Bandwidth Management

Internet Society Technology Roundtable Series

Introduction

Bandwidth management is clearly a complex and evolving concern for most, if not all, network operators around the globe. A recent survey of over 400 European network operators found that blocking of voice-over-IP and peer-to-peer traffic is commonplace today. This survey also identified a very wide variety of traffic management practices, implementation methods and policy justifications. The rapid introduction of new Internet applications and services makes it hard for network operators to predict or robustly categorize uses of the network. Meanwhile, users want their Internet service to 'just work'.

Building on an earlier briefing panel and associated report into the evolving landscape of Internet bandwidth, the Internet Society convened an invitational roundtable meeting of network operators, technologists, researchers and public policy experts with an interest in broadband regulation and deployment. This report documents the Internet Society Technology Roundtable meeting on the topic of bandwidth management that took place on October 11th and 12th 2012 in London, England.

Executive summary

At the macro scale there is a good story to tell about new fiber deployments, and more open and competitive cable consortia. However, dark fiber availability is a local concern for some as fiber infrastructure operators seek to provide higher margin services.

Content Delivery Networks are part of a larger architectural shift toward localized content delivery and content aggregation services. This shift yields benefits for network users in terms of quality of experience. Content aggregators are helping to meet the growing demand for Internet content but there still remain huge spikes in traffic related to new releases of software and other kinds of popular content that are creating new challenges for network capacity planners and performance engineers.

Latency management is the key to understanding the causes and solutions to many of the performance problems witnessed on today's broadband access networks.

Latency management is the key to understanding the causes and solutions to many of the performance problems witnessed on today's broadband access networks. There are opportunities for development and research stemming from this fact in terms of new algorithms, an improved access network router ecosystem and greater consumer and network operator awareness of the long-term costs of poor engineering and cheap products in this critical region of the end-to-end network path.

Bandwidth at the macro scale

- Latency management (by which we mean controlling for and minimizing unnecessary network transmission delays) is the root of many of the problems and evolutionary trends we are seeing on the Internet today. Queue latency (where data packets are

delayed in the network) is often the dominant determining factor in perceived Internet performance rather than raw bandwidth.
- Even in the middle of a financial crisis and amidst falling prices there are still investments in new undersea cable capacity being made.
- Investment in bandwidth is not going to solve performance problems unless latency management issues are also addressed.

Network operator challenges

- Constant or increasing rates of traffic growth are commonplace – growth that is not always visible from publicly available sources.
- Concerns are beginning to be raised about the scalability of the current Internet architecture given the linear growth of core networking interface speeds.
- Traffic growth rates are not tailing off when looking at individual operator statistics, however, per-user consumption curves are starting to flatten out at the very high bandwidth speed tiers now coming on the market as the bandwidth bottleneck moves to the content provider end of the path.
- The benefits of middle-box solutions to bandwidth management challenges in mobile networks are short-term and will have trouble scaling as mobile broadband Internet and fixed broadband Internet service definitions converge.

Future directions

- We can't build infinite bandwidth to the edge everywhere. It is therefore necessary to engineer and deploy mechanisms that allow performance to degrade much more gracefully than is typically the case today.
- Better architectural support for technical collaboration between network and content aggregators and application providers to deliver more network and application management data is a high priority.
- Novel solutions for both per-customer bandwidth management and per-flow bandwidth management are on the horizon.
- Specifying a latency-under-load test is one very concrete step that will help root out network paths that have bufferbloat-related problems.
- Separating software from hardware in the Customer Premises Equipment (CPE) business is highly desirable. Creating open and modular CPE software platforms, like OpenWRT[1] , will help to remove the CPE software as a bottleneck on progress and innovation. Openly specified testing with network operator input, coupled with certification programs could help to improve CPE quality.
- Growing awareness of latency management problems and promising, partial, solutions give some hope but deployment of major technological upgrades will always be a challenge for a large distributed system like the Internet. Deployment challenges for such new mechanisms are severe and a focused effort to radically improve the quality of the network edge seems necessary to make improvements in this regard.

