Security Now! #512 - 06-16-15

Mozilla's Tracking Protection

This week on Security Now!

- The LastPass network breach
- More bad news from the Office of Personnel Management
- Did China & Russia obtain and decrypt Snowden's document cache?
- Bye bye browser FTP?
- When BGP goes very wrong
- Wikipedia goes all-HTTPS
- A big SSD reliability report
- The state of tracking in today's websites

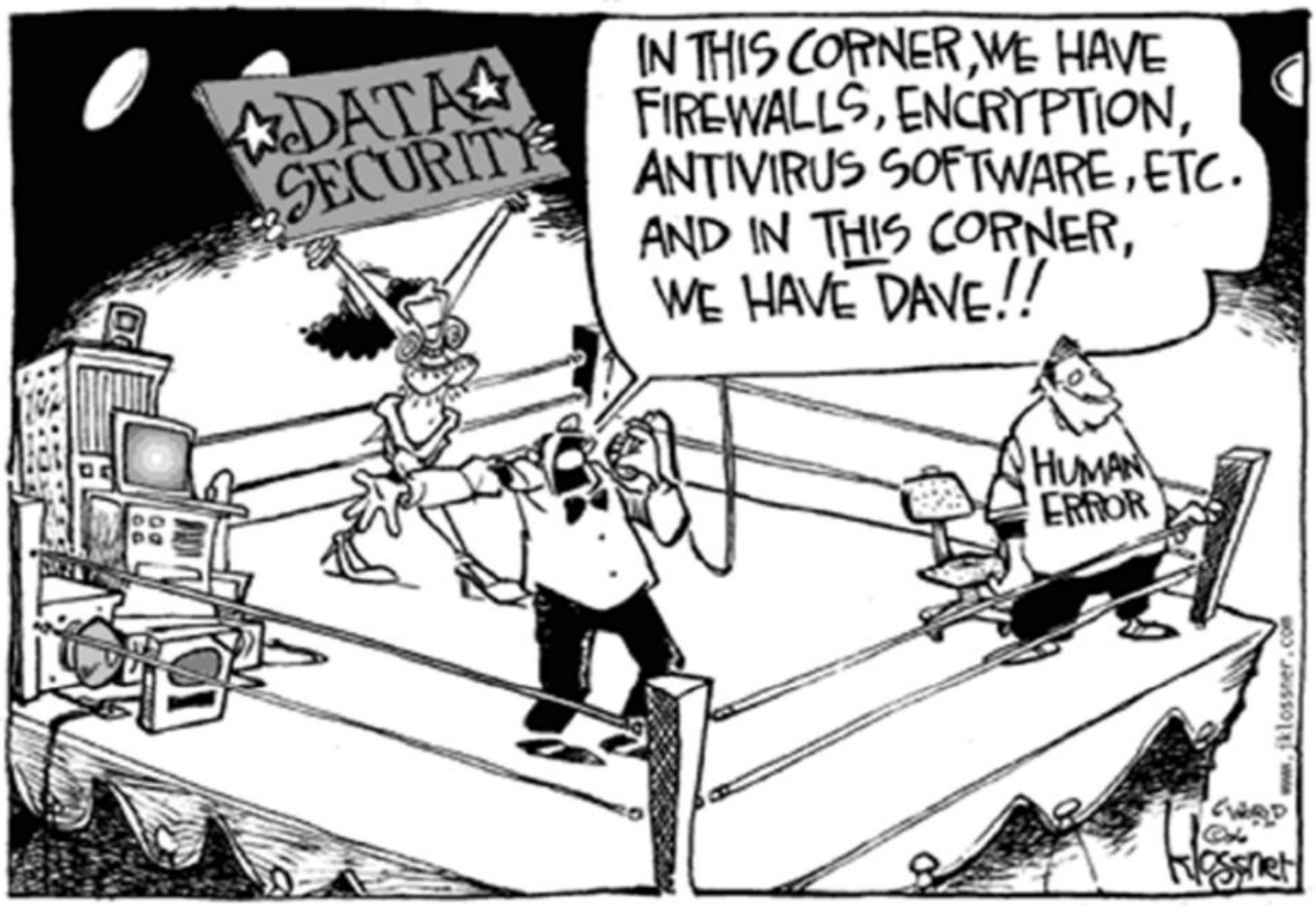

Sunday's TWiT: ***Merlin Mann is a fabulous guest!***

Security News

LastPass Network Breach:

