

Traffic Microbursts and their Effect on Internet Measurement

Measurement Lab
May 25, 2016

Table of Contents

[Abstract](#)
[Traffic Microbursts and Switch Discards](#)
[Remediation and Next Steps](#)
[Affected Data Identification](#)
[Demarcated Data Availability](#)
[Validating Previous Research](#)
[Implications for Internet Performance](#)
[Disclosure Policy and Contact](#)
[Summary of Actions Taken](#)

Abstract

In August 2015, M-Lab was notified of potential degradation of site performance by a measurement partner based on discrepancies compared to results for their own servers. After a full investigation these patterns were found to have been caused by the unique confluence of several specific conditions. Interim remediation measures were taken in early October 2015, and the resolution of the degradation was confirmed by the partner and others. Due to these administrative actions, the episode, which we are calling the “switch discard issue,” has not affected testing conducted in the United States (the region impacted by this problem) since October 11, 2015, and thus measurements after this period are not affected by the incident. M-Lab has also conducted an evaluation of data collected during the time period in which the issue occurred, and has taken steps to remove affected measurements from its dataset. This incident will not affect use of its dataset, past or present, as a result.

M-Lab is committed to open data and operations, and this includes transparently disclosing issues it discovers. This document was drafted to help users, partners, and stakeholders better understand an issue that M-Lab believes is relevant to the operation of network infrastructure generally, about which we have seen little research.

This document does the following:

- outlines the technical causes behind instances of calibration anomalies and their implications for Internet infrastructure management overall;
- describes measures taken by M-Lab to prevent further issues;

- discusses several strong forward-looking techniques to provide ongoing calibration monitoring;
- describes measures taken to prepare M-Lab infrastructure for the next generation of consumer networks; and,
- reproduces past research conducted by M-Lab, using unaffected data, in order to provide updated analysis and confirm the continued validity of our previous findings

Traffic Microbursts and Switch Discards "

In the past, slower average broadband connections have effectively acted as traffic shapers, preventing large traffic bursts from creating issues with performance. The increasing presence of high capacity connections and edge connectivity presents a shift in network conditions that generally impacts all network infrastructure, and has broad implications for the Internet. This is the account of how M-Lab came to this knowledge, and the means by which M-Lab researchers detected, confirmed, and remediated the impact of these changing conditions on its measurement infrastructure.

On August 13, 2015, Measurement Lab became aware of reports of inconsistent performance for a selection of sites, initially believed to apply to only one location and a limited number of service providers. Historically, patterns of degradation have been caused by interconnection-related congestion and abnormal concentrations of concurrent measurements against sites, which were resolved through changes to server selection and scheduling regimes. M-Lab began to collaborate with measurement partners and sourced data from other records to reconstruct the utilization of sites and site capacity. A test client placed within the same metro location as an M-Lab site demonstrated that the identified sites were able to achieve throughput approaching the 1 Gbps upstream capacity standard to M-Lab's platform – i.e. that they were provisioned properly. Additionally, while perceptible differences in measured performance were seen across IPv4 vs. IPv6, and based on other network topology factors, reports of degradation over multiple access ISPs indicated that the performance trend was not the result of transit ISP failure and not exclusively due to intermediary network segment issues impacting one or another ISP. M-Lab was able to achieve the high-throughput measurements that the platform was designed to support.

At the beginning of September 2015, M-Lab came to understand that the reported issues were not limited to one location, but occurred to a lesser extent on other geographically-dispersed sites. The sites in question were each connected via different transit ISPs and were deployed at different times. These findings greatly reduced the possibility that the issues detected were due to connectivity or hardware-related causes. Further testing on alternative operating systems, run on M-Lab hardware, showed similar patterns of performance, and therefore indicated that the performance characteristics were not solely the result of software or device drivers.

Working with measurement partners to gather data about the impact of the issue on testing regimes, Measurement Lab attempted to reduce a sampled set of sites (the servers and switch deployed at a particular location) to their minimum components – removing as many variables as possible in order to better isolate the cause of the problem. As a part of this process, for a short period of time, impacted M-Lab sites were pared down to one server (reduced from the standard three production servers), running CentOS and none of the other M-Lab software. Under these conditions, a more expansive testing regime

was conducted against the bare-bones sites, with the variables in play reduced to the single server, its hardware, the site's switch, and the upstream connectivity. The results of these tests produced throughput measurements that were stable and within normal, expected levels, further indicating that the upstream connectivity and hardware could support stable, high-throughput transfers, and thus were not the source of the inconsistency. Continuing this testing regime, the experimental CentOS server was replaced by an operational M-Lab server – reintroducing the full M-Lab software suite as a potential variable. Testing under these conditions revealed the same, consistent results, revealing that the M-Lab software was not responsible for the issue. Only when additional servers were turned back on, and hosted experiments began to test normally again did we see degradation recur. Therefore, the introduced loss was not the result of upstream connectivity, server hardware or software issues, but the devices between M-Lab's server and upstream connectivity. This clearly implicated M-Lab's switch hardware.