Deployment challenges for new mechanisms are severe and a focused effort to radically improve the quality of the network edge seems necessary.

[1] https://openwrt.org

Background

"Days in which the technological reasons were impacting the performance of the residential Internet users are coming to an end and in the near future, the performance offered to the end users will be predominantly shaped by the economic factors rather than physical bottlenecks."[2]

Bandwidth management is clearly a complex and evolving concern for most, if not all, network operators around the globe. A recent survey of over 400 European network operators found that blocking of voice-over-IP and peer-to-peer traffic is commonplace today.[3] This survey also identified a very wide variety of traffic management practices, implementation methods and policy justifications. The rapid introduction of new Internet applications and services makes it hard for network operators to predict or robustly categorize uses of the network. Meanwhile, users want their Internet service to 'just work'.

The rapid introduction of new Internet applications and services makes it hard for network operators to predict or robustly categorize uses of the network. Meanwhile, users want their Internet service to 'just work'.

Building on an earlier briefing panel and associated report into the evolving landscape of Internet bandwidth, the Internet Society convened an invitational roundtable meeting of network operators, technologists, researchers and public policy experts with an interest in broadband regulation and deployment. The findings of the earlier work were:

- The growth of Internet bandwidth globally is not about to cause global problems. International and inter-carrier links are not, in general, unable to cope with the demands of growing bandwidth consumption.
- While gross Internet capacity is meeting demand today, new capacity will be required in the long term.
- Adding more capacity to address access network constraints may or may not resolve network congestion issues, and should be considered in light of the relative cost of alternative solutions.
- Increasingly expensive and complicated network-specific bandwidth management techniques do not address the problems arising from broadband deployment.
- Solutions emerging today are a better fit to the Internet architecture than complex bandwidth management functions in the network and are cheaper to deploy as well.
- It is not sufficient to solve a problem for one network; we need global solutions for a global network.

The agenda for the latest roundtable was deliberately broad and included talks and discussion of the physical network layer (optical transmission state of the art and trajectory), Internet Exchange Points (IXPs) and scaling interconnection, the structure and state of the IP interconnection market, access ISP traffic growth and future trends and challenges, and development and deployment of new mechanisms and algorithms to address the evolving bandwidth management challenge.

This was a closed meeting, with invited talks followed by general discussion of the issues raised. The Internet Society presents this report of the meeting (with the consent of those present) to help further distribute the information and ideas discussed and the conclusions reached.

---

[2] R. Stanojevic, N. Laoutaris, P. Rodriguez. "On Economic Heavy Hitters: Shapley value analysis of the 95th-percentile pricing", Proceedings of IMC 2010

[3] 'A view of traffic management and other practices resulting in restrictions to the open Internet in Europe. Findings from BEREC's and the European Commission's joint investigation', 29 May 2012, https://ec.europa.eu/digital-agenda/sites/digital-agenda/files/Traffic%20Management%20Investigation%20BEREC_2.pdf

Internet bandwidth – history, evolution and diversity

Latency management is the root of many of the problems and evolutionary trends we are seeing on the Internet today.

- Latency management is the root of many of the problems and evolutionary trends we are seeing on the Internet today.
- Even in the middle of a financial crisis and amidst falling prices there are still investments in new undersea cable capacity being made.
- Investment in bandwidth is not going to solve performance problems unless latency management issues are also addressed.

This session explored the history and evolution of Internet bandwidth and identified the diversity of current bandwidth provision around the globe. The discussion also addressed the concern raised in some quarters that optical physics research is now lagging behind the bandwidth demand curve.

The first strongly supported observation was that latency management is the root of many of the problems and evolutionary trends we are seeing on the Internet today. In one example, a multi-million dollar cable deployment is being justified on the grounds of communication latency improvements on the order of milliseconds. The challenges posed by growing use of high-bandwidth video applications and interactive applications can also be characterized in terms of latency management (specifically queue latency).

