



- 3 Congestion is a *capacity phenomena* in terms of Megabit per second use, not "excess" use of volume quantity in Gigabytes over time.
- 4 Comcast controls congestion by arbitrary, non-neutral restriction on *volume use by particular subscriber, content or protocol* of Gigabyte use over time, ignoring excess use of *capacity from all sources of peak usage*.
- 5 Comcast uses vague contract language in its terms of service designed to *avoid specific commitments of dedicated, non-transferable connection capacity* to individual subscribers *after* capacity has been marketed and sold to its subscribers as fully available.
- 6 Comcast has blamed incorrectly, specific subscribers, content or protocol as the *exclusive* "cause" of congestion in peak periods simply for using the connection capacity *as marketed, encouraged, sold, maintained and made explicitly available by Comcast*.
- 7 Reasonable network management of congestion should be a routine and essential role for *capacity availability*, administered in net neutral fashion.
- 8 Congestion management does not require, and should not include management of users, applications, content, protocol or volume flow quantities over time *by identity or size*, which constitutes violations of net neutrality.
- 9 Conventional flat monthly rates for specific amounts of dedicated, non-transferable tiers of connection capacity which do not exceed network capacity *would avoid chronic congestion by definition in the absence of oversold capacity*.
- 10 Comcast should admit it has oversold network capacity and be required to downgrade it at the connection level in net-neutral proportion to that oversold to match total network capacity.

11 If Comcast is allowed to continue overselling non-dedicated capacity which causes congestion, it should be required to offer net neutral *peak and off-peak* rates or undertake non-price neutral rationing of capacity by the wide variety of neutral technical means available to control congestion in peak periods.

### Network Congestion

Chronic network congestion is a *predictable* consequence of insufficient bandwidth speed at the network level necessary to provide "firm" service to internet subscribers during peak periods, measured in Megabits per second.<sup>1</sup>

Chronic congestion on a node-wide or network-wide basis appears as degraded or interrupted *capacity* service to *all subscribers during peak periods*.

Chronic congestion is the combined result of Comcast overselling (overstating) its "firm" available capacity to subscribers at the level of individual connections, which used by subscribers *as sold, maintained and made explicitly available by Comcast*, exceeds the capacity of the network to provide "firm" service.

When sufficient levels of bandwidth use occur at available connection capacity *as allowed by Comcast*, congestion occurs on the network *because there is not enough capacity at the network level to meet and service that simultaneous use in peak periods as measured in Megabits per second*.

Contrary to claims by Comcast, network congestion is not "caused" exclusively by a small number of large content users, labeled as

---

<sup>1</sup> Comcast definitions of service are vague at best, allowing for wide variations in bandwidth speed and service quality at its discretion. In this context, "firm" service is generally whatever subscribers have come to expect as "normal" service in terms of bandwidth capacity, latency and jitter, from which "congestion" constitutes a degradation, particularly in the form of lower bandwidth speed.

Chronic "predictable" congestion refers to those degradations or service interruptions not associated with emergencies or maintenance, instead caused by regular subscriber use in the aggregate.

"internet hogs", *in isolation from other subscriber use during peak periods.*

Further, congestion is not "caused" by creative use of network capacity through multiple paths among end users, *who cannot possibly use more total network capacity than is made available by Comcast at the connection level.*

Comcast responds to congestion by blaming it on *particular* subscribers, content or protocol, claiming that the *particular* traffic must be shaped, interrupted or otherwise controlled *for the claimed benefit of all (remaining) subscribers, despite that congestion results from all sources of peak use.*<sup>2</sup>

### Reasonable Network Management of Congestion

Because of oversold network capacity, individual subscribers do not have service in terms of *dedicated increments of clarified, firm, non-transferable capacity.*

Comcast should admit it has oversold capacity and be required to correct the problem by downgrading and dedicating connection capacity in proportionate, neutral fashion until it matches network capacity.

Absent this correction, attempts to control congestion by pricing connection capacity will be useless by definition and will continue to encourage congestion as subscribers attempt to use connection capacity as purchased *and made explicitly available by Comcast.*<sup>3</sup>

If the network is allowed to remain in a chronic state of oversold capacity, the only net-neutral solutions to control the resulting

---

<sup>2</sup> The incentive and actions taken to block competitor content by Comcast are not specifically addressed in these comments beyond the broader context of doing so to "control" congestion as claimed by Comcast. Comcast also controls congestion in other non-neutral ways. Congestion *per se* has become a major, general problem for net neutrality and reasonable network management, and requires far more correction than mere case-by-case restraint of manipulating specific content on a pretext of controlling congestion.

