Cookies

What is a Cookie?

A cookie[1] is (usually) a small amount of information sent by a Web server to a Web browser with the expectation that, when the user visits the server (website) again, the browser will send the information back to the server. Web servers use cookies to store user-specific information on the user's computer. For example, a cookie may be used to store the user's website id on his/her computer so that the user does not have to supply the id every time when visiting the website. In successive visits to the website, the browser will send the cookie with the id to the server, thereby identifying the user.

A browser will send a server only cookies that were previously sent to it by that server, or by a related server. (A website can have multiple servers.) A server cannot download arbitrary cookies from a user's computer.

Some cookies have lifetimes longer than a single "session" (set of visits before the server decides that the user has gone). Upon termination, a browser stores long-lived cookies in a file on the hard disk.

Cookies are implemented using specific request and response headers in HTTP (Hypertext Transfer Protocol), the underlying Web protocol.

Why Cookies?

To understand the need for cookies, it helps to understand how HTTP works. In the simplest case, a browser sends a request message to a Web server. The server sends a response back to the browser. The server then forgets everything about the request. The server is thus ``state-less." From the server's point of view, each request from a browser is independent of every other request.

However, many Web applications require some enduring information, known as state information, to be remembered from one user request to another. For example, in a shopping application the server needs to maintain state somewhere. Otherwise, it will not be able to maintain a ``shopping basket." One place to store the shopping basket is on the user's computer. The server, in each response to the browser, sends a cookie representing the contents of the shopping basket. The server does not have to remember the basket's contents, because the browser remembers the contents for it.

---

[1] The term ``cookie'' has long been used in computer circles to mean an arbitrary token, something that gets passed around between components of a software system. Lou Montulli of Netscape coined this use of cookies to describe information sent by a server to a browser for storage on a user's computer.

© Silicon Press

When the user shopping on the website initiates the next request to the server, the browser sends the shopping basket cookie back to the server.

Cookie Alerts

New versions of browsers can be configured to alert a user whenever a website sends a cookie. This may result in lots of alerts, because many websites now use cookies for many purposes. One common use of cookies is to track the path of an individual user through a website. The website can do the tracking without knowing the user's identity and it can determine the set of pages the user (presumably) found interesting because the user visited them. Because each request to a Web server is independent, this kind of tracking would be hard to do without cookies. Browsers that provide more selective control over which cookies to accept can reduce the number of alerts.

Cookies and Personal Information

Servers can use cookies to collect personal information, but only if a user volunteers information that can be associated with a cookie. One example of volunteering information is registering at a website. Such a website will be able to associate the user's registration information with the user's visits to the website. The user's trajectory around the website may provide the website with a profile of the user's interests which could then be used, for example, for targeted advertising.

What Are Some Relevant Standards?

The Internet Engineering Task Force (IETF) hosted a standardization effort for cookies that led to the specification RFC 2965, titled *HTTP State Management*. RFC 2965 set out to do two things:

- Define more precisely the way browser and server applications use cookies, to ensure interoperability.
- Provide greater user control over cookies.

RFC 2965 contains a section on privacy that calls for greater user notification and control than supported by many current browsers. Specifically, it says users should at least be able to

- disable sending and saving of cookies,
- determine when cookies are being used, and
- control the use of cookies based on the website sending them.

Where Can I Find More Information?

- HTTP State Management (Cookie Specification), RFC 2965 (ftp://ftp.isi.edu/in-notes/rfc2965.txt)
- HTTP Specification (ftp://ftp.isi.edu/in-notes/rfc2616.txt)
- Netscape's Original Cookie Specification (http://www.netscape.com/newsref/std/cookie_spec.html)

Author: David M. Kristol, info@silicon-press.com

SILICON PRESS www.silicon-press.com
Email: info@silicon-press.com
Tel: 908-273-8919 Fax: 908-273-6149
© Silicon Press