

Network Working Group  
Internet-Draft  
Intended status: Informational  
Expires: January 7, 2013

G. White  
J-F. Mule  
D. Rice  
CableLabs  
July 6, 2012

**Analysis of SPDY and TCP Initcwnd  
draft-white-httpbis-spdy-analysis-00**

**Abstract**

Making the Internet faster is the goal of many product and service companies. Many strategies exist, from data caching to packet switching, performance improvements in Internet browser clients and server software, without forgetting the transport network scaling from the fiber core to the access edges.

This report investigates two proposed protocol enhancements aimed at making the web browsing user experience better by optimizing the transport of web objects and thereby reducing the page load time. The two enhancements are SPDY, a replacement of the HyperText Transfer Protocol (HTTP) requiring client and server changes, and a TCP tune-up, an increase in the TCP initial congestion window accomplished by a setting change on the web servers. Both show promise, but there are some caveats, particularly with SPDY. SPDY is a candidate for standardization as HTTP/2.0.

In addition to a discussion of the two enhancements, this report provides the results of laboratory testing on SPDY version 2 and the proposed increase in the TCP initial congestion window in a variety of simulated conditions, comparing the page load time when using the proposed enhancement to the page load time for the default case of HTTPS and current initial congestion window settings.

The proposed enhancements generate mixed results: web page load times were reduced in some scenarios but increased significantly in others. The performance improvement (or degradation) varied depending on the number of servers, configuration of the initial TCP congestion window, and especially any network packet loss. The following results were obtained across all scenarios comparing SPDY and congestion window enhancements to standard HTTPS.

- o Average reduction in page load time was 29%
- o Best improvement was over 78% reduction in page load time

- o Worst cases showed a negative impact, resulting in a 3.3x increase in page load time

These results lead us to the following conclusions:

- o The SPDY protocol is currently a moving target, and thus it would be challenging to realize a return on investment for general-purpose usage in web servers.
- o Protocol improvements, standardization in the IETF and wider adoption by the client/server software may warrant a second look at SPDY.
- o Some applications in controlled environments may gain by leveraging SPDY. SPDY might be a valuable tool where the a single entity provides the servers, the client software, and the web content.
- o If SPDY were adopted very widely it may have some secondary benefits for network operators through improved infrastructure scalability due to a significant reduction in concurrent TCP sessions, as well as a reduction in Packets Per Second.
- o The proposed increase in the TCP initial congestion window is straightforward, requires no client modifications, and on its own provides consistent (albeit modest) performance gains.

This report is available in a somewhat more detailed form in [\[SPDY-ANALYSIS\]](#).

#### Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 7, 2013.

#### Copyright Notice



Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.



## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction . . . . .</a>	<a href="#">5</a>
<a href="#">2.</a>	<a href="#">SPDY . . . . .</a>	<a href="#">6</a>
<a href="#">2.1.</a>	<a href="#">The SPDY Protocol . . . . .</a>	<a href="#">7</a>
<a href="#">2.2.</a>	<a href="#">SPDY Server Implementations . . . . .</a>	<a href="#">9</a>
<a href="#">2.3.</a>	<a href="#">SPDY Client Implementations . . . . .</a>	<a href="#">9</a>
<a href="#">2.4.</a>	<a href="#">SPDY in the Wild . . . . .</a>	<a href="#">9</a>
<a href="#">2.5.</a>	<a href="#">SPDY Protocol Development and IETF Standardization . . . . .</a>	<a href="#">9</a>
<a href="#">3.</a>	<a href="#">TCP Initial Congestion Window . . . . .</a>	<a href="#">10</a>
<a href="#">3.1.</a>	<a href="#">Historical Settings for initcwnd . . . . .</a>	<a href="#">10</a>
<a href="#">4.</a>	<a href="#">The Current Web Environment . . . . .</a>	<a href="#">11</a>
<a href="#">4.1.</a>	<a href="#">Content Environment . . . . .</a>	<a href="#">11</a>
<a href="#">4.2.</a>	<a href="#">Network Environment . . . . .</a>	<a href="#">13</a>
<a href="#">4.2.1.</a>	<a href="#">Round-Trip Time (RTT) . . . . .</a>	<a href="#">13</a>
<a href="#">4.2.2.</a>	<a href="#">Access Network Data Rate . . . . .</a>	<a href="#">15</a>
<a href="#">4.2.3.</a>	<a href="#">Packet Loss . . . . .</a>	<a href="#">15</a>
<a href="#">5.</a>	<a href="#">Laboratory Testing and Results . . . . .</a>	<a href="#">16</a>
<a href="#">5.1.</a>	<a href="#">Goals . . . . .</a>	<a href="#">16</a>
<a href="#">5.2.</a>	<a href="#">Laboratory Environment . . . . .</a>	<a href="#">16</a>
<a href="#">5.2.1.</a>	<a href="#">Web Server Configuration . . . . .</a>	<a href="#">16</a>
<a href="#">5.2.2.</a>	<a href="#">Client Configurations . . . . .</a>	<a href="#">17</a>
<a href="#">5.3.</a>	<a href="#">Test Conditions . . . . .</a>	<a href="#">17</a>
<a href="#">5.3.1.</a>	<a href="#">Protocol Options . . . . .</a>	<a href="#">17</a>
<a href="#">5.3.2.</a>	<a href="#">Web Sites . . . . .</a>	<a href="#">18</a>
<a href="#">5.4.</a>	<a href="#">Test Results . . . . .</a>	<a href="#">20</a>
<a href="#">5.4.1.</a>	<a href="#">CASE 1: SPDY vs. HTTPS . . . . .</a>	<a href="#">21</a>
<a href="#">5.4.2.</a>	<a href="#">CASE 2: initcwnd=10 vs. initcwnd=3 . . . . .</a>	<a href="#">24</a>
<a href="#">5.4.3.</a>	<a href="#">CASE 3: SPDY+initcwnd=10 vs. HTTPS+initcwnd=3 . . . . .</a>	<a href="#">27</a>
<a href="#">5.4.4.</a>	<a href="#">Results Summary . . . . .</a>	<a href="#">29</a>
<a href="#">6.</a>	<a href="#">Recommendations and Conclusions . . . . .</a>	<a href="#">30</a>
<a href="#">7.</a>	<a href="#">IANA Considerations . . . . .</a>	<a href="#">31</a>
<a href="#">8.</a>	<a href="#">Security Considerations . . . . .</a>	<a href="#">31</a>
<a href="#">9.</a>	<a href="#">Acknowledgements . . . . .</a>	<a href="#">31</a>
<a href="#">10.</a>	<a href="#">References . . . . .</a>	<a href="#">32</a>
	<a href="#">Authors' Addresses . . . . .</a>	<a href="#">33</a>



## 1. Introduction

The current method of transporting web objects from a server to a client utilizes Hypertext Transfer Protocol version 1.1 (HTTP/1.1), running atop the Transport Control Protocol (TCP) [[RFC0793](#)]. HTTP/1.1 was published as [[RFC2616](#)] in 1999, and has since become the most widely used application-layer protocol on the Internet. TCP pre-dates HTTP/1.1 by about 18 years, and even though TCP has evolved over time, it is fundamentally unchanged, and runs atop IPv4 and IPv6.