Based on traffic records and experimentation, M-Lab staff identified that specific patterns in the use of the platform correlated to short periods of underperformance, which was only evident when a number of particular conditions were met simultaneously. Based on this information, M-Lab staff hypothesized that the switches could incur packet loss when multiple high-throughput transfers were concurrently conducted against different servers at the same site in a manner that filled the buffer space allocated by switch firmware to the uplink port. This would explain why the patterns of performance only manifested when multiple servers were up and receiving tests at a given site. This led to the hypotheses that the underperformance was a result of the combination of the following conditions:

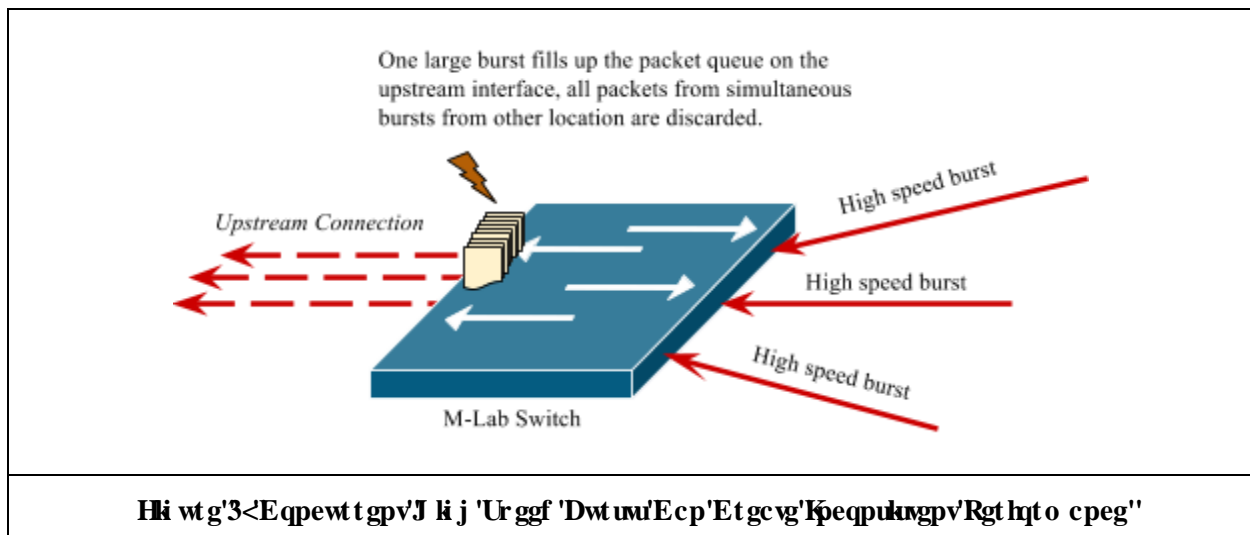
- A standard TCP implementation which is prone to sending line-rate microbursts, because it does not pace TCP packet sending;
- Switch queue buffer sizes below a particular bound; and,
- Simultaneous tests running from high-throughput test clients to different M-Lab servers at the same site.

Broadly, these factors occurring at the same time would result in brief periods of network underutilization – on the order of seconds – caused by the M-Lab measurement servers concurrently sending high-throughput microbursts (milliseconds or microseconds) beyond the capacity of the switch, resulting in packet loss. When a high-throughput client conducts a test, the M-Lab server (under its current kernel version¹) will send line-rate packet trains (up to 1 Gbps) with a large TCP transmission window as a part of the server-to-client (download throughput) measurement.² The upstream connectivity is a shared resource between all of the servers at an M-Lab site, requiring the switch to properly manage concurrent traffic flows to clients. As the sizes of concurrent flows sent from multiple servers behind the switch increased, due to requests from faster measurement clients, microbursts that could previously co-reside in the switch's queue memory grew to exceed the capacity of the device. When the switch queue buffer memory became full, head-of-line issues in the switch would cause an arriving packet train sent from the M-Lab servers to experience substantial loss. Since switch behavior is a product of vendor design and