Bandwidth growth on the Internet has been exceptionally fast for such a complex system. Compound annual growth rates in excess of 50% have delivered eightfold growth of global Internet traffic over the last five years and a CAGR of around 30% is predicted to deliver threefold growth over the coming five years.[4] Growth is much faster in emerging economies and varies greatly by country and by geographic region. Video streaming and download, cloud services, and mobile data services are driving bandwidth demand. Peer-to-peer traffic and voice traffic are proportionately less significant contributors to growth rates.

As far as the interconnection market is concerned, prices for transit have declined consistently by an annual average of 38%, while the savings afforded by peering have remained fairly constant. A new wave of investments in undersea cables is being driven by requirements for greater route diversity, improved quality (lower latency), and greater competition for underserved regions. Africa in particular is witnessing enormous levels of investment in new undersea cable capacity (see Figure 1). Even in the middle of a financial crisis and amidst falling prices there are still investments being made.

[4] Cisco Visual Networking Index: Forecast and Methodology, 2011-2016

Figure 1 - New cable routes around Africa

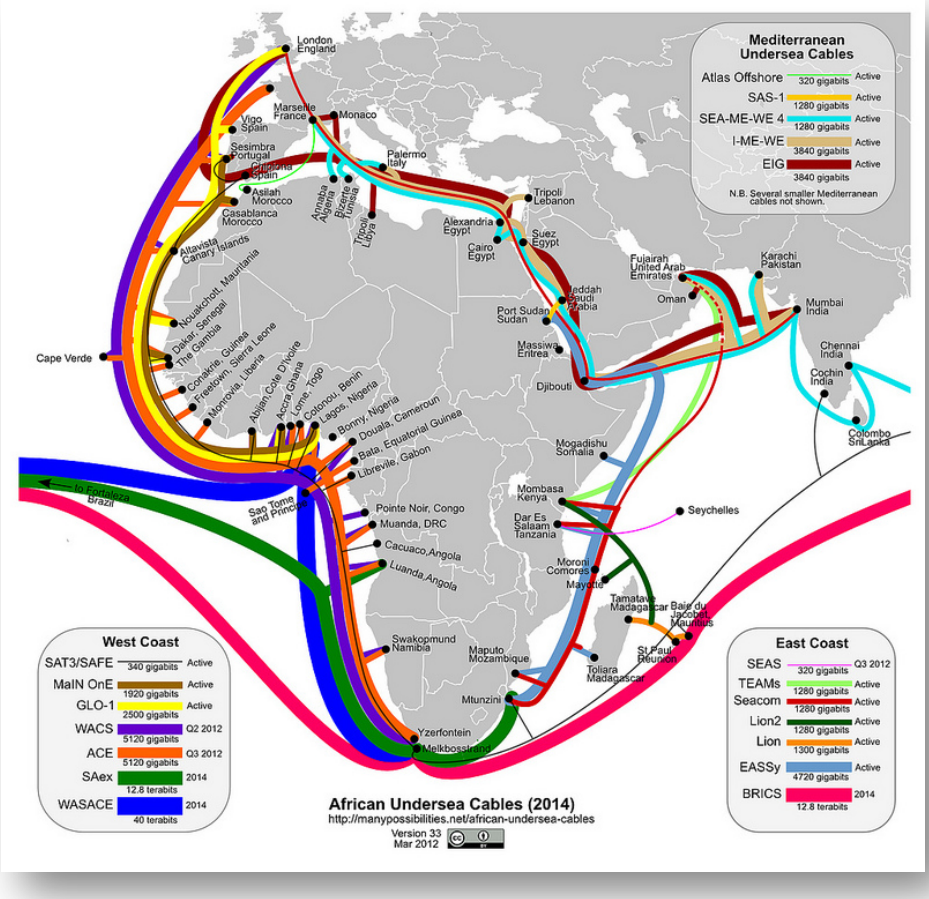

Growth in network interface speeds has declined from the earlier exponential trend. A combination of reduced levels of research investment and difficulties of running higher speeds over anything other than very short distances have combined to encourage aggregations of lower link speeds in the short term.