Since the advent of the World Wide Web 15+ years ago, access network bandwidths as well as server and client CPU horsepower have increased by a factor of approximately 1.5x year-over-year, yet the apparent speed of the web (from the web browsing user's perspective) has grown much more slowly. In part, this can be explained by a concomitant increase in web page complexity, as measured by any number of attributes including total page size, number of linked resources, amount of server-side code, and lines of JavaScript. However, there is a view that the underlying protocols used to transport web resources are becoming increasingly out-of-date with the network and computing resources that are in use today, and that significant improvements in performance (generally page load time) can be achieved by revisiting these fundamental technologies.

Bolstering this view is the fact that while network bandwidths have been improving rapidly, network latencies have not, and there is little hope of achieving significant reductions in network latency, as it is dominated by propagation delay. The result is that a fundamental network parameter, the Bandwidth-Delay Product (BDP), is much larger today than in the early days of the Internet. It is becoming clear that HTTP/1.1 and TCP are not particularly optimized for networks with high bandwidth-delay product.

Furthermore, in order to maximize performance using the existing tools of HTTP/1.1 and TCP, web architects (browser, server and site developers) have employed work-arounds that have some negative implications. The most significant implications result from the use of multiple simultaneous TCP connections. Modern browsers will open up to six simultaneous TCP connections to each server from which they need to retrieve resources, and optimized websites will compound this by spreading content over multiple servers, a practice known as "domain sharding". The result is that a single browser may have 20 or more simultaneous TCP connections to the hosts of a particular site while it is downloading a single web page. The parallel nature of this approach is intended to reduce the page download time (compared to the alternative of a single, non-pipelined TCP connection), and in many cases it succeeds. However, a web browser





with 20 or more simultaneous TCP sessions can drown out other applications that use a small number of TCP sessions (e.g., one). This is due to the fact that the TCP congestion avoidance algorithm provides approximately fair sharing of the bandwidth on the bottleneck link on a per-TCP-session basis. Additionally, some browsers do not reuse TCP connections, so each web resource is retrieved using a separate TCP connection. The result is that the vast majority of sessions never reach the congestion avoidance phase, and network resources can therefore either be underutilized or excessively congested. The network overhead for establishing and tearing down all of the TCP connections can also be a concern.

A number of efforts have sprung up to develop ways to improve page download time and consequently improve the web browsing user experience. This paper discusses two proposed enhancements, both originally proposed by Google as part of their "Let's Make the Web Faster" project [[Faster](#)]. The first is a replacement for HTTP 1.1, called SPDY [[SPDY](#)], which aims to make more efficient use of the network in common web browsing scenarios. The second is an incremental enhancement to the TCP protocol in the form of an increase in the server's initial congestion window (`initcwnd`).

In addition to a discussion of the two proposed enhancements, this paper presents the results of laboratory testing conducted by CableLabs on both technologies independently, and on the combination of the two technologies.

## **[2.](#) SPDY**

SPDY [[SPDY](#)] is an experimental protocol developed by Mike Belshe and Roberto Peon of Google in late 2009 to replace HTTP/1.1 communication for transporting web content between a client and a server. It is deployed in production on many Google servers but requires a compatible browser such as Google Chrome. The stated goals of SPDY are:

- o 50% reduction in page load time
- o Minimize deployment complexity
- o Avoid changes to content by web authors
- o Solve this collaboratively via open-source software

At the time of this study, the current implemented version of SPDY is version 2 (SPDY/2). Version 3 is currently being developed.



### **2.1. The SPDY Protocol**

SPDY/2 retains the HTTP/1.1 metadata, but replaces the transport aspects of HTTP with a more streamlined approach. SPDY has three basic features that are used to accelerate the loading of a page:

- o Multiplexed streams - The fundamental enhancement of SPDY is that multiple resources can be retrieved via a single TCP connection. The expectation (and current implementation in Chrome) is for the client to open a single TCP connection to each server, and to request all resources of the server over that single connection. The use of a single TCP connection in this way allows the congestion avoidance algorithm in TCP to more effectively manage data flow across the network. Also, the client can include multiple requests in a single message, thereby reducing overhead and the number of round-trip times necessary to begin file transfer for requests beyond the first. While this is similar to HTTP pipelining [[HTTP Pipelining](#)], the difference introduced by SPDY is that the server can transfer all of the resources in parallel (multiplexed) without the "head-of-line blocking" problem that can occur with HTTP pipelining.
- o Request prioritization - While SPDY may serve to decrease total page load time, one side-effect of stream multiplexing is that a critical page resource might be delivered more slowly due to the concurrent delivery of a less critical resource. To prevent this, SPDY provides a way for the client to indicate relative priorities for the requested resources. For example, if a JavaScript library is required in order to generate URLs for additional page resources, the client can request that the library be delivered with higher priority so that it isn't holding up those subsequent requests.
- o HTTP header compression - SPDY mandates that HTTP headers be compressed to reduce the number of bytes transferred. HTTP Header compression has been an Internet standard but it is not widely used. Google SPDY makes it mandatory. This might be a fairly small gain on the response headers, but it can provide a significant benefit on the requests (which in many cases are entirely headers). For example, each request from a particular client to a particular server includes an identical user-agent string (which can be in excess of 100 bytes), and in some cases the same cookie is sent in each request. Both Chromium and Firefox developers have reported approximately 90% header compression using the proposed zlib compression. This could allow a lot more requests to be packed into each request packet.

SPDY also has two advanced features that can be used in certain cases



to further accelerate the web user experience:

- o Server Push - SPDY allows the server to create a stream to the client, and via this stream send web resources that were not explicitly requested by the client. The expectation is that the client will cache the pushed resource, and then upon needing it, will retrieve it from local cache. This function could potentially be used by the server to push objects (of which the client isn't yet aware) that are necessary for rendering the current page. Additionally, this function could be used to push resources for related pages that the user may be likely to request. The heuristics used by the server to decide when/if to push objects and what those objects are is left to the server implementer. This feature is somewhat controversial, but the authors defend it by pointing out that it is better than the practice of "in-lining" resources into the html page, since it allows the pushed resources to be cached for multiple future usages.
- o Server Hint - SPDY defines a new header that the server can use to suggest additional resources that the client should request. This is a less forceful approach than the Server Push, and allows the client to be involved in the decision whether or not a particular resource is delivered. The client might, for example, examine its own cache and only request resources that are not resident in local cache. On the other hand, Server Hint theoretically requires one additional round trip that would be eliminated by Server Push.

