


IN THE SUMMER of 2013, I spent my days sifting through the most extensive archive of top-secret files that had ever reached the hands of an American journalist. In a spectacular act of transgression against the National Security Agency, where he worked as a contractor, Edward Snowden had transmitted tens of thousands of classified documents to me, the columnist Glenn Greenwald, and the documentary filmmaker Laura Poitras.





I decided to delve more deeply. The public debate was missing important information. It occurred to me that I did not even know what the records looked like. At first I imagined them in the form of a simple, if gargantuan, list. I assumed that the NSA cleaned up the list—date goes here, call duration there—and converted it to the agency’s preferred “atomic sigint data format.” Otherwise I thought of the records as inert. During a conversation at the Aspen Security Forum that July, six weeks after Snowden’s first disclosure and three months after the Boston Marathon bombing, Admiral Dennis Blair, the former director of national intelligence, assured me that the records were “stored,” untouched, until the next Boston bomber came along.

Even by that account, the scale of collection brought to mind an evocative phrase from legal scholar Paul Ohm. Any information in sufficient volume, he wrote, amounted to a “database of ruin.” It held personal secrets that “if revealed, would cause more than embarrassment or shame; it would lead to serious, concrete, devastating harm.” Nearly anyone in the developed world, he wrote, “can be linked to at least one fact in a computer database that an adversary could use for blackmail, discrimination, harassment, or financial or identity theft.” Revelations of “past conduct, health, or family shame,” for example, could cost a person their marriage, career, legal residence, or physical safety.

Mere creation of such a database, especially in secret, profoundly changed the balance of power between government and governed. This was the Dark Mirror embodied, one side of the glass transparent and the other blacked out. If the power implications do not seem convincing, try inverting the relationship in your mind: What if a small group of citizens had secret access to the telephone logs and social networks of government officials? How might that privileged knowledge affect their power to shape events? How might their interactions change if they possessed the means to humiliate and destroy the careers of the persons in power? Capability matters, always, regardless of whether it is used. An unfired gun is no less lethal before it is drawn. And in fact, in history, capabilities do not go unused in the long term. Chekhov’s famous admonition to playwrights is apt not only in drama, but in the lived experience of humankind. The gun on display in the first act—nuclear warheads, weaponized disease, Orwellian cameras tracking faces on every street—must be fired in the last. The latent power of new inventions, no matter how repellent at first, does not lie forever dormant in government armories.

Memorial Day Sale. Memorial Day Sale.
Don't miss the future. Get 1 year for ~~\$10~~ \$5.

Subscribe Now

These could be cast as abstract concerns, but I thought them quite real. By September of that year, it dawned on me that there were also concrete questions that I had not sufficiently

explored. Where in the innards of the NSA did the phone records live? What happened to them there? The Snowden archive did not answer those questions directly, but there were clues.

I stumbled across the first clue later that month. I had become interested in the NSA's internal conversation about "bulk collection," the acquisition of high-volume data sets in their entirety. Phone records were one of several kinds. The agency had grown more and more adept, brilliantly creative in fact, at finding and swallowing other people's information whole. Lately the NSA had begun to see that it consumed too much to digest. Midlevel managers and engineers sounded notes of alarm in briefings prepared for their chains of command. The cover page of one presentation asked "Is It the End of the SIGINT World as We Have Come to Know It?" The authors tried for a jaunty tone but had no sure answer. The surveillance infrastructure was laboring under serious strain.

One name caught my eye on a chart that listed systems at highest risk: Mainway. I knew that one. NSA engineers had built Mainway in urgent haste after September 11, 2001. Vice President Dick Cheney's office had drafted orders, signed by President George W. Bush, to do something the NSA had never done before. The assignment, forbidden by statute, was to track telephone calls made and received by Americans on American soil. The resulting operation was the lawless precursor of the broader one that I was looking at now.

MAINWAY CAME TO life alongside Stellarwind, the domestic surveillance program created by Cheney in the first frantic weeks after al Qaeda flew passenger airplanes into the Pentagon and World Trade Center. Stellarwind defined the operation; Mainway was a tool to carry it out.

At the time, the NSA knew how to do this sort of thing with foreign telephone calls, but it did not have the machinery to do it at home.