Partly as a consequence, port density limitations and IXP scalability challenges are increasing. Average port utilization is increasing as bandwidth demands increase and the challenges of scaling continue. As router slot capacity is growing less rapidly than bandwidth demand and link capacity increases, this has implications for port density, power consumption and cooling requirements. Router bit density is not growing as fast as traffic due to the limits of chip I/O throughput where growth has been linear. Modern multi-chassis routers can handle huge volumes of traffic, but could leave the operator with 24kW of heat to dissipate.

Modern multi-chassis routers can handle huge volumes of traffic, but could leave the operator with 24kW of heat to dissipate.

Transmission systems today are capable of carrying tens of 100G channels. New multiplexing techniques could increase this by an order of magnitude. New transmission systems entirely, such as hollow core photonic crystal fiber, would have a more dramatic effect on capacity, but to be useful to the Internet layer, this bandwidth will still require termination on router blades, which in turn will require power and appropriate levels of cooling. October 2011 saw AMS-IX offer its first 100Gb/s ports and in January 2012, Atrato became the first exchange point member to deploy 100Gb/s connectivity.

There was discussion of whether the pressures arising from scalability challenges, interface speed declines, power dissipation and so on mean that it is now time to rethink the Internet

architecture. Content distribution networks (CDNs) are providing an additive, rather than a subtractive, approach to delivering higher quality services and lower latency communications. The CDN market is growing rapidly and is delivering structural change in the broader industry and value chains. Access networks can profit from their ability to attract content resources into their networks, delivering an improved quality of experience to the end user.

It was proposed that many of the questions raised in the net neutrality debate (such as *How should Internet content be delivered to end users?* and *What resources should be deployed to improve quality, where, by whom and on what terms?*) are already being answered in the market for IP interconnection agreements, while new questions are emerging from the shift in content distribution architecture.

Bandwidth challenges faced by operators today

- Constant or increasing rates of traffic growth are commonplace – growth that is not always visible from publicly available sources.
- We can't build infinite bandwidth to the edge everywhere. It is therefore necessary to engineer and deploy mechanisms that allow performance to degrade much more gracefully than is typically the case today.
- Better architectural support for technical collaboration between network and content aggregators and application providers to deliver more network and application management data is a high priority.

We can't build infinite bandwidth to the edge everywhere. It is therefore necessary to engineer and deploy mechanisms that allow performance to degrade much more gracefully than is typically the case today.

What are the operational challenges faced by major ISPs today as they seek to manage efficient use of their networks, plan for network expansion and stay profitable?

Network operators are experiencing constant or increasing rates of traffic growth on their networks. One network operator cited over 50% traffic growth in 2011/12 and predicted growth in 2012/13 of 72%. It is unclear whether provisioning is tracking demand or vice-versa given the adaptive nature of streaming video applications (and their users) that make up the bulk of the growth.

P2P traffic volume is decreasing as a share of the overall traffic volume. In one example, the share of P2P traffic fell from 11% to 8% in the period 2010 to 2012 although this may in part be related to throttled bandwidth remaining constant while overall capacity increased. During the same timeframe, video-related traffic increased in share from 36% to 49%. Content now comes from a few large content aggregators: 5 aggregators contribute up to 60% of the traffic. Statistical multiplexing assumptions are being challenged by the dominance of streaming video. These changes in uses of the network are putting pressure on network operator upgrade cycles, making capacity planning harder in the presence of unpredictable spikes in demand and leading to new requirements for better latency management to allow for smooth performance of interactive applications and better apportioning of network resources between users.

