Web Browsers

Web browsers are software on your machine that communicate with servers or hosts on the Internet. Using a web browser causes data to be stored on your computer and logs to be stored on the web servers you visit, and frequently transmits unencrypted information.

Until you have understood the mechanisms by which this occurs — and taken steps to prevent them — it is best to assume that anything you do with a web browser could be recorded by your own machine, by the web servers you're communicating with, or by any adversary that is able to monitor your network connection.

Controlling and Limiting the Logs Kept by Your Browser

Web browsers often retain a large amount of information about the way they are used. A browser typically keeps a history of the web pages it visits. Browsers also often retain cached copies of the pages you've visited, information about which accounts you log into on web servers, names and other data you enter into web forms, and cookies that record preferences and link your browser to records on third party web servers. Fortunately, browsers also include features for managing these records. In general, the features are getting better, so it's getting easier to control browser records.

For example, here are the stored data privacy settings pages for Firefox, the free web browser:

For each type of information your browser stores, you can either set it to not collect it at all, set it to delete within a certain span of days, set it to delete when you quit the browser, or press "clear" to manually erase the data. Or you can "clear all" of the info — all the data your browser's been keeping on you.

Apple's Safari browser also has an easy one-click option to clear everything. Just select "Reset Safari" from the "Safari" pull-down menu and you'll get this option:

Controlling and Limiting the Logs Kept By Web Servers

Web servers usually see and retain a large amount of information about what you do when you surf to them. For instance, if you type any information into a form on a web page (such as a search engine), the server will record not only what you sent it, but also information that might identify you: your IP address, the browser and operating system you are using, whether you followed a link from another web page to get to the page, what that previous site/page was, your account if you are logged in to the site, and cookies that were created when you previously looked at pages on the site.

Web Privacy is Hard

If you use a particular website a lot, the chances are that it is going to end up retaining a huge amount of information about you. To get a sense of the kinds of information, and what needs to be done to prevent them from being aggregated, read our white paper on search privacy. Although that document primarily discusses search engines, the issues to consider for other kinds of sites are similar.

Cookies

Cookies are pieces of information that a web site can send to your browser. If your browser "accepts" them, they will be sent back to the site every time the browser accepts a page, image or script from the site. A cookie set by the page/site you're visiting is a "second party" cookie. A cookie set by another site that's just providing an image or script (an advertiser, for instance), is called a "third party" cookie.

Cookies are the most common mechanisms used to record the fact that a particular visitor has logged in to an account on a site, and to track the state of a multi-step transaction such as a reservation or shopping cart purchase. As a result, it is not possible to block all cookies without losing the ability to log into many sites and perform transactions with others.

Unfortunately, cookies are also used for other purposes that are less clearly in users' interests, such as recording their usage of a site over a long period of time, or even tracking and correlating their visits to many separate sites (via cookies associated with advertisements, for instance).

With recent browsers, the cookie setting that offers users the most pragmatic tradeoff between cookie-dependent functionality and privacy is to only allow cookies to persist until the user quits the browser (also known as only allowing "session cookies").

You can enable this in the "Privacy" tab of Firefox 3's "Preferences" pane:

