

LARGE-SCALE INTERNET MEASUREMENTS FOR DATA-DRIVEN PUBLIC POLICY

Henning Schulzrinne

(+ Walter Johnston & James Miller)

FCC & Columbia University

Overview

- Why measure?
- Results of FCC MBA 2011 and 2012
- The challenge of mobile measurements
- Baking measurements into the infrastructure

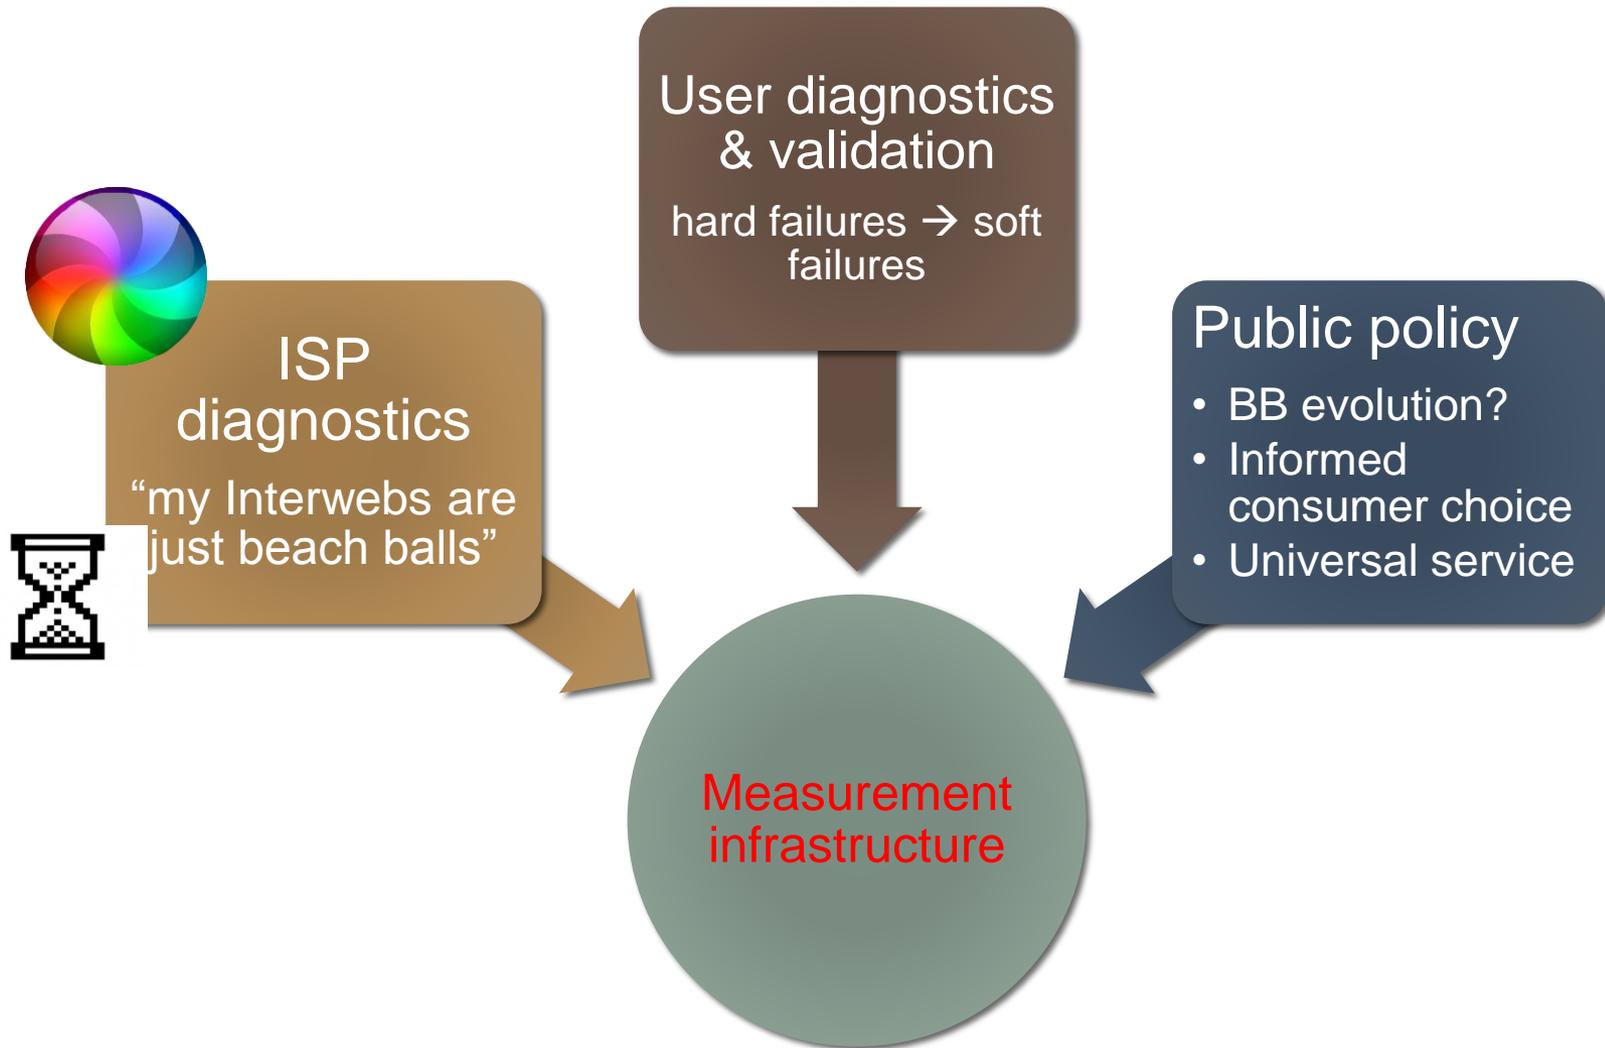
Why measure?

- Whole conferences on measurement
 - often driven by lost-key method
- Long history of IETF measurement work
 - RFC 2330: 1998
 - focused on specific metrics, not so much overall architecture
- Performance measurements don't work well
 - Individual consumers:
 - not statistically reliable (geek-heavy)
 - no ground truth data
 - confounding factors (e.g., home networks)
 - data generally not published → hard to compare metrics
 - Service providers:
 - limited infrastructure → hard to scale

FCC measurement history

- FCC has acquired and analyze data on legacy PSTN
- More recent and evolving broadband interest
 - Section 706 of 1996 Telecommunications Act → annual report on availability of advanced telecommunications services to all Americans
 - Resulted in information on deployment of broadband technology (“Form 477”)
 - but not its performance
 - FCC’s National Broadband Plan – March 2010
 - Proposed performance measurements of broadband services delivered to consumer households
 - Work plan evolved from recommendations of National Broadband Plan

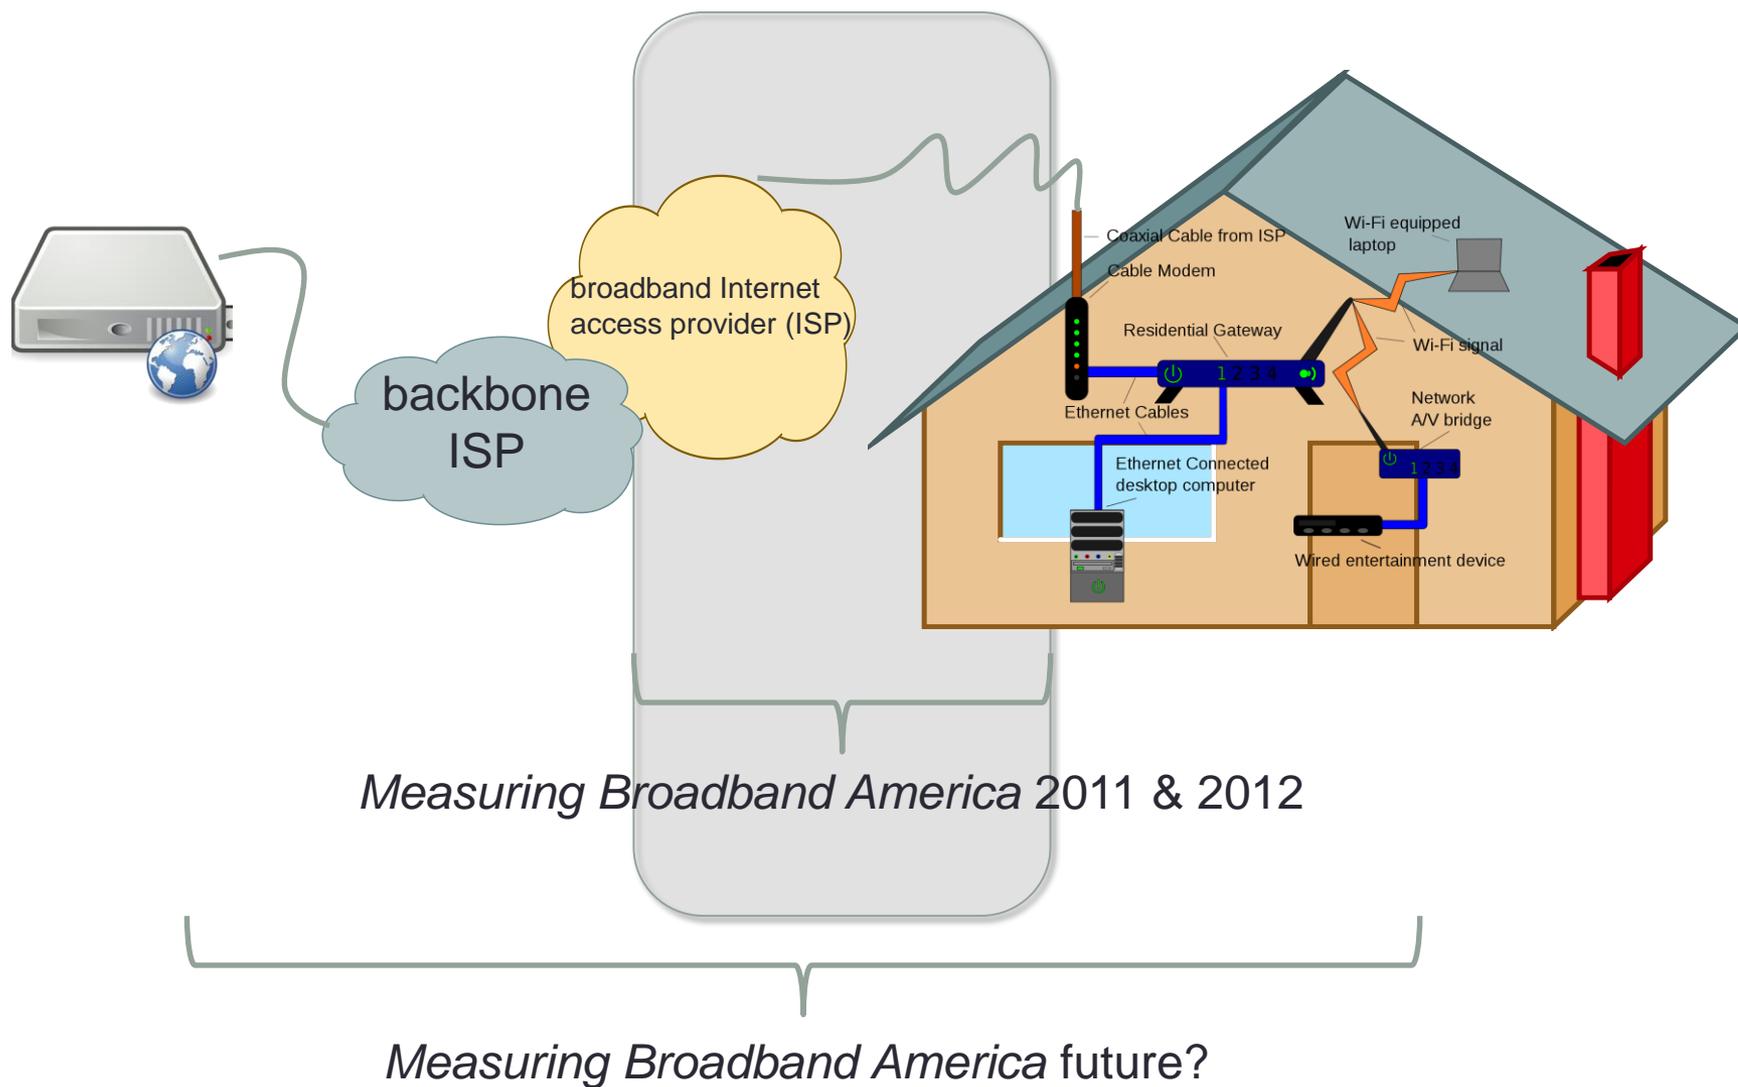
The role of network measurements



Principles

- The FCC Measuring Broadband America program is based on principles of openness, transparency and partnership with diverse stakeholders.
- We are committed to:
 - Ensuring that commonly accepted principles of *scientific research*, *good engineering practices*, and *transparency* guide the program;
 - Encouraging *collaboration* of industry, academia and government;
 - ***Publishing the comprehensive technical methodology*** used to collect the data, including the source code for the tests as open source;
 - ***Releasing data*** used to produce each report coincident with the report's release, and releasing all data for each collection cycle within one year of collection.