While it isn't fundamentally required for the protocol to function, in practice SPDY/2 runs only over an encrypted (TLS) connection. The rationale for this is three-fold:

- o TLS involves a client-server handshake to negotiate cipher-suite. The authors of SPDY extend this handshake to negotiate whether SPDY can be used. This Next Protocol Negotiation (NPN) [[I-D.agl-tls-nextprotoneg](#)] allows the use of the https:// URI, rather than a new spdy:// URI, and as a result a single html page can work for clients that support SPDY and clients that don't.
- o TLS passes through firewalls and bypasses intermediaries. Many HTTP requests today are processed (and modified) by transparent proxies without the knowledge of the end-user. It would create a tremendous barrier to adoption of SPDY if it were necessary for intermediaries to be upgraded to handle the new protocol.
- o Encryption is good. The authors of SPDY state a philosophical belief that all HTTP traffic should be encrypted. They cite the



relative ease by which traffic (particularly Wi-Fi traffic) can be snooped.

## **2.2. SPDY Server Implementations**

There are a number of SPDY server implementations that are in various stages of development as open source projects. The Chromium project maintains a list of the implementations available [[SPDY](#)]. At the time this study was initiated, many of these implementations supported a subset of SPDY functionality and/or supported SPDY version 1. At the time this study began, only two implementations appeared to support SPDY/2 in a reliable way. The first is the Chromium "FLIP Server" (FLIP was an early internal code-name for SPDY within Google). This server is built off of the immense Chromium source tree. The second is a Javascript implementation called "node-spdy" which is built on the node.js server framework. There is also an Apache module for SPDY, but at the time this study was initiated it was incomplete.

## **2.3. SPDY Client Implementations**

On the client side, both the Chrome browser and the Mozilla Firefox browser support SPDY/2. Chrome also includes a built-in diagnostics function that allows the user to examine the active SPDY sessions `<chrome://net-internals/#spdy>`. Finally, the Silk browser in the Amazon Kindle Fire tablet computer purportedly utilizes SPDY/2 for its connection to the Amazon EC2 cloud when performing web acceleration.

## **2.4. SPDY in the Wild**

In terms of live web content, the most significant, publicly available sites that serve content via SPDY are run by Google. Many of the Google properties have SPDY/2 enabled, including Gmail, Google Docs, Picasa, Google+, and Google Encrypted Search. All of these sites utilize only the basic SPDY features; there are no known live instances of Server Push or Server Hint. In addition, the Twitter and Wordpress websites utilize SPDY, and the web acceleration companies Cotendo (acquired by Akamai in Dec. 2011) and Strangeloop indicate that they have deployed SPDY in some capacity.

## **2.5. SPDY Protocol Development and IETF Standardization**

Google has actively solicited input on enhancements to the SPDY/2 protocol. Up until very recently, development and discussion of SPDY/3 has taken place on an open, Google-managed forum. However, the SPDY/3 draft was submitted to the IETF HTTPbis working group on February 23, 2012, for comments [[I-D.mbelshe-httpbis-spdy](#)].





### 3. TCP Initial Congestion Window

The TCP initial congestion window (`initcwnd`) is used at the start of a TCP connection. In the context of an HTTP session, the server's `initcwnd` setting controls how many data packets will be sent in the first burst of data from the server. It is a standard protocol parameter that can be changed on Linux servers via a simple command line.

Absent packet loss or receiver window limits, the TCP slow start operation looks like:

Table 1. TCP Slow Start Operation

Round-Trip #	Client	Server
1	TCP SYN	TCP SYN/ACK
2	TCP ACK and HTTP GET <url>	<code>initcwnd</code> data packets
3	TCP ACKs	$2 * \text{initcwnd}$ data packets
4	TCP ACKs	$4 * \text{initcwnd}$ data packets
		and so on....

#### 3.1. Historical Settings for `initcwnd`

A larger value for `initcwnd` will clearly result in fewer Round-Trip Times (RTTs) to deliver a file. However, the downside to an excessively large `initcwnd` is that there is an increased risk of overflowing a router buffer on an intermediate hop, resulting in an increase in latency from packet loss and retransmissions. A value of `initcwnd` that is greater than the Bandwidth-Delay Product (BDP) of the network path between the server and client has an increased likelihood of causing packet loss that may lead to poor performance. As network bandwidths have increased over time, the BDP has increased, and as a result the defined value for `initcwnd` has been adjusted. The early specifications for TCP ([RFC1122], [RFC2001]) required that a TCP implementation set `initcwnd` to 1 packet. Starting in 2002 ([RFC2414], [RFC3390]), the `initcwnd` value was raised to 4000 bytes (effectively 3 packets in most networks).

Google has submitted an IETF draft [I-D.ietf-tcpm-initcwnd] proposing that `initcwnd` now be increased to at least 10 packets, based on a series of tests performed using production Google servers



[[INITCWND-PAPER](#)]. In Google's tests, a value of 10-16 packets resulted in the minimum average latency for delivery of web resources.

In November 2011, CDN Planet [[CDNPLANET](#)] performed an assessment of CDNs to see what value is currently used for initcwnd. Many CDNs have already increased their initcwnd beyond 3 packets, some as high as 10-16 packets.

This increase from a value of 3 to a value of 10 will result in the elimination of up to 2 RTTs for each file transfer. The maximum gain will be seen for files that take 9 packets to deliver (i.e., files around 11 kB in size). Files of that size would take 4 RTTs to deliver using the default initcwnd, but can be delivered in 2 RTTs with the proposed increase, a 50% improvement. Files smaller than about 3 kB will not experience any acceleration.

## **4. The Current Web Environment**

### **4.1. Content Environment**

The construction of web page content plays a very important role in determining the page load time. In order to validate the two proposed enhancements, we will create a test bed that consists of web servers, web clients and web content. We performed a survey of some popular (and resource heavy) websites in order to understand the content environment. The sites used in the survey were selected as examples of sites that would stand to benefit from an acceleration technology such as SPDY. Many of these sites are composed of a fairly large number of resources, the content changes frequently (reducing local cache acceleration), and they would presumably be interested in page load performance.

The results shown in Table 2 indicate, for each site's home page, the number of servers from which content was drawn, the number of resources requested, the total transferred size of all page resources, and the page load time. The reported page load time is the total time to download all of the resources for the page. It is possible, and even likely, that one or more of these pages are largely complete (perhaps only missing some non-critical resources) in significantly less time than is reported here.

Additionally, for each webpage we evaluated the percentage of total resources that were requested from the top N servers for the page, for values of N between 1 and 5. For example, Engadget receives 34% of the web objects from the most used server, 82% of the objects from 3 most used servers and 87% of them come from the 5 servers with the



highest count. These servers may include the host servers for the website, CDN servers for images and media content, advertising servers for ad content, and analytics collection, among other types of resources.

