Democratize the Internet Now!

"Some buddies of mine need your help with a database," my friend said.

"I'm good at databases," I replied. "And I like to help."

So I found myself in a large conference room with two nervous men who wanted to know exactly how exposed they were by the Ashley Madison leak. They wanted me to look inside the leaked data to see if I could find any traces of their exploits.

"I can do this," I said.

This was several leaks ago, so a refresher: Ashley Madison is a web site that helps you have affairs. You enter your personal information and the site lets you look at other people who have also entered their information. Then you can make arrangements to have sex with these other people. And because it is digital, it felt anonymous, which meant you could structure the social interactions so that no one in your immediate vicinity (spouse, children) would ever know.

Ashley Madison is a heavily advertised digital product, and until the leak it was doing pretty well. Roughly 35 million people had signed up, and presumably more than one of them had sexual intercourse without the knowledge of their spouses. By which I mean wives—the users were overwhelmingly men.

Then it all went pear-shaped. One day last summer, an individual or individuals known as "The Impact Team" determined that all of the information in the Ashley Madison database should be made public. This was a lot of data: tens of millions of names, addresses, profiles, and credit card transactions. The database, in essence, *was* Ashley Madison.

It felt as if a nuclear bomb had gone off in the datasphere. The files were spread through the BitTorrent network, which meant that lots of people could easily download them; it also meant that the files were difficult to suppress, because they were so widely distributed.

I downloaded the database in anticipation of my new friends' arrival. And while I was unimpressed by the database itself—the typical mess of MySQL fields, with functional but hardly exemplary data modeling—I was staggered by its scope, the fact that one big database could serve all these people looking for affairs. It was the encapsulation of so much human desire. I poked around. Then, seeing a name I recognized, I stopped poking around. It wasn't worth knowing.

One of the visitors was sweating as he handed me a list of names written in pencil on the back of an envelope. I'd never seen that before: a man in a cold, air-conditioned room sweating, slightly wild-eyed.

"I just need to know," he said. "How bad is it?"

I started with last names.

"That's not me," said the other man. "That's my brother. I'll have to talk with him."

We found some profiles, and I read them. Walks on the beach. Long nights, fine wines. Nothing cruel or strange. We also found records of credit card transactions. They'd signed up, left a trail, and it was still there. One of them had a credit card connected to his home address. His friend shook his head. Poor bastard.

"OK," the man said. "It's better to know."

Once there was a time when if you wanted to have an affair, you had to take charge of keeping it a secret. Wink slyly at a prospective lover and receive a subtle nod in return. Leave a scented note in a mailbox. Meet at a motel off the main route. Pay the bill in cash; never call your lover's house. Ashley Madison's innovation was that it took care of all that for you. The entire come-on of the site was that it would make seamless—*frictionless*—something that had previously been difficult and time-consuming. All while reducing the risk.

It seemed safe and secure, but it wasn't. Ashley Madison knew what you told it: name, email address, sometimes a home address. It knew your credit card information, provided so that you could pursue conversations and thus sex. It promised anonymity, but the service it delivered was sitting right in the middle of a world of transactional processing. Thirty-five million individuals had placed responsibility for the continuance of their marriages and relationships in the hands of a single company. Ashley Madison was a massive, centralized agglomeration of indiscretion.

The internet was once a highly decentralized system. In the earliest days, there were no large corporations or service providers like Ashley Madison or Facebook or Twitter, or behemoth databases to house your information. If you wanted to join up, you plugged in a computer and found a connection through a service provider, and that was

basically it. You were online. Your computer was a "peer" of the other computers. It was a computocracy.

When the web came along, it was the same. You wanted to say something, so you ran a web server on a computer. You put some web pages in a folder. Your web server waited, night and day, for other computers to ask it for pages and files, and then sent those files back over the network. The servers were still off on their own, but now they could talk to each other.