- https://blog.lastpass.com/2015/06/lastpass-security-notice.html/
- Once again... correct and responsible disclosure.
- Late last week they detected some anomalous traffic on their network, between systems, when no one was working to account for it... and that set off alarm bells.
    - (Reading between the lines: An employee got hacked.)

- The investigation has shown, however, that LastPass account:
    - email addresses
    - password reminders
    - server per user salts
    - authentication hashes... were compromised.

- Consequences:
    - PBKDF2 (eMail + Passphrase) x 100,000 *after* what is done in the user's browser first.
    - Brute forcing would be theoretically possible.
    - Something like 8 thousand password guesses tests per second.
    - If your LastPass master password was weak: trouble.
    - And if you used *IT* anywhere else: possible trouble.
    - 2FA would not prevent this, since 2FA is orthogonal to and secondary to this process.

- User response:
    - Change your LastPass master password.
    - LP locked out any logons from new devices or IPs, requiring confirmation eMail.
    - (LastPass servers were overwhelmed by the demand.)
    - (They had to do 100,000 iterations of PBKDF2, too!)

- (This is some of the many things SQRL prevents.)
    - No centralized cloud-based database in the sky.
    - No need to synchronize independent copies of password databases.
    - (No more password databases.)
    - Last week Chat room: "Why is that better than browser remembers password?"

Update on the OPM (office of personnel management) troubles:

- http://bigstory.ap.org/article/d842d757851b4a59aca2aecf2f31995a/union-says-all-federal-workers-fell-victim-hackers
- Apparently there was a second intrusion involving a great many more pieces of even more sensitive data.
- <Liberally paraphrased from the Associated Press coverage>
  Hackers linked to China ((American officials have said the cybertheft originated in China and that they suspect espionage by the Chinese government, which has denied any involvement.)) gained access to the sensitive background information submitted by

intelligence and military personnel for security clearances, in a cyberbreach of federal records dramatically worse than was first acknowledged.

The forms, which authorities believe may have been stolen en masse, known as Standard Form 86, require applicants to fill out deeply personal information about mental illnesses, drug and alcohol use, past arrests and bankruptcies. They also require the listing of contacts and relatives, potentially exposing any foreign relatives of U.S. intelligence employees to coercion. Both the applicant's Social Security number [AND] that of his or her cohabitant is required. Beyond Social Security numbers, the data include military records and veterans' status information, addresses, birth dates, job and pay histories; health insurance, life insurance and pension information; and age, gender and race data.

In a statement, the White House said that on June 8, investigators concluded there was "a high degree of confidence that ... systems containing information related to the background investigations of current, former and prospective federal government employees, and those for whom a federal background investigation was conducted, may have been exfiltrated."

Joel Brenner, a former top U.S. counterintelligence official said: "This tells the Chinese the identity of almost everyone who has a United States security clearance. That makes it very hard for any of those people to function as an intelligence officer. The database also tells the Chinese an enormous amount of information about almost everyone with a security clearance. That's a gold mine. It helps you approach and recruit spies."

The White House statement said the hack into the security clearance database was separate from the breach of federal personnel data previously announced — a breach that is itself appearing far worse than first believed. Nearly all of the millions of security clearance holders, including some CIA, National Security Agency and military special operations personnel, are potentially exposed in the security clearance breach, the officials said. More than 4 million people had been investigated for a security clearance as of October 2014, according to government records.