Table 2. Website Survey

Page	servers	GETs	total page size (KB)	total time (s)	% of GETs from top N servers				
					N=1	N=2	N=3	N=4	N=5
Engadget	26	278	1500	23.0	34%	67%	82%	85%	87%
NY Times	27	148	1500	13.6	33%	49%	59%	67%	73%
Gizmodo	25	116	3900	13.2	34%	60%	69%	76%	78%
CNN	22	158	1180	12.6	59%	70%	75%	80%	83%
comcast.net	12	55	606	12.0	51%	67%	75%	80%	84%
ESPN	26	144	5400	11.6	40%	51%	60%	67%	70%
Amazon	15	139	1100	11.0	34%	58%	81%	86%	90%
RoadRunner	30	128	85	9.50	27%	48%	61%	65%	68%
Wired	32	130	1200	8.90	48%	56%	61%	65%	68%
eBay	14	53	520	8.60	23%	40%	55%	68%	75%
LinkedIn	9	75	457	7.17	48%	80%	88%	92%	95%
Hulu	12	193	2000	4.96	60%	79%	87%	95%	96%
Yahoo	6	51	423	4.60	88%	92%	94%	96%	98%
YouTube	6	37	2400	4.16	41%	68%	81%	92%	97%
Google   shopping	3	20	73	2.45	45%	90%	100%	100%	100%
Average	18	115	1490	9.82	44%	65%	75%	81%	84%



## **4.2. Network Environment**

The network environment between the servers and the client plays an equally large role in determining page load time. The most important parameters that drive performance are the round-trip time, the bottleneck link bandwidth, and the packet loss rate.

### **4.2.1. Round-Trip Time (RTT)**

In the absence of congestion, typical Internet round-trip times can range from as low as 10 milliseconds when accessing content served in close proximity, to as high as 500 milliseconds or more when accessing servers across continents. For wireline broadband customers in the U.S., the most popular websites can generally be reached with RTTs ranging from 15 ms to 150 ms. For the websites included in the survey in [Section 4.1](#), a sample TCP RTT to each of the top-five servers (in terms of number of resources requested) from CableLabs headquarters in Louisville, Colorado, is shown (in milliseconds) in Table 3 along with a summary of the minimum, mean, and max RTT of those top-five servers. The RTT was calculated by doing a TCP "ACK RTT" analysis of a packet capture of each page load using the Wireshark tool. The ISP connection used during this testing consisted of two load-balanced 100 Mbps duplex links with a mean RTT of less than 2 ms.





Table 3. Round-Trip Times to Website Servers in Milliseconds

Page	server 1	server 2	server 3	server 4	server 5	min	average	max
Engadget	63*	62*	64*	62*	38	38	57.8	64
NY Times	64*	64*	62*	64*	64*	62	63.6	64
Gizmodo	37*	39*	134	25	22	22	51.4	134
CNN	46*	53	69*	32	45	32	49.0	69
Comcast.net	65*	69*	94	28	71	28	65.4	94
ESPN	64*	33	65*	64*	64*	33	58.0	65
Amazon	35	72*	16*	98	33	16	50.8	98
RoadRunner	56	58	29*	52*	61*	29	51.2	61
Wired	70*	114	66*	24*	85	24	71.8	114
eBay	62*	75	64*	62*	62	62	65.0	75
LinkedIn	50*	72*	78	54	63*	50	63.4	78
Hulu	65*	67*	66*	64*	22	22	56.8	67
Yahoo	28	27	64	29	61	27	41.8	64
YouTube	23	25	23	22	22	22	23.0	25
Google shopping	34	42	43	22	-	22	35.3	43

\* denotes a host run by a CDN provider.

When network links are congested, a phenomenon referred to as "bufferbloat" can result in an increase in RTT on the order of hundreds of milliseconds. Bufferbloat arises due to the fact that many network elements support much more buffering than is needed to maintain full link utilization. These oversized buffers will be kept full by the TCP (particularly when there are multiple simultaneous sessions), thereby causing an increase in latency. In recent updates to the DOCSIS 3.0 cable modem specification, a new feature has been



added to mitigate this effect by proper sizing of the cable modem's upstream buffer. At this time, 14 cable modem models from 10 manufacturers have been certified as supporting this feature. For more detail on bufferbloat and the DOCSIS Buffer Control feature, see: [[BUFFERCONTROL](#)].

#### **4.2.2. Access Network Data Rate**

Residential access network data rates have steadily risen over time, with a 10 Mbps downstream rate and 1 Mbps upstream rate being fairly common in North America, and 20 Mbps x 2 Mbps rates becoming increasingly available.

As a result, for purposes of our investigation of the two proposed technologies, we will utilize the 10x1 and 20x2 data rate configurations.

Some access networks, in particular cable modem access networks, are able to provide higher burst data rates that go well beyond the sustained traffic rate for the connection. In this report we will not attempt to study the impact that these "Powerboost" type features have on the two proposed technologies. In addition, many operators offer higher speed tiers than those mentioned above, with speeds in Europe up to 360 Mbps on the very high end, but they currently have fairly low take rates. Finally, some access networks may have lower speed limitations placed on the services such as public and outdoor Wi-Fi networks.

#### **4.2.3. Packet Loss**

It is important to note that packet loss due to router (or switch) buffer overflow is expected behavior in networks, particularly when TCP is utilized. In fact, packet loss due to buffer overflow is fundamental to the congestion avoidance algorithm in TCP; it is the only signal that the TCP has (absent the Explicit Congestion Notification field in IP packets) that it has saturated the link and needs to reduce its transmission rate. As a result, it is not a concern in this testing. The concern here is random packet loss due to noise, interference or network faults that will effectively send TCP an erroneous signal to reduce its transmission rate.

In wired networks, packet loss due to noise is fairly uncommon, for example with packet loss rates of approximately  $10^{-5}$  being typical for DOCSIS cable modem networks. In wireless networks, packet loss can vary dramatically as a result of interference and fluctuations in carrier-to-noise ratio.



## **5. Laboratory Testing and Results**

### **5.1. Goals**

The goals of the testing are to investigate the page load performance impacts of the two proposed technologies. In particular, we examine Google SPDY with an eye toward evaluating SPDY's ability to achieve the stated performance goal and to understand what aspects of SPDY provide acceleration and in which scenarios. Testing is limited to the SPDY "basic" features; no testing of server-push or server-hint is included in this study, primarily due to the fact that the performance enhancement gained by using them would be heavily dependent on how they are configured and utilized by the website administrator. Some benchmarks of server-push and server-hint have been published by external sources; see [[GOOGLE-BENCHMARKS](#)]. The results have not shown a statistically significant gain.

As stated previously, the TCP initcwnd increase is tested both independently (using HTTPS) and in conjunction with SPDY.