While it is hard to make concrete predictions about the implications of these observed trends for the long term, today we can observe that per-user consumption curves are starting to flatten out at the very high bandwidth speed tiers now coming on the market as the bandwidth bottleneck moves to the content provider end of the path. Better tools to help diagnose the root causes of poor subscriber experience will become increasingly important as these new, higher speed tiers become more widely deployed.

This growing traffic demand is not evenly distributed: unlimited users on fiber access products are driving the majority of network growth. Event-driven traffic spikes are creating operational issues in some cases. For example, the recent release of the iOS6 operating system update from Apple created a new bandwidth demand record on some networks. The network operators present also highlighted that some applications are incredibly bandwidth inefficient. Commonly used software libraries don't always do the right thing, or provide developers with enough obvious ways to do the right thing with regards to bandwidth consumption.

For mobile networks, typical business-case driven approaches to bandwidth management include deploying devices that break the Internet - middleboxes, 'intelligent' caching proxies, etc. The benefits of these middle-box solutions are short-term and will have trouble scaling as mobile broadband Internet and fixed broadband Internet service definitions converge. Many mobile providers have deployed TCP proxies that make speedtest services go 10% faster, for example. Alternative approaches including usage based pricing, fiber to the tower, layer two edge QoS, and CDN deployments are considered long-term solutions that are consistent for wired and wireless networks.

A fundamental observation for both fixed and mobile networks is that we can't build infinite bandwidth to the edge everywhere. It is therefore necessary to engineer and deploy mechanisms that allow performance to degrade much more gracefully than is typically the case today. Congestion is a normal operating condition at the edge of the network, guaranteed by the TCP design that attempts to run as fast as possible. Congestion is not abnormal in any way. TCP cannot work correctly in the face of large packet buffers: the design requires *timely* notification of possible congestion by packet drop (given ECN is not deployed). Excessively large packet buffers (so called 'bufferbloat') create a performance 'cliff' whereby, when capacity is reached, performance as seen by most interactive applications drops dramatically. If the presence of large buffers is not mitigated, then deploying ever more bandwidth all the way to the network edge is the only alternative to maintain performance.

Mobile networks suffer particularly badly from the bufferbloat problem.

Mobile networks suffer particularly badly from the bufferbloat problem. The LTE network stack includes acknowledgements at all levels (from L2) guaranteeing 100% packet delivery, which can lead to round-trip latencies of tens of seconds. This results in users timing out before the applications do.

For operators and for content providers alike, quality of experience for the end user is a very high priority. But capacity planning is an ongoing challenge, and resilience and reliability are increasing concerns given the greater dynamic range of the consumption profile. Better architectural support for technical collaboration between network and content aggregators and application providers to deliver more network and application management data is a high priority. Regulation and pressure from competitive markets for services are necessary to keep bandwidth controls to a minimum. Most network operator bottlenecks will remain in their internal infrastructure, not in their connections to external networks, since they have very diverse peering and transit connectivity already.

Delivering big content - now and in the future

- Delivering big content requires collaboration with access network operators to succeed and there is scope for improvement in this area.
- The long-term scalability of this approach to application delivery is questionable.

For major Internet content providers, managing distribution of their content on a global basis is a massive operation. Networks of tens of PoPs, peering with access networks in nearly a hundred IXPs and deploying caching infrastructure into operator networks in over a hundred countries helps to guarantee low latency and resilient end-user experiences. Content providers derive tremendous benefits from on-net caches and CDNs. For example, rebuffering rates of streaming video services can be reduced four-fold. For some content providers over half of their traffic is served from inside 'eyeball' networks, and this is trending upwards. However, better technical coordination between content providers and access network operators could yield further benefits, for example by reliably scheduling background transfers in (locally-determined) off-peak hours or providing information about the kind of network a specific user is connected to. This coordination could be accomplished with new signaling protocols or in-band signaling.