Measurement architecture



The MBA project - logistics

- Enlisted cooperation:
 - 13 ISPs covering 86% of US population
 - vendors, trade groups, universities and consumer groups
- Reached agreement reached on what to measure and how to measure it
- Enrolled roughly 9,000 consumers as participants
 - 6,800 (7,782) active during March 2011 (April 2012)
 - A total of 9,000 active over the data collection period

What was measured

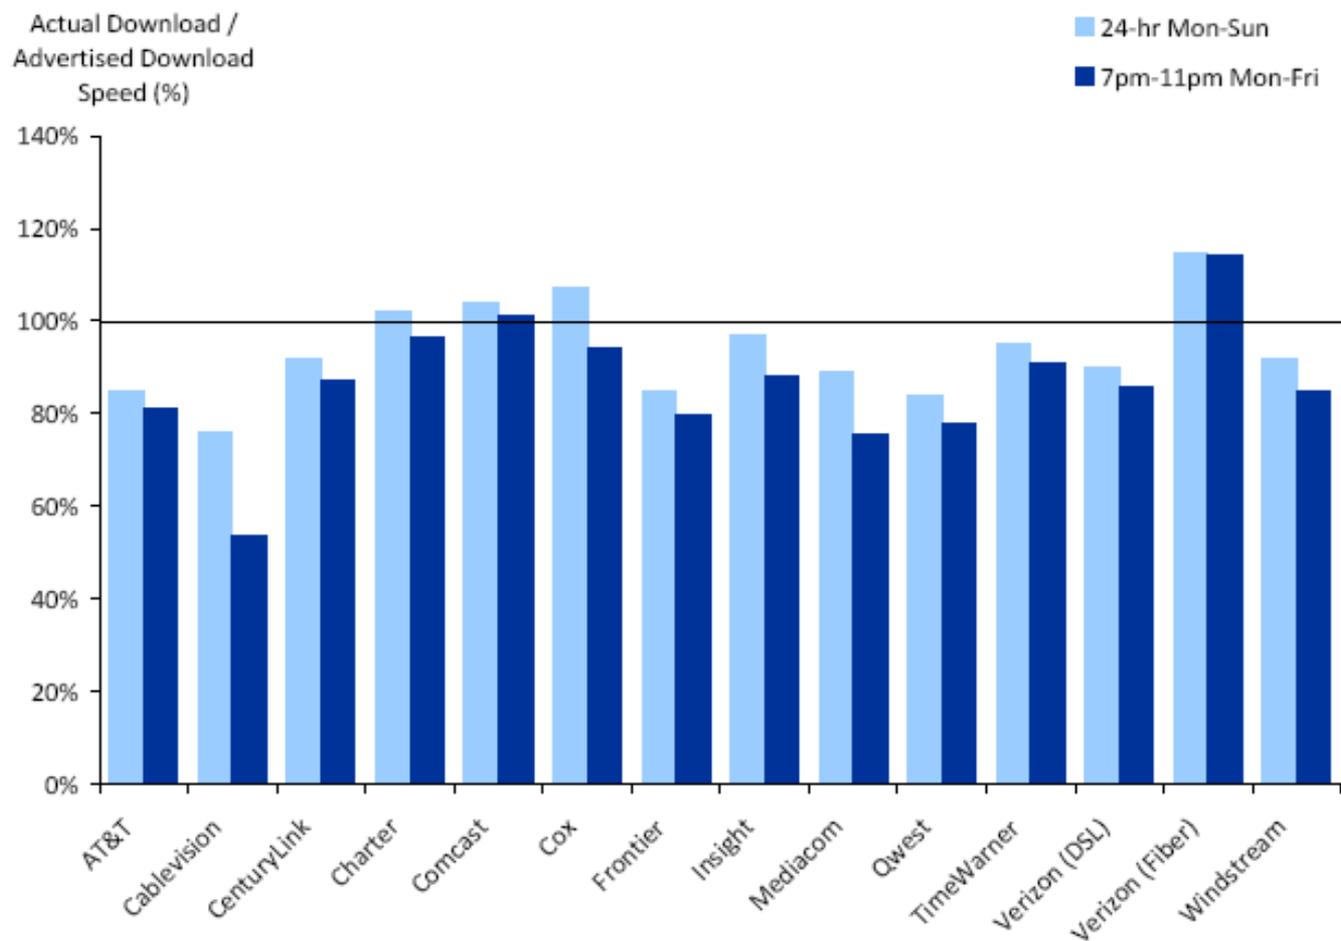
| | |
|------------------------|-------------------------|
| Sustained Download | Burst Download |
| Sustained Upload | Burst Upload |
| Web Browsing Download | UDP Latency |
| UDP Packet Loss | Video Streaming Measure |
| VoIP Measure | DNS Resolution |
| DNS Failures | ICMP Latency |
| ICMP Packet Loss | Latency Under Load |
| Total Bytes Downloaded | Total Bytes Uploaded |

What was released

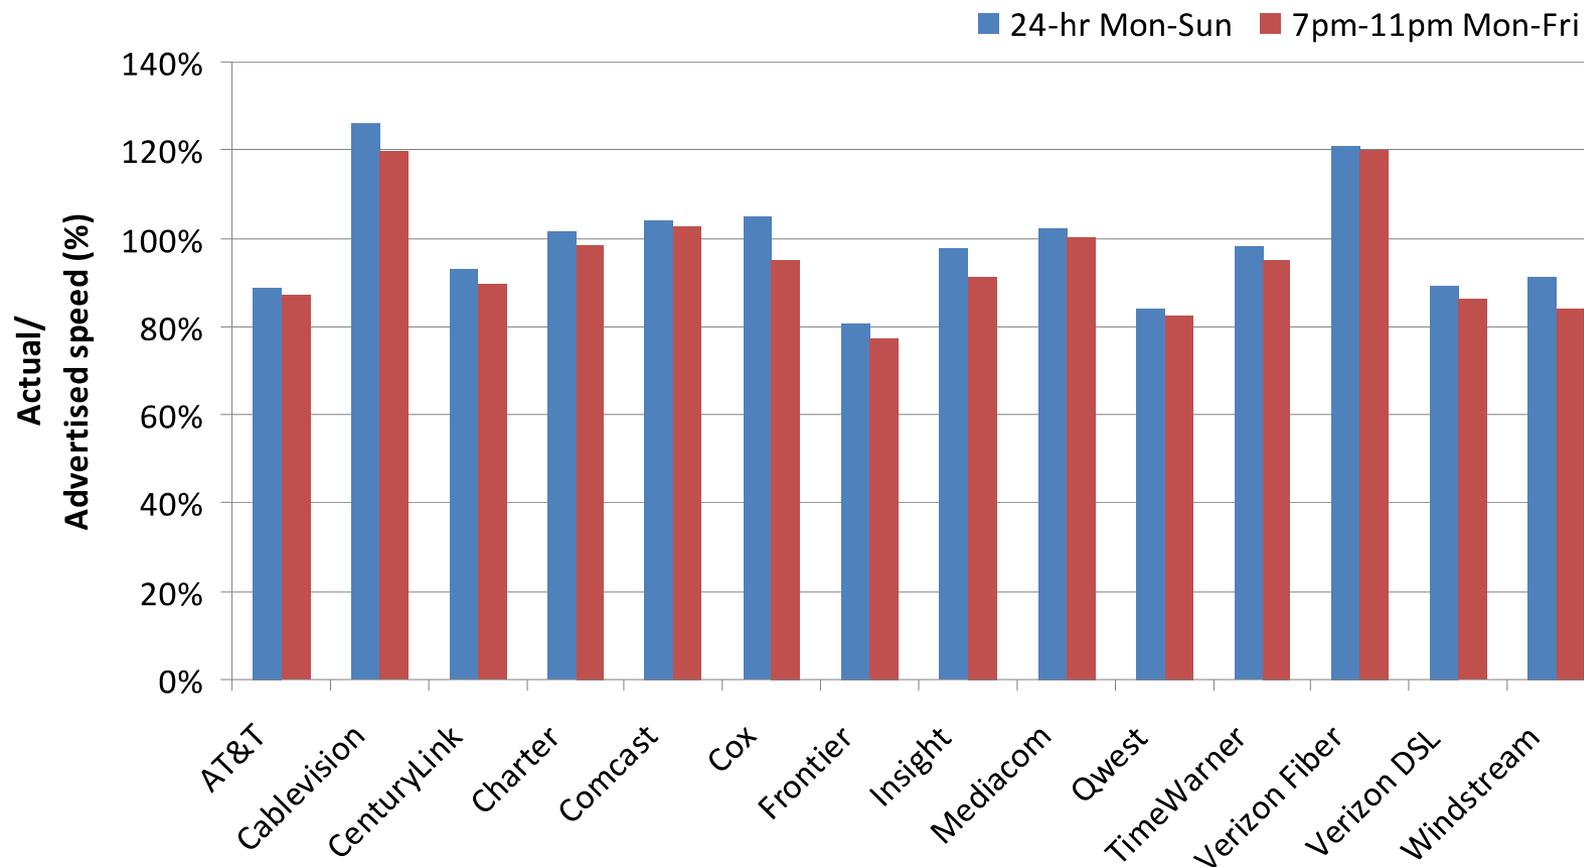
- *Measuring Broadband America* reports
 - Main section describing conclusions and major results
 - Technical appendix describing tests and survey methodology
- Spreadsheet providing standard statistical measures of all tests for all ISPs and speed tiers measured
- Report period data set with 4B data elements from over 100M tests
 - Data set presented as used with anomalies removed
 - Documentation provided on how data set was processed
 - All data, as recorded
- Geocoded data on test points recently released
- Information available at <http://www.fcc.gov/measuring-broadband-america>

2011: Most ISPs deliver close to advertised during peak hours

Chart 1: Average peak period and 24-hour sustained download speeds as a percentage of advertised, by provider