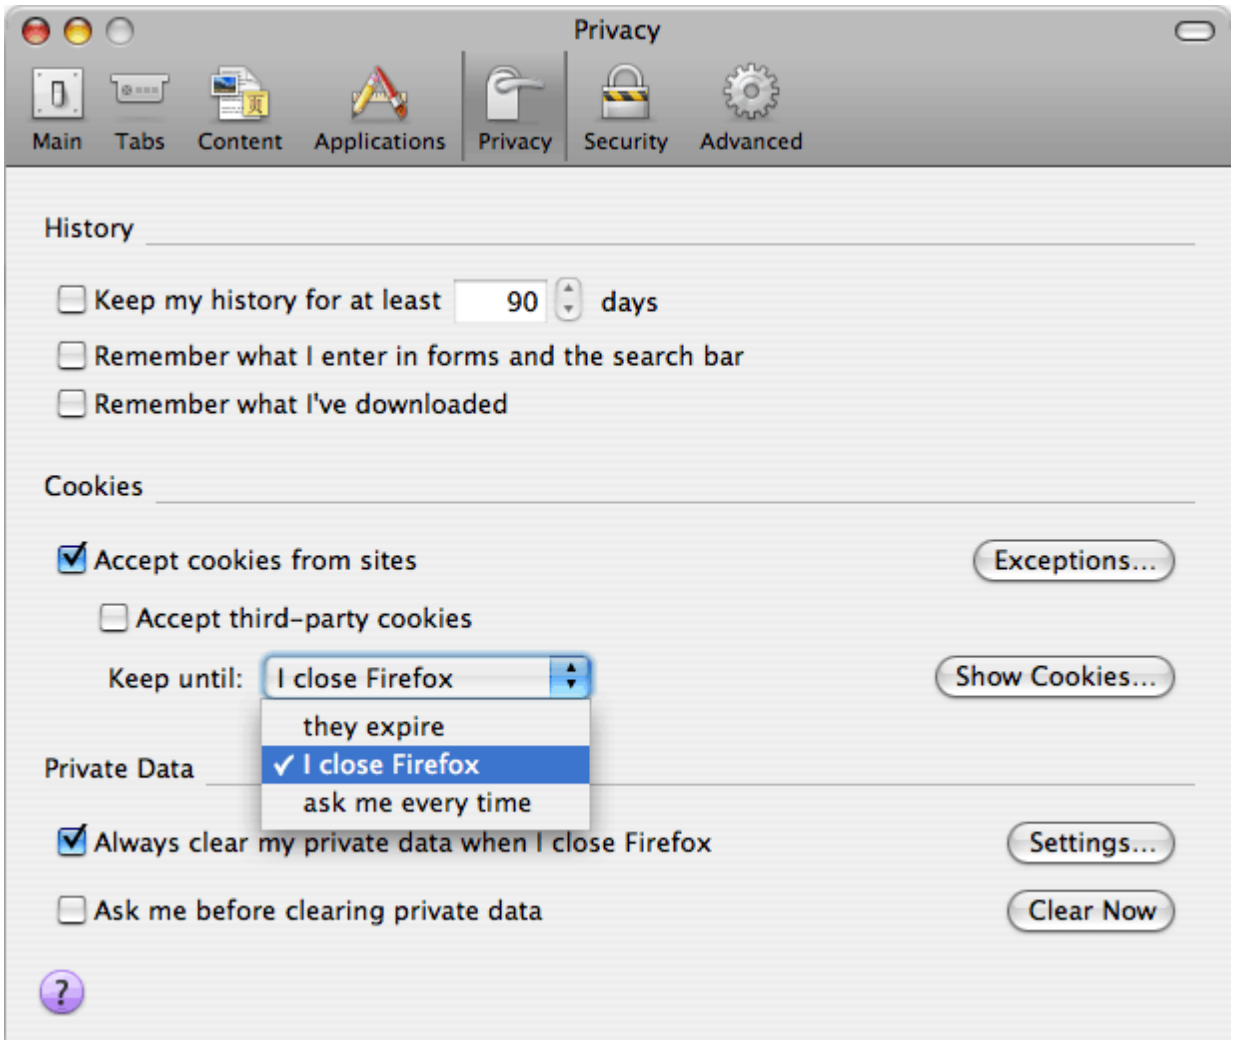

Unfortunately, if you only quit your browser entirely once every week or two, web sites will still collect a huge amount of information about your habits, such as the IP addresses you use at home, at work, at friends' houses and at Internet cafes. However, the "Incognito" mode offered by Google's Chrome browser and the "InPrivate" mode offered by Internet Explorer 8 are signs that in future browsers may offer more convenient ways to limit cookie tracking.

Sophisticated users can configure their browser to manually decide whether each site they visit is allowed to set cookies. This may have good privacy outcomes, such as allowing session cookies for sites the user logs in to or purchases things from, but not any other sites. But it requires a lot of work. A certain amount of debugging may also be required for situations where sites are poorly designed and fail to function without certain third-party cookies.

Recent Cookie-Like "Features" in Web Browsers

In addition to the regular cookies that web browsers send and receive, and which users have begun to be aware of and manage for privacy, companies have continued to implement new "features" which behave like cookies but which aren't managed in the same way. Adobe has created "Local Stored Objects" (also known as "Flash Cookies") as a part of its Flash plug-ins; Mozilla has incorporated a feature called "DOM storage" in recent versions of Firefox. Web sites could use either or both of these in addition to cookies to track visitors. We recommend that users take steps to prevent this.