That's really all there is to it. It is, at its core, a wonderfully autonomous, independent, and decentralized arrangement. Anyone can set up a web site and point to all the other web pages. Everyone is a publisher. Everyone is a peer. That's why it's called a web. Individuals knit themselves together by linking to one another. Everyone tends his or her own little epistemological garden, growing ideas from seed and sharing them with anyone who comes by.

Yet as the web grew, the problem of *finding* information arose. Search engines were needed that could crawl across the web, indexing the words in web pages; this way someone could type a word in a box and the machine could consult its index and list the pages that matched. But once you do that successfully, you have created something that appeals to larger forces. The search engine has power over other pages—you're no longer a peer.

Imagine you had a huge bread machine and an enormous bag of flour. You made so much bread that you gave some away. And people came to eat the free bread, and they liked it and wanted more. They told their friends. Free bread! People just kept coming—ten people; 100; 100,000; 100,000 million. A googol. To keep up with demand, you find yourself in need of not just more flour, but more bread machines. Fortunately, there are companies that are willing to pay you—not for your bread, but for the right to say, "This bread was brought to you by…" All because you've done the work of getting a lot of people in one place to eat free bread. Eventually you turn out enough loaves that you're designing wearable technology and self-driving cars.

There were other technical demands that chipped away at the decentralized nature of the internet. All those files spread across the web, it turned out, are a chore to manage. You need to update the copyright statement at the bottom of every page you've published, because you're in a new year. Or you need to wipe the CEO's name from all the pages after he resigned in the wake of a sexual harassment scandal. Doing that one page at a time would be a real pain. At the same time, it became possible to rent access to a database on a server somewhere. To solve this too-many-pages problem, people began to put their "content" into databases, and then publish everything through consistent, replicable "templates." As a result, every page on your web site—and everyone else's—eventually came to look roughly the same. Data went into the database via forms and came out via templates. Content was thusly managed. To change the copyright notice across all 100,000 pages of your retro sneaker site, you only needed to change a single line of one template. The CEO's photo could be briskly removed and replaced by the photo of the new, interim CEO. This was obviously a better state of affairs.

Soon the home page, which had enabled individual expressions of interest in [mycology](#) or *Star Trek* or bondage, was subsumed by the blog, which brought form and chronological order to the universe of web content. Tools like [Movable Type](#), [Blogger](#), and [Typepad](#) emerged, which "hosted" your content in their databases. No longer did people tend their own digital gardens. The gardens were tended for them.

Freed from the need to build and manage their own web sites, people could do more social things with their computers. They could talk to each other, start conversations, argue endlessly. They could leave private messages. Many found a community. And the companies that hosted the databases found a business model. Make the messages short, and adapt the database to manage millions of "friends" and "followers" ([Friendster](#), then Twitter). Make a blogging engine that allows you to post short updates and keep track of your friends ([MySpace](#), then Facebook). The computocracy was now something else—a Googlopoly.

The technology that let people make web sites never went away. You can still set up a site as if it were 1995. But culture changes, as do expectations. It takes a certain set of skills to create your own web site, populate it with cool stuff, set up a web server, and publish your own cool-stuff web pages. I would argue that those skills should be a basic part of living in a transparent and open culture where individuals are able to communicate on an equal field of play. Some fellow nerds would argue the same. But most everyone else, statistically, just uses Facebook and plays along.

There's an obvious connection between a decentralized internet, in which individuals create and oversee their own digital identities, and a functioning democracy, in which we make informed choices about who rules us and how we are ruled. Yet too few people make that link. We live in a world in which sensitive information of every conceivable sort—financial, sexual, medical, legal, familial, governmental—is now kept, and presumably guarded, online. It's guarded in gigantic treasure chests labeled "important data here." So many plums for hackers to pluck.

If you don't take care of yourself online, someone else will. That someone is likely not a peer but a megacorporation that is tracking and selling your preferences in a silent auction, a government surveilling your movements and religious affiliations, or a hacker collective that feels entitled to publish your sexual indelicacies. That someone probably already is.

