

The Right to Delete

Chris Conley¹

ACLU of Northern California
39 Drumm St., San Francisco, CA 94111
cconley@aclunc.org

Abstract

Most of us have incidents in our past that we'd rather leave there, but that is increasingly difficult in a world teeming with devices and services that capture our words and actions and record them indefinitely. Instead of organically being forgotten, records of the past increasingly persist in digital storage unless and until they are deleted by someone.

Should an individual have the right to demand that records about her be deleted? Does it matter who holds these records or what forms the records take? And even if this right would be socially beneficial, can it be implemented?

In this paper, we argue that an individual should have the right to delete information about her that is held by others, and sketch out frameworks of how such a right might work. We suggest methods of implementing this right using technical tools, legal regulation, and/or social norms and market forces. Even without the legal component, we believe that collective action has the potential to give individuals greater control over their own personal information by establishing a widely (if not universally) accepted right to delete.

Introduction: Do We Need a Right to Delete?

For human beings, forgetting is easy and remembering is hard. While this can be a challenge, it is also in many ways a boon: we can distill the past into a few simple memories rather than reliving it verbatim over and over again; we can “forgive and forget;” we can grow and change without being forever linked to our past (Mayer-Schönberger 2009).

Modern technology changes this paradigm. With computers and electronic devices, remembering, rather than forgetting, is increasingly the default. Search engines record every search to improve their performance, and have only recently begun to delete the oldest records on their servers. Social networks take the transient “tweets” and “status updates” of millions of users and turn them into

permanent records. Cell phones and email services generate logs of our conversations, however mundane and “forgettable” they might be.

These records of the past can have consequences long after the event they record is forgotten to the human mind. A teacher's career may be ruined by a picture of her holding a drink at a party long ago. Candidates for a job may be evaluated based on their student activism years prior. Off-color remarks or former relationships may reappear in unexpected contexts.

The preservation of even innocuous information can have disturbing consequences when that information is aggregated. Companies and governments can increasingly infer a great deal about our private lives from records of our “public” activities. For example, recent research has shown that sexual orientation can be predicted solely on the basis of a person's network of friends on Facebook. (Jernigan and Mistree 2008). With a far richer and deeper set of data available on the Internet, our intimate secrets may well be exposed to the world.

Deletion and Privacy

The concerns here should seem familiar, as they are the fundamental concerns that underlie arguments for privacy rights: without control over our own information, we are vulnerable to external forces—and this vulnerability affects the way we think, behave, and grow. Privacy, understood broadly, protects us from abuses of power and allows us to maintain our individuality and liberty (Schneier 2006).

Indeed, observers have noted the risks of persistent data retention within the context of privacy. The Supreme Court recognized that we should not be forever under the shadow of past events or mistakes when it suggested that some information, even public information, should remain in “practical obscurity” (*DOJ v. Reporters for Freedom of the Press*, 1989). But due to today's vast data retention and search capabilities, we cannot simply rely on our personal information remaining hidden while we move on with our lives. As one scholar stated, “technologies are making the past easily and eternally present” (Allen 2008).

Can we retain our privacy, and the freedoms that come with it, absent a counterbalance to this trend? Without some mechanism to delete records that have escaped our

¹ Chris Conley is the Technology and Civil Liberties Fellow of the ACLU of Northern California. The views expressed in this paper are those of the author and should not be ascribed to the ACLU-NC. The author is grateful for the assistance of Tamar Gubins and Hari O'Connell in preparing this paper.

immediate control, we have no remedy for information that is inadvertently made public or revelations of intimate details about us that are derived from public information, and the benefits of privacy may be lost.

Deletion and User Expectations

In addition, establishing some mechanism to delete records provides protections for individuals who fail to understand the extent of current record collection or their rights concerning those records. Users of new technologies may not fully realize the extent to which the content they create and records of their activities are shared and stored. Even those who are aware that their personal information is held by third parties may not understand what rights (if any) they retain over this information. For example, in a 2009 national telephone survey conducted by the University of Virginia and University of California-Berkeley, 54% of American adults falsely believed that, if a website has a privacy policy, the site must comply with a request to delete information about a user by that user (Turow et al. 2009).