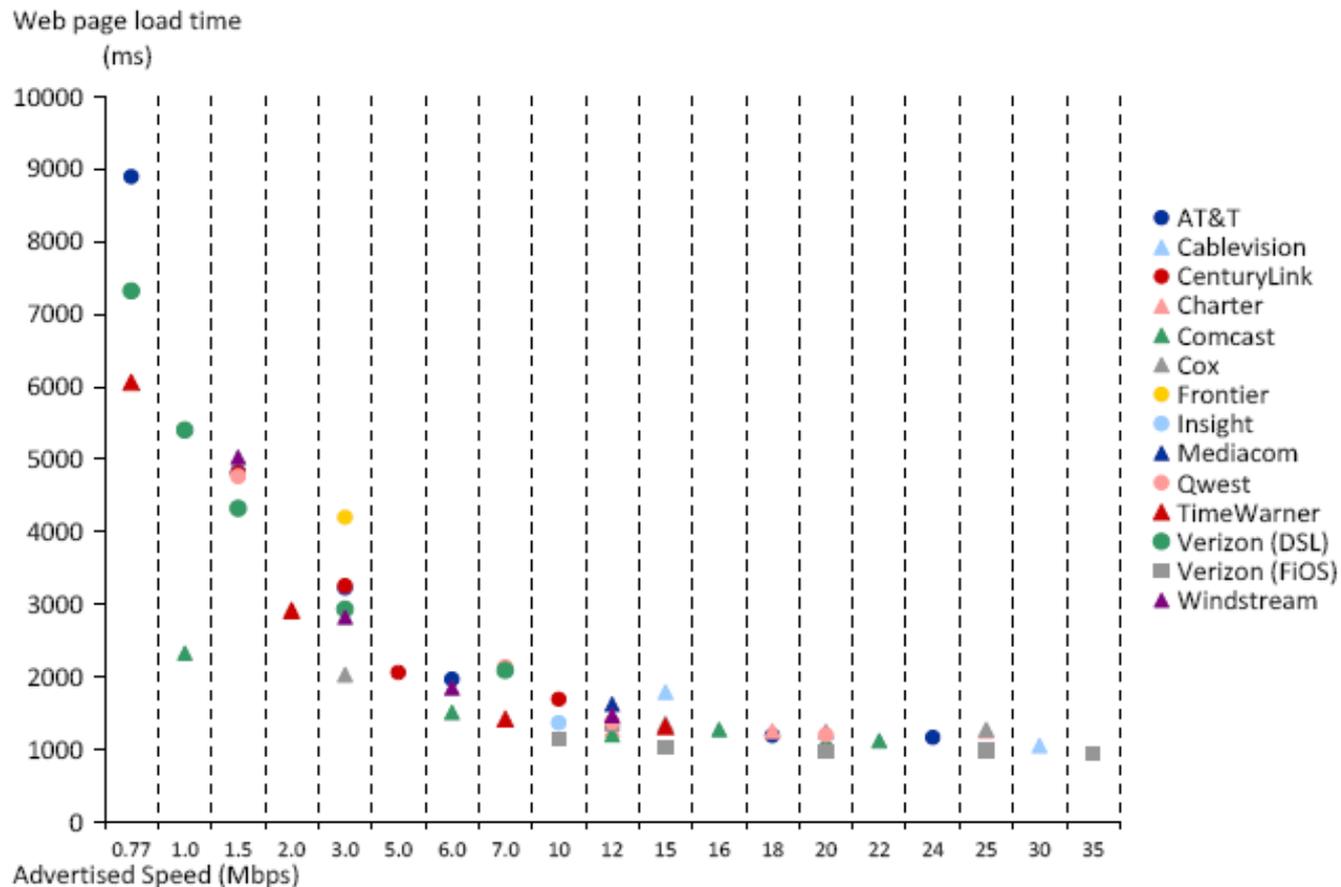


2012: You improve what you measure...



Web page downloading

Chart 10: Web loading time by advertised speed, by technology



The Internet is not a series of (fixed-width) tubes

- Some cable companies advertise burst speed
 - Quota based technique providing temporary speed increase of < 15 seconds
 - Also affected by other household activity
 - Can't be applied generally to DSL where sync rate often limiting factor
 - Marginal value to fiber where each subscriber has potentially available 37 Mb/s to 75 Mb/s provisioned bandwidth
 - → **Links are no longer constant-size bit pipes**
- Measured both burst and sustained speed



Broadband 2012

- Deployment
 - USF: Connect America Fund
- Performance
 - Measuring Broadband America
 - mobile project announced
- Significant progress:
 - wider availability of 100 Mb/s
 - fiber available to 46 million homes (FiOS, U-verse)
 - community/non-traditional broadband (Chattanooga, Kansas City)
 - LTE networks

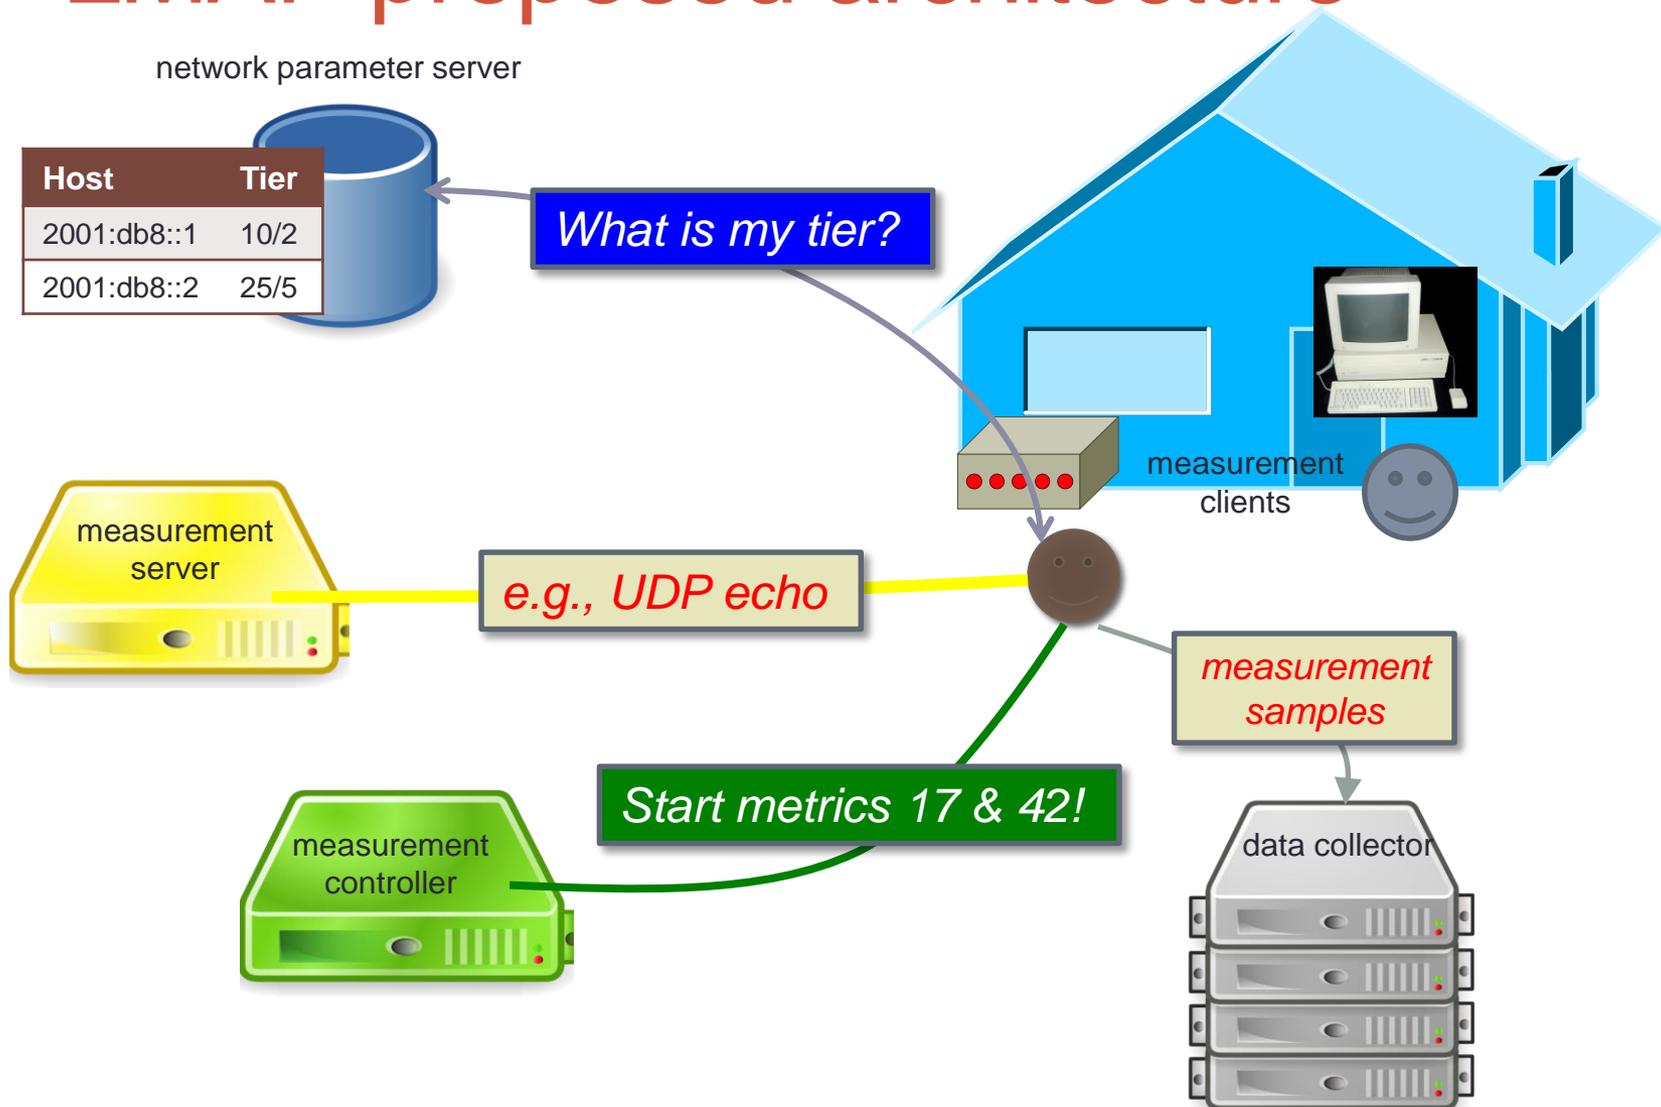
Mobile performance

- Announced effort with 4 largest wireless providers
- Options available
 - structured drive testing → expensive
 - semi-formal drive testing → snapshot, limited coverage
 - passive measurements → privacy concerns, limited sampling
- → **Smartphone app + volunteers + existing infrastructure**
- Challenges:
 - how to capture variation in time, space and device?
 - how to ensure location privacy?
 - impact of bandwidth usage for metered & capped plans?

What can't we measure?

- Small providers
 - may be more infrastructure- & engineering-challenged
- Anchor institutions
- Network reliability
 - large-scale outages gathered by “Part 4” outage reporting, but not public
 - small-scale outages?
- Network features
 - Who has access to DNSsec? IPv6? Which ports are being blocked?
- Which country has the cheapest broadband?
 - And why?
- What drives consumer adoption of higher speeds & new applications?
 - Speed → applications OR applications → speed?
- Why does 1/3 of US not use the Internet?
 - Relative importance of availability, affordability, relevance?