When NSA director Mike Hayden received the execution order on October 4, 2001 for "the vice president's special program," NSA engineers assembled a system from bare metal and borrowed code within a matter of days, a stupendous achievement under pressure. They commandeered 50 state-of-the-art computer servers from Dell, which was about to ship

them to another customer, and lashed them into a quick and dirty but powerful cluster. Hayden cleared out space in a specially restricted wing of OPS 2B, an inner sanctum of the gleaming, mirrored headquarters complex at Fort Meade, Maryland. When the cluster expanded, incorporating some 200 machines, Mainway spilled into an annex in the Tordella Supercomputer Facility nearby. Trusted lieutenants began calling in a small group of analysts, programmers, and mathematicians on October 6 and 7.

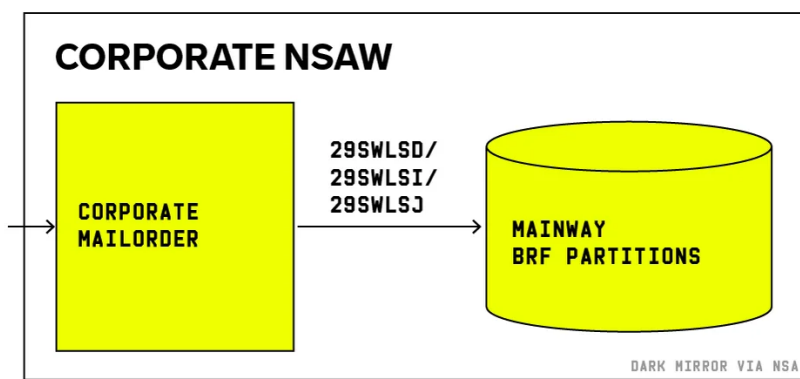
On Columbus Day, October 8, Hayden briefed them on their new jobs in a specially compartmented new operation. That day he called it Starburst. The Stellarwind cryptonym replaced it soon afterward. During the same holiday weekend, Hayden dispatched personnel from Special Source Operations to negotiate the secret purchase of telephone data in bulk from companies including AT&T and Verizon. The price would surpass \$102 million in the coming five years.

It was impossible to hide the hubbub from other NSA personnel, who saw new equipment arriving under armed escort at a furious pace, but even among top clearance holders hardly anyone knew what was going on. Stellarwind was designated as ECI, “exceptionally controlled information,” the most closely held classification of all. From his West Wing office, Cheney ordered that Stellarwind be concealed from the judges of the FISA Court and from members of the intelligence committees in Congress.

According to my sources and the documents I worked through in the fall of 2013, Mainway soon became the NSA's most important tool for mapping social networks—an anchor of what the agency called Large Access Exploitation. “Large” is not an adjective in casual use at Fort Meade. Mainway was built for operations at stupendous scale. Other systems parsed the contents of intercepted communications: voice, video, email and chat text, attachments, pager messages, and so on. Mainway was queen of metadata, foreign and domestic, designed to find patterns that content did not reveal. Beyond that, Mainway was a prototype for still more ambitious plans. Next-generation systems, their planners wrote, could amplify the power of surveillance by moving “from the more traditional analysis of what is collected to the analysis of what to collect.” Patterns gleaned from call records would identify targets in email or location databases, and vice versa. Metadata was the key to the NSA's plan to “identify, track, store, manipulate and update relationships” across all

forms of intercepted content. An integrated map, presented graphically, would eventually allow the NSA to display nearly anyone's movements and communications on a global scale. In their first mission statement, planners gave the project the unironic name "the Big Awesome Graph." Inevitably it acquired a breezy acronym, "the BAG."

The crucial discovery on this subject turned up at the bottom right corner of a large network diagram prepared in 2012. A little box in that corner, reproduced below, finally answered my question about where the NSA stashed the telephone records that Blair and I talked about. *The records lived in Mainway.* The implications were startling.



The diagram as a whole, too large to display in full, traced a "metadata flow sourced from billing records" at AT&T as they wended through a maze of intermediate stops along the way to Fort Meade. Mailorder, the next to last stop, was an electronic traffic cop, a file sorting and forwarding system. The ultimate destination was Mainway. The "BRF Partitions" in the network diagram were named for Business Records FISA orders, among them a dozen signed in 2009 that poured the logs of hundreds of billions of phone calls into Mainway.