But in this newly revealed hack of standard personnel records announced last week, two people briefed on the investigation disclosed Friday that as many as 14 million current and former civilian U.S. government employees have had their information exposed to hackers... records going back to the 1980s. Since there are about 2.6 million executive branch civilians, the majority of the records exposed relate to former employees. Contractor information also has been stolen, officials said.

The personnel records would provide a foreign government an extraordinary roadmap to blackmail, impersonate or otherwise exploit federal employees in an effort to gain access to U.S. secrets —or entry into government computer networks.

Outside experts were pointing to the breaches as a blistering indictment of the U.S. government's ability to secure its own data two years after a National Security Agency contractor, Edward Snowden, was able to steal tens of thousands of the agency's most sensitive documents.

After the Snowden revelations about government surveillance, it became more difficult for the federal government to hire talented younger people into sensitive jobs, particularly at intelligence agencies, said Evan Lesser, managing director of ClearanceJobs.com, a website that matches security-clearance holders to available slots.

"Now, if you get a job with the government, your own personal information may not be secure," he said. "This is going to multiply the government's hiring problems many times."

- https://www.techdirt.com/articles/20150612/16334231330/second-opm-hack-revealed-even-worse-than-first.shtml
- Mike Masnick: And yet... this is the same federal government telling us that it wants more access to everyone else's data to "protect" us from "cybersecurity threats" -- and that encryption is bad? Yikes.

Did China & Russia obtain and decrypt Snowden's document cache?

- Seems much less than likely
- Glenn Greenwald's take:
  - https://firstlook.org/theintercept/2015/06/14/sunday-times-report-snowden-files-journalism-worst-also-filled-falsehoods/
- Craig Murray:
  - https://www.craigmurray.org.uk/archives/2015/06/five-reasons-the-mi6-story-is-a-lie/
- Factual troubles:
  - <quote> It is not clear whether Russia and China stole Snowden's data, or whether he voluntarily handed over his secret documents in order to remain at liberty in Hong Kong and Moscow.

    David Miranda, the boyfriend of the Guardian journalist Glenn Greenwald, **was seized at Heathrow in 2013 in possession of 58,000 "highly classified" intelligence documents after visiting Snowden in Moscow.**

  - What's the problem with that *Sunday Times* passage? It's an utter lie. David did not visit Snowden in Moscow before being detained. As of the time he was detained in Heathrow, David had never been to Moscow and had never met Snowden. The only city David visited on that trip before being detained was Berlin, where he stayed in the apartment of Laura Poitras.

- Everything we know about Edward Snowden suggests that he knows how to protect his sensitive information.

Firefox and Chrome considering dropping native FTP support

- https://bugzilla.mozilla.org/show_bug.cgi?id=1174462
  - Bug 1174462 - Remove built-in support for FTP
- https://code.google.com/p/chromium/issues/detail?id=333943
- <quote> We should consider removing built-in support for FTP from Chrome and move it out to an app.

  Over a 7-day period, only .1-.2% of users end up navigating to any FTP URL (with slightly

higher numbers amongst Linux desktop users). This has been fairly stable over the last year, so it doesn't look there are trends for FTP to disappear altogether.

With the combination of the sockets API and the downloads API it may be possible to construct a Chrome App which handles this well. Also would need a way to be able to register an app/extension to handle a particular URL scheme so that navigations would be seamless for users of FTP apps.

This isn't urgent priority, but might be a nice code cleanup for a little-used feature.

Massive route leak caused Internet slowdown

- http://www.bgpmon.net/massive-route-leak-cause-internet-slowdown/
- Earlier today a massive route leak initiated by Telekom Malaysia (AS4788) caused significant network problems for the entire global Internet routing system. Primarily affected was Level3 (AS3549 – formerly known as Global Crossing) and their customers.

  Starting June 12th at 08:43 UTC, AS4788 Telekom Malaysia started to announce about 179,000 routing prefixes to Level3 (AS3549, the Global crossing AS), who in turn accepted these and propagated them to their peers and customers. Since Telekom Malaysia had thus inserted itself between these thousands of prefixes and Level3 it was now responsible for delivering these packets to the intended destinations.

