Managing Broadband Networks: A Policymaker's Guide

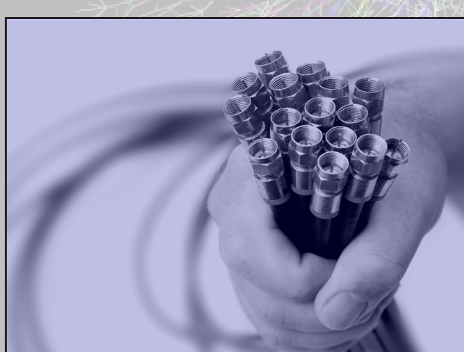

George Ou

The Information Technology and Innovation Foundation

Executive Summary

To make progress to a ubiquitous digital world, bigger pipes are not sufficient. We need not just expanded network capacity but networks that are better and more intelligently managed.

The Internet has changed the face of communications, commerce, and indeed the world. And over time the Internet itself has changed too. Until recently, most Americans at home accessed the Internet using telephone dial-up connections rather than today's faster broadband connections. With slower connections, home users limited themselves to a few basic online activities, such as email and web browsing, which perform passably well even on a slow network. In this environment, the need for Internet service providers (ISPs) to manage their networks to ensure the best possible experience for their customers was limited.

Today most Americans connect to the Internet over broadband connections that are in some cases 400 times faster than the dial-up connections of the late 1990s. But it is precisely because of these new bigger "pipes" that ISPs are finding that they need to more actively manage their networks. Broadband networks have enabled the rise of new applications, including those that need to be managed if they are to work effectively (e.g., voice over Internet Protocol, online gaming, video conferencing, and Internet Protocol-based TV) and those that can cause other applications to fail on an unmanaged network (e.g., many peer-to-peer (P2P) applications).

With this exciting transformation of the Internet into the universal communication platform of the future, network engineers face an array of daunting challenges. Specifically, to provide customers a good Internet service and operate their networks efficiently, ISPs must be able to do two very important things: 1) allocate limited bandwidth fairly among users; and, 2) apply network management tools to shape traffic from multiple applications. ISPs can and should do these things in a fair and nondiscriminatory manner. Thus, they should strive to ensure that customers who pay for the same tier of service get roughly the same bandwidth at a given level of usage, eliminate harmful variations of delay (i.e. jitter), make consumers' broadband service more conducive to using multiple applications simultaneously, while at the same time treating other applications and content fairly.

Unfortunately, network management solutions have come under heavy criticism from many advocates of "net neutrality." The issue of network management came to the fore when Comcast limited the ability of peer-to-peer (P2P) users to operate in upload-only mode whenever P2P traffic exceeded 50 percent of total upstream capacity of the entire neighborhood. More generally, the issue of network management refers to whether and to what extent ISPs can manage their networks to ensure quality of service for the majority of their customers.

Strong advocates of net neutrality argue that ISPs should have little flexibility to manage their networks and that the solution to any kinds of network congestion or other network performance challenges can and should be solved by simply adding more network capacity—primarily in the form of "bigger pipes." Indeed, they fear that using efficient network management techniques may enable network operators to abuse their power, thereby stifling free speech and civic expression and erecting unfair barriers to other companies seeking to distribute digital content or applications. Moreover, some proponents of net neutrality fear that any improvement in the efficiency of the Internet will eliminate the motivation of ISPs to expand network capacity by "building bigger pipes." As we transition to a ubiquitous digital world, bigger pipes are necessary – and public policy should support their deployment – but they are not a substitute for network management. We need to not just expand network capacity, but also build networks that are better and more intelligently managed.

Many if not most of the fears of the proponents of net neutrality stem from a lack of understanding of the history of the Internet, the economics of the ISP industry, and the science of network engineering. This guide is intended to help policymakers better understand how broadband networks and the applications that run on them work, and calls for a balanced approach to the regulation of broadband network management. A balanced approach should be based on reality: both the economic realities of building broadband networks and the scientific realities of network engineering. In addition, it should provide ISPs the flexibility they need to manage complex networks while also ensuring oversight to insure that network management practices are not being applied in anti-competitive ways.