In-network caching and CDN delivery also provide benefits to network operators, enabling some applications to be functional where previously they had not. New application functionality makes access networks more attractive to users thereby increasing subscriber numbers. New services can drive up consumption-related revenues as well, even in an environment of declining prices.

Concerns were raised about the scalability and costs of the distributed content approach. Rapid growth in real-time video services will present a challenge to content delivery mechanisms that currently rely on the amenability of popular content to caching. Managing literally hundreds of small datacenters and interconnection partnerships is extremely challenging and expensive compared with a more centralized approach, but is currently the best way to deliver a resilient and responsive application to a global audience.

New directions in bandwidth management – network and host support

- Novel solutions for both per-customer bandwidth management and per-flow bandwidth management are on the horizon.
- Deployment challenges for such new mechanisms are severe and a focused effort to radically improve the quality of the network edge seems necessary to make improvements in this regard.
- Specifying a latency-under-load test is one very concrete step that will help root out network paths that have bufferbloat-related problems.
- More attention to cross-layer optimizations is also desirable.

Specifying a latency-under-load test is one very concrete step that will help root out network paths that have bufferbloat-related problems.

The need for network support to enable something more than simple best effort traffic forwarding and flow-rate fairness was acknowledged at least as far back as the original work on congestion control for TCP on the Internet. In the intervening period there has been considerable work on TCP variants, congestion signaling protocols and other technologies to address perceived problems. The latest thinking in this area was discussed with a view to uncovering any emerging consensus on what is needed now.

The objective for congestion management is to be responsive to dissimilar customer application demands. Different approaches have been tried and some clearly do not work:

- Deep Packet Inspection (DPI) everywhere raises serious privacy implications, creates single points of failure in the network if inline, requires forklift upgrades, is frequently inaccurate or misleading, and can be cripplingly expensive;

- Protocol-specific blocking or throttling motivates users to obfuscate and encrypt their communications, drives traffic to port 80, and encourages an arms race with application developers that network operators are not flexible enough to win.

In 2006 Comcast deployed an automated throttling solution for P2P traffic and as a consequence customer complaints about the quality of VoIP calls went away. At the time Comcast thought they were dealing with a congestion issue, but it now seems clear that part of the problem was related to bufferbloat. This raises an interesting question of how systemic problems in the Internet that get triggered in this way can be correctly diagnosed in future. Comcast now use two priority levels for traffic over their DOCSIS network: priority best effort and best effort. In times of congestion the heaviest users are given best effort while all other users remain on priority best effort. This congestion management system impacts less than 1% of customers and is described in RFC6057.

One of the downsides of a mechanism like that deployed by Comcast that works on timescales of several minutes is that it will penalize subscribers even when they are utilizing applications that make use of scavenger transport protocols like Ledbat.

A long-term solution to the challenges of bandwidth management on shared access networks must include a number of different elements. It must provide the best possible network experience for the broadest set of subscribers. It must enable subscribers to control their own network experience. It must enable continued Internet evolution, avoiding ongoing ‘cat and mouse’ games between application developers and network operators. And it must support a reasonable network capacity upgrade schedule that will require greater technical collaboration between content aggregators and network operators.

Fairness between users takes place at different timescales and in different parts of the network. The Comcast scheme delivers per-user fairness in a timeframe of minutes. Alternatively, an Active Queue Management (AQM) scheme like Random Early Detection (RED) provides fairness on millisecond timescales.

Oversized and under-managed buffers in parts of the access ecosystem (applications, drivers, hardware, CPE) mean that performance under load is today very sub-optimal. As already mentioned above, buffers in mobile networks are essentially infinite to guarantee no packet loss, which is a highly undesirable optimization.

So buffers may be too big, but picking a smaller size isn’t the right answer either. Fixed size buffering can never be expected to work given the huge variations in access network bandwidths and path delays. One size will never fit all and AQM is therefore essential for good performance of interactive real-time flows in the presence of sustained, throughput-maximizing background flows.