LMAP proposed architecture

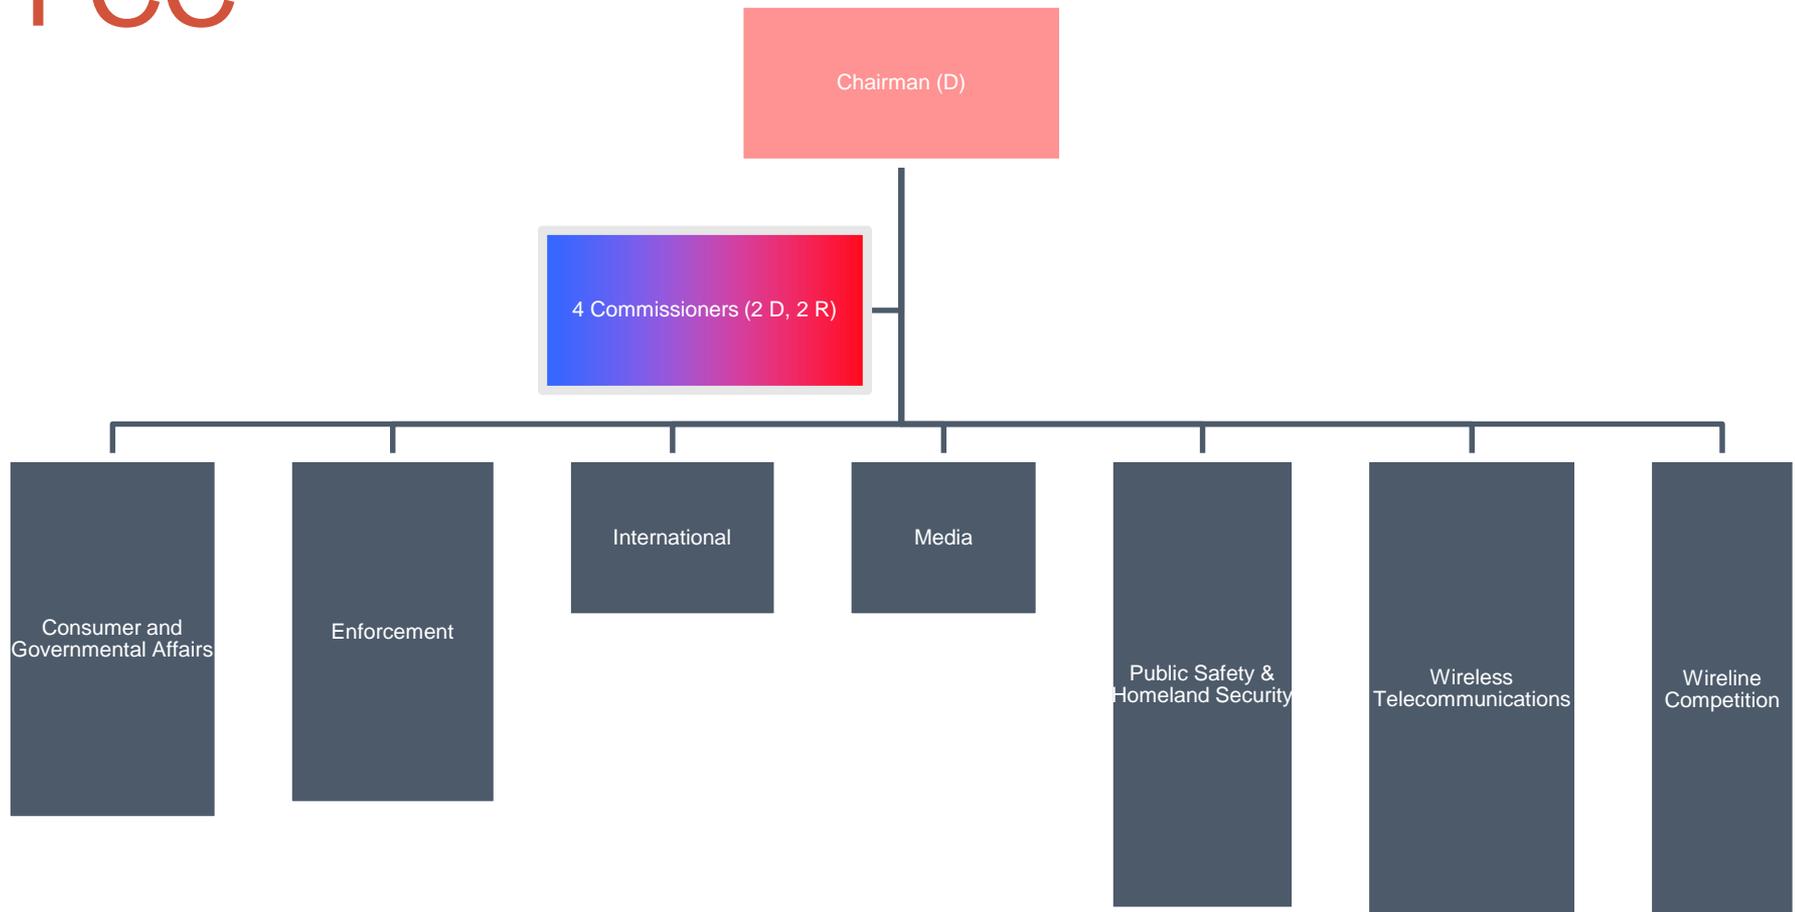


Conclusion

- Measurements: from sometime experiment and LAN to built-in capability
- Good telecom policy needs good data
 - not just counting lines
 - PSTN transition to IP → there is no second network
- Re-use measurement for three purposes:
 - ISP diagnostics and planning
 - Consumer diagnostics
 - Public policy data gathering
- Dear IETF: We need your help!

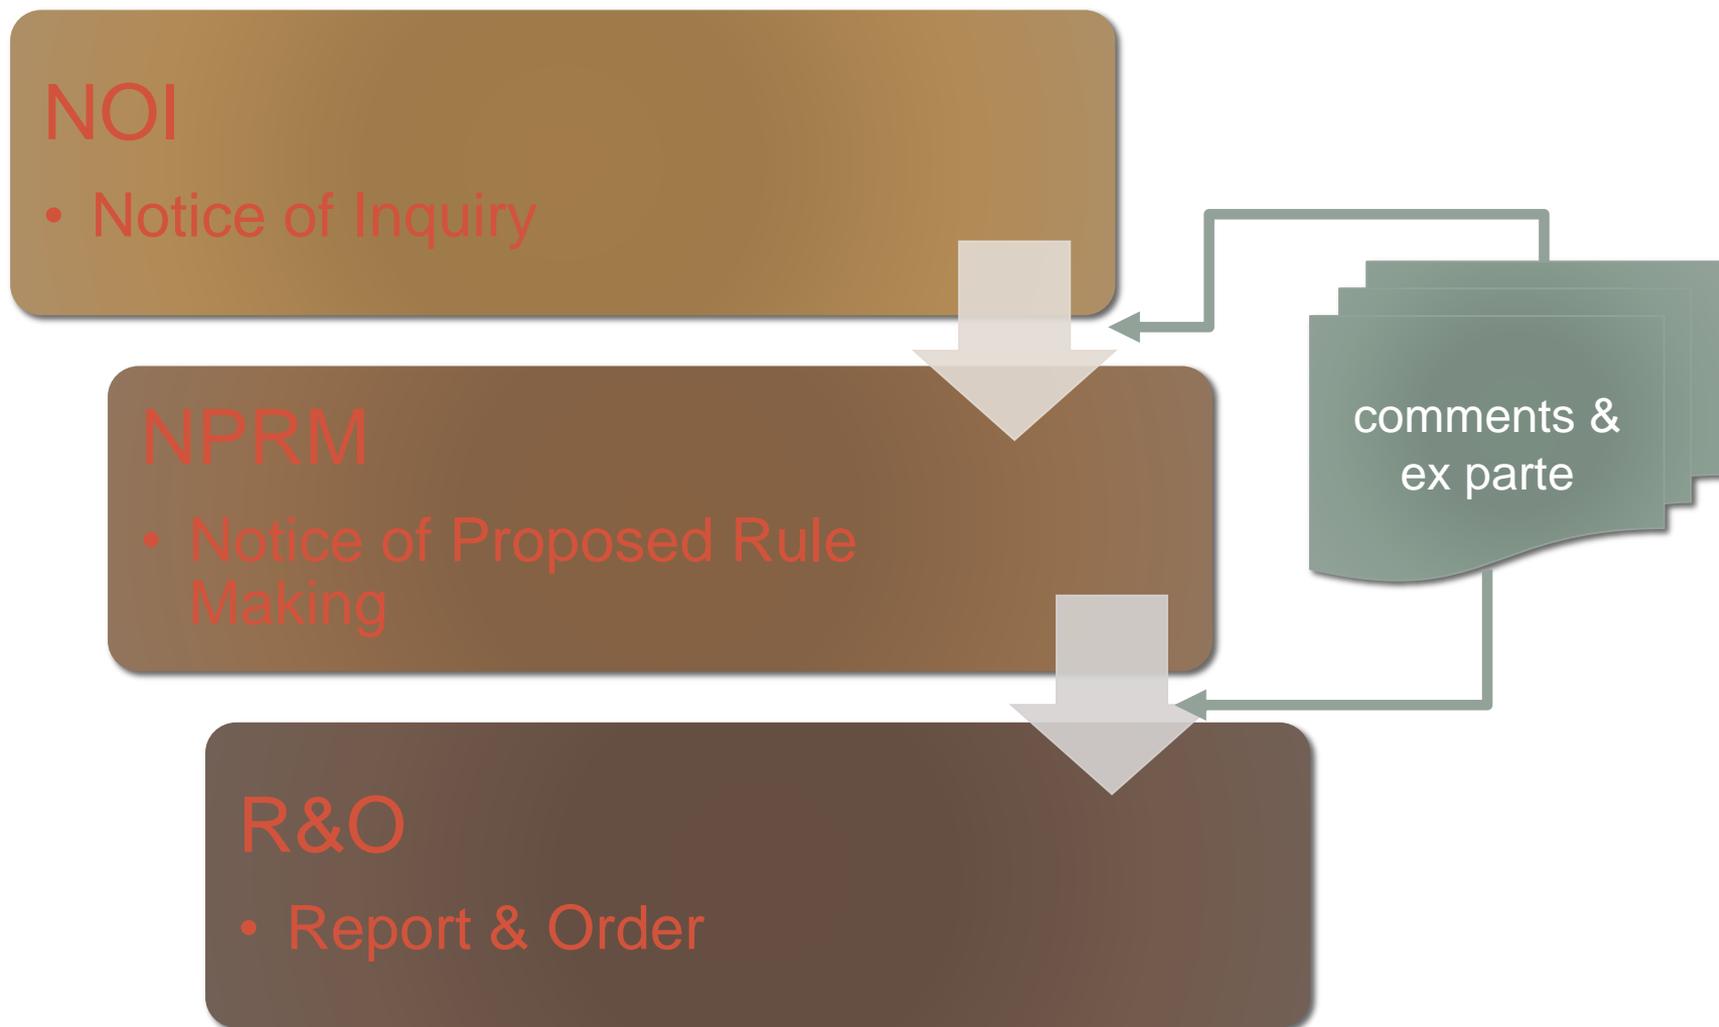
BACKUP

FCC



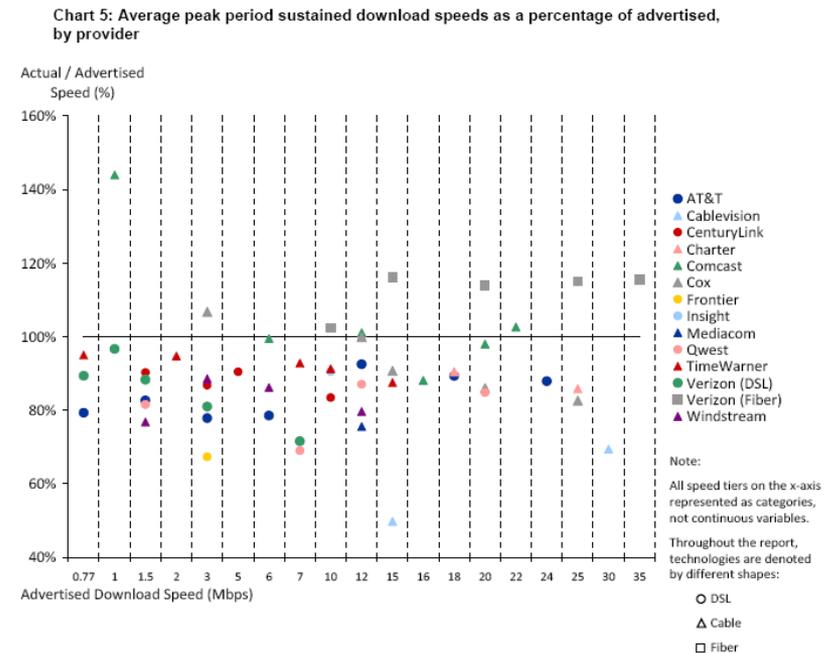
- Independent federal agency
- About 1,600 employees