Effective policy in this area must be based on facts. Unfortunately much of the debate over broadband network management to date has been informed more by rhetoric and emotion than by an actual examination of how advanced networks and the applications that run on them work. By providing policymakers with this guide, ITIF hopes to better inform this debate.

Key Findings and Conclusions:

- **Packet-switched networks, like the Internet, have advantages, but also disadvantages.** Packet-switched networks like the Internet were invented for their flexibility and efficiency, characteristics which are optimum for data applications. But they have two key deficiencies in the absence of network management: 1) inability to equitably allocate bandwidth; and 2) high jitter, which are essentially micro-congestion storms that last tens or hundreds of milliseconds, and which can disrupt real-time applications such as VoIP, online gaming, video conferencing, and IPTV.

- **The Internet and its predecessor ARPANET became the first adopter of packet-switching networks because it was more efficient and flexible than the circuit-switching telephone network.** Unlike telephone networks which only connected a small percentage of users at any given time, packet-switched networks allow everyone to be on the network at the same time and dynamically divide up the resources among the active users. If few users are on the network, then those users get a lot of resources allocated to them. If many users are on the network, then each user gets fewer resources but no user is locked out. This dynamic expansion and contraction of bandwidth makes packet switching networks very efficient but the allocation of bandwidth can become disproportionate whenever applications like P2P resist reallocations of bandwidth. Network management can balance the allocation of bandwidth such that each customer in the same service tier gets an equitable share of the total bandwidth.

- **Network management techniques, such as quality of service (QoS) mechanisms, make a packet-switched network more conducive to simultaneous application usage.** Network management techniques such as QoS essentially carve out virtual circuits within a packet-switched network by providing the necessary resources and performance characteristics that real-time applications need. This gives a packet-switched network the real-time characteristics of a circuit-switched network while maintaining the robustness and flexibility of a packet-switched network.

- **Even since its early days, the Internet has been a managed network.** The Internet has had basic network management mechanisms built into it since its inception, although these mechanisms have undergone and continue to undergo much refinement as usage patterns on the Internet change. Since 1987, for example, computers have used a revised version of the transmission control protocol (TCP) that includes a network congestion control mechanism developed by computer scientist Van Jacobson to slow down endpoints and prevent network meltdown.

- **Peer-to-peer (P2P) applications pose special challenges to broadband networks.** P2P users on unmanaged networks can use a disproportionately high amount of bandwidth and cause network congestion. In Japan, for example, P2P users represent 10 percent of the total broadband population but account for 65 to 90 percent of traffic on the network. By running multiple TCP flows (i.e. connections) per file transfer, P2P applications can effectively circumvent the Jacobson algorithm intended to allocate bandwidth. As a result, P2P applications can maximize the use of available bandwidth, sometimes at the expense of other applications, such as VoIP and video conferencing, which require low latency and jitter.

- **An ISP that dynamically allocates its network capacity can always offer its customers far more unguaranteed bandwidth than its guaranteed minimum level of service.** Because broadband networks are shared, it is more efficient to give consumers access to speeds that can increase when there is less congestion. Since only 1 to 10 percent of network users are active at any point in time, packet switching networks can dynamically allocate 10 to 100 times more bandwidth to each active user. If a network can be built to guarantee 1 megabit per second (Mbps) of performance for each user, for example, it can just as easily offer the customer 1 Mbps of guaranteed performance and up to 20 Mbps of unguaranteed performance. But building a network that provided a guaranteed performance of 20Mbps for example, would be much more expensive and require much higher monthly costs for the consumer.