One size will never fit all and AQM is therefore essential for good performance of interactive real-time flows in the presence of sustained, throughput-maximizing background flows.

RED, dating from 1993, is one candidate mechanism for AQM that requires considerable tuning to deliver useful results. More recently, simpler alternative AQM solutions have been proposed including CoDel (Controlled Delay) and BCP (Bottleneck Congestion Policer). CoDel has a published algorithm and is implemented in recent Linux kernels as both codel and fq_codel queuing discipline variants. The fq_codel implementation combines CoDel with Stochastic Fair Queuing (SFQ) that provides most of the benefits of explicit classification without having to do the work of explicit classification. Some concern was expressed in the discussion that CoDel’s fixed parameters will inevitably be inadequate for higher speed links. This needs further exploration and testing.

The problems of per-customer traffic management (that Comcast's mechanism is trying to address) are separate to the issues of in-home traffic management (apps vs. apps, or kids vs. parents) that fq_codel tries to address. While CoDel seems well suited to help deliver bounded queue latencies at the edge of the network in CPE, BCP promises to deliver short queues and per-subscriber fairness at bottlenecks inside the broadband aggregation network.

When the discussion turned to deployment of these AQM mechanisms, the group observed that most home routers use extremely out-of-date software and provide little or no functionality for a user or operator-initiated upgrade. Tackling the problem of poor software in Customer Premises Equipment (CPE) should be a high priority for many reasons, not just dealing with poor queue management.

Separating software from hardware in the CPE business is therefore highly desirable. Creating open and modular CPE software platforms, like OpenWRT, will help to remove the CPE software as a bottleneck on progress and innovation. Openly specified testing with network operator input, coupled with certification programs could help to improve CPE quality.

Latency under load tests are potentially a very useful indicator of bufferbloat problems in a network path. There is a need to have an openly specified metric for latency under load, potentially in the IPPM Working Group at IETF. By providing operators and users with information we can support market forces and achieve incremental improvements in overall performance.

We also need better cross-layer coordination of optimization (every network layer tries to optimize independently of other layers leading to interference and sub-optimal results when viewed from a more holistic perspective).

Opportunities for collaboration

Discussion focused next on attempting to identify requirements for cross-industry collaboration on new mechanisms, practices or technology deployments to enable the continued growth and evolution of the Internet and the global deployment of innovative applications in the future.

Provide easily understandable tools to get network performance measurement results in front of consumers and into operators' minds to help shift the debate away from speed and bandwidth as the only, or even the key metric in determining Internet performance.

An overarching observation was the large number of different pieces that go to make up the Internet making it difficult to change. Figuring out how to make a focused change is hard, and even when it is clear how to proceed it can still seem overwhelming. For example, content providers and network operators don't need to make any special agreements to get connected, but there remains an unmet need for a more optimized approach. System-wide optimizations are hard given the nature of this particular highly distributed system.

Concrete next steps proposed were to:

- Provide easily understandable tools to get network performance measurement results in front of consumers and into operators' minds to help shift the debate away from speed and bandwidth as the only, or even the key metric in determining Internet performance;
- Develop a better and broader understanding of the buffer bloat problem amongst networking technologists of all kinds, and develop long-term solutions to the problem. Build awareness of bufferbloat issues amongst engineering management;
- Develop test tools for consumers, engineers, and network operators to root out bufferbloat issues, specifically including an openly specified latency under load test.