<sup>3</sup> Technically, Comcast could simply cap individual connection capacity to avoid congestion, but doing so would betray that it was oversold in the first place.

congestion are to ration capacity use for *all users - not just "internet hogs"* in peak periods with non-price technical means or administer explicit pricing of peak versus off-peak use of capacity.

Current connection rates are already designed to control congestion. Typical retail internet service is currently sold in bandwidth tiers to subscribers in terms of maximum flow rates, or maximum capacity in terms of Megabits per second for a flat monthly rate.

*Absent oversold capacity, existing flat rates effectively control congestion because they also serve as maximum peak rates for dedicated capacity.* Under this condition, total capacity use at the connection level cannot exceed total network capacity by definition.

Where typical terms and conditions of internet service should relegate "congestion" to limited exceptions under "good faith" and "best effort" clauses, Comcast has abused these exceptions to "firm" service by expanding, encouraging and allowing congestion *among all subscribers during peak periods* through deceptive marketing practices.

After the capacity is sold, then Comcast assigns dynamically, available *unused* capacity from *some* subscribers to *other* subscribers in arbitrary fashion, particularly an excess of *new* subscribers encouraged by a policy of oversold capacity on the front end combined with a policy of severe use restriction on the back end.

While technically legal under vague contract language, this practice can no longer offset the problem of emerging, routine, network-wide congestion caused in general by oversold capacity.

As the use per connection increases in peak periods, Comcast has been forced to intervene and elected to prevent node-wide or network-wide service degradation or interruption with discriminatory restrictions of peak use for selected subscribers, content or protocol.<sup>4</sup>

---

<sup>4</sup> The presumption is Comcast is controlling congestion, whether or not it is also blocking competitor content in the process.

By simultaneously refusing to commit to specific amounts of sold, available connection capacity while insisting on managing an artificially created shortage of total available network capacity, Comcast places itself in a *de facto* position of operating a network *intentionally designed* to create congestion on an ongoing basis.

The Myth of "Internet Hogs": Congestion is neither Caused nor Controlled by Particular Volume Use of Gigabytes over Time

"Internet hog" is a propaganda term used to undermine net neutrality with claims that "excessive use" cannot be managed under net neutral conditions to the detriment of "ordinary" users.

If fact, just the opposite is the case. Because "hog use" is defined as large *volume use* rather than large *capacity use*, controlling and limiting "hog" use is a convenient means of systematically overselling capacity at the expense of imposing peak congestion on *all* subscribers, large and small.

Specifically, when subscribers use capacity as sold *in net neutral fashion to generate "large" amounts of Gigabytes, Comcast declares the use as "excessive" to justify forced reductions of volume use, service degradation or interruption, or cancellation of service altogether.*

In other words, a practice of bait-and-switch allows Comcast to market and sell *explicit* amounts of *connection capacity* in terms of Megabits per second, then *reduce it substantially* by limiting greatly the volume of Gigabytes it could provide over time to subscribers.

Comcast applies the concept of "hog use" in its terms of service by using *average Gigabyte volume used across subscribers* as a standard against which *excessive use* is penalized.

On its face, this flatly violates the stated *available maximum volume use* marketed and sold to subscribers in terms of the total potential Gigabytes that *could* be provided by the respective connection at the particular tier of rated capacity in terms of bandwidth speed.

For example, if a connection of 4 Megabits per second *as sold, maintained and made explicitly available by Comcast* is used *at maximum capacity* continuously on a 24/7 basis for one month, then total volume use would be 1,296 Gigabytes.<sup>5</sup>

*Yet Comcast warns subscribers in its terms of service that total use over 2 Gigabytes is "excessive", which is less than 1/4 of one percent of the potential total Gigabytes provided on a sustained basis by a connection capacity of 4 Megabits per second.*

This compares to driving a rental car for one mile a day instead of a hundred miles, using one gallon of water a day instead of a hundred gallons, using 1 kilowatt hour of electricity a day instead of a hundred kilowatt hours and so on.

In each case, the amount of *unused idle connection* capacity is absurdly high, compared to the amount actually used over time due to severe restrictions on total *volume* use.