  This event resulted in significant packet loss and Internet slow down in all parts of the world. The Level3 network in particular suffered from severe service degradation between the Asia pacific region and the rest of their network. The following graph shows the packet loss -- often hitting 100% -- as measured by OpenDNS between London over Level3 and Hong Kong. The same loss patterns were visible from other Level3 locations globally to for example Singapore, Hong Kong and Sydney.

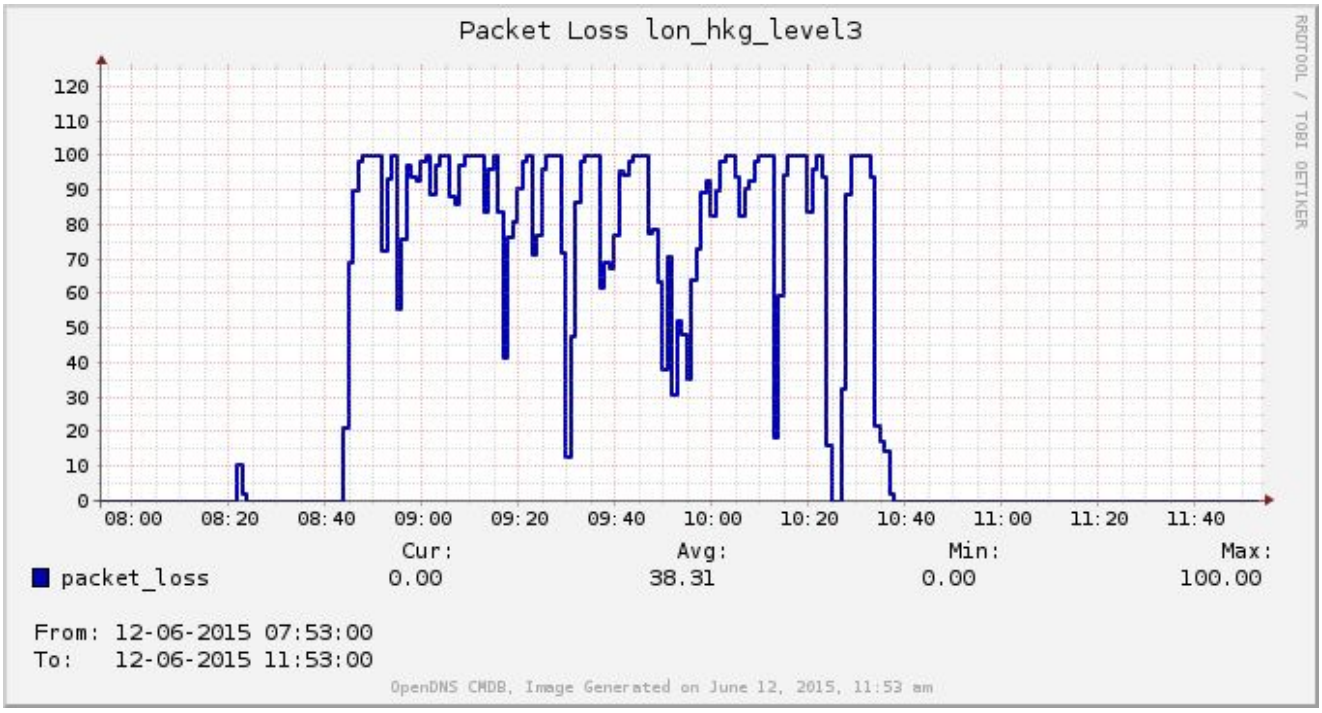

Wikipedia goes all-HTTPS, starting immediately!

- https://blog.wikimedia.org/2015/06/12/securing-wikimedia-sites-with-https/
- The Foundation blog posted:
  To be truly free, access to knowledge must be secure and uncensored. At the Wikimedia Foundation, we believe that you should be able to use Wikipedia and the Wikimedia sites without sacrificing privacy or safety.