### **5.2. Laboratory Environment**

The entire test network was virtualized on a dedicated VMWare ESX Virtual Server. The test network consisted of 7 virtual hosts, each having two network interfaces. A single network interface on each host was connected to a virtual LAN within the ESX environment. The second network interface on each host was connected to the corporate network. During the experiments, the CPU load on the virtualization server was checked to ensure that it was not impacting the results.

#### **5.2.1. Web Server Configuration**

Four of the hosts were configured identically as web servers, with the following software:

- o OS: Ubuntu 11.04
- o Web Server: "Express-SPDY" written in JavaScript for node.js

The "Express-SPDY" server is the melding of a compact and efficient open-source HTTPS server (known as "Express") with an open-source SPDY/2 implementation called "node-spdy". The implementation is written in JavaScript using the node.js server environment. This implementation was chosen over the Chromium package published by Google due to the fact that it was lightweight and could be deployed relatively easily. Unfortunately, the version of Express-SPDY available at the time of this testing was found to not be stable, and required some work to debug and fix prior to beginning any testing.



Since our testing commenced, the author of node-spdy rewrote a large part of the package.

### **5.2.2. Client Configurations**

The remaining three hosts were configured as client machines. The primary client ran Windows 7, and the other two clients (used for spot testing) ran Windows XP and Ubuntu 11.04. All clients used the Chrome 16 browser, with the SPDY Benchmarking plugin, and were installed with Wireshark to capture network traffic. The Dummynet [[DUMMYNET](#)] network simulator was used to simulate the various network conditions.

For each test condition, Dummynet was configured with the desired upstream and downstream rates, with the desired RTT split equally between the two links. In cases where packet loss was being tested, the same packet loss rate was configured on both links. Both links were left with the default buffering configuration of 50 packets.

Note that two of the hosts (Win7 and Ubuntu) support the TCP Window Scaling Option by default, while the third (WinXP) does not.

## **5.3. Test Conditions**

### **5.3.1. Protocol Options**

As previously stated, the testing is intended to compare the SPDY protocol to the HTTPS protocol, and to investigate the use of the increased TCP Initial Congestion Window (`initcwnd`) (both with HTTPS and in conjunction with SPDY). As a result, we measured the time it took to load a web page using four different protocol options as shown in Table 4. One option defines the baseline scenario or "Base Case": with HTTPS for transport and the default `initcwnd` setting of 3. The results for the three cases involving a proposed enhancement are then compared to the "Base Case".

Table 4. Matrix of Protocol Options

	+-----+-----+
	HTTPS         SPDY
+-----+	+-----+
initcwnd = 3	Base Case         CASE 1
+-----+	+-----+
initcwnd = 10	CASE 2         CASE 3
+-----+	+-----+





### **5.3.2. Web Sites**

Nine different "websites" were used to evaluate the impact that the content and server configuration would have on the test results. The nine websites were formed by the 3x3 matrix of three different server configurations and three different web pages as follows.

#### **5.3.2.1. Server Configurations**

The server configuration parameter determined how all of the web resources for a page were served.

- o Server Case 1, all resources were served by a single server.
- o Server Case 2, the resources were divided equally across two servers.
- o Server Case 3, the resources were divided equally across four servers.

##### **5.3.2.1.1. Web Pages**

Three web pages were used in testing. These pages were chosen to span a range of possible page configurations. In particular, Page B was modeled based on the data collected in the website survey ([Section 4.1](#)), while Pages A and C represent more extreme situations. Page A is the smallest of the three pages, and is populated by a number of small image files. Page C is the largest, and is populated by a small number of very large image files. It could be argued that a page like Page C would be fairly rare on the Internet compared to pages like Page A or Page B.

Page A consisted of 102 total resources, with a total page size of 369 KB. It was composed of the following resources:

- o 1 HTML file: 12.8 KB
- o 1 JavaScript Library (jquery.js): 94.5 KB
- o 100 JPEG images: min/mean/max/stddev size = 1.3/2.6/5.5/1.2 KB

Page B consisted of 101 total resources, with a total page size of 1.4 MB. It was composed of the following resources:

- o 1 HTML file: 8.2 KB
- o 100 JPEG/PNG/GIF images:



- \* min/mean/max/stddev size = 46B/14KB/103KB/24KB
- \* Approximately log-uniform size distribution

Page C consisted of 11 total resources, with a total page size of 3.0 MB. It was composed of the following resources:

- o 1 HTML file: 0.8 KB
- o 10 JPEG images: min/mean/max/stddev size = 298/302/310/4 KB

#### **5.3.2.2. Channel Conditions**

Eight different channel conditions were used in the testing. These eight cases were formed from the 2x2x2 matrix of the following parameters:

Configured Data Rate:

- 10 Mbps downstream, 1 Mbps upstream
- 20 Mbps downstream, 2 Mbps upstream

Round Trip Time:

- 20 ms
- 100 ms

Packet Loss Rate:

- 0%
- 1%

The two Configured Data Rates were chosen to match two commonly used residential high-speed data service configurations. The Round Trip Times correspond to a fairly short RTT that models a user accessing a site that is located fairly close to them (or is accelerated by a CDN), as well as a longer RTT which models a U.S. coast-to-coast page load or a case where there is increased upstream latency due to bufferbloat. Performance in higher RTT cases, such as an overseas connection, is not considered in this study. The two packet loss rates correspond to an idealistic case that might be approached by an end-to-end wired network connection and a case with packet loss that might be more typical of a network connection that includes a Wi-Fi link.



#### **5.3.2.3. Test Execution**

The test execution will be performed in the following fashion:

- o Run all 288 test conditions on Win 7
- o Run spot tests on WinXP and Ubuntu
- o For each test condition, perform 20 page loads (clearing cache and resetting connections before each load) and calculate median load time.

#### **5.4. Test Results**

The result of the spot tests on WinXP and Ubuntu were largely consistent with the results for the Win7 testing. However, there was a subset of the test cases with WinXP where SPDY showed degraded performance compared to the other OSes. This occurred with the cases that used a single server, 100 ms RTT, and 0% packet loss. This is most likely a result of the lack of TCP Window Scaling in the WinXP client, which limits the data rate of each TCP connection to 5 Mbps in these conditions. Since SPDY is using a single TCP connection (compared to six for HTTP), SPDY will see a performance hit due to its inability to use the entire 10 Mbps or 20 Mbps pipe.

The raw data for all test cases (including the WinXP and Ubuntu test cases) are provided in [[SPDY-ANALYSIS](#)]. The remainder of this report will focus on the Win7 test results.

As described in [Section 5.3.1](#), the 2x2 matrix of protocol options will be presented as three different "cases", where each case represents the gain achieved by using one of the three proposed protocol enhancement options compared to the base case. This section is broken into three subsections that correspond to the three cases.