Process



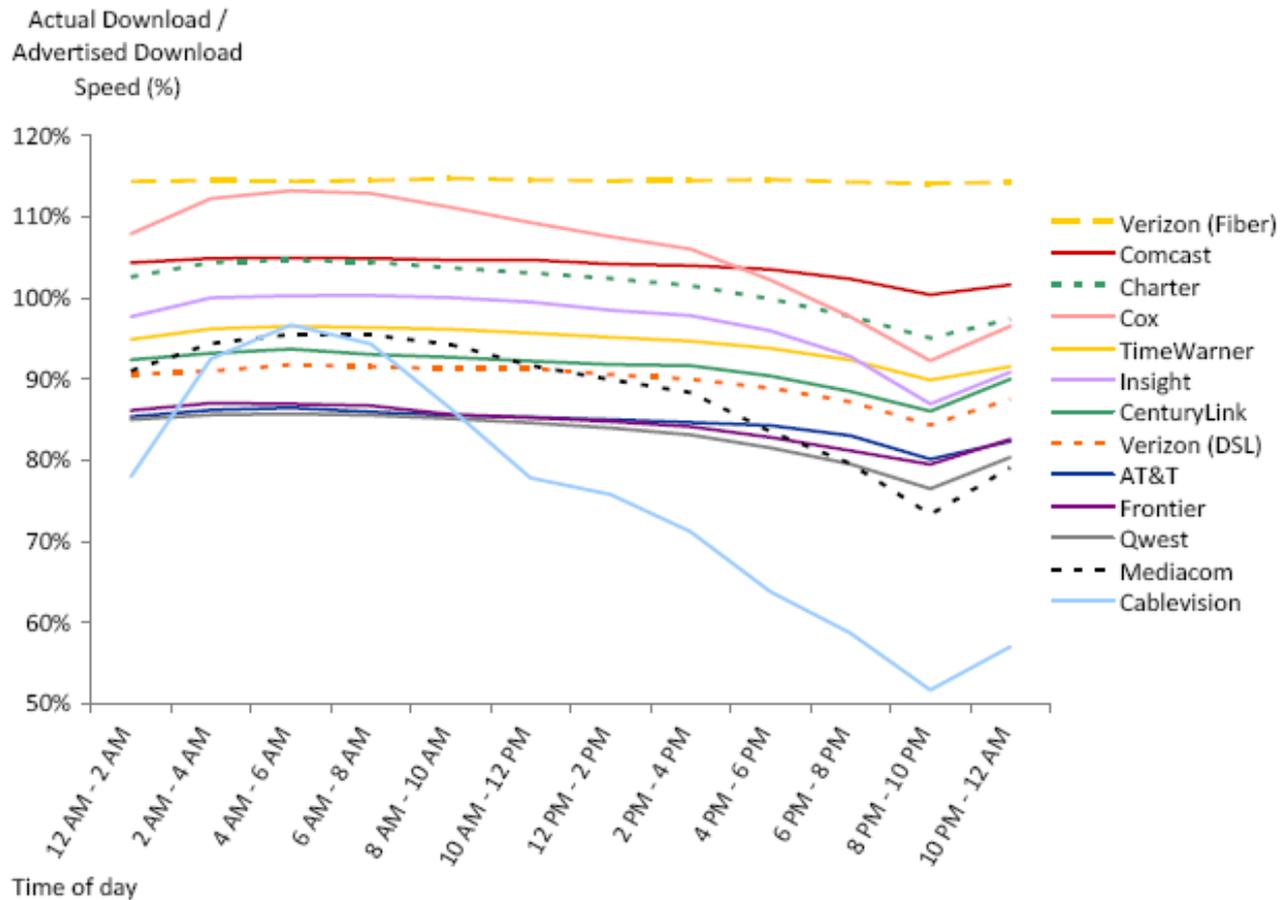
Performance varies

- ISPs seem to impose network wide performance standards
- However, there can be exceptions by speed tier

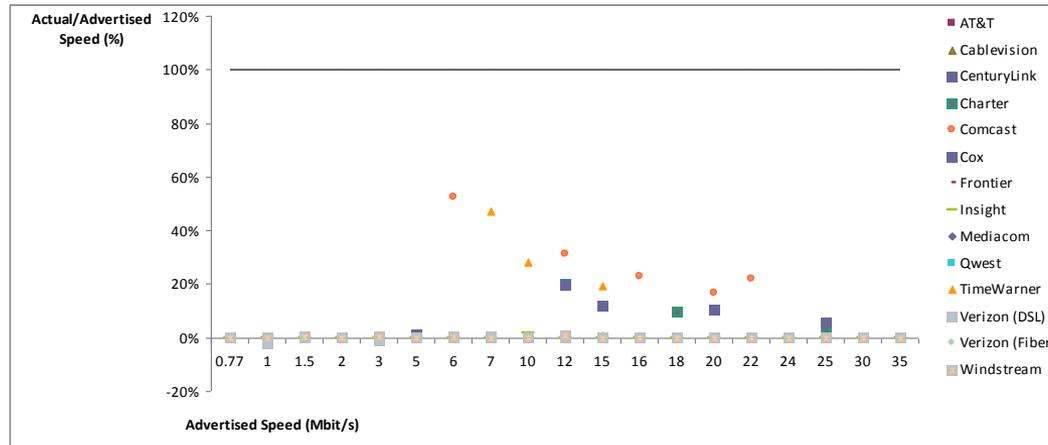


2011: some don't

Chart 11: Average sustained download speeds as a percentage of advertised over a 24-hour period, by provider