Managing Mozilla/Firefox DOM Storage Privacy. If you use a Mozilla browser, you can disable DOM Storage pseudo-cookies by typing **about:config** into the URL bar. That will bring up an extensive list of internal browser configuration options. Type "storage" into the filter box, and press return. You should see an option called **dom.storage.enabled**. Change it to "false".

Managing Adobe Flash Privacy. Adobe lists advice on how to disable Flash cookies here. There are some problems with the options Adobe offers (for instance, there is no "session only" option), so it's probably best to globally set Local Stored Object space to 0 and only change that for sites which you're willing to have tracking you. On the Linux version of Adobe's Flash plugin there doesn't seem to be a way set the limit to 0 for all sites — consider donating or contributing to the Gnash project to give users an alternative to Adobe's privacy-unfriendly design decisions.

Aside from being an annoying medium for advertising, Flash poses other kinds of privacy and security risks. Some people choose not to use Flash at all (using other tools like youtube-dl for watching Youtube videos). Others install a Flash management browser plugin like FlashBlocker. Unfortunately, while FlashBlocker makes surfing the web a more peaceful experience, it does not protect you from being tracked by Flash cookies or from exposure to other flash-based security risks.

IP Addresses

Whenever your browser fetches a page, image or script from a website, you should expect the website to record the IP address of the computer you're using. Your ISP, or anybody with the power to subpoena your ISP, could tie those records to the Internet account subscription you

are connected through. Use Tor (or a proxy server, which is faster but less secure) if you wish to prevent these records from being created.

Privacy on the wire

HTTPS

Most sites on the web are accessed using the unencrypted HTTP protocol. HTTP is susceptible to eavesdropping, and even to intermediaries that might set out to modify the pages a browser is fetching.

HTTPS is a more secure alternative to HTTP. HTTPS encrypts pages, and attempts to ensure three things: (1) that third parties cannot see the contents of the page; (2) that the page cannot be modified by third parties; (3) that the page was really sent by the web server listed in the URL bar.

Unfortunately, a web server must be configured to support HTTPS properly before you can use it. If there is a site you were planning to send sensitive information to, ensure that you are using HTTPS. If a site doesn't support HTTPS, don't send sensitive information to it.

Some Notes on Using HTTPS

Check three indicators to ensure that you're at an HTTPS page: (1) the URL begins with `https://`; (2) there is a lock icon in the corner of the browser; and (3) the URL/location bar is colored.

If you receive a warning about certificates, or a see broken lock icon, you should assume that any of the security properties of the page could be broken. Contact the site's webmaster and have them fix the problem before sending any sensitive information to the site.

Blocking Javascript for Browser Security and Privacy?

Javascript is a simple programming language which is part of modern web browsers. Unlike HTML, javascript allows a page to make the browser perform complicated and conditional calculations in determining what a page will look like and how it will function.

Javascript has many uses. Sometimes it is simply used to make webpages look flashier by having them respond as the mouse moves around or change themselves continually. In other cases, javascript adds significantly to a page's functionality, allowing it to respond to user

interactions without the need to click on a "submit" button and wait for the web server to send back a new page in response.

Unfortunately, javascript also contributes to many security and privacy problems with the web. If a malicious party can find a way to have their javascript included in a page, they can use it for all kinds of evil: making links change as the user clicks them; sending usernames and passwords to the wrong places; reporting lots of information about the users browser back to a site. Javascript is frequently a part of schemes to track people across the web, or worse, to install malware on people's computers.

For this reason, sophisticated users with strict security and privacy requirements may wish to consider selectively blocking javascript in their browser. There is a Mozilla/Firefox plugin called NoScript which is very useful for this purpose. Noscript (1) allows you to see the sources of any javascript in a page (many pages include javascript from third parties); (2) blocks javascript by default and (3) allows javascript from particular sources to be temporarily or permanently reenabled. Surfing the web with NoScript is more work (because when you visit new sites, you may have to enable some javascript sources to make them work properly), but surfing the web with NoScript is also much more secure.