- **One goal of network management is to fairly allocate bandwidth between paying customers.** Fairness dictates that customers who are paying for the same tier of broadband service from a broadband provider should get roughly the same bandwidth at a given level of usage. Fair bandwidth allocation shouldn't just measure instantaneous bandwidth usage, duration should also be factored in to the equitable distribution of bandwidth. If one application or one customer uses the network hundreds or thousands of times more frequently than another application or customer, it isn't unreasonable to let the short duration application or customer get a short boost in bandwidth over the long duration application or customer.

- **To achieve fair bandwidth allocations, protocol-agnostic schemes are the best solution.** ISPs can use protocol-agnostic network management systems (systems that measure the aggregate bandwidth consumption of each customer and not what protocols they are using) to ensure that bandwidth is shared fairly between customers. Early network management systems that used less accurate protocol-specific schemes to allocate bandwidth between customers worked well most of the time but experienced occasional problems. These protocol agnostic solutions are being evaluated by broadband providers. A key downside of protocol-agnostic network management systems is that they are often too expensive for smaller ISPs to deploy.

- **Another goal of network management is to better share network resources between many different applications.** Different types of applications have different network requirements. Real-time applications (e.g., VoIP) are most sensitive to network jitter. Video streaming applications (e.g., YouTube) have moderate fixed bandwidth requirements and moderate jitter tolerance. Interactive applications (e.g., web browsing) have brief bursts in bandwidth that could disrupt real-time or streaming applications. Background applications (e.g., P2P applications) are designed to be unattended with no one waiting for an instant response.

- **Packets should be ordered logically with priority given to real-time applications first, streaming applications second, interactive applications third, and background applications last.** In order for all applications efficiently and fairly share an Internet connection, those with higher duration and higher bandwidth consumption (e.g., P2P) are given lower priority than applications with lower duration and lower bandwidth consumption (e.g., VoIP applications). This does not mean P2P applications are being mistreated because they still receive the highest average bandwidth from the network.

- **To better enable multiple applications to share an Internet connection, protocol-specific schemes are necessary.** Application protocols that require low packet delay must be identified and must be protected against high variations in packet delay (e.g., jitter) and Quality of Service network management techniques are the mechanism that provides that protection.

- **Wireless networks require more management than wired networks.** Wireless networks require more network management than wired networks because they have less bandwidth available and it must be shared more frequently. Furthermore, multiple radio transmitters sharing the same wireless frequency in the same geographic location results in a high probability of radio interference which can bring networks to a halt. These unique challenges of wireless networks require the most elaborate network management system of all in the form of a centralized scheduler which coordinates the transmission slots for network users as tightly and efficiently as possible without collision.

- **Wireless network management enables innovation.** Intelligent wireless networks will ultimately spur more adoption and usage of wireless broadband, which facilitates more mobile e-commerce and enables more innovation and generation of wealth.

Responding to Common Misperceptions About Network Management:

- **Network management techniques, such as QoS, do not put low priority applications on a "dirt road."** QoS gives higher prioritization to applications that have lower bandwidth, lower duration, and higher sensitivity to packet delay. In spite of this, applications that are given the least priority still end up receiving the highest average bandwidth from the network. But with this logical prioritization scheme in place, low priority applications like P2P applications interfere less with other applications sharing the same network. This in turn allows P2P applications to operate freely without any artificial constraints on when to use them or how much bandwidth to allocate to them which are commonly used on unmanaged networks.

- **Building more bandwidth, while desirable, does not eliminate the need for network management.** Advancing the digital economy requires higher speed broadband. However, higher speed networks will not preclude the need for network management. First, as network capacity grows, network demand also grows, as new kinds of applications emerge to take advantage of the capacity. Second, networks with plenty of spare unused capacity on average can still suffer instantaneous shortages at peak times of the day. Third, networks operating at low utilization levels can still suffer packet delay in the form of jitter.

- **Metered pricing and usage caps alone will not solve the problem of network congestion.** Metered pricing and bandwidth usage caps are legitimate tools for ensuring the efficient use of networks, but they cannot control instantaneous bursts in demand nor can they deal with the problem of jitter and the inability of dumb networks to gracefully support multiple applications. Only advanced network management techniques like quality of service can deal with these challenges.