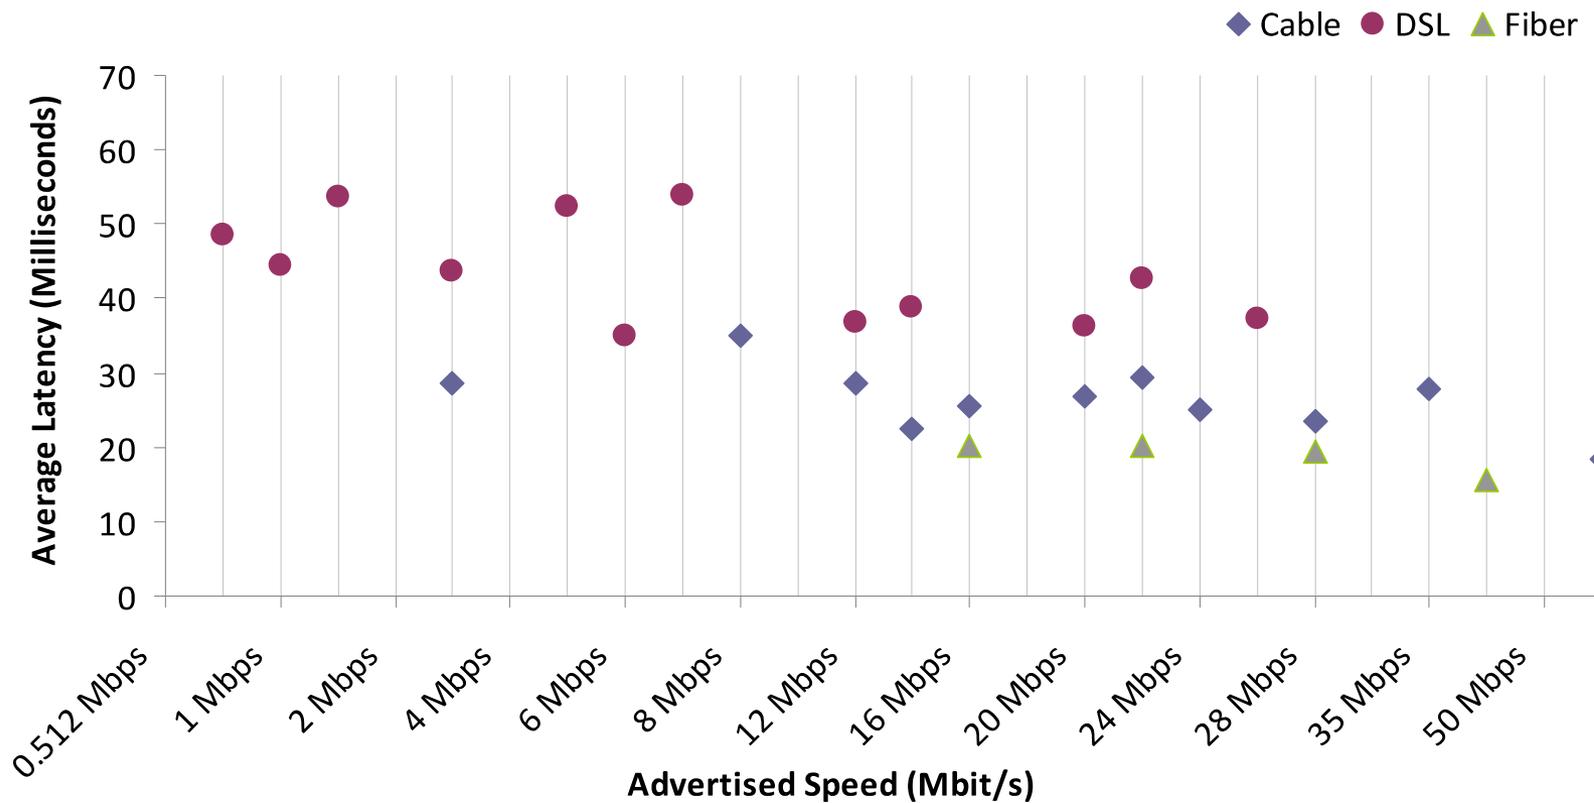


Burst speed increase

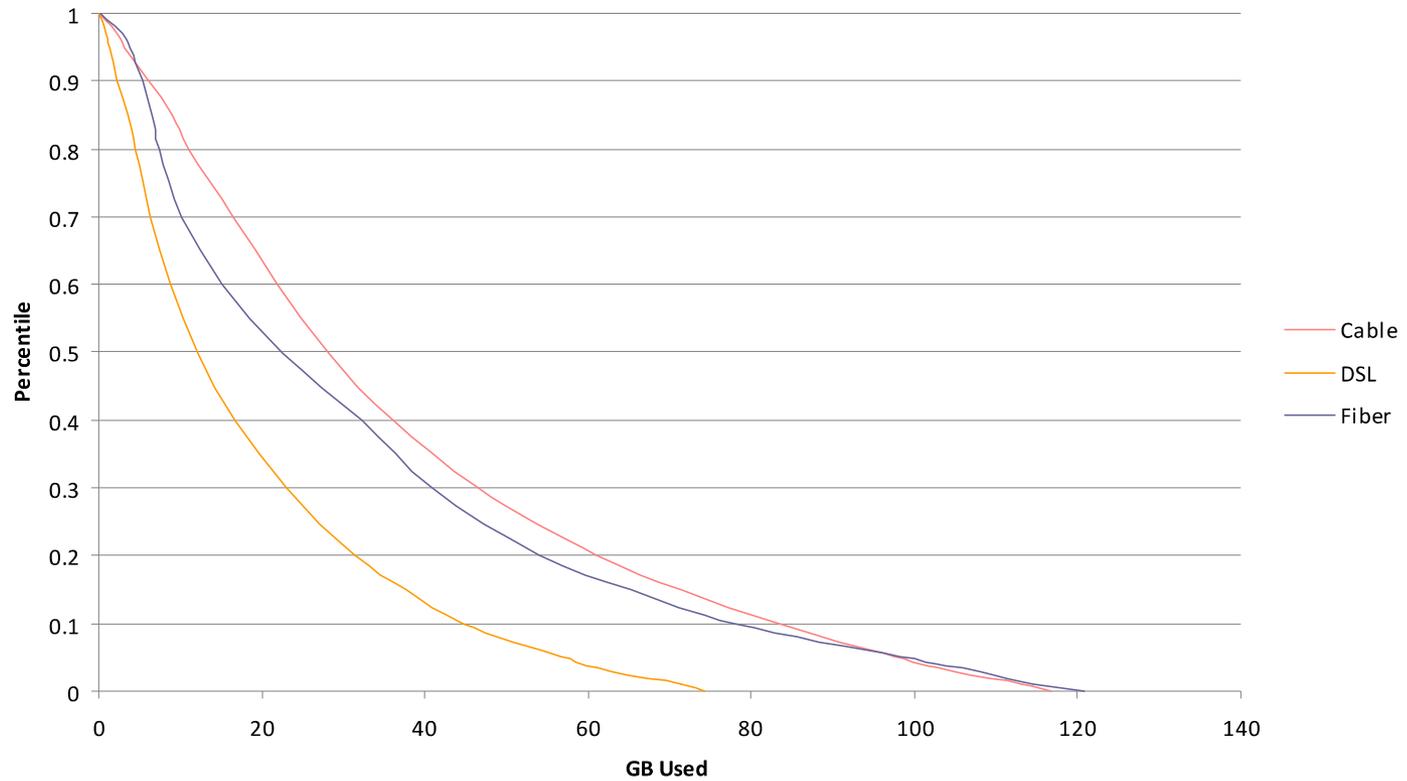


- Most impact of burst speed seen between 6 and 12 Mb/s
- Note: This chart not in report and shows calculated difference between burst and sustained performance

Latency by technology



Data usage



Reliability

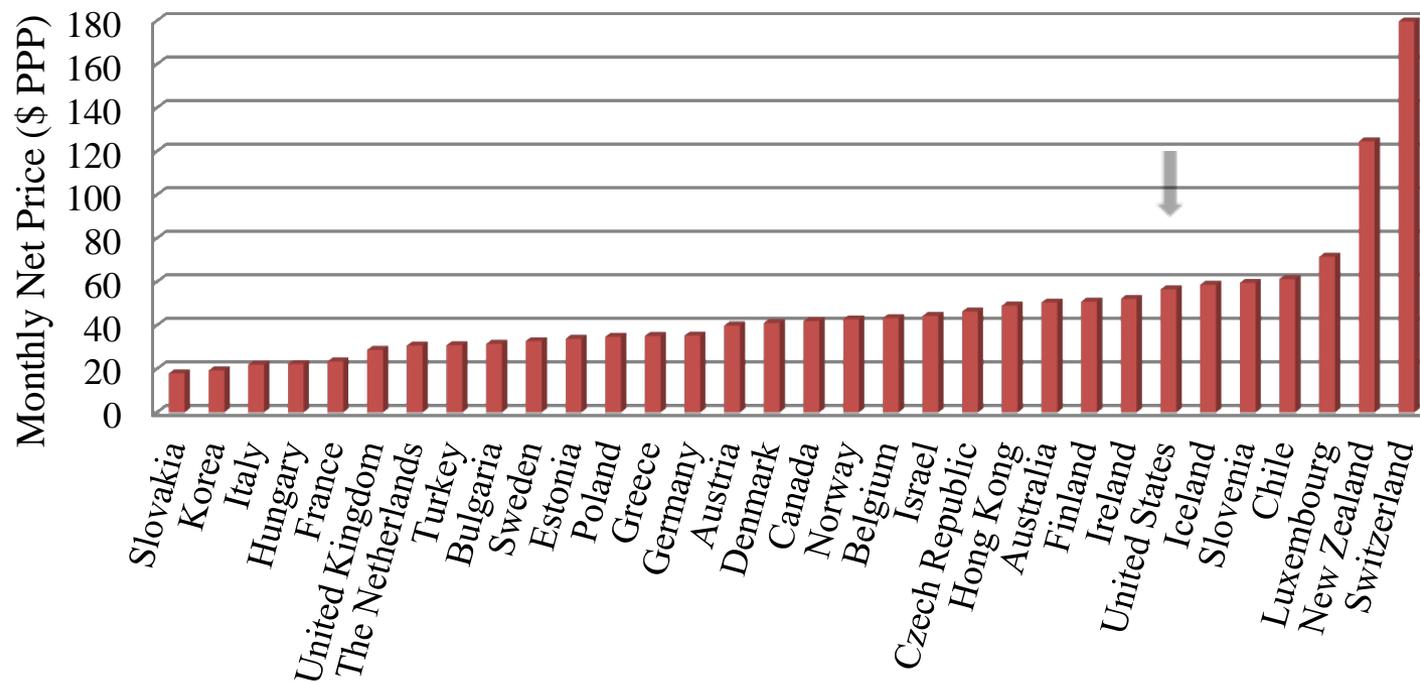
- Packet loss rate < 1%
- Correlation between peak periods and packet loss
 - Higher loss during peak hours
- Most companies during peak experience < .4% packet loss
- Worst case seen during March: .8%
- Data from other periods may have numbers in excess of 1% (Georgia Tech)
- 1% packet loss often cited as video threshold

Web page downloading: Canary in the Coal Mine?

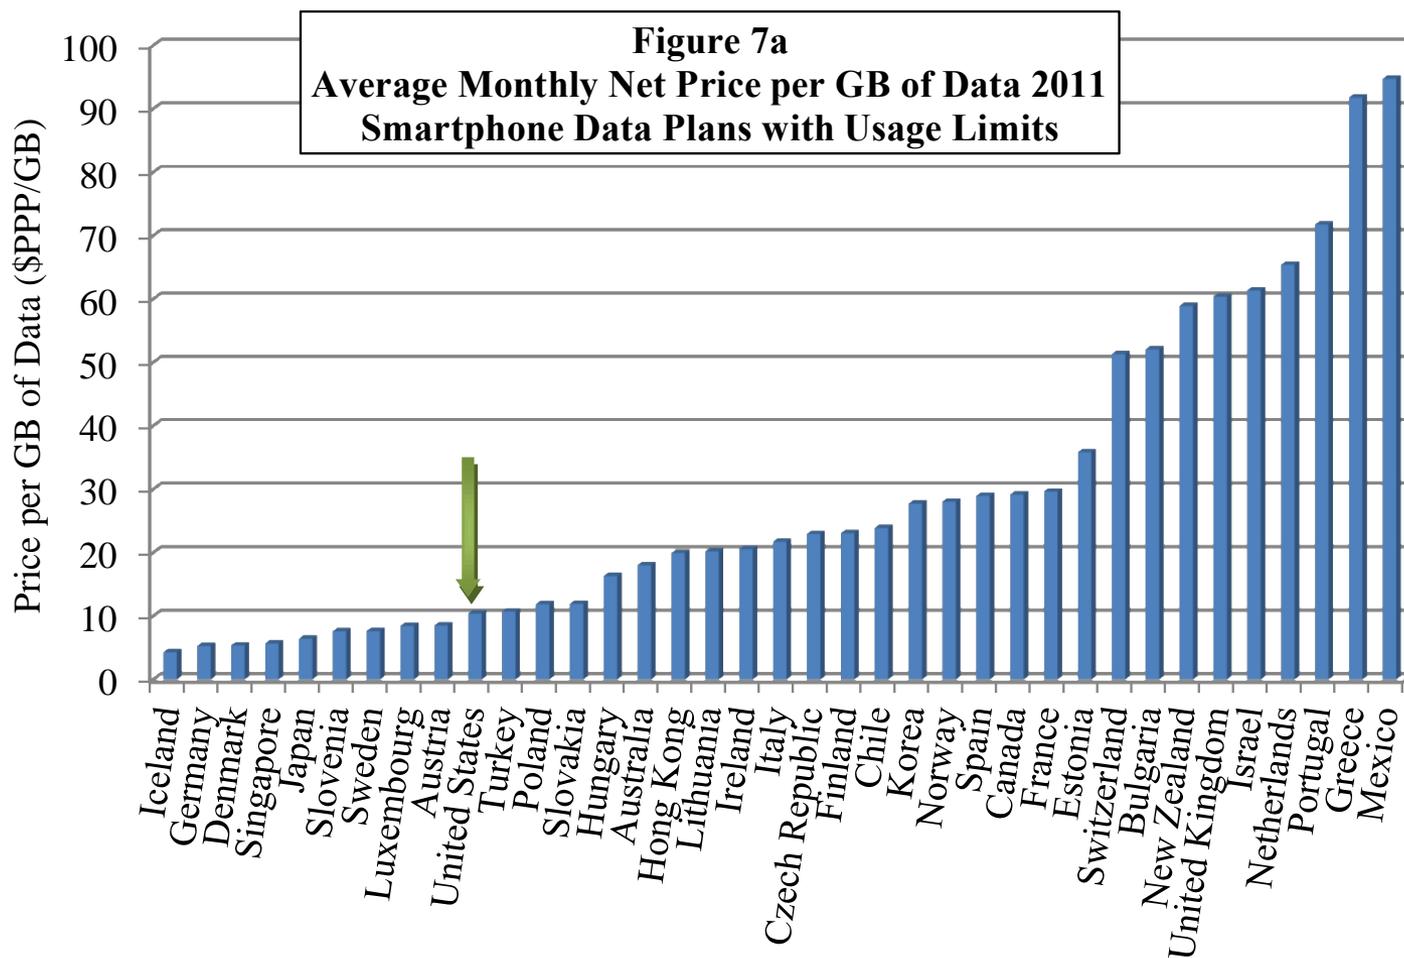
- Performance seems to top out after 10 Mb/s
- Many possible explanations
 - Latency, TCP issues, server loading, household platform limitations, ...
- However, discussions with Georgia Tech indicate that they have seen similar performance issues
- Discussion with Ofcom and others suggest that globally, full benefits of higher line rates not being realized
- Higher ISP speed may challenge industry to examine performance bottlenecks
- More data needed

International comparison: fixed

Figure 2c
Average Monthly Net Price (\$ PPP) of Residential (Fixed) Standalone
Broadband 2011
15-25 Mbps of Download Speed