  Today, we're happy to announce that we are in the process of implementing HTTPS to encrypt all Wikimedia traffic. We will also use HTTP Strict Transport Security (HSTS) to protect against efforts to 'break' HTTPS and intercept traffic. With this change, the nearly half a billion people who rely on Wikipedia and its sister projects every month will be able to share in the world's knowledge more securely.

  HTTPS is not new to Wikimedia sites. Since 2011, we have been working on establishing the infrastructure and technical requirements, and understanding the policy and community implications of HTTPS for all Wikimedia traffic, with the ultimate goal of making it available to all users. In fact, for the past four years, Wikimedia users could access our sites with HTTPS manually, through HTTPS Everywhere, and when directed to our sites from major search engines. Additionally, all logged in users have been accessing via HTTPS since 2013.

  Over the last few years, increasing concerns about government surveillance prompted members of the Wikimedia community to push for more broad protection through HTTPS. We agreed, and made this transition a priority for our policy and engineering teams.

- SSL Labs:
  - TLS only.
  - PFS.
  - Very nice cipher-suite ordering.
  - OCSP stapling.
  - HSTS (but only 259200 - short)

Misscellany:

- Steve Needs Help! -- "Tip Subjects" for TNSS snippets!
- SyFy: "Dark Matter" sadly... looks awful. Awful acting, of awful writing.
- AMC: "Humans"
- Twitter to allow 10K character DM's.
  - https://twittercommunity.com/t/removing-the-140-character-limit-from-direct-messages/41348

SpinRite -&- SSD's

Carnegie Mellon & Facebook -- latest research on SSDs
http://users.ece.cmu.edu/~omutlu/pub/flash-memory-failures-in-the-field-at-facebook_sigmetrics15.pdf

ABSTRACT
Servers use flash memory based solid state drives (SSDs) as a high-performance alternative to hard disk drives to store persistent data. Unfortunately, recent increases in flash density have also brought about decreases in chip-level reliability. In a data center environment, flash-based SSD failures can lead to downtime and, in the worst case, data loss. As a result, it is important to understand flash memory reliability characteristics over flash lifetime in a realistic production data center environment running modern applications and system software.

This paper presents the first large-scale study of flash-based SSD reliability in the field. We analyze data collected across a majority of flash-based solid state drives at Facebook data centers over nearly four years and many millions of operational hours in order to understand failure properties and trends of flash-based SSDs. Our study considers a variety of SSD characteristics, including: the amount of data written to and read from flash chips; how data is mapped within the SSD address space; the amount of data copied, erased, and discarded by the flash controller; and flash board temperature and bus power.

http://www.zdnet.com/article/facebooks-ssd-experience/?tag=nl.e539&s_cid=e539&ttag=e539&ftag=TRE17cfd61

What they found:
The good news: some issues that worry people, aren't issues.
The bad news: there's other stuff to worry about.

Temperature
SSDs are sensitive to temperature - more so than hard drives. When they get hot, the SSD may throttle back performance. Experiencing unexplained slowdowns on some servers?: check temperature.

First generation SSDs failed more often as temp rose, possibly due to a lack of throttling.
Some second generation SSDs throttle aggressively enough to reduce failure rates, while others kept the failure curve flat.

Bus power
SSDs are thirsty. PCIe v2 SSDs ran anywhere from 8 to 14.5 watts, a high and surprisingly wide range. The team found that as power consumption rose, so did failure rates.

Write fatigue
The team found that the level of system write activity correlated with SSD failure, probably because flash writes require a lot of power. Disks could be a better choice for heavy write applications such as logging.

SSD failures
SSD failures - i.e. UREs - are relatively common: 4.2 to 34.1 percent of the SSDs reported uncorrectable errors. In fact, 99.8 percent of the SSDs reporting an error in one week reported another error in the next week.