For each case, the results comprise a five-dimensional matrix of test conditions. In order to present the results in a compact way, each subsection below will first examine the impact that website configuration has on the achieved gain, then second will examine the impact of the channel conditions.

The results provide a comparison between the median web page download time achieved using the proposed protocol enhancement option and that which is achieved in the base case.



#### 5.4.1. CASE 1: SPDY vs. HTTPS

Case 1 compares the median web page download time achieved using SPDY/2 to the median time achieved using HTTPS. The comparison is made for each website configuration and channel condition combination. The comparison is reported as a "gain", which is calculated as the ratio of median page load time using HTTPS to the median page load time using SPDY. As a result, gain values greater than 1 indicate that SPDY/2 provides an advantage over HTTPS. As an additional point of reference, a gain of 2 corresponds to a 50% reduction in page load time, the goal to which SPDY aspires.

As stated above, the impact that the website configuration has on the achieved gain will be examined first. Table 5 shows the 3x3 matrix of website configurations. For each website, the maximum, minimum and mean gain values (across all channel conditions) are provided, in the manner shown below.

```
+-----+
|           max |
|           mean |
|           min  |
+-----+
```

Additionally, row-wise, column-wise, and overall statistics are calculated and shown around the periphery of the 3x3 matrix.

Table 5. Website Impact on SPDY Gain

	1 server	2 servers	4 servers	average
Page A	4.1 2.4 1.6	1.8 1.3 1.1	1.8 1.3 1.1	4.1 1.7 1.1
Page B	3.3 1.7 0.8	1.4 1.0 0.6	1.4 1.0 0.7	3.3 1.3 0.6
Page C	1.0 0.5 0.3	1.0 0.7 0.4	1.1 0.9 0.5	1.1 0.7 0.3
average	4.1 1.6 0.3	1.8 1.0 0.4	1.8 1.1 0.5	4.1 1.2 0.3





These results show that SPDY worked well for Page A (many smaller images), showing gains across all test conditions and an average gain of 1.7. For Page B (many images more typical size), the results were a bit more hit-or-miss, with SPDY not always resulting in improved performance. Nonetheless, on average a gain of 1.3 was achieved. Page C (small number of large images), on the other hand, resulted in almost universally worse performance with SPDY than with HTTPS. In the worst case, SPDY resulted in a 3.3x increase in page load time (gain of 0.3). On average SPDY took 1.4x longer (gain of 0.7) to download Page C than traditional HTTPS.

In general, the SPDY impact (positive or negative) diminished as more servers were utilized. This is the result of the increased parallelism and decreased number of objects requested per SPDY session, which if taken to the extreme would result in a pattern of requests that is similar to the HTTPS case.

The results shown in Table 6 examine the impact that channel conditions have on SPDY performance (relative to HTTPS). Since the channel conditions were formed from a 2x2x2 matrix of data-rate, RTT, and packet-loss rate (PLR), the results are depicted as two 2x2 matrices of data-rate and RTT, one corresponding to the 0% PLR test cases, and the other corresponding to the 1% PLR test cases. For each channel condition, the results for all nine websites are summarized via the maximum, mean, and minimum gain achieved. Similar to the row and column statistics provided in the website impact analysis in Table 5, summary statistics for all three dimensions are provided around the periphery of the 2x2x2 cube (i.e., in the right-most column, the last row, and the third table).

Table 6a. Channel Impact on SPDY Gain (0% PLR)

0% PLR	20ms	100ms	average
10/1	2.1 1.4 0.7	4.1 1.7 0.8	4.1 1.6 0.7
20/2	1.6 1.2 0.5	3.4 1.6 0.7	3.4 1.4 0.5
average	2.1 1.3 0.5	4.1 1.6 0.7	4.1 1.5 0.5



Table 6b. Channel Impact on SPDY Gain (1% PLR)

1% PLR	20ms	100ms	average
10/1	2.3 1.1 0.4	2.2 0.9 0.3	2.3 1.0 0.3
20/2	2.0 1.0 0.3	1.7 0.8 0.3	2.0 0.9 0.3
average	2.3 1.0 0.3	2.2 0.8 0.3	2.3 0.9 0.3

Table 6c. Channel Impact on SPDY Gain (all PLR)

all PLR	20ms	100ms	average
10/1	2.3 1.3 0.4	4.1 1.3 0.3	4.1 1.3 0.3
20/2	2.0 1.1 0.3	3.4 1.2 0.3	3.4 1.1 0.3
average	2.3 1.2 0.3	4.1 1.2 0.3	4.1 1.2 0.3

The most interesting observation here is the comparison between the two PLR tests. SPDY provided significant gains when the PLR was 0% (average gain 1.5), but showed worse average performance than HTTPS when packet loss was 1% (average gain 0.9). This effect was more pronounced in the large RTT cases as compared to the small RTT. This result points to a weakness in the way that SPDY utilizes TCP. With a typical HTTPS download of a web page, the browser will open a large number of simultaneous TCP connections. In this case, a random packet loss will cause the one affected connection to temporarily reduce its congestion window (and hence the effective data rate), but the other connections will be unaffected, and in fact may be able to opportunistically make use of the bandwidth made available by the



affected connection. The result for HTTPS is that random packet loss has only a minor impact on page download time. In the case of SPDY, the number of parallel TCP connections is dramatically reduced (by as much as a factor of six), so that random packet loss has a much bigger impact on the overall throughput.

In the absence of packet loss, SPDY provides better gains as the RTT increases. This is the result of reducing the number of round trips needed to fetch all of the page resources. Nonetheless, there were test cases with 0 packet loss where SPDY performed worse than HTTPS. This only occurred with Page C; SPDY always provided a benefit for Pages A and B when there was no packet loss. The root cause of the degradation is unknown, but it may be the result of the buffering configuration of the network simulator Dummynet. As stated previously, testing was performed with the default configuration of 50 packet buffers. This configuration is not reflective of real networks (which in some cases may have more than double that amount of buffering), and could bias the results somewhat against applications (such as SPDY) that use a small number of TCP sessions.

Additionally, SPDY generally shows more gain in the lower Data Rate cases.

The overall average gain of 1.2 seen in our experiments aligns well with the performance gains that Google is seeing in their live deployments. In results presented in a December 8, 2011, Google TechTalk [[TECHTALK](#)], they report a 15.4% improvement in page load time (equivalent to a gain of 1.18).

#### **5.4.2. CASE 2: initcwnd=10 vs. initcwnd=3**

The initcwnd increase provides very modest gain across the majority of test cases. In a few cases there was a marginal degradation of performance compared to the default initcwnd case. Across all test conditions an average gain of 1.1 (or 9% reduction in page load time) was seen, as indicated in Table 7 and Table 8.