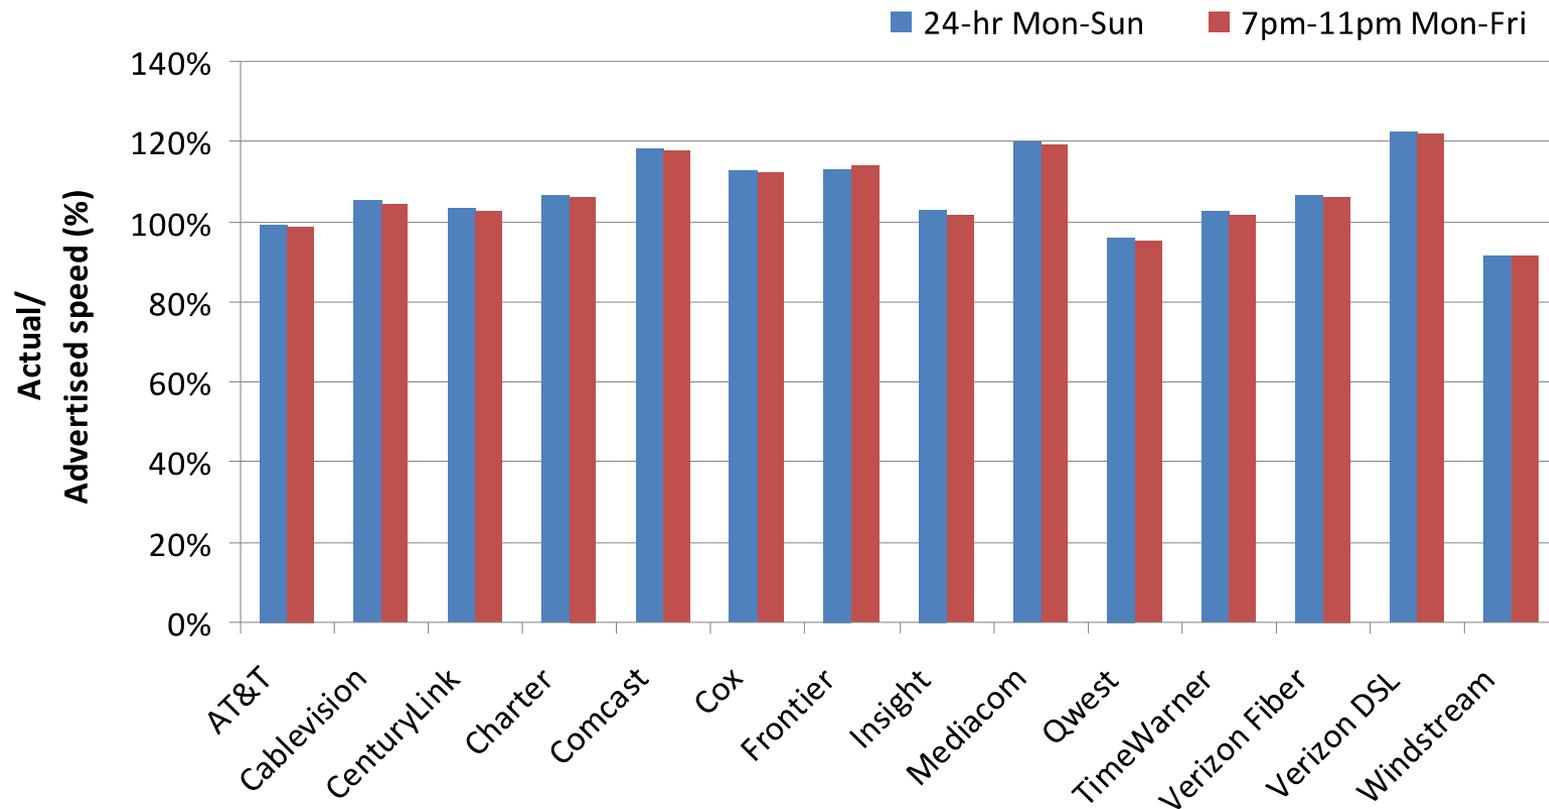


International comparison: mobile



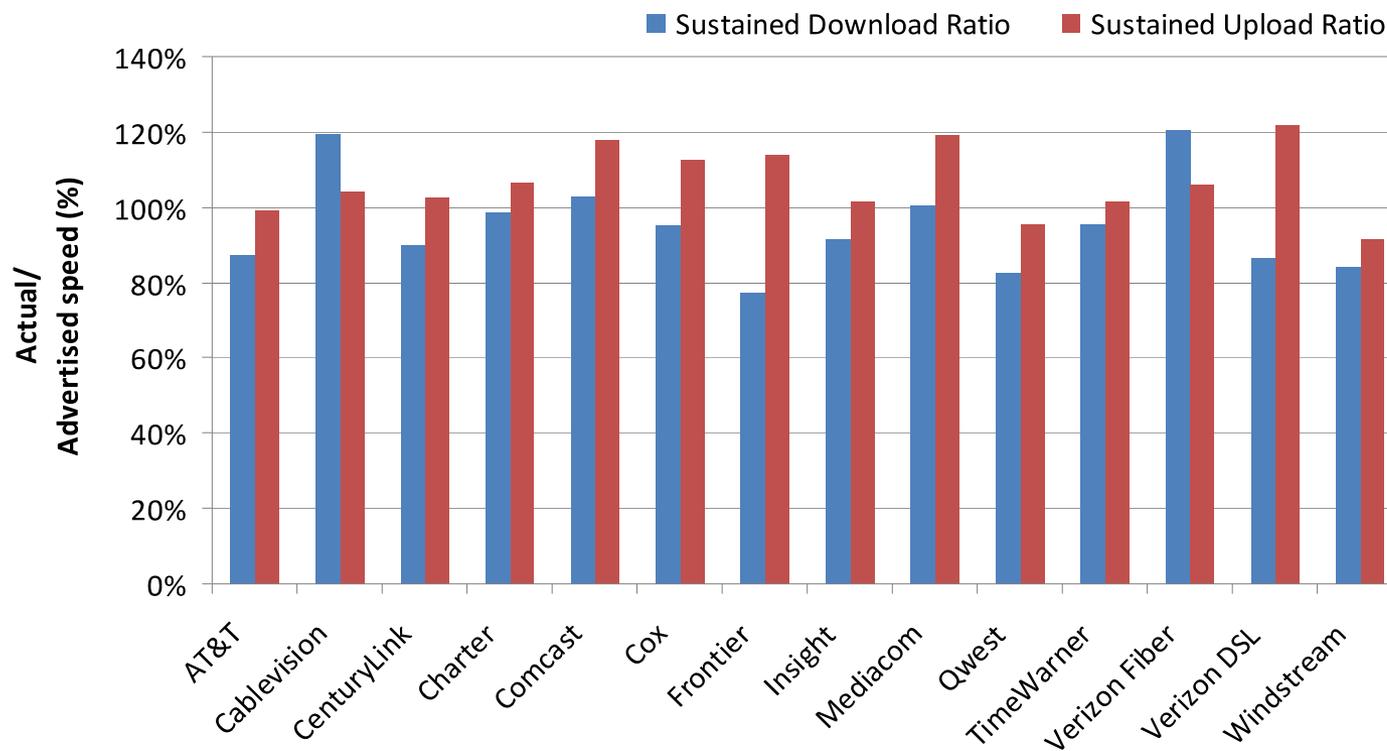
2012 Upload speeds

Chart 2: Average Peak Period and 24-Hour Sustained Upload Speeds as a Percentage of Advertised, by Provider—April 2012 Test Data



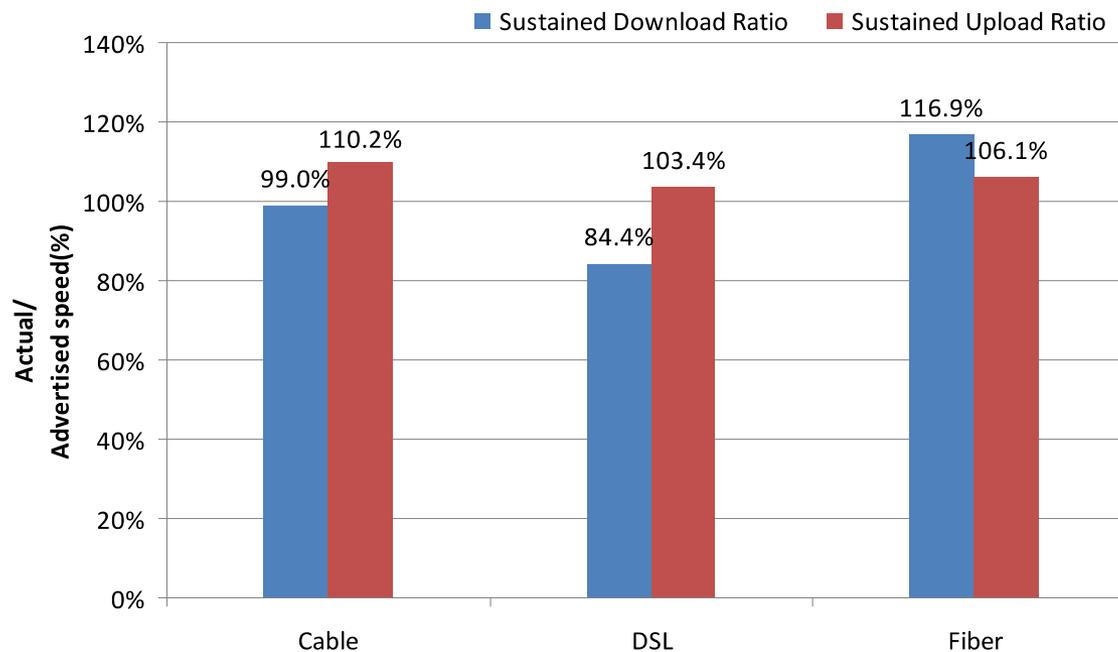
2012 peak period actual vs. advertised

Chart 3: Average Peak Period Sustained Download and Upload Speeds as a Percentage of Advertised, by Provider—April 2012 Test Data



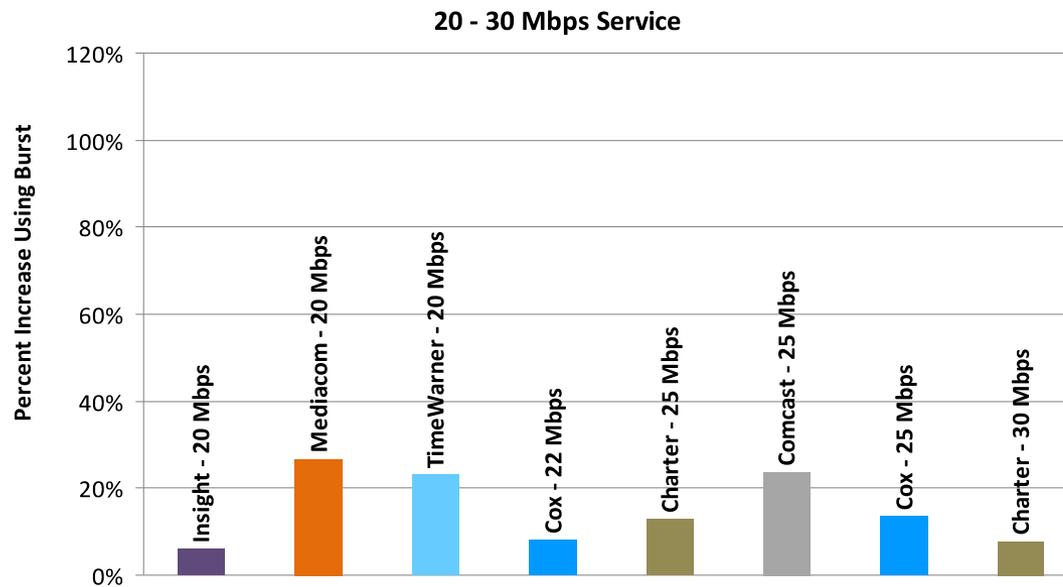
Advertised vs. actual

Chart 4: Average Peak Period Sustained Download and Upload Speeds as a Percentage of Advertised, by Technology—April 2012 Test Data