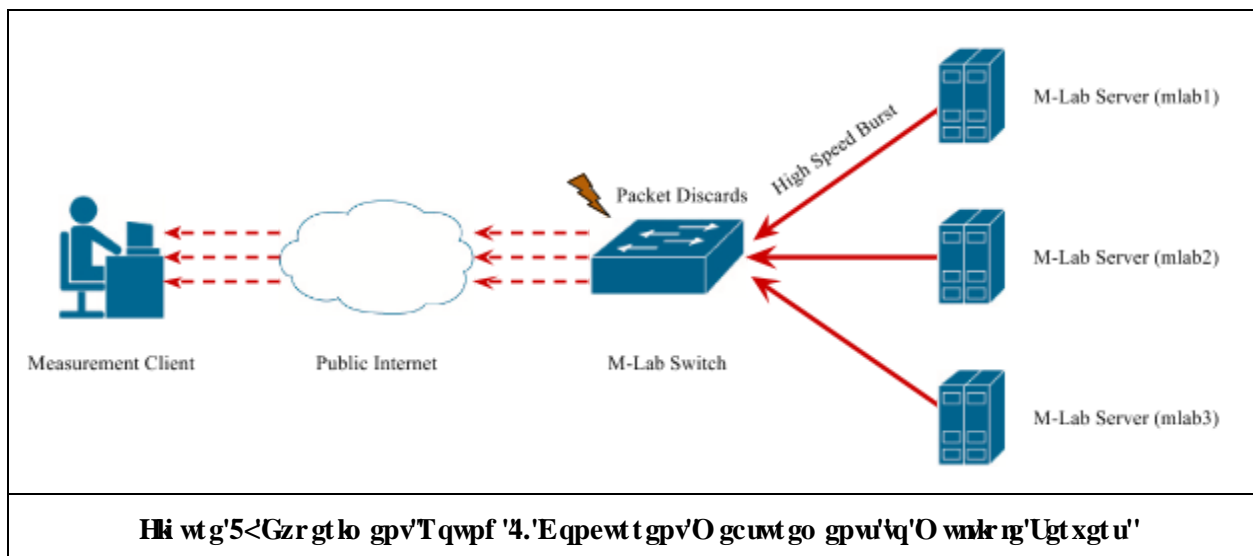
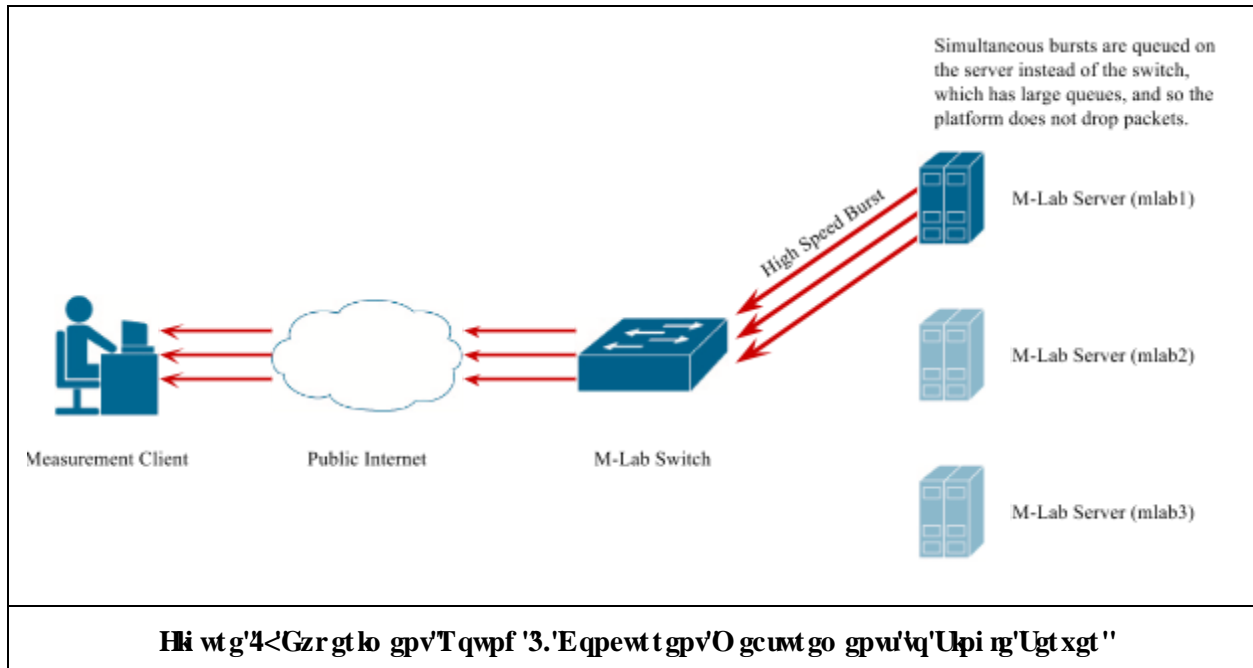
¹ As of May 23, 2016, Linux kernel version 2.6.32-131.

² Pacing was added in Linux 3.12, such that the packets can be spread out across the window. <https://lwn.net/Articles/564825/>; For more on the underlying mechanisms, <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=1146410>

firmware, features to avoid blocking and to smooth out microbursts can differ between models in an opaque manner. This loss would cause the TCP implementation to back off under the rules of standard congestion control. The overly-large backoff due to correlated loss would cause underutilization of the network, as traffic from all servers would back off. When the necessary conditions were met, the end effect would be a decrease in measured throughput performance to rates below the other bottlenecks in the path for server-to-client tests for brief intervals of time. The increase of high-capacity clients would affect the probability of observing switch discards, but not the severity of the degradation. At lower latencies between test client and M-Lab server, TCP's standard congestion control would recover from the loss more quickly, explaining why the geographically-proximate test clients were less impacted by the discard events, and why the initial testing did not show evidence of performance impacts.

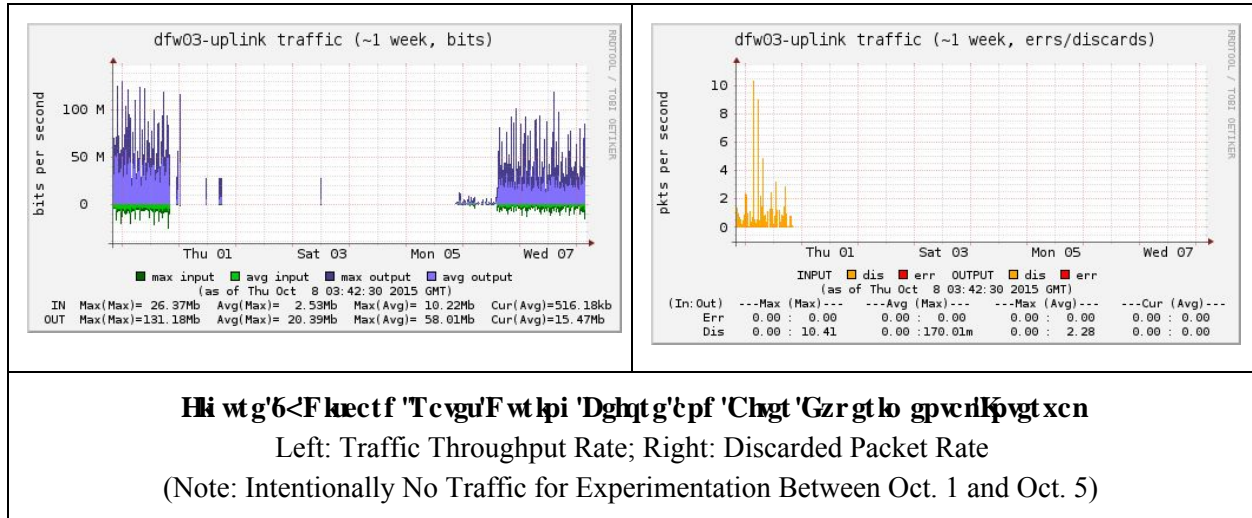


In order to prove this hypothesis, we constructed a controlled experiment emulating measurements from high-throughput test clients at latencies more commonplace on consumer broadband connections. In the first round, three concurrent tests were run against one server at a single site. In the second round, three otherwise identical tests were run against three different servers at the same site. With three concurrent requests to one server (Round 1), the aggregate test traffic was consistently able to approach the site's upstream capacity without any increased packet loss (i.e. no degradation was detected). Where the requests were spread across different servers at the same site (Round 2), aggregate performance was substantially less than the upstream capacity, and the switch exhibited high packet loss caused by buffer overruns at the switch. Both controlled experiments and trials conducted by partners using real testing clients validated the hypothesis that the discarding of packets by the switch was the source of the observed behavior. Once we could easily reproduce the problem, we were able to study all of its symptoms and evaluate techniques for remediation.