One could, of course, simply endorse a mindset of “caveat user” and put the burden of understanding the system of data collection on the user. But even savvy users have trouble predicting the extent to which new technologies will expose information about their past from seemingly innocuous sources. Major companies such as AOL and NetFlix have released “anonymized” records to the public for research purposes, only to discover that these records could be de-anonymized and linked back to specific individuals (Ohm 2010). If sophisticated online companies cannot accurately understand the consequences of sharing data, it is unrealistic to expect that a typical user will do better. One way to address the risks this presents is to provide some mechanism for deleting older records.

Why a “Right” to Delete?

So how do we protect privacy and self-determination interests in the context of permanent memory? One possibility is to restore the previous status quo by building a technological “self-destruct” mechanism into records (Mayer-Schönberger 2009; Geambasu 2009). This approach, however, has one considerable issue: either the “timer” must be the same for all data, ensuring that data with long-term value is subject to the same destruction mechanism as sensitive data whose survival does more harm than good, or the timer must be set by the creator of the record when the record is created, who may not have either the foresight or the incentive to ensure that potentially harmful records are destroyed in a timely manner. (It also suffers from similar problems as DRM for copyrighted materials, namely hindering the distribution and use of valuable material and risking circumvention of the self-destruct mechanism—but these problems are pervasive in any technological approach to content control, which is why technological solutions alone are insufficient.)

A second approach, of course, is to simply allow technology to eliminate forgetting entirely. Some proponents of “lifelogging” suggest that the benefits of recording and remembering every facet of one’s life will far outweigh any negatives from doing so (Bell and Gemmell 2009). However, even if this were true for individuals, ubiquitous permanent recording threatens serious harm on the societal scale: fewer will be willing to adopt unorthodox views and challenge the status quo if all ideas and comments are recorded and made permanently available. It is precisely these unconventional ideas that spur societal growth, change, and innovation.

We suggest, therefore, a different approach: a right to delete certain records from any permanent repository where they might reside. By making deletion a possibility, we hope to preserve the right to privacy and the social breathing space it enables; by making it manual rather than automatic, we hope to empower individuals to control their own information so that they can choose what to retain and what to delete. We also hope to build in processes to balance this right with competing interests, such as freedom of expression, freedom of contract, preservation of socially valuable information, and mutual parties with interests in the same record.

As we move into a society built around information, we need to expand our conception of “person” to match. We need to recognize that records about an individual are actually a part of that individual’s digital identity, and that the individual has the right to control her identity rather than have it controlled by others who hold these records. And an essential part of this right to self-determination is the right to delete.

Framing a Right to Delete

Of course, the idea of a right to delete is far more complex than a simple assertion that an individual should have control of their own past. This section attempts to define what a right to delete might encompass, balancing interests in privacy and self-determination with competing expressive, economic, and social goods. We start by considering existing legal frameworks, namely property law and privacy torts, as potential models; we then proceed to the details by addressing the following questions: (1) what is the scope of the right, i.e., which content or records does it cover; (2) who are the parties to the right—who can exercise the right and who is subject to it; (3) what does the right actually entail, and what burdens or duties do the various parties have; and (4) what limitations should be placed on the right or the exercise thereof?

Models for the Right to Delete

Before diving into specifics, we want to address the broader framework of a right to delete. We examine two existing frameworks as models for such a right: property law and privacy torts. Although privacy torts provide the closest current approximation of a right to delete in the

U.S., we find the tort model inadequate to address some of the particular challenges of a right to delete. A framework based on the affirmative right to control records about one's self, without requiring objectively demonstrable harm, seems preferable. (However, as discussed below, the enforcement mechanisms generally associated with property law may not be suited to the right to delete.)

Property Rights. Property law concerns ownership, which amounts to a right to exercise control over that property, including the right to exclude others from using the property, and even a right to destroy the property. A variety of intellectual property regimes, including copyright, trademark, and patent law, extend property rights to nonrivalrous, intangible property, including digital property. Copyright law, which provides authors with a bundle of rights concerning production and distribution of copies of the original work, is the closest analogy to a right to delete.