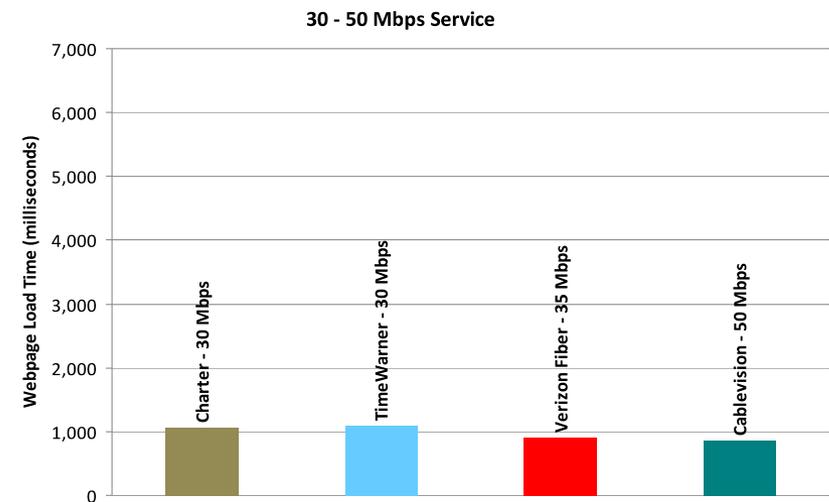
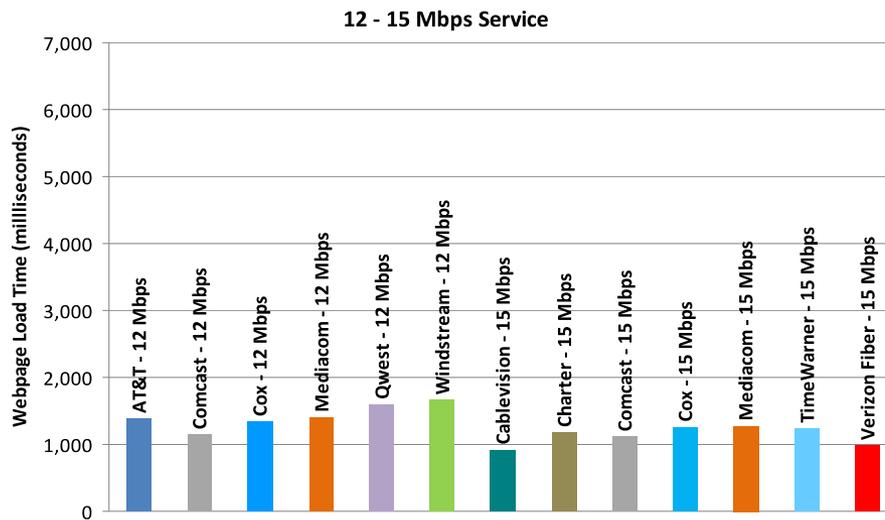
Burst download vs. sustained download

Chart 7.2: Average Peak Period Burst Download Speeds as a Percentage Increase over Sustained Download Speeds, by Provider (20-30 Mbps)—April 2012 Test Data



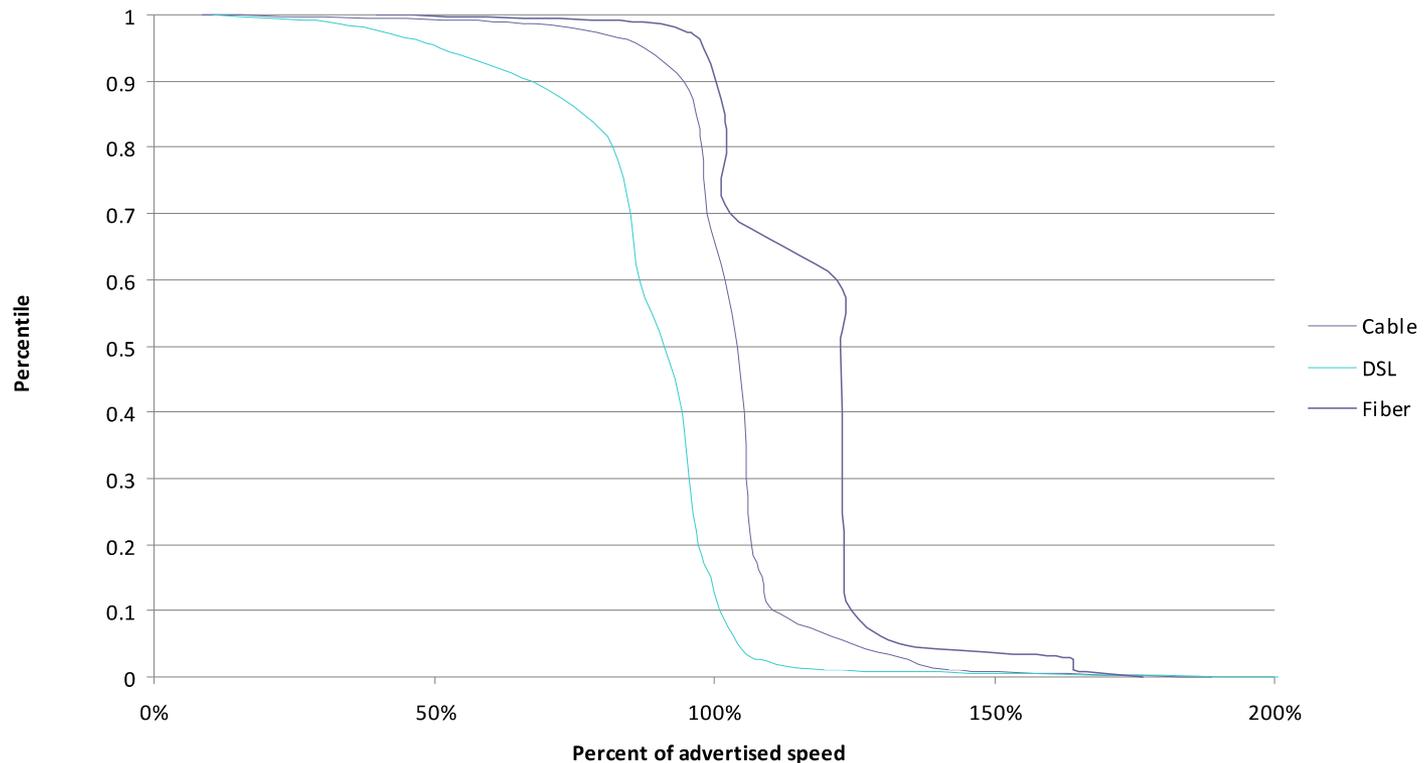
Web page loading time

Chart 11.3: Web Loading Time by Advertised Speed, by Technology (12-15 Mbps Tier)—April 2012 Test Data



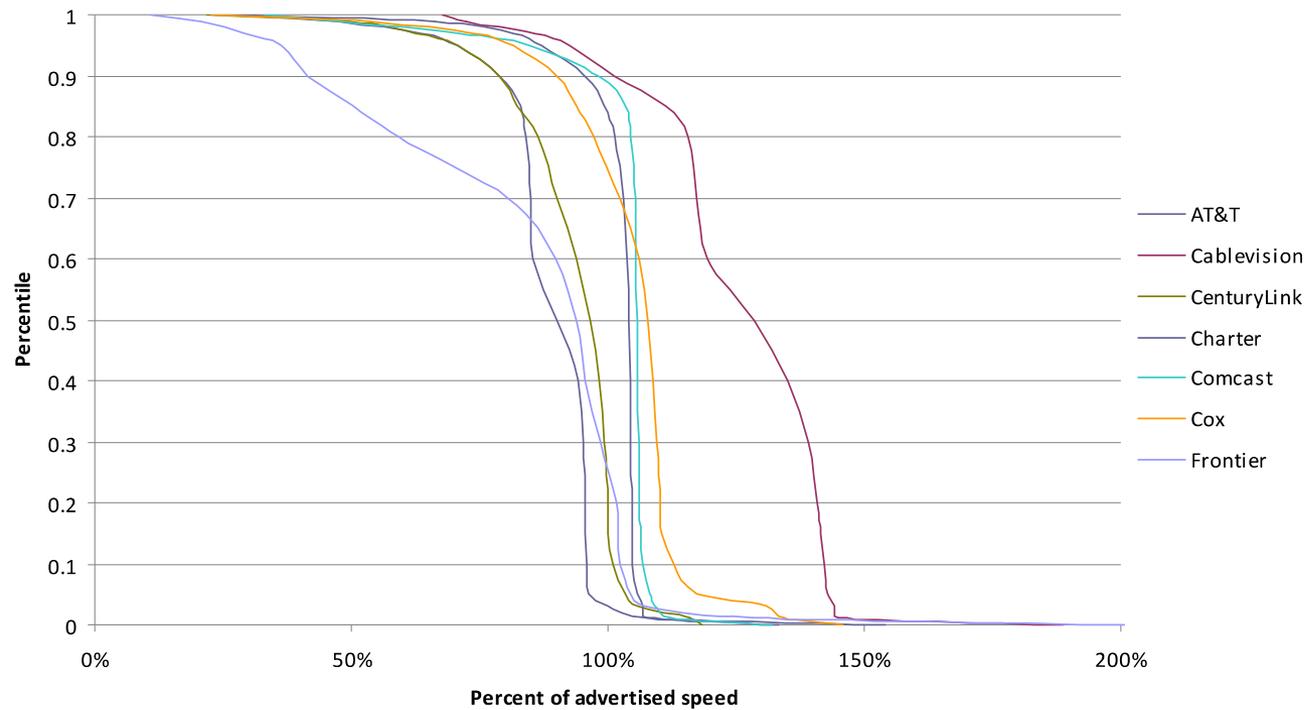
Cumulative distribution of download speeds

Chart 14: Cumulative Distribution of Sustained Download Speeds as a Percentage of Advertised Speed, by Technology—April 2012 Test Data



CDF of sustained download speed

Chart 15.1: Cumulative Distribution of Sustained Download Speeds as a Percentage of Advertised Speed, by Provider (7 Providers)—April 2012 Test Data



Throughput predictability

Figure 1: Percentage of Sustained Advertised Download Speed Delivered During Peak Period, by Provider

| | 70th Percentile | 80th Percentile | 90th Percentile | 95th Percentile |
|----------------------|-----------------|-----------------|-----------------|-----------------|
| AT&T | 85% | 84% | 79% | 71% |
| Cablevision | 117% | 116% | 101% | 93% |
| CenturyLink | 90% | 86% | 79% | 71% |
| Charter | 103% | 101% | 96% | 87% |
| Comcast | 105% | 105% | 98% | 85% |
| Cox | 102% | 97% | 90% | 82% |
| Frontier | 81% | 59% | 42% | 37% |
| Insight | 98% | 97% | 92% | 88% |
| Mediacom | 104% | 102% | 97% | 88% |
| TimeWarner | 97% | 97% | 95% | 91% |
| Verizon Fiber | 122% | 106% | 102% | 99% |
| Verizon DSL | 86% | 77% | 60% | 53% |
| Windstream | 87% | 84% | 78% | 67% |

Actual download speeds: 2011 vs. 2012

Figure 2: Year by Year Comparison of Sustained Actual Download Speed as a Percentage of Advertised Speed (2011/2012)

| | Year 2011 | Year 2012 |
|-----------------|-----------|-----------|
| AT&T | 81% | 87% |
| Cablevision | 54% | 120% |
| CenturyLink | 87% | 89% |
| Charter | 96% | 98% |
| Comcast | 101% | 103% |
| Cox | 95% | 95% |
| Frontier | 81% | 79% |
| Insight | 89% | 92% |
| Mediacom | 75% | 100% |
| Qwest | 77% | 83% |
| TimeWarner | 91% | 96% |
| Verizon (DSL) | 86% | 87% |
| Verizon (Fiber) | 114% | 120% |
| Windstream | 85% | 84% |

Chart 18: Normalized Average User Traffic—April 2012 Test Data