Table 7. Website Impact on Initcwnd Gain

	1 server	2 servers	4 servers	average
Page A	1.1 1.0 0.9	1.1 1.0 1.0	1.1 1.0 0.9	1.1 1.0 0.9
Page B	1.1 1.0 1.0	1.3 1.1 1.0	1.2 1.0 0.9	1.3 1.0 0.9
Page C	1.2 1.1 1.0	1.3 1.1 1.0	1.5 1.2 1.0	1.5 1.1 1.0
average	1.2 1.1 0.9	1.3 1.1 1.0	1.5 1.1 0.9	1.5 1.1 0.9

Overall the Page C test cases showed a slightly higher gain over the other two pages. Since Page C involves large files, it isn't surprising that the effect is slightly more pronounced than in Page A. In Page A, the majority of resources can be delivered in 3 packets or less anyway, so the increase in initcwnd doesn't reduce the number of round-trips.

Table 8a. Channel Impact on Initcwnd Gain (0% PLR)

0% PLR	20ms	100ms	average
10/1	1.1 1.0 1.0	1.2 1.1 1.0	1.2 1.0 1.0
20/2	1.1 1.0 1.0	1.1 1.0 0.9	1.1 1.0 0.9
average	1.1 1.0 1.0	1.2 1.0 0.9	1.2 1.0 0.9





Table 8b. Channel Impact on Initcwnd Gain (1% PLR)

1% PLR	20ms	100ms	average
10/1	1.5 1.1 0.9	1.2 1.1 1.0	1.5 1.1 0.9
20/2	1.4 1.1 0.9	1.3 1.1 0.9	1.4 1.1 0.9
average	1.5 1.1 0.9	1.3 1.1 0.9	1.5 1.1 0.9

Table 8c. Channel Impact on Initcwnd Gain (all PLR)

all PLR	20ms	100ms	average
10/1	1.5 1.1 0.9	1.2 1.1 1.0	1.5 1.1 0.9
20/2	1.4 1.1 0.9	1.3 1.1 0.9	1.4 1.1 0.9
average	1.5 1.1 0.9	1.3 1.1 0.9	1.5 1.1 0.9

When viewing the results broken down by channel condition, we don't see a significant difference in gain from one channel condition to the next. Notably, even in the high RTT cases we don't see much change in the gain. This is likely due to the fact that the majority of resources had sizes that either fell below the 3 packet default initcwnd, or were much larger (e.g., over 210 packets each for the images in Page C) so that the initcwnd change might save a single RTT for a transfer that took 9 RTTs using the default initcwnd.



#### 5.4.3. CASE 3: SPDY+initcwnd=10 vs. HTTPS+initcwnd=3

The combination of SPDY and the initcwnd increase resulted in a benefit that was more than the product of the two individual gains. This is a result of the fact that, due to SPDY, all of the TCP sessions involved multiple round-trips, so the acceleration of initial data rate provided by the initcwnd increase resulted in tangible benefits. In fact, one case (Page A, single server) saw an astounding gain of 4.7. By examining the raw data in [\[SPDY-ANALYSIS\]](#), it can be seen that this combination reduced the median page load time of 8.7 seconds down to 1.9 seconds. Unfortunately, this result is not typical, and Page C again experienced worse performance when the new protocol options are used. The overall average gain seen was 1.4 across all test cases as shown in Table 9 and Table 10.

Table 9. Website Impact on SPDY+Initcwnd Gain

	1 server	2 servers	4 servers	average
Page A	4.7 3.2 2.2	2.3 1.5 1.2	1.5 1.4 1.2	4.7 2.0 1.2
Page B	3.3 1.9 0.9	1.6 1.1 0.6	1.5 1.1 0.7	3.3 1.4 0.6
Page C	1.0 0.6 0.3	1.0 0.7 0.4	1.1 0.9 0.6	1.1 0.7 0.3
average	4.7 1.9 0.3	2.3 1.1 0.4	1.5 1.1 0.6	4.7 1.4 0.3



Table 10a. Channel Impact on SPDY+Initcwnd Gain (0% PLR)

0% PLR	20ms	100ms	average
10/1	3.2	4.7	4.7
	1.7	1.8	1.8
	1.0	0.9	0.9
20/2	3.1	3.9	3.9
	1.6	1.7	1.6
	0.7	0.7	0.7
average	3.2	4.7	4.7
	1.7	1.7	1.7
	0.7	0.7	0.7

Table 10b. Channel Impact on SPDY+Initcwnd Gain (1% PLR)

1% PLR	20ms	100ms	average
10/1	3.1	2.4	3.1
	1.2	0.9	1.1
	0.4	0.3	0.3
20/2	2.8	2.2	2.8
	1.2	1.0	1.1
	0.3	0.3	0.3
average	3.1	2.4	3.1
	1.2	1.0	1.1
	0.3	0.3	0.3



Table 10c. Channel Impact on SPDY+Initcwnd Gain (all PLR)

all PLR	20ms	100ms	average
10/1	3.2	4.7	4.7
	1.5	1.4	1.4
	0.4	0.3	0.3
20/2	3.1	3.9	3.9
	1.4	1.3	1.3
	0.3	0.3	0.3
average	3.2	4.7	4.7
	1.4	1.3	1.4
	0.3	0.3	0.3

In the packet loss test cases, there was a fairly significant performance loss with SPDY, as reported in earlier test cases. This is particularly true in the high RTT cases.

#### 5.4.4. Results Summary

The results are summarized further in Table 11. A caveat on the average results presented here is worth noting: the 72 test cases were selected to test the protocol changes in a fairly wide range of conditions in order to understand the impact that individual factors have on the performance gains. As a result, the 72 test cases likely do not comprise a statistically accurate sampling of real-world sites, and so the average results presented here may not accurately reflect the average performance gain that would be seen in the real world. However, as noted previously, the average results do appear to align well with the average gains seen by Google via their live deployments with users utilizing the Chrome browser. In addition, it is interesting to examine the average gain achieved in a particular subset of the test conditions. For that we select the subset consisting of pages A and B operating on a wireline network (PLR=0%). For that subset, we see that the combination of SPDY/2 and the initcwnd increase achieves an average gain of 2.1 (52% reduction in page load time).





Table 11. Summary of Results

	Case 1	Case 2	Case 3
	SPDY	initcwnd	SPDY+initcwnd
Best Gain	4.1	1.5	4.7
(all test conditions)	(A,1,10,100,0)	(C,4,10,20,1)	(A,1,10,100,0)
Average Gain			
(all test conditions)	1.2	1.1	1.4
Worst Gain	0.3	0.9	0.3
(all test conditions)	(C,1,x,x,1)	(A,1,20,100,0)	(C,1,20,100,1)
# test cases achieving gain >= 1	40 of 72 (56%)	58 of 72 (81%)	44 of 72 (61%)
# test cases achieving gain >= 2	8 of 72 (11%)	0 of 72 (0%)	14 of 72 (19%)
Average Gain (Pages A and B with PLR=0%)	1.8	1.0	2.1

## 6. Recommendations and Conclusions

The results presented here indicate that SPDY/2 in conjunction with the TCP Initial Congestion Window increase has the potential to improve or impair web page download performance depending on a number of factors.