Privacy Torts. Tort law is based not on ownership of property or personality but on the right to restitution for harm. There are already several privacy torts that bear some semblance to the right to delete. In particular, the tort of public disclosure of private facts—which is recognized in 36 states—provides a cause of action following public disclosure of information if the disclosure would be highly offensive to a reasonable person.

Comparison. The tort law model is difficult to expand into a broader right to delete. Tort law requires harms that are definite and measurable; this standard may be difficult to satisfy in the case of information privately held, as it would require a plaintiff to both know with certainty that a given individual held a record and show that she was harmed by the fact that they held that record. Property law provides a better model, giving individuals affirmative rights without any need to demonstrate harm.

Conceiving a right to delete as an individual's affirmative right does not necessarily entail individual enforcement, however. As we discuss below, a regulatory scheme for ensuring that data providers comply with requests to delete information may be preferable to an individually enforceable right to delete.

Specific Framework: Scope, Parties, Rights, and Limitations

Within the general framework of a property- or ownership-based right to delete records about one's self, however, there are still several specific questions that need to be addressed. First, what kind of records or content is covered by the right? Second, who possesses the right, and who is subject to it? Third, what does the right actually entail, and what duties does it impose on others? And, finally, what are the limitations on the right that keep it in balance with other individual or societal interests?

Scope. We begin by addressing the issue of scope. Unlike copyright, which covers only expressive works, a functional right to delete needs to cover content that is

informational in nature, such as records of attendance or logs of search queries.

However, content can be both expressive and informational; a photograph of a political rally may be both a creative piece and documentary evidence that a given person attended the rally. In order to protect freedom of expression, we suggest that a right to delete should be limited to records that are primarily informational (which might include some photographs such as security camera footage) or that can be separated from any associated expressive content (such as tags connecting the above photograph with a specific person's name, but perhaps excluding the from address in an email). Obviously, this distinction does limit the efficacy of the right to delete, but we do not believe that proposing a right that substantially burdens freedom of expression is appropriate. Furthermore, particularly egregious uses of private information are still subject to the tort of public disclosure of private facts, as described above.

Beyond this limitation, the right to delete should be applicable to any record that is associated with a specific individual. This association may be direct, as in an email address or a tag associated with a photograph; it may also be indirect, such as an IP address in combination with a timestamp, or a description that demonstrably applies to a specific person.

Parties. As defined above, a right to delete would naturally belong to the individual who is referenced in a given record. Because this right is grounded in concepts of self-determination and individual liberty, we limit it to real persons and see no reason to extend a right to delete to corporations. In addition, to exercise the right, the rightsholder must be authenticated in an appropriate fashion to prevent abuse of the delete mechanism by malicious parties. (While we acknowledge that this authentication is no trivial matter, we are not able to address it in this paper.)

There are two possible ways to define those subject to the right. We could apply a right to delete to anyone who possesses a record within the scope of the right, without regard for the nature of their possession or use of the record. Alternately, we could limit those subject to the right to anyone using the record for commercial purposes. Narrowing the group subject to the right is appealing in many ways; it seems to get to the root of the current threats that the lack of a right to delete poses, and it avoids many of the conflicts with other individual rights. However, as data mining and web crawling technologies advance, the threat from non-commercially held information may come to pose as great a threat. In addition, a broad definition would not necessarily subject corporations and individuals to the same duties and burdens, as discussed below.

Note that, by extending obligations to the "possessor" of a record, we do include intermediaries who did not create the record but are storing it on behalf of another (who may or may not be the subject of the record). Rather than exclude these parties from the right entirely, we instead define their duties differently below.

Rights & Duties. As the name implies, the right involved here is a right to delete—but to delete what? The identifying information in a record? The entire record? Copies of records sent to third parties? Archived copies of the offending record in offline storage?

To promote a balance between the interests of privacy and other interests, we suggest that the right to delete should generally encompass the deletion only of any association with a given record, not necessarily the entire record itself. Of course, in many cases, the record is inherently identifiable—a person’s face captured by a security camera, for example—such that the record cannot be de-associated without deleting it entirely. In addition, this de-association must be sufficient to resist expected attempts to re-establish the link between record and individual. It is increasingly clear that simply replacing identifying marks with pseudonymous data or similar efforts to camouflage identity is insufficient. (Ohm 2010).