This test scenario authentically predicted the increased rate of discarded packets observed in our monitoring system during high-intensity testing periods whenever high performance test clients were present, test clients such as those often associated with measurement programs conducted by partners. The switch would discard the packets itself due its own limitations in receiving the microbursts from the multiple M-Lab servers. This discard information is collected as a part of monitoring from all M-Lab sites. However, these microbursts and subsequent switch discards were partially concealed by the monitoring system, which reports throughput rates and discards on an aggregated basis averaged over an interval of five minutes. Because of this, short intervals of extremely high throughput, or discard rates for several seconds (the common manifestation of the incident), would be averaged out to a seemingly

acceptable level over a longer interval, a consideration that has prompted improvements to M-Lab’s monitoring systems.³ Figure 4 graphs a real example of the experiment on a Dallas site, a location initially identified by partners as experiencing inconsistent performance. Until October 1, the site operated with three servers and a full set of clients testing against all three. After about 4 days of inactivity (the gap in the middle of the graph on the left), testing was resumed against only one server. Whereas in the three-server configuration the site experienced packet discards, after the change, discards did not occur (see the right side of the graph on the right) – even when the site was subject to a similar volume of test client traffic.



The incident did not uniformly impact sites as a result of differing traffic patterns, however, the causes of the problem and M-Lab’s ability to identify affected intervals of time remain the same. While M-Lab provides a directory service to measurement partners for selecting a server for each test, those making use of the M-Lab platform or its tools are free to choose their own selection method. M-Lab’s directory services works by selecting the site that is geographically closest to a given test client, and directing the client to test against a server within that site. Where multiple M-Lab sites are in place in a given region (a common practice that allows M-Lab to provide comparative measurements across different network segments and spread out load), those using M-Lab’s directory service will be sent randomly to one or another site in the region. Other approaches to server selection used by some measurement partners focus on different properties of the network, such as determining which sites to test against by checking the latency between a test client and a multiple sites, then selecting the site that is fastest in terms of latency. Such test clients are more likely to use a more limited portion of the platform and concentrate traffic on particular sites – those with lower latencies. As a result, the degree to which particular sites were subject to discard events differs across sites and time.

³ M-Lab has begun to collect its normal behavior parameters at more fine-grained intervals to allow more precise monitoring, and has expanded to retention other indicators of site status, such as system load over time. These measures will be documented in later posts, and M-Lab will be releasing this data where such disclosures can confidently be understood as not posing a risk to the platform.

As a result of the interrogation of our hypothesis, developed through a rigorous experimental regime, M-Lab has confirmed that test clients would have intermittently failed to operate within expected levels when other high-throughput connections concurrently tested against the platform due to clumped packet loss as a result of limitations with the current switch hardware. Further testing strengthened the conclusion that lower throughput streams would not trigger performance degradation, even at relatively high aggregate loads across multiple servers at a given site, aligning with reports that the situation became more apparent with the introduction of new, higher-capacity test clients. This investigation additionally confirms that the measures taken by the FCC to exclude collected data from a limited set of test clients in the Measuring Broadband America report released in December 2015 was an appropriate measure, and that the M-Lab platform is able to support a full range of test clients in the upcoming MBA analysis. The measures outlined in this report will provide further calibration testing and ensure data quality for forthcoming measurement programs conducted by partners.

Remediation and Next Steps''

As described previously, the switch discard condition required the confluence of multiple factors. Only once this set of causes was identified could the issue be remedied through removal of any one of the contributing factors. This presented at least three mitigation strategies:

- Disabling of all servers except one at each M-Lab site, so that the switch's queues could not be overwhelmed by competing line-rate traffic;
- Upgrading switch hardware so that simultaneous tests could not overwhelm the switch's queues; or,
- Updating the Linux kernel of the servers to a version with pacing, such that TCP could be configured not to produce line-rate packet trains.

Based on this understanding, and upon consultation with measurement partners, M-Lab turned a selection of locations to single-server sites and verified that performance stabilized following that move. M-Lab also increased monitoring of resource utilization and found that the individual servers did not suffer from an increase of load under this configuration. This approach was preferable in consideration of the data collection needs of measurement partners (i.e. the ongoing Measuring Broadband America program), since the change could be made immediately. The mitigation strategy was validated after several days of data collection in collaboration with measurement partners, and M-Lab made the temporary transition for each site used within current MBA testing period as a short-term resolution. Additionally, the directory service provided by M-Lab (mlab-ns) was updated to only utilize one available server per site, to reduce potentially competing flows at sites where the other servers were not turned down.⁴

Measurement Lab has also identified additional, longer-term changes to replace the single-server site configuration and improve platform resiliency overall. While the switches had been chosen due to their tested reliability under similar past circumstances,⁵ the increased demand for buffer memory created by

⁴ Outside of the United States, where there is currently little indication of significant increases of discards.