Deployment of SPDY/2 for general-purpose web servers should be considered in light of the following concerns:

- o While the Chrome and Firefox browsers are important (approximately 29% and 22% market share respectively as of June 2012), client support isn't ubiquitous.
- o Some web page downloads were significantly impaired by SPDY.
- o Work is underway on a revision to the protocol, and with standardization in the IETF a possibility in the near future, waiting may be worthwhile to see what develops in this space.

On the other hand, for a controlled environment where a single entity



provides the servers, the client software, and the web content, SPDY might be a valuable tool. An example application might be a remote user interface (RUI) that is delivered via HTTP to a CPE device.

In addition to potential uses for SPDY by network operators, another area of interest is the impact that SPDY would have on network traffic if it were to be widely adopted. In general the story here is good, by reducing the number of simultaneous TCP sessions and extending the duration of many of the sessions, other applications could see improved performance as a result of TCP congestion avoidance being invoked far more often. Secondly, the expectation would be to see a slight reduction in data usage, due to the greater efficiency that SPDY provides (fewer TCP control packets), as well as a slight increase in average packet size, due to the multiplexing of HTTP responses. Both of these factors will serve to reduce the packets per second rate of the network, which improves scalability of CMTSSs, routers, DPI boxes, etc. Finally, equipment implementing Carrier Grade NAT (which will be a key element of the IPv6 transition in some networks) could see improved scalability as the number of simultaneous TCP connections is reduced.

In regards to the increase in the TCP Initial Congestion Window, while we see only marginal gains resulting from this enhancement, the change can be made simply by changing a server operating system parameter. It requires no client modifications. As a result, we see no reason not to set the server `initcwnd` at the proposed value of 10.

## **7. IANA Considerations**

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

## **8. Security Considerations**

See [[I-D.mbelshe-httpbis-spdy](#)] and [[I-D.ietf-tcpm-initcwnd](#)] for Security Considerations.

## **9. Acknowledgements**

The authors would like to thank Marc Weaver, Darshak Thakore, Eric Winkelman, Robin Sam Ku, Scott Maize, and Rob Moon for their assistance in developing and configuring the test bed, and collecting the data presented in this report. The authors would also like to



gratefully express their appreciation to Chris Donley, Joan Strosin, Ken Barringer and Christie Poland for manuscript preparation assistance, and to Mike Belshe and Roberto Peon for their review and very insightful comments.

## 10. References

### [BUFFERCONTROL]

CableLabs, "Cable Modem Buffer Control", <<http://www.cablelabs.com/specifications/CM-GL-Buffer-V01-110915.pdf>>.

### [CDNPLANET]

"www.cdnplanet.com", <<http://www.cdnplanet.com>>.

### [DUMMYNET]

"Dummysnet", <<http://info.iet.unipi.it/~luigi/dummysnet/>>.

### [Faster]

Google, "Let's Make the Web Faster", <<http://code.google.com/speed>>.

### [GOOGLE-BENCHMARKS]

Lloyd, M., "SPDY and Server Push: Analysis and Experiments", June 2010, <[https://docs.google.com/View?id=d446246\\_0cc6c6dkr](https://docs.google.com/View?id=d446246_0cc6c6dkr)>.

### [HTTP\_Pipelining]

Wikipedia, "HTTP Pipelining", <[http://en.wikipedia.org/wiki/HTTP\\_pipelining](http://en.wikipedia.org/wiki/HTTP_pipelining)>.

### [I-D.agl-tls-nextprotoneg]

Langley, A., "Transport Layer Security (TLS) Next Protocol Negotiation Extension", [draft-agl-tls-nextprotoneg-03](#) (work in progress), April 2012.

### [I-D.ietf-tcpm-initcwnd]

Chu, J., Dukkupati, N., Cheng, Y., and M. Mathis, "Increasing TCP's Initial Window", [draft-ietf-tcpm-initcwnd-03](#) (work in progress), February 2012.

### [I-D.mbelshe-httpbis-spdy]

Belshe, M. and R. Peon, "SPDY Protocol", [draft-mbelshe-httpbis-spdy-00](#) (work in progress), February 2012.

### [INITCWND-PAPER]



Dukkipati, N., Refice, T., Cheng, Y., Chu, J., Sutin, N., Agarwal, A., Herbert, T., and A. Jain, "An Argument for Increasing TCP's Initial Congestion Window", <[http://code.google.com/speed/articles/tcp\\_initcwnd\\_paper.pdf](http://code.google.com/speed/articles/tcp_initcwnd_paper.pdf)>.

- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), September 1981.
- [RFC1122] Braden, R., "Requirements for Internet Hosts - Communication Layers", STD 3, [RFC 1122](#), October 1989.
- [RFC2001] Stevens, W., "TCP Slow Start, Congestion Avoidance, Fast Retransmit, and Fast Recovery Algorithms", [RFC 2001](#), January 1997.
- [RFC2414] Allman, M., Floyd, S., and C. Partridge, "Increasing TCP's Initial Window", [RFC 2414](#), September 1998.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", [RFC 2616](#), June 1999.
- [RFC3390] Allman, M., Floyd, S., and C. Partridge, "Increasing TCP's Initial Window", [RFC 3390](#), October 2002.
- [SPDY] "SPDY - The Chromium Projects", <<http://www.chromium.org/spdy>>.
- [SPDY-ANALYSIS] White, G., "Analysis of Google Spdy and TCP initcwnd", May 2012, <[http://www.cablelabs.com/downloads/pubs/Analysis\\_of\\_Google\\_SPDY\\_TCP.pdf](http://www.cablelabs.com/downloads/pubs/Analysis_of_Google_SPDY_TCP.pdf)>.
- [TECHTALK] "http://www.cnx-software.com/2012/01/27/spdy-aims-to-make-the-web-faster-and-replace-http/", <<http://www.cnx-software.com/2012/01/page/2/>>.





Authors' Addresses

Greg White  
CableLabs  
858 Coal Creek Circle  
Louisville, CO 80027-9750  
USA

Email: [g.white@cablelabs.com](mailto:g.white@cablelabs.com)

Jean-Francois Mule  
CableLabs  
180 Montgomery St  
Suite 2480  
San Francisco, CA 94104-4203  
USA

Email: [jf.mule@cablelabs.com](mailto:jf.mule@cablelabs.com)

Dan Rice  
CableLabs  
858 Coal Creek Circle  
Louisville, CO 80027-9750  
USA

Email: [d.rice@cablelabs.com](mailto:d.rice@cablelabs.com)