The duty to delete should be satisfied whenever the possessor of a record makes a reasonable effort to find and delete the indicated record. The definition of “reasonable” may be very different for different entities: a company that uses records commercially is likely to have a far greater ability to locate records associated with an individual than a private person with a disorganized collection of files on her hard drive, and could breach this duty if it does not have deletion capabilities that meet industry standards.

For possessors of a record who are merely hosting or storing it for a third party, the duty to delete could be replaced with a duty to forward the deletion request to the author or owner of the record and to delete the record if the author consents or if no response is received within a reasonable period of time. This would allow an author the opportunity to challenge a deletion request if the record falls into a given exception, rather than forcing the hosting service to arbitrate any requests on behalf of its users. (We note that this model presents difficulties in the context of anonymous hosting services; see below.)

To be effective, this right must also follow any copies of the original record that remain accessible, including those sent to third parties. This could be handled in one of two ways: the party receiving the delete request could either pass on the delete request to any subsequent recipient, or it could notify the rightsholder of the identities of subsequent recipients. The latter option divides the burdens between the two parties: the possessor of an associated record has the obligation to keep track of potential recipients of a given record, while the identified individual retains the burden of enforcing her rights against any subsequent possessors of her records. In addition, as with the duty to delete, we limit this duty by a standard of reasonableness: corporations with extensive record-keeping should be able to determine who has or might have been provided with a copy of the record, while individuals would not be required to keep archives of every interaction in order to comply with this requirement.

Balancing the Right to Delete. Finally, having set out a fairly broad framework for a right to delete, we come to the

question of limitations. When should the right be limited so as to balance properly with other rights and interests? We identify four separate areas where the right might be limited or a process to resolve disputes could be established: conflicts with freedoms of speech and of the press; interactions with the right to contract; records associated with multiple individuals; and situations where deletion is impossible, infeasible, or socially harmful.

By espousing a right to delete records that are specifically about a given person, we threaten to infringe upon the freedom of expression, giving one person (the subject of the record) the right to silence another (the creator of the content). Limiting the right to delete to non-expressive content goes some distance towards

The more difficult conflict between a right to delete and free expression comes in the context of anonymous speech. Our proposal would allow individuals to demand that a hosting company, such as a blog service, delete records if it is unable to contact the creator of those records—something that is to be expected when the creator is anonymous. Unfortunately, we do not see a simple approach to remedying conflicts of this sort. Negating the right in this context could lead to abuse, but leaving the authority to consent or contest a deletion demand with the host may not sufficiently protect anonymous speech. Ideally, hosting servers will be willing to stand up for their anonymous clients’ rights and resist demands to delete truly expressive content, but that ideal remains largely untested.

In addition, a right to delete may have additional impact on freedom of the press, in the sense that a right to delete incriminating records from one’s past reduces or eliminates the press’s ability to (re)publish these incidents if they return to relevance. Providing an exception for “newsworthy” facts and records would seem to protect press freedom—but defining the contours of that exception may be difficult, especially as it pertains to information that has no apparent newsworthiness at present but may be of public interest in the future.

A second concern with a right to delete is how it will interact with the right to contract. If the right to delete can be waived simply by agreeing to a Web site’s Terms of Service, it is likely to have no practical effect whatsoever; on the other hand, if it can never be waived at all, it may hinder beneficial projects that involve long-term collection and use of data. We suggest that the best approach is to allow waiver of the right to delete only with specific, express, and actual informed consent. Furthermore, this waiver should apply only to the uses of the record specified when the consent is given, and not to any additional uses or to any third parties who obtain a copy.

Most records associated with multiple individuals present little difficulty for a right to delete: simply deleting the references and identifiers for the given individual (assuming the record falls into the scope of the right) satisfies that person’s rights without impacting anyone else. But what of the situation where the record is specifically *about* the associations between individuals, such as a friend connections on a social network? In the modern age,

deleting such records is often seen as a necessary part of fully ending a relationship. Should the right to delete trump any other person's interest in retaining a record specifically about a relationship?