So what is the alternative? For starters: In a utopian vision of a better, devolved-but-more-human internet, I would never post to your database. There'd be zillions of personal data sets, and every individual would have the technical capacity and social resolve to share only what they wanted, plus the power to revoke information from the commons. It's much easier to load my thoughts into someone else's little box and hit "Submit" (perhaps the most well-chosen interface word of all time). But submission comes at a price. My personal information, my finances, my family connections, my ideas—all are now in the hands of those to whom I have submitted.

The temptations of centralization are powerful. With a few employees you can make something worth a billion dollars, as Instagram did. You don't need to worry about advertising. You just create a situation where a larger company sees an opportunity to insert lots of ads. People are desperate to buy places to insert lots of ads so they can resell those places, which is what Facebook did when it bought Instagram. It bought future ad inventory. No one is paying attention to individuals online, at least not any more. There's no money in it. What they are looking for is tens of millions of people all in one place, moving in one direction. If you're Facebook, you need to get in front of that mass of humanity; you need to define their destiny. You need centralization.

Standing against this tide of centralization is the indie web movement. (And hackers, the black hat kind and otherwise.) Perhaps "movement" is too strong—it's more an aesthetic of independence and decentralization. The IndieWebCamp web page states: "When you post something on the web, it should belong to you, not a corporation." You should own your information and profit from it. You should have your own servers. Your destiny, which you signed over to Facebook in order to avoid learning a few lines of code, would once again be your own.

But an affair? That's trickier. A decentralized dating system would end up being a lot like Bitcoin—confusing, briefly exciting, and overpopulated by desperate, libertarian men. And the sites would be vulnerable, hackable. But big decentralized systems have many points of failure, rather than just one. The government could still spy on you, but the damage would be limited. The data wouldn't be in one place.

I'm not proposing some sort of digital back-to-the-land, communal-living, anti-regulation paradise that does away with food-delivery apps and secure online banking. I am an avid self-publisher and web site geek, but I also make a living as a paid client of centralizers.

Oddly, the people most excited about peer-to-peer technologies are not hackers but bankers. "Decentralized applications will someday surpass the world's largest software corporations in utility, user-base, and network valuation," writes David Johnston, managing director of the Dapps Fund, which helps bankroll decentralized consumer apps. There will always be money to be made in big and few. But there is also money—lots of it—in small and many.

Think of Bitcoin, which pioneered a block chain model of financial transactions that has been used by millions. Or Ethereum, which raised $18 million in a single crowd-funding campaign for its secure, peer-to-peer platform for consumer transactions. Or all those little apps on our phones—those incredible pocket supercomputers—talking to billions of other little apps. This is how file-sharing networks already work. BitTorrent uses "trackers" to keep, well, track of the files that people are sharing; the software functions as a tiny server. "I am here," it says. "I have these files, some in completeness and others in parts. I seek parts of some files as well." And that signal goes out to one or more trackers, and then to other clients, and in this way files are distributed. This was how the centralized Ashley Madison database was ultimately decentralized, by force. Someone took the pirated data, zipped it up, and made it available as one big torrent. An older, smaller internet protocol brought down a newer, larger, corporatized database.

Making a shift to a more democratized internet won't be easy. Once you start to rally your energies toward a more open future, you will be shocked by the forces arrayed against you; the intransigence of the people who want to buy and sell your information; the amorality of the hackers who play with millions of people for sport; the cold, endemic corruption of intellectual property and patent law; the infinite protections for copyright. It can get a person down.

We could still live in that decentralized world, if we wanted to. Despite the rise of the all-seeing database, the core of the internet remains profoundly open. I can host it from my apartment, on a machine that costs $35. You can link to me from your site. Just the two of us. This is an age of great enterprise, no time to think small. Yet whatever enormous explosion tears through our digital world next will come from exactly that: an individual recognizing the potential of the small, where others see only scale.

Paul Ford is a contributing editor at the *New Republic*. His book about web pages will be published in 2016 by Farrar, Straus and Giroux.