Life cycle and failures
The SSD failure profile differs from disks, where disks exhibit infant mortality, then they enjoy a few years of good reliability, before age catches up with them. SSDs have an early period of UREs as faulty cells are identified, increasing reliability, until cell wear-out leads to increasing read failures.

The data layout surprise
Disk drives aren't much affected by data layout - unless it involves lots of random seeks.
But SSDs are very different.
Sparse logical data layouts - non-contiguous data - lead to higher SSD failure rates as do very dense data structures.
Such behavior is potentially due to the fact that sparse data allocation can correspond to access patterns that write small amounts of non-contiguous data, causing the SSD controller to more frequently erase and copy data compared to writing contiguous data.

(What we learn about today's Internet thanks to) Mozilla's Tracking Protection

http://monica-at-mozilla.blogspot.ca/2015/05/tracking-protection-for-firefox-at-web.html
http://monica-at-mozilla.blogspot.com/2015/05/advertising-sustainable-utopia.html
http://ieee-security.org/TC/SPW2015/W2SP/papers/W2SP_2015_submission_32.pdf
http://venturebeat.com/2015/05/24/firefoxs-optional-tracking-protection-reduces-load-time-for-top-news-sites-by-44/

Disconnect -- https://disconnect.me/
Headline: Online protection, simplified.
Subhead: Join over 10 million people who actively use our software to protect themselves from hackers and trackers.

Founded in 2011 by Google people.

- "Disconnect" is the new default search provider on the Tor browser.

Monica Chew:
Posted May 21st, 2015:
Tracking Protection for Firefox at Web 2.0 Security and Privacy 2015 Conference
<quote> This paper is the last artifact of my work at Mozilla, since I left employment there at the beginning of April. I believe that Mozilla can make progress in privacy, but leadership needs to recognize that current advertising practices that enable "free" content are in direct conflict with security, privacy, stability, and performance concerns -- and that Firefox is first and foremost a user-agent, not an industry-agent.

Advertising does not make content free. It merely externalizes the costs in a way that incentivizes malicious or incompetent players to build things like Superfish, infect 1 in 20 machines with ad injection malware, and create sites that require unsafe plugins, take twice as many resources to load... all quite expensive in terms of bandwidth, power, and stability.

It will take a major force to disrupt this ecosystem and motivate alternative revenue models. I hope that Mozilla can be that force.

Posted May 28th, 2015: (heavily snipped, edited, and paraphrased)

[Online Internet] Advertising generates $50 billion annually in the US alone, but how much of that figure reflects real value? Approximately 1/3rd of click traffic is fraudulent, leading to $10 billion in [fraud and] wasted spending annually.

Even when ads are displayed to real people, they often create little to no value for the ad creator.

According to Google, half of ads are never viewable, not even for a second.

[Presumably they are "below the fold" and require vertical scrolling that never occurs]

In addition, adblocking usage grew by 70% last year, and 41% of people between 18-29 use an adblocker.

The advertising industry responds to these trends by making ads increasingly distracting (requiring large amounts of resources and unsafe plugins to run), collecting increasingly large amounts of data, and creating more opportunities for abuse by government agencies and other malicious actors. As [Mozilla Foundation and Corporation Chairwoman] Mitchell Baker put it: Do we want to live in a house or a fish bowl?

There has to be a better way. Why can’t a person buy and blank out all of the ad space on sites they visit at a deep discount, since targeting machinery would no longer be relevant? Why aren’t subscriptions available as bundle deals, like in streaming video? Solutions like these are hypothetical and will remain so as long we maintain the fiction that the current advertising revenue model is a sustainable utopia.

Steve's Experience & Position:

- Malware IS a problem.
- Scripting IS a problem.
- Tracking/profiling... doesn't really worry ME much, but I FULLY respect the rights and sentiments of those who ARE concerned over the practice.
- Intrusion into my browsing IS a problem:
    - I strongly dislike cover-up ads that require me to take action to get them out of the way are annoying.
    - The other day: "Change your device orientation to proceed."
        - (Chrome's scripting was disabled.)
- I *do* pay a few dollars to turn ads off in apps I want to use.