⁵ While extensive performance testing was conducted prior to the deployment of the switches currently in use, these trials did not assume the same conditions that lead to the inconsistent performance.

higher tiers of access drove them out of their tested range of reliability. It is important to note that the problems documented by M-Lab happened across multiple similar switches manufactured by different vendors, pointing to broader challenge than that of a peculiarly underperforming model of equipment. As a result, M-Lab will be replacing all switches across its fleet, taking the occasion as an opportunity to move to hardware capable of servicing upstream connectivity in excess of 1 Gbps, in order to support future efforts to measure next generation broadband access. Switches and uplinks that support 10 Gbps links will not produce such behavior on 1 Gbps flows, as their buffers are designed to support higher capacity throughput, even if the TCP implementation of the M-Lab platform servers continues to send line-rate packet trains. An over-provisioned switch will also ensure that similar behaviors do not occur at higher throughput rates, such as discrete 1 Gbps links per server, and improvements to monitoring will inform M-Lab staff of any further issues with switch capacity. Additionally, M-Lab had already planned to update the platform in the medium term to run newer kernels that could support features such as pacing. These software updates are anticipated to incorporate more recent changes to Linux's TCP implementation that do not produce the same packet sending behavior, which should further ensure that the problem does not arise later when higher throughput measurements are conducted. Once these changes have been made, M-Lab will restore the additional servers at the sites and return to its redundant provisioning.

In addition to the steps detailed above, M-Lab has taken additional measures to greatly improve the platform's overall calibration and calibration testing, including the expansion of a controlled testing lab, where test clients and servers can be evaluated under a wider range of emulated network environments (including stress tests). M-Lab has also introduced fine-grained resource monitoring telemetry for switch buffer overflows and other resource contention. The buffering behaviors of the switches and their operating system, such as behavior under memory stress, are typically proprietary information and not well documented, resulting in M-Lab emphasizing monitoring as a mechanism to ensure future calibration and accountability. M-Lab is currently working to integrate platform-wide resource telemetry information into the data collection pipeline. The resource telemetry will provide fine-grained signals useful to researchers and to M-Lab in planning future platform buildout, and are also considering additional techniques that might be used to validate platform calibration in situ.

Affected Data Identification

The behavior that led to inconsistent performance would appear to the M-Lab site as a period when multiple high-throughput test clients conducted measurements concurrently, with increased packet retransmissions as a result of discards at the switch. Fortunately, M-Lab retains a fine-grain accounting of the operations of its infrastructure through the public SideStream dataset. SideStream collects statistics about the TCP connections used by all measurements running on the M-Lab platform, including metrics on packet retransmission and throughput, which would indicate periods where discards could occur. This allows M-Lab to analyze the historical operations of the platform and identify measurements occurring during potential windows of inconsistent performance.

During February 2016, in cooperation with measurement partners, M-Lab recreated the conditions resembling normal activity during Fall 2015 at a site in Dallas that initially experienced the discard

episodes (dfw02), doing this through concurrent testing from multiple test clients to all three servers at the site. This was done to purposefully induce switch discards for diagnostic purposes. While the testing was underway, M-Lab collected the usual SideStream logs as well as new, high-resolution packet discard statistics directly from the switch. Simultaneously, at multiple sites still running three servers and experiencing normal load and non-zero – but acceptably low level – switch discards, M-Lab performed the same data collection from SideStream and switch discards, creating a comparative baseline. Researchers then combined the M-Lab SideStream records with the high-resolution switch traffic statistics to perform a LASSO regression analysis, a machine learning technique.⁶ The result of the analysis is a linear model that predicts when switch discard events are more or less likely, based only on SideStream data. The resulting model can be used to estimate fine-grained switch discard data at sites and times for which our only fine-grained traffic data is from SideStream.

For every 10 second period, i , if a discard was possible, the Discard Score for that period is calculated as:

$$\begin{aligned} \text{Discard Score}_i = & -32.584 + \\ & 0.073 * \text{SUM}(\text{SegsRetrans})_i + \\ & 2.256 * \text{COUNT}(\text{host_peak_burst_rate_gte_20mbps})_i + \\ & 1.585 * \text{COUNT}(\text{host_peak_burst_rate_gte_100mbps})_i + \\ & 0.000000213 * \text{MOVING_AVERAGE}_i(\text{SndLimBytesCwnd}, 6) \end{aligned}$$

Each variable in the Discard Score equation is derived directly from historical SideStream data. The model was trained and validated using segregated subsets of the measurement data preserved for the purpose, in accordance with standard statistical techniques. Based on the Discard Score equation, each time period (and all tests occurring during that time period) is considered *affected* (a score greater than zero) or *unaffected* (a score less than or equal to zero) by a switch discard. This allows researchers to consider the potential for inconsistent performance in their analysis of the M-Lab results based on the labelling of the data.

Demarcated Data Availability

The Discard Score was applied to the entirety of M-Lab NDT dataset in order to determine all tests that may have been conducted when the platform was potentially operating outside of normal levels due to packet discards at the switch. In applying this model, M-Lab has taken an extremely cautious approach in order to ensure that the representations made in the dataset remain as accurate as possible. In the next quarter, M-Lab will mark this information in BigQuery. In the now standard “fast tables” version⁷ of the M-Lab dataset, data impacted by the switch discards will be omitted to ensure that tools querying the data will not inadvertently use affected data. In the monthly tables, which provide more fine-grained information on measurements, these results will be marked in an additional column rather than excluded.