Finally, a right to delete needs to exempt records that cannot feasibly be deleted. Records that an entity is legally required to retain, records of financial transactions needed for verification or accounting purposes, archives of a database in secure offline storage that cannot be altered without losing their archival character, and records that are the subject of an existing search warrant or wiretap order should not be subject to a deletion demand. The possessor of a record should always be able to claim that the deletion request is impossible or impracticable, and such claims should be judged on their merits. However, this exception should be narrowly interpreted: it should apply only where the record holder is required for some reason to retain the information in the record, not simply where doing so suits the holder's business model or database structure.

Implementing a Right and Duty to Delete

We have framed the concept of the right to delete on the basis of law, specifically property law—but in doing so, we do not mean to suggest that the right to delete should be a strictly legal construct, or that it must be implemented similarly to copyright. Instead, there are a number of approaches to implementing a right to delete: technical measures that give individuals more control over personal records; laws and regulations that require the compliance of certain actors; and market pressure and social norms that encourage the recognition of a “right” to delete whether or not it has any legal status.

Enabling Through Automation and Standardization

Several projects have attempted to enable the automatic or manual destruction of records. By building an expiration date into the content that we create, we could indeed address some of the privacy concerns that persistence presents (Mayer-Schönberger 2009). The challenge with this approach is twofold. First, it requires some mechanism to track and delete all instances of a record or document—a problem that is particularly difficult in the context of informational records which are easily integrated into other records and shared with other parties. But the greater limitation with this type of tool is that it places the timer in the hands of the record's creator—and, where the creator is not the subject, that creator may have very little incentive to ensure that the record is destroyed in a timely fashion. Thus, while this may be a valuable tool for controlling content that we create even if held by others, it seems to have limited applicability to records about us that are created by others.

Instead, we advocate for a manual ability to delete records. From the user side, the key technology needed to enable the framework we envision is a “deletion manager”:

a tool that can automate the process of deleting records by identifying and interacting with record-holding parties, track the flow of records from one party to another, and—most importantly—provide an interface that gives an individual the information she needs to decide when or whether to delete records without being overwhelmed by the volume of records that capture her life in otherwise-permanent storage. Standardizing and automating the process of deleting one's history, rather than being forced to utilize idiosyncratic interfaces and options to attempt to delete records one by one, would make exercising a right to delete feasible.

Some progress in this direction has already been done. Google's Dashboard, for example, gives users the ability to delete selected records from a wide range of products and services that Google offers (Gross 2009). And products like the “Web 2.0 Suicide Machine” allow users to erase their records and profiles from multiple social networking sites using a single tool (Schonfeld 2009). By combining the functionalities of these products to allow fine-grained control over a wide range of information, and tying in discovery mechanisms that can locate and interact with record holders, a usable “deletion manager” could be produced.

Enforcing Through Legal Requirements

The framework of the right to delete spells out the basic duties of those subject to a right to delete: taking appropriate steps to identify and delete requested records and to identify and disclose any other party who may have a copy of the deleted records. There are a number of possible legal mechanisms to require compliance with these duties, including private or class-action lawsuits to penalize anyone who fails to comply. However, these mechanisms are often quite onerous for individuals with limited resources to hire lawyers, and may limit the effectiveness of the right.

An alternative approach would be to appoint an administrative agency, such as the Federal Trade Commission (FTC), to establish required practices for all parties subject to the right to delete, or to approve self-regulation by industry parties. In either case, these regulations should include procedures to ensure that requests to delete are either complied with or challenged in a timely manner and the adoption of relevant interfaces and other standards that allow individuals to use tools such as those described above to submit and monitor requests for deletion.

If any legal enforcement of a right to delete is implemented, it will require the establishment of legal standards for evaluating any of the disputes identified above: whether a record is expressive or informational, authentication of the requesting party, whether the individual has waived her right to delete, and so forth. Exceptions that are common to many parties, such as deletion requests for financial transaction records, should be defined and applied broadly; exceptions for specific parties should be evaluated individually. Finally, there may

be a need for a panel with the authority to effectively adjudicate specific disputes between rightsholders and recordholders.

