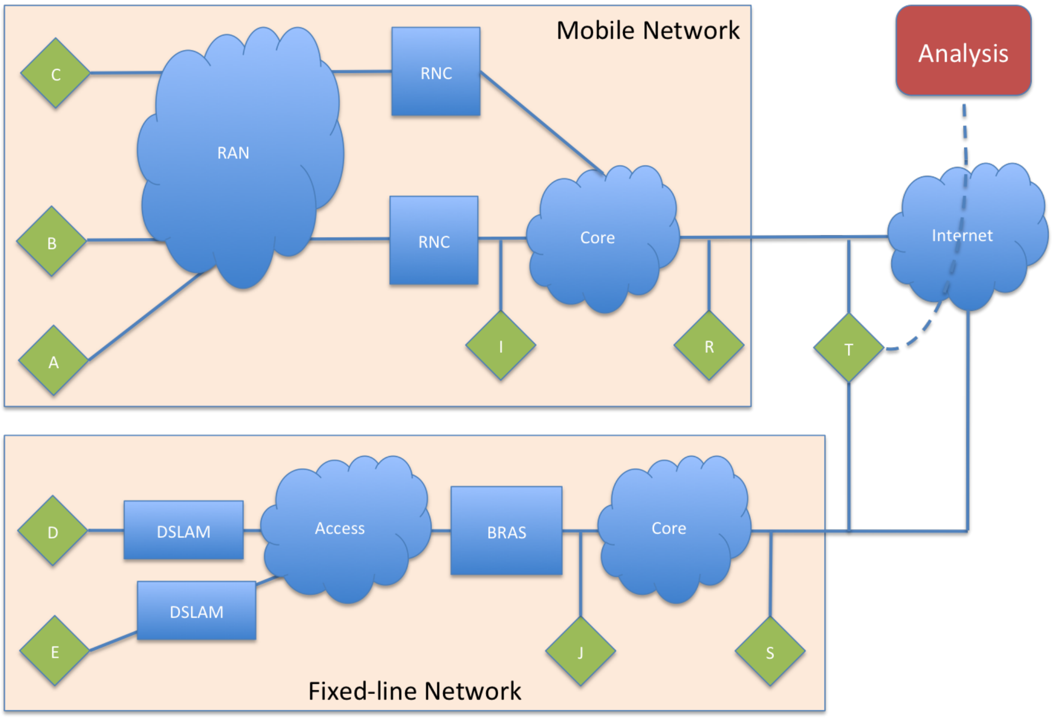

Figure 1: Typical Probe Locations

3 Decomposing Quality Attenuation

Following the transit of packets between multiple probes provides a measure of how packets are lost and delayed in their end-to-end journey, which is called 'quality attenuation' (ΔQ). ΔQ is the result of multiple factors in the network, and to make proper sense of the measurements, it can usefully be decomposed into three principle components for each network path:

G: The constant effect of geography and network topology, for example due to the speed of light and distance;

S: The effect of serialising packets onto links of various speeds and characteristics along the network path; this is a function of packet size, and;

V: The variability in transport quality due to the sharing of resources in the network; the effect of contention.

Examining each of these components in detail reveals particular aspects of the network performance. They can be thought of as a composable 'basis set' for network performance.

'G' captures issues relating to topological and physical distances between network elements, the irreducible cost of processing a packet, and the delay in gaining access to the transmission medium (for example in 802.11). It can be thought of as the time taken for a zero length packet to be transported between two measurement points, given that the rest of the network is completely idle. When measuring this at multiple locations there should normally be a simple relationship between distance and the measured value; where this is not the case there is an anomaly worth investigating, typically by placing additional probes to further isolate the portion of the network concerned. It is important to isolate this factor as topology can dominate the absolute measured delay. Once geography dominates, the only option is to change the speed of light! Once G is known, it can be discounted when comparing measurements from two different locations. This allows other features of the data transport quality to be revealed.

The effect of packet size, that is removed from G, is captured in 'S'. S can be thought of as the relationship between packet size and the transport time of a packet across an otherwise idle network. S should not vary between locations that are connected at the same access link speed; any such variation would imply an imbalance in the network topology. In principle S should be a simple constant factor, but packet fragmentation can produce substantial variation from this ideal. Isolating and resolving such fragmentation can have a substantial effect on the quality of experience (QoE) delivered to end users.

Networks are of course not idle, and as a packet traverses network elements in its end-to-end journey it is will compete with other packets for access to resources. Those resources can be onward transmission capacity, buffer space, CPU cycles for packet processing, as well as share of capacity over a radio link. The contention for such resources is reflected in the 'V' component of the measurements. It is not commercially feasible to make V disappear, its presence is a natural consequence of using statistical multiplexing. Excessive and/or rapid variation of V can indicate poor management of contention and is a major factor in the variability of delivered quality of experience. Understanding the causes, and hence the effective management, of V is essential for the network operator to achieve flow and resource efficiency, as discussed in section 4 below.