⁶ [https://en.wikipedia.org/wiki/Lasso_\(statistics\)](https://en.wikipedia.org/wiki/Lasso_(statistics))

⁷ <http://www.measurementlab.net/blog/new-tables-release/>

Validating Previous Research''

In October 2014, Measurement Lab published the report “ISP Interconnection and its Impact on Consumer Internet Performance” (Interconnection Study), which demonstrated how interconnection disputes between network providers had resulted in degraded broadband performance in the United States using measurement from the Network Diagnostic Tool (NDT). The techniques and tools developed for the report were further used in the June 2015 blogpost “New Opportunities for Test Deployment and Continued Analysis of Interconnection Performance,” which extended the interconnection research to data collected from more recent M-Lab sites connected by other transit networks. This increased perspective into network topologies broadened the scope of the observed degradation. Since both reports used the download throughput measurement from NDT clients toward the M-Lab platform, the switch discard effect would be potentially relevant to the outcome of the research. Through applying the Discard Score analysis to the NDT dataset, M-Lab is able to reassess the interconnection degradation episodes identified in this research using only “unaffected” NDT tests, and in the end confirm that the original findings of the Interconnection Study remain correct.

The research methods employed in the Interconnection Study focused on relative differences over the course of diurnal cycles and longitudinal periods of time. The report found that during periods of congestion, overall performance of consumer broadband for users testing across certain interconnection relationships would substantially decline during periods of peak use, in particular during evening hours. While the exclusion of potentially-affected measurements could change the appearance of the graphs, the overall findings of the interconnection research did not rely on absolute numbers and were less susceptible to false positives as a result of the initially cautious approach. The research also considered baselines of historical performance from alternative transit ISPs and access ISPs, in order to identify whether the source of degradation was potentially the result of failures that were solely within particular networks. The interconnection relationships highlighted in the research typically had comparative alternatives that performed within expectation, such as comparing Verizon over Cogent (congested) to Verizon over Internap (normal), or to Cablevision over Cogent (normal). Since the switch discard behavior would affect test clients and measurements similarly, independent of provider or technology, the methodology also reduced the likelihood that the switch behavior biased performance in favor of particular access ISPs. For example, had the behavior impacted all measurements against a site, such as Cogent in New York, then M-Lab would not have been able to identify consistent baselines such as Cablevision, as all clients would have encountered degradation.

M-Lab has used a temporary database of the results of the Discard Score equation for each NDT measurement in order to generate updated information and graphs related to the Interconnection Study. The most notable impact to the original report is the slight modification of graphs as a result of the exclusion of potentially-affected data. In both the interconnection research and the Internet Observatory, M-Lab set a threshold of thirty measurements per unit of time as a confidence level prior to consideration of datasets. Analysis and visualization of this data was then based on median values of the measurements. The exclusion of certain measurements from the pool of considered data may reduce the number of samples for certain graphs that are now below this standard threshold, and may change the median value

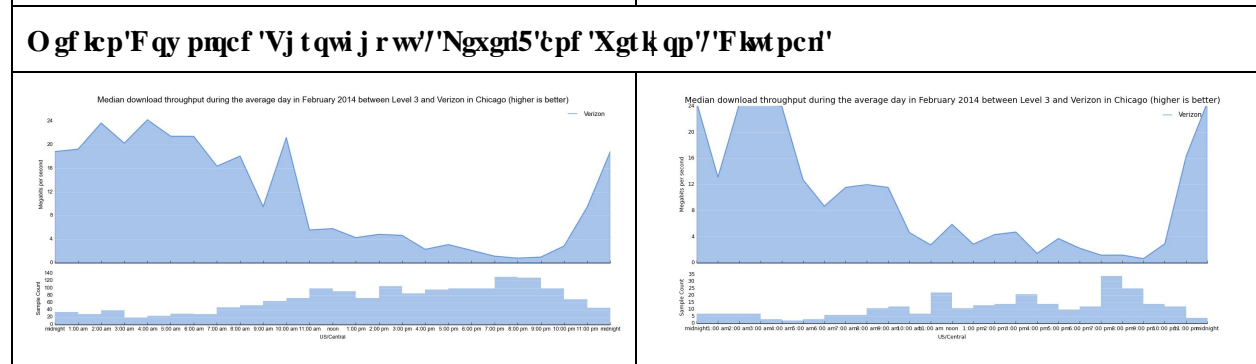
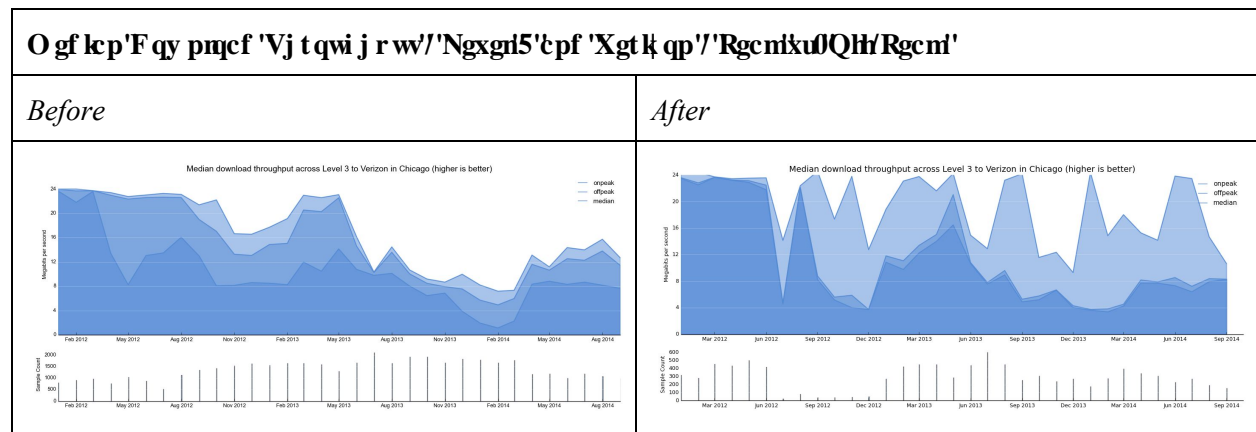
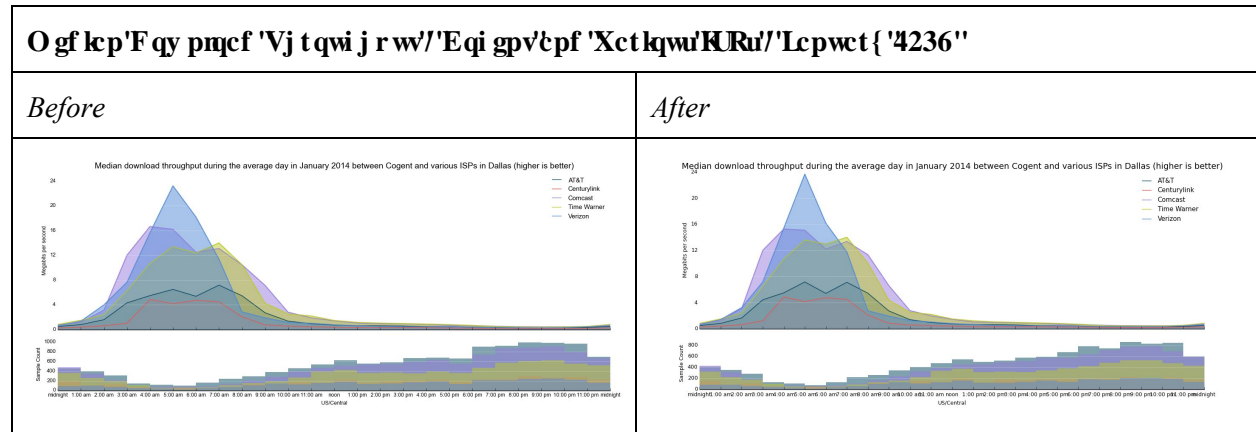
for a given timeframe. In cases where the number of samples per unit of time was already low, namely in early morning hours, the median may be subject to outlier values, resulting in less smooth graphs.