Policy Implications:

- **Legislation and regulations should not limit efforts by ISPs to fairly use network management to overcome technical challenges and maintain a high quality Internet service for their customers.** As described in this report, ISPs face many technical challenges to manage network congestion and support various online applications. Network management is a necessary and important component of broadband networks, and policymakers should support its use. However, this freedom to manage the network is not a license for ISPs to behave in anti-competitive ways such as blocking legitimate websites or unreasonably degrading services that users have paid to access. Neither should ISPs unreasonably discriminate against any content or service on the open Internet.

- **Policymakers should be cognizant of the effects of certain proposed legislation on the use of network management.** Some proposed net neutrality bills ban differentiated pricing for enhanced QoS and would have undesirable and unintended consequences. One intent of these bills is to facilitate more open Internet bandwidth for broadband consumers, but the result may be just the opposite. Not allowing network operators to prioritize their own IPTV content above other Internet content, for example, will simply push those cable TV-like services onto private circuits that share the same physical network. That would result in less Internet bandwidth being available on a permanent basis for broadband consumers even when they are not using their IPTV service.

- **The federal government has a key role to ensure openness and fair play on the Internet.** However, it should do this with sensible rules. Policies should strive to prevent any potential abuse without eliminating the ability of ISPs to manage their networks in ways that produce the best possible user experience for the largest number of users, and without eliminating incentives to build the next generation broadband network. Toward that end the FCC should oversee broadband providers and ensure that they ISP network management practices are open, transparent and not harmful to competition. And the ISP industry should continue its efforts to develop and abide by industry codes of good conduct regarding network management that include, but are not limited to, fuller and more transparent disclosure to consumers of network management practices.

Conclusion:

The Internet in all its glory has never had a perfect architecture. There have always been conflicts between users and applications competing for scarce network resources. Network management is necessary to fairly allocate bandwidth between customers and seamlessly support multiple applications on shared network connections.

The Internet and broadband technology are continuing to evolve at a fairly rapid rate, and neither shows any signs of maturing. Network engineers continue to find new solutions to improve the Internet experience for all users. This situation makes it very difficult, if not impossible, to predict where the market and technology will evolve. The Internet is so valuable precisely because it is open to anyone, for any use, and for any business model, but participation has always required varying levels of payment for varying levels of service between willing parties. Given this environment, it is best for policymakers not to issue blanket prohibitions on network management technology and existing business models. Instead, policies should focus on creating better transparency for all Internet companies along with FCC oversight to ensure that broadband providers are managing networks in ways that are not unfair or anticompetitive.

ABOUT THE AUTHOR

George Ou is a Senior Analyst at the Information Technology and Innovation Foundation and works out of Silicon Valley. Outside of ITIF, Mr. Ou is an Information Technology and CISSP Security Consultant who was the founder of ForMortals.com. He recently served two years as Technical Director and Editor at Large for TechRepublic and ZDNet (both property of CNET Networks) doing in-depth coverage on IT and technology topics. Before journalism, he worked as an IT professional who designed and built wired network, wireless network, Internet, storage, security, and server infrastructure for various Fortune 100 companies.

ABOUT THE INFORMATION TECHNOLOGY AND INNOVATION FOUNDATION

The Information Technology and Innovation Foundation (ITIF) is a nonprofit, non-partisan public policy think tank committed to articulating and advancing a pro-productivity, pro-innovation and pro-technology public policy agenda internationally, in Washington and in the states. Through its research, policy proposals, and commentary, ITIF is working to advance and support public policies that boost innovation, e-transformation and productivity.

ACKNOWLEDGMENTS

The author would like to thank Bob Briscoe, Brett Glass, and Richard Bennett for their technical insights. In addition, the author would like to thank Daniel Castro, Kerry Kemp, Priscilla Jang, and Rob Atkinson for editorial efforts.