Enabling Through Market Pressure and Social Norms

Of course, binding legal requirements are politically difficult to implement, and may be difficult to enforce, as the recording industry has seen in ongoing efforts to stem copyright infringement. Societal buy-in may be both necessary and sufficient to establish a right to delete.

If members of society can agree that individuals deserve the right to own their own digital persona, including records that are held by third parties, then a right to delete can be established absent any legal change. Private individuals will find a negotiated means of complying with requests to delete information (assuming that they even retain information in searchable form), and market actors will adapt to changing consumer expectations.

Of course, actors that rely on monetizing records about individuals may resist this trend, but these same actors are often susceptible to collective action and public pressure. Facebook, for example, reversed changes to its Terms of Service and reassured users that they retained the right to delete content that they post to the social network (Stelter 2009). With more transparency and increasing public demand for control over personal information, we remain hopeful that these actors will give users the rights they demand and deserve, including the right to delete.

Conclusion

As we move into a world where memory is perfect and permanent, we should consider whether we need some mechanism to replace the ability to forget. Without such a mechanism, we may lose our ability to invent and reinvent ourselves, and instead find ourselves constrained by actual records of our past or feared records in our future. The right to privacy, a right many consider fundamental to our society, may be rendered impotent if our private actions can be reconstructed from countless permanent records.

We propose that the best way to address this concern is to create a right to delete that gives individuals the ability to control their own history and thus escape it. This right comes from the idea that records are not just about a person; in our modern world, they are functionally part of our digital persona, and thus should be under our control whether we create them or not. By establishing a right to delete, and balancing that right with other concerns, we believe that we can reap the benefits of our ever-expanding technological capacities without leaving privacy behind.

We envision this right as a combination of technical tools, legal regulation, and social norms and market pressure that will work in combination with other laws and technologies to promote individual control of personal information. And we have already seen some progress as companies have launched products and altered terms of

service to give users the right and ability to delete their own records. Although establishing this right remains challenging, we believe it can be done.

References

- Allen, A.L. 2008. Dredging up the Past: Lifelogging, Memory, and Surveillance. *University of Chicago Law Review* 75: 47.
- Bell, G., and Gemmell, J. 2009. *Total Recall: How the E-Memory Revolution Will Change Everything*. Boston, Mass.: Dutton .
- DOJ v. Reporters Committee for Freedom of the Press*, 489 U.S. 749 (1989).
- Geambasu, R., et al. 2009. Vanish: Increasing Data Privacy with Self-Destructing Data. In *Proceedings of the Eighteenth USENIX Security Symposium*, 299–316. Montreal, Quebec, Canada: The USENIX Ass’n.
- Gross, D. 2009. Google Releases Dashboard Privacy Tool. *CNN.com*.
- Jernigan, C. and Mistree, B. 2008. Gaydar: Facebook Friendships Expose Sexual Orientation Student Paper, Department of Computer Science, Massachusetts Institute of Technology, Cambridge, Mass.
- Mayer-Schönberger, V. 2009. *Delete: The Virtue of Forgetting in the Digital Age*. Princeton, N.J.: Princeton University Press.
- Ohm, P. 2010. Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization. *UCLA Law Review* 57 (forthcoming).
- Samuelson, P. 2000. Privacy as Intellectual Property. *Stanford Law Review* 52: 1125.
- Schneier, B. 2006. The Eternal Value of Privacy. *Wired News*.
- Schonfeld, E. Wipe the Slate Clean for 2010, Commit Web 2.0 Suicide. *TechCrunch*.
- Stelter, B. Facebook’s Users Ask Who Owns Information. *The New York Times* Feb. 17, 2009: B3.
- Turow, J., et al. 2009. Americans Reject Tailored Advertising. *Social Sciences Research Network*.
- Werro, F. 2009. The Right to Inform v. the Right to Be Forgotten: A Transatlantic Clash. In *Liability in the Third Millennium*, 285–300. Baden-Baden, F.R.G.: Nomos.