The analysis and graphs produced after the exclusion of potentially-affected measurements strongly reflect those of the original Interconnection Study, and confirm the original findings of sustained congestion across the United States. Only in the case of the peak versus off-peak graphs for the three sites hosted on the Level 3 network and the XO site in D.C., do we see a reduction in samples that could lead to exclusion of certain charts. As noted earlier, certain sites received more traffic than others based on their physical location, or the performance of their network for certain measurement clients. As the only sites in those regions at that time, the Level 3 and XO sites were amongst the most popular (until the deployment of more sites across the United States) and were therefore more impacted by the switch discard issue due to higher volumes of traffic. As a result of M-Lab's cautious approach in the Discard Score analysis, more measurements were excluded from these sites than from others, resulting in reduced insight in certain aspects of performance at particular times. Despite the potential changes in certain graphs leading to a less consistent appearance, diurnal trends for all sites, including Level 3 and XO in DC, still indicate the presence of the same level of performance degradation, and our original findings remain the same.

An archive of updated graphs from the Interconnection Study is available at:

<https://console.cloud.google.com/storage/browser/m-lab/interconnection-study-2014/>

Ecug'Uwvf { "	Ɔ rcev'
Cablevision, Comcast, TWC, and Verizon - Cogent and Internap - New York City	<i>No Change</i>
AT&T, CenturyLink, Comcast, Time Warner Cable, and Verizon - Cogent - Dallas, Los Angeles, and Seattle	<i>No Change</i>
Comcast, Time Warner Cable, and Verizon - Level 3 - Atlanta and Chicago	Same Diurnal Result, More Fluctuations in Peak vs. Off-Peak Graphs
Comcast and Time Warner Cable - XO - Washington, D.C.	Same Diurnal Result, Less Clear Peak vs. Off-Peak Graphs
AT&T over GTT in Atlanta, Chicago, and Los Angeles	<i>No Change</i>
Centurylink over Tata in Seattle	<i>No Change</i>
Comcast, Time Warner Cable, and Verizon over GTT and Tata	<i>No Change</i>
Verizon over GTT in Washington D.C.	<i>No Change</i>



Implications for Internet Performance

Measurement Lab’s infrastructure is supported by hardware and software that are commonly deployed in datacenters serving content to broadband users, reflecting the performance under real world conditions. Performance measurement currently requires sustained bursts of traffic that seeks to saturate the end user’s connectivity and reach congestion on the path, a behavior that approximates the quality of experience when accessing commonplace Internet services. Tests deployed on the M-Lab platform that attempt to emulate applications, such as streaming media performance, were more likely to incur discards as a result of their burstiness. As a result, the switch discard behaviors documented by M-Lab have

broader implications for Internet performance and protocol implementation requirements that warrant attention in measurement and networking communities.

Microbursts of traffic from high-throughput clients present a potential problem for core Internet hardware overall as access speeds increase. Previously, edge hardware and limits on performance effectively acted as traffic shapers, preventing large traffic bursts from meeting in the core of the network without being spread out in time. In the case of the Measurement Lab platform, concurrent high-throughput streams overloaded switch queues and caused auto-correlated packet drops, leading to path underutilization and preventing bursts from being experienced upstream. Modern edge hardware, such as next generation switches, can increasingly transmit traffic bursts that could overwhelm other queues in hardware that are in the path, including those nodes central to the network. The potential effect of that failure to handle overwhelming microbursts by spreading out flows is underperformance.