Further, restrictions on volume use *have no particular causal relation to peak use and congestion* beyond coincidence. Whether in *small or large "amounts of content"*, *volume use of any size can cause congestion in peak periods by definition*, just as in off-peak periods, *volume use causes no congestion regardless of size.*

The proposition that "internet hogs" could possibly *cause* congestion *in isolation from other users during peak periods is a myth.*

If a large number of water, electric or internet subscribers used their service *simultaneously for ten minutes* in a peak period, it could result in congestion from *"small use" per subscriber in terms of volume but "high use" in terms of peak aggregate network capacity.*

By singling out and controlling particular sources of high volume use to control congestion claimed "caused" by these sources, Comcast is

---

<sup>5</sup> One byte consist of 8 bits, where capacity is usually reported in bits and volume usually reported in bytes. For example, a 4 Megabit per second connection is roughly equivalent to a .5 Megabyte per second connection, which if used at maximum on a sustained basis provides 1,296 Gigabytes (.5 Megabyte x 60 seconds x 60 minutes x 24 hours x 30 days = 1,296,000 Megabytes = 1,296 Gigabytes)

giving a free ride to all other users during peak periods who contribute *equally* to congestion in terms of Megabit per second use.

When Comcast uses *high volume* identity to control congestion, it is controlling congestion *de facto* by *content discrimination* and in the process *overpenalizing* that particular content source for congestion caused in relation to other peak users.

The current pattern of heavy *volume* use by a relatively small number of subscribers in no way justifies actions under practices of net neutral, reasonable network management by Comcast to restrict only that particular use to control congestion.

Comcast is exploiting the coincidence of high volume use by these particular subscribers, content or protocol with peak use and congestion, particularly given its technically, observable prominence within aggregate peak use.

In context of the example above for peak use during a ten-minute period, Comcast would dip into aggregate peak use and single out *only* say, "large" users of two hours or four hours over the peak period without disturbing "smaller" users who would cause *identical* congestion during the ten-minute peak period at an equivalent bandwidth speed.

Restricting *only use by volumetric means* sets exactly the wrong, non-neutral standard going forward for increased use per subscriber by *all* subscribers *in peak periods up to the levels of assigned connection capacity*.<sup>6</sup>

Comcast has reached a tipping point of uncontrollable congestion caused by its flawed policy to oversell capacity and can no longer control it absent a major overhaul of its "reasonable management practices", absent increasingly flagrant violations of net neutrality.

---

<sup>6</sup> This emphasizes the obvious question raised earlier of why Comcast refuses to cap individual connection capacity in order to control congestion rather than explicitly making it available and then penalizing selected subscribers for using "too much of it" in terms of Gigabyte volume.

Any rulemaking by the Federal Communications Commission on the matter of reasonable network management practice of broadband practices should include the following:

- Establishment of clear definitions and sources of congestion as a capacity phenomena;
- Practices required or not allowed in the control of predictable congestion from any source in net neutral fashion;
- Clear language that specifically excludes the possibility of controlling or managing congestion by identity or size of subscribers, users, content, protocol or application;
- In the affirmative, congestion should be controlled through bandwidth *at or below the packet* level by net neutral technical rationing or pricing means, independent of Gigabyte volume;<sup>7</sup>
- Specific solutions for congestion caused by oversold capacity, including whether oversold capacity should be allowed at all;
- Clear terms and conditions of service designed to match what service is marketed and sold in terms of available capacity in Megabits per second versus available use of Gigabyte volume over time;.
- For volume restrictions applied to Gigabyte use, committed, capped, maximum amounts by time interval, tied to specific amounts of available connection capacity;<sup>8</sup>

---

<sup>7</sup> This would not restrict technical methods designed to discard certain packets which remain in the network for lack of retrieval after a designated amount of time, since they can contribute to congestion of more recent arrived packets. Controlling packet congestion a time-identity basis rather than a volume-identity basis would be consistent with net neutrality.

<sup>8</sup> Technically on a stand-alone basis, metered pricing of volume use over time in uniform fashion among all customers would be consistent with conventional standards of net neutrality to the extent that specific content is not subject to separate treatment *within each pricing package*. The problem arises across the pricing packages for which the price per Gigabyte can vary greatly and by definition, favor certain content over others by size alone.

- Clear definitions and distinctions for terms and conditions of service between "firm", best-effort capacity and all other versions of degraded, interruptible or transferable service, committed to in marketing and at the point of sale.<sup>9</sup>

### About the Author

Dr. Payne has worked in public utility regulation for twenty years including state commissions and agencies in Florida, Minnesota, Illinois, North Dakota and Indiana.

He has sponsored substantial expert testimony in these state jurisdictions and also worked as a senior economist at the Federal Communications Commission and taught telecommunications at the University of Ohio.

---

<sup>9</sup> Includes "premium firm service" where latency and jitter is held to a higher standard for certain types of services.