Analysing the measurements into G, S, and V is a key component of the service to a network operator, making different types of packet flow performance impairment visible. Examples in practice are given in section 5.

4 Flow Efficiency vs Resource Efficiency

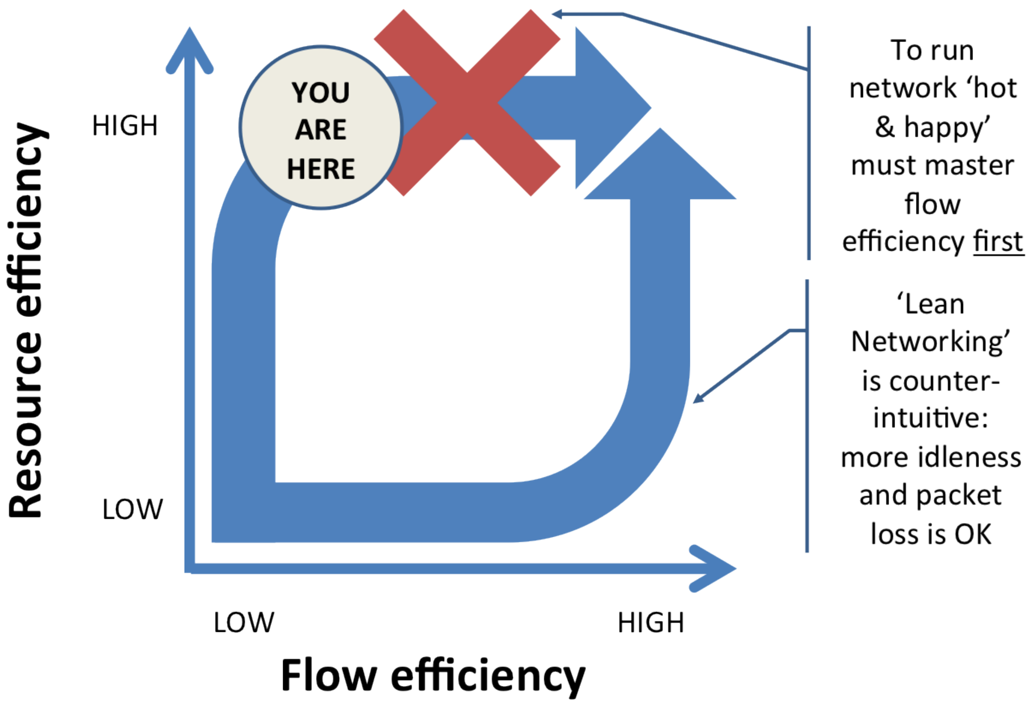

Figure 2: The Path to Lean Networking

Measuring quality attenuation provides a way to focus on the 'flow efficiency' of the network, i.e. its effectiveness in delivering end-user outcomes and hence appropriate QoE. This enables a path to 'lean networking', which is both effective and efficient in its use of resources; as shown in figure 2, it is essential to optimise flow efficiency first before trying to optimise resource usage[1]. This is because a network that has been optimised for resource efficiency but fails to deliver satisfactory end-user outcomes cannot be changed without becoming less efficient, whereas one that delivers acceptable performance outcomes can be made more efficient without damaging those outcomes (provided appropriate measurements are in place).

[1] See, for example http://www.futureofcomms.com/blog/2012/5/16/lean-networking-2.html.

5 Typical Issues Uncovered

Detailed analysis of how the different components of ΔQ behave reveals performance impairments that are often easily corrected.

5.1 G-Based Issues

These reveal routing and topology bottlenecks.

Observation: In an MNO network, some NodeBs connected to an RNC displayed much larger values of G than others.

- **Deduction:** ATM IMA over multiple E1s was being used to connected the NodeBs, and one of the E1s in the group was routed differently to the others. The action of IMA caused every packet to experience the maximum delay of the group.
- **Recommendation:** Improve the routing of the E1s.
- **Benefit:** Improved subscriber QoE at minimal cost.

Observation: Subscribers of a fixed-line operator in a major city experienced a larger G than those in the surrounding region.

- **Deduction:** To deal with the subscriber numbers in the city, an extra network layer had been added.
- **Recommendation:** Re-balance the network topology.
- **Benefit:** Improved customer satisfaction leading to better competitiveness in the key market. Rebalancing the architecture freed up some equipment which could be re-purposed.

Observation: MNO traffic that was, geographically, assumed to travel just 20km had a value of G equivalent to nearly 1000km, 19ms delay instead of <2ms. Note that this constituted 25% of the end-user end-to-end delay!

- **Deduction:** The Tier-1 connectivity supplier to the MNO was delivering extremely sub optimal routing; nothing in the supply contract constrained this.
- **Recommendation:** As the MNO was moving to a new supplier, ensure that the new contact terms include sensible routing paths.
- **Benefit:** More effective supplier management, leading to increased end-user QoE at no extra cost.