Raise ISP executives awareness of how much customer experience is impacted by the quality of CPE

Being able to identify whether latency problems are being caused inside the home or outside could be a useful first step;

- Clean up the CPE mess (bad/old software, etc.) – something akin to energy star ratings to indicate levels of conformance to a consensus-based test specification might help here;
- Develop more advanced and readily updated home routers. An industry consortium or open source CPE alliance could be constituted to tackle this problem;
- Raise ISP executives awareness of how much customer experience is impacted by the quality of CPE;
- Explore new ways to link the economic motivations of ISPs and application and content providers;
- Motivate application developers to build more bandwidth efficient applications. Encourage better application design to reduce peak load;
- Support a focused trial of relevant new protocols emerging from IETF, e.g. MPTCP, DCCP, Ledbat, Conex, to improve the situation with regard to lack of adoption;
- Document CoDel and fq_codel;
- Continue researching queuing theory;
- Encourage collaboration between network and transport experts, possibly in the form of an IETF cross-layer working group producing guidelines, gateway requirements and host requirements;
- Build measurement (instrumentation) into the components (as opposed to products) that make up the Internet as a matter of routine;
- Test against generic problematic traffic, not just current applications. It's not just about the applications running on the Internet today, but the applications that will be running on the Internet in future – we need forward-facing solutions.

Concluding remarks

It is remarkable that Internet technologies and protocols have kept working and stayed relevant despite the huge advances in link-layer technology and increases in core and access link bandwidths that we have witnessed over four decades of operation. Novel mechanisms to deliver popular, high-bandwidth content to a global audience of billions are now being widely deployed and are starting to raise new questions concerning scalability. It is clear that as the ratio of access link bandwidth to core link bandwidth approaches unity, new opportunities for latency and bandwidth management will arise. In this report we have identified a number of concrete actions that will serve to improve the quality of experience for Internet users, relieve some of the pressures on network operators, and ensure that the Internet continues to thrive as a platform for innovation and global communication on an unprecedented scale for the next forty years.

Attendees

Mike Blanche, Google

Bob Briscoe, BT

Cameron Byrne, T-Mobile

Stuart Cheshire, Apple

Alissa Cooper, Center for Democracy and Technology

Sam Crawford, SamKnows

Leslie Daigle, Internet Society

Lars Eggert, Netapp

Rob Evans, Janet

Nick Feamster, Georgia Institute of Technology

Mat Ford, Internet Society

Jim Gettys, Alcatel-Lucent Bell Labs

Jason Livingood, Comcast

Michael Menth, University of Tübingen

Remco van Mook, Equinix

Andrei Robachevsky, Internet Society

Phil Roberts, Internet Society

Andrea Soppera, BT

Nicholas Weaver, ICSI

Dennis Weller, Navigant Economics

Magnus Westerlund, Ericsson

About the Internet Society Technology Roundtable Series

The Internet Society Technology Roundtable Series are recurring events on various topics that the Internet Society organizes to facilitate shared understanding of different perspectives on Internet evolution and the development and adoption of technologies that are beneficial for the growth of the Internet as a whole. These roundtables offer a venue for the open discussion of technology issues among concerned parties to engender collaboration in ways that might not be feasible in other venues. Our experience has shown that these meetings provide an extremely valuable means of meeting and connecting with individuals from other companies who are deeply involved in the technology issues under discussion, and aid understanding of complex issues from a range of different perspectives.

About the Internet Society

The Internet Society is the trusted independent source for Internet information and thought leadership from around the world. We are the administrative home of the Internet Engineering Task Force (IETF), a standards-setting body, and the Internet Architecture Board (IAB), which provides technical advice to policymakers and various technology communities. For more than 20 years the Internet Society has supported the evolution of the Internet through technical expertise, capacity building, and innovation. Our mission is that we are dedicated to ensuring the open development, evolution, and use of the Internet for the benefit of people around the world.

Internet Society
1775 Wiehle Avenue, Suite 201
Reston, VA
20190-5108 U.S.A
Tel: +1-703-439-2120
Fax: +1-703-326-9881
http://www.internetsociety.org

Internet Society
Galerie Jean-Malbuisson 15
CH-1204 Geneva
Switzerland
Tel: +41 22 807 1444
Fax: +41 22 807 1445
Email: info@isoc.org