Traditional TCP is self-clocked, as it dynamically adapts the rate of packet transmission based on the speed of the network and client, in response to its peer sending acknowledgements (ACKs). Under most conditions data packets sent on the forward path are a reflection of the ACKs arriving on the return path, setting the pace of sending. Likewise the ACKs sent on the return path are a reflection of the data packets arriving at the receiver. In the absence of something to disturb packet spacings, each round of the TCP protocol conversation is an echo of the previous round, preserving the pattern of bursts and silences from one round to the next. When a connection starts (or restarts after some types of loss recovery episodes or pauses in application data) the slowstart algorithm is used to re-establish the self-clock. During slowstart, each ACK causes the sender to emit an additional number of packets, as much as twice as many as those that arrived at the receiver in the previous round, but at the spacing determined by the dominant bottleneck in the path, and returning ACKs. This implies that the sender will eventually create a queue at the bottleneck because it is sending data into the network faster than the data is leaving the network. The queue at the bottleneck drains at the full network speed in consecutive order. Since slowstart creates a queue, it tends to close any gaps that might have been caused by cross traffic at the bottleneck in a previous round trip.

The maximum rate for the data bursts is typically determined by the bottleneck, the slowest link on the connection. If the path to the client has a remote bottleneck, say for example a 10 Mbps modem, then slowstart consists of 15 or 20 Mbps data bursts that will become queued at the subscriber access link, sustained by the ACKs generated by the data arriving at the client at 10 Mbps. If there is no remote bottleneck (e.g., where the end-to-end path from the user to the server has a capacity of 1 Gbps), then the queue is formed within at the outbound interface of the M-Lab server, and slowstart consists of a sequence of full 1 Gbps line rate bursts. However, the sending rate is also not consistent for applications. Fixed rate streaming video applications are commonly implemented by periodically writing large chunks of data to the TCP socket buffer. For example streaming media at 4 Mbps might be accomplished by sending 1,000,000 Bytes of data every 2 seconds. To the servers on the M-Lab platform, each chunk of data causes a new slowstart, which consists of a sequence of microbursts at either some rate determined by a bottleneck modem or at full interface line rate. This pattern of repeated bursts on a single stream greatly increases the probability of a full queue on the switch upload port. When that queue is full the switch will drop packets from one or more outgoing streams, even if the queue fullness is an extremely

transient condition. It is these drops, caused by the transiently full queue, which can cause the path to be systemically underutilized.

There are several mechanisms that drive this phenomena, and some that help to mitigate it. Mainline Linux now supports "paced TCP," where packets can be scheduled on fine-grained timers reducing the significance of the bursts by adding space between outgoing packets so that other streams may interleave with the burst in the switch queue. Although pacing is present in mainline Linux, it is a relatively new feature, not on by default, and has a mixed track record in research literature.⁸ However, pacing provides a strong antidote to the self-preserving bursts caused by slowstart and the traditional TCP self-clock. Pacing can be used to mitigate the effects of microbursts by spreading out the packets at the sender, such that they do not generally cause queues in the network, except when the average sending rate exceeds the bottleneck rate.⁹

Disclosure Policy and Contact"

Measurement Lab seeks to maintain a high-level of confidence with its partners and the research community regarding the integrity of the data collected and stored via the M-Lab platform. M-Lab is a complex system, requiring the maintenance of measurement tools, the administration of server infrastructure and the provision of a data extraction pipeline that processes large volumes of data on a daily basis. M-Lab has developed and deployed systems for monitoring the platform and validation of measurement tools, which it seeks to make publicly-available where possible. This scrutiny does not preclude the possibility of errors, as anyone who has managed a complex system can attest. When M-Lab becomes aware of platform issues, it makes every effort to resolve issues in a timely and transparent manner.

Where M-Lab determines that an external factor could have impacted the collection and analysis of any data, it commits to:

- A description of the incident posted publicly, including the identification of technical causes behind the episode, any potential impacts on the data, steps taken to resolve the matter, and measures to prevent future issues;
- The notification of the M-Lab community through the M-Lab discussion list;¹⁰ and
- If needed, the demarcation of potentially impacted data in BigQuery, where necessary and possible.

Summary of Actions Taken

Upon notification of inconsistent performance to the platform in mid-August 2015, M-Lab conducted a thorough investigation of the incident and identified the underlying causes. This led to immediate steps to

⁸ The advantages and disadvantages of pacing are out of scope for this document, and the literature has addressed a number of aspects of this issue. <https://www.usenix.org/system/files/conference/atc12/atc12-final236.pdf>

⁹ The maximum achievable performance at the slowest link across the end-to-end path from the user to M-Lab.

¹⁰ <https://groups.google.com/a/measurementlab.net/forum/#!forum/discuss>

mitigate the root cause of the behavior. As a result, the inconsistent performance has not affected M-Lab sites in the United States since October 11, 2015. M-Lab takes potential problems with platform seriously, and have taken the following actions to address this issue:

- Reduced sites temporarily to one active server in order to prevent switch buffer overruns;
- Defined a Discard Score metric based on a machine learning process in order to delineate which measurement could have been conducted during windows of inconsistent performance;
- Posted this disclosure and notified M-Lab researchers of our data via the discuss@measurementlab.net mailing list; and,
- Published amended graphs from the M-Lab 2014 Interconnection Study and subsequent blogpost on Google Cloud Storage;

Additionally, in order to improve consistency and resiliency in the long term, M-Lab:

- Has established new monitoring processes to collect site performance information at a finer granularity;
- Has reprocessed the affected measurements in the BigQuery dataset to mark potentially affected data based on the Discard Score metric;
- Is evaluating next generation switches, which will also provide an opportunity to support upstream connectivity above 1 Gbps.
- Has built a more expansive testing lab, where servers and test clients can be evaluated under a wider range of emulated network environments, including stress testing;
- Is deploying a new resource telemetry pipeline, which will be integrated into the data collection pipeline for labeling tests that might have experienced resource contention, such as packet drops caused by switch buffer contention; and,
- Is evaluating additional measures to improve measurement calibration.

For any questions or concerns about the integrity of Measurement Lab dataset, the potential impact of previous performance incidents or the operations of measurement clients, please contact: support@measurementlab.net