The scenario above can occur in many situations of international transit. It can be a significant factor in the delivered QoE to roaming customers for MNOs, which in-turn affects an important revenue source.

Observation: In a large scale corporate VoIP system outpost-to-outpost calls had issues. Measurements showed that G was 30ms higher than expected at each end, rendering VoIP unusable.

- **Deduction:** ADSL provider had optimised throughput by turning on interleaving, leading to increased delay.
- **Recommendation:** Renegotiate ADSL provision to trade lower peak bandwidth for lower delay.
- **Benefit:** VoIP issues resolved without extra cost. VoIP supplier contract was no longer in breach of its SLA.

5.2 S-Based Issues

These are frequently caused by transiting multiple layer-2 domains, excessive layer-3 routing or inappropriately configured other layer-3 equipment such as firewalls and traffic management equipment.

Observation: Examining the variation of S with packet size for a major MNO showed sub-optimal behaviour.

Deduction: Layered tunnelling protocols were introducing fragmentation and overheads that reduced performance, particularly in the presence of diverse routing.

Recommendation: Eliminate fragmentation either through increasing maximum L2 frame size or through appropriate configuration at edges.

Benefit: 100% improvement in the end-user web download KPI.

Such layering is common in MNOs (it is in the standards), and also occurs where VPNs or other tunnelling approaches are used, typically by large corporate customers.

Observation: A major international research institute routed traffic through an intermediary with a rate limit. Despite shaping traffic to this limit, high packet loss occurred.

Deduction: The intermediate rate limit was applied to packets encapsulated in MPLS over Ethernet.

Recommendation: Take these overheads into account when shaping traffic at the IP layer.

Benefit: Packet loss eliminated; headline-making scientific discoveries facilitated.

Observation: In an MNO network, S was found to have multiple values.

Deduction: Packets may be taking multiple heterogeneous routes, or RNC configuration is incorrect.

Recommendation: Change the granularity on which packets are dispersed to avoid out-of sequence delivery; check RNC configuration.

Benefit: A serious mis-configuration of the RNCs was uncovered; correcting this improved performance at no extra cost.

5.3 Dynamic Properties and V-Based Issues

G and S should be constant; issues relating to un-expected G and S tend to be fixed issues to do with static properties of the network (routing, transmission paths, topology, aspects of equipment configuration). Variation due to contention, V, should be bounded provided the statistical multiplexing is well managed.

Observation: With both a military customer and a DSL broadband provider, G, S and V were all seen to vary dramatically during changing environmental conditions.

Deduction: The transmission equipment reconfigured itself in response to lower available link throughput (e.g. rain fade on wireless links, local environmental issues on DSL).

Recommendation: New profiles of operation for the equipment using those particular links.

Benefit: Transitions to lower throughput operation now just delivered lower capacity instead of causing the equipment to reset, resulting in temporary loss of service.

Observation: Working with a major network wholesaler, changes in V during different loading conditions were seen, and at some loads G began to wander causing distributed clock synchronisation to be lost with consequent loss of service on certain equipment such as NodeB, femto-cells etc, where accurate distributed time is essential to operation.

Deduction: The predictable region of operation of the network element/sub-network was being exceeded.

Recommendation: Changes to planning rules to take account of the limitations of the installed equipment, including its scheduling algorithms and configurations. Re-configuration of scheduling in key network elements to reduce wander of G for key timing traffic to acceptable levels.

Benefit: Reduction in costs (better overall utilisation of equipment before upgrading) combined with higher overall operational stability by removing a key failure mode.

Observation: In an MNO network, much of G,S, and V accrued on the customer's network and within the Network Operators management domain.

Deduction: The QoE impairment was largely due to the practices of the customer (the intra-subscriber contention) as compared to the practices of the network operator (the inter-subscriber contention).

Recommendation: How the network operator could aid the customer achieve its goals, and how the data transport portion of the SLA could be re written.

Benefit: Rapid dispute resolution due to access to reliable quantitative measurements. Happier customers (they knew what they had to do) and happier Network Operator (less resources being spent in dispute resolution).

Observation: Working with application service providers, the G,S, and V being delivered from the same location by different Network Operators (fixed, mobile and wireless) are found to differ substantially.

Deduction: Checking the relative quality attenuation costs of each network provider.

Recommendation: The likely effects of the differences on different uses by the subscriber - the relative fitness for purpose for applications of the different providers.

Benefit: Quantitative competitive analysis, enabling the value of each competing service to be assessed. Such evidence has a strong scientific basis of both capture and analysis and, as such, could form the basis of a claim by, for example, an advertising standards authority.

6 Conclusion

NetHealthCheck is a service based on a profound analysis of the fundamentals of network operation, which can expose and isolate the causes of performance impairments. In many cases the solutions to these problems are straightforward and inexpensive, and produce significant improvements in end user QoE, network operating costs, or both.