- Paywalls don't work for me.
    - I'm not at the Financial Times or New York Times enough to "subscribe".
    - Limited pages per month seems an equitable solution.

The big problem: Perverse commercial incentives.

- Examples of commercial incentives that don't work:
    - Military-industrial complex.
    - HealthCare-industrial complex.
    - Educational-industrial complex.
    - Prison-industrial complex.

Zero-cost to the publisher to "host" ads on their site.

- I'm AFRAID to put ads on GRC... because the slope is SO slippery.

How bad has it become??

Georgios Kontaxis @ Columbia University
Monica Chew @ Mozilla Corporation

- Online advertising has a symbiotic relationship with the Internet ecosystem. Advertisers pay content publishers, i.e., websites, to embed promotional material in the content they generate. Publishers in turn use that revenue to mitigate the need for users to directly purchase the content they consume. In 2013, revenue from US online advertising reached $43 billion, supporting the vast majority of publishers [1]. A prime reason for its accelerated growth is the rise of efficient markets for targeted promotions.

  Unfortunately, the technology to personalize advertisements and other site content relies extensively on tracking online activities of users. On top of that, governments misuse these technologies to facilitate warrantless surveillance without the knowledge or consent of the users or the original tracking services themselves.

  The rise of online advertising has led to a corresponding increase in ad-blocking and privacy-oriented software. In 2013 ad-blocking grew nearly 70%, with nearly 41% of 18-29 year olds reporting usage [2]. Even though typical users may not be aware of the broad reach of tracking technologies, they feel that the advertising industry is insufficiently regulated and lack confidence protecting their data online [3], [4].

  In response to this clarion call from users, we designed and implemented Tracking Protection in Mozilla Firefox 35. We demonstrate a 67.5% reduction in the number of HTTP cookies set during a crawl of the Alexa top 200 news sites. The reduction in cookies corresponds to blocking 11 tracking elements on 50% of these sites. Since Firefox does not download and render content from tracking domains, Tracking Protection also enjoys performance benefits of a 44% median reduction in page load time and 39% reduction in data usage in the Alexa top 200 news sites.

Note: this is NOT ad blocking. I still see ad-laced pages.

What is it?

- In order to cover the entire array of tracking (more than just 3rd-party cookies), Firefox blocks ALL traffic between the browser and previously identified tracking domains.
- An URL-based API based on Google's Safe Browsing was implemented.
- Using a subset of "Disconnect's" privacy-oriented blocklist.
- ~1500 domains are maintained & updated every 45 minutes.

A "Shield" icon is shown to the left of the browser's URL.

- Sites can be individually white-listed if desired.

My own test at USA Today:

- With blocking:139 requests for 3,662.34 KB
- Without blocking: 162 requests for 4,530.11 KB
- 23 tracking domains and nearly another megabyte of data.

Cookies Set by the Alexa Top 200 News sites:

- Default 4006 0.0%
- Tracking Protection 1300 - reduction: 67.5%
- Adblock Plus 2398 - reduction: 40.1%

What's being blocked?

- SCRIPT 1,890 79.2%
- IMG 334 14.0%
- IFRAME 149 6.3%
- OBJECT 13 0.5%

A real-world example using www.weather.com

- Page loads in 3.5 seconds with Tracking Protection versus 6.3 seconds without.
- Results in data usage of 2.8 MB (98 HTTP requests) versus 4.3 MB (219 HTTP requests), respectively.
- Though Tracking Protection prevents initial requests for only 4 HTML <script> elements, without Tracking Protection, an additional 45 domains are contacted.
- Of the additional resources downloaded without Tracking Protection enabled, 57% are JavaScript (as identified by the content-type HTTP header) and, 27% are images. The largest elements appear to be JavaScript libraries with advertisement-related names, each on the order of 10 or 100 KB.