To a first-time reader of network maps, Mainway's cylindrical icon might suggest a storage tank. It is not. The cylinder is a standard symbol for a database, an analytic service that runs on the hardware. Mainway was not a container for data at rest. The NSA has names for those. They are called data marts and data warehouses. If the agency merely stored the US telephone records, it would have left them in a system called Fascia II, the "call detail record warehouse" that feeds Mainway. Mainway's mission, laid out in its first fiscal year, was to "enable NSA ... to dominate the global communications infrastructure, and the

targets that currently operate anonymously within it.” The way the system accomplished that task had huge implications for American privacy.

For reasons that will become apparent soon, I want to reproduce the entry for Mainway in the *SSO Dictionary*, a classified NSA reference document:

(TS//SI//REL) Mainway, or the Mainway Precomputed Contact Chaining Service, is an analytic tool for contact chaining. It's helping analysts do target discovery by enabling them to quickly and easily navigate the increasing volumes of global communications metadata. Mainway attacks the volume problem of analyzing the global communications network.

There were three noteworthy terms in that short passage: volume problem, contact chaining, and precomputed. The last two, in combination, turned my understanding of the call records program upside down. Before we get to them, a note on the *volume problem*.

The NSA has many volume problems, actually. Too much information moving too fast across global networks. Too much to ingest, too much to store, too much to retrieve through available pipes from distant collection points. Too much noise drowning too little signal. In the passage I just quoted, however, the volume problem referred to something else—something deeper inside the guts of the surveillance machine. It was the strain of an unbounded appetite on the NSA's digestive tract. Collection systems were closing their jaws on more data than they could chew. Processing, not storage, was the problem.

For a long time, intelligence officials explained away the call records database by quoting a remark from President Bush. “It seems like to me that if somebody is talking to al Qaeda, we want to know why,” he had said.

In fact, that was not at all the way the NSA used the call records. The program was designed to find out whether, not why, US callers had some tie to a terrorist conspiracy—and to do so, it searched us all. Working through the FBI, the NSA assembled a five-year inventory of phone calls from every account it could touch. Trillions of calls. Nothing like that was needed to find the numbers on a bad guy's telephone bill.

This is where *contact chaining* came in. The phrase is used to describe a sophisticated form of analysis that looks for hidden, indirect relationships in very large data sets. Contact chaining began with a target telephone number, such as Boston bomber Dzhokhar Tsarnaev's, and progressively widened the lens to ask whom Tsarnaev's contacts were talking to, and whom those people were talking to, and so on.

Software tools mapped the call records as "nodes" and "edges" on a grid so large that the human mind, unaided, could not encompass it. Nodes were dots on the map, each representing a telephone number. Edges were lines drawn between the nodes, each representing a call. A related tool called MapReduce condensed the trillions of data points into summary form that a human analyst could grasp.

Network theory called this map a social graph. It modeled the relationships and groups that defined each person's interaction with the world. The size of the graph grew exponentially as contact chaining progressed. The whole point of chaining was to push outward from a target's immediate contacts to the contacts of contacts, then contacts of contacts of contacts. Each step in that process was called a hop.

Double a penny once a day and you reach \$1 million in less than a month. That is what exponential growth looks like with a base of two. As contact chaining steps through its hops, the social graph grows much faster. If the average person calls or is called by 10 other people a year, then each hop produces a tenfold increase in the population of the NSA's contact map. Most of us talk on the phone with a lot more than 10 others. Whatever that number, dozens or hundreds, you multiply it by itself to measure the growth at each hop.

Former NSA deputy director John C. Inglis testified to Congress in 2013 that NSA analysts typically "go out two or three hops" when they chain through the call database. For context, data scientists estimated decades ago that it would take no more than six hops to trace a path between any two people on Earth. Their finding made its way into popular culture in *Six Degrees of Separation*, the play by John Guare (which subsequently was adapted into a film). Three students at Albright College refashioned the film as a parlor game, "Six Degrees of Kevin Bacon." The game then inspired a website, The Oracle of Bacon, that calculates the shortest path from the *Footloose* star to any of his Hollywood peers. The site

is still live as I write this, and it makes for an entertaining guide on hops and where they can take you.

Bacon shared screen credits with a long list of actors. Those were his direct links, one hop from Bacon himself. Actors who never worked alongside him, but appeared in a film with someone who had, were two hops away from Bacon. Scarlett Johansson never worked with Bacon, but each of them had starred alongside Mickey Rourke: Bacon in *Diner*, Johansson in *Iron Man 2*. Two hops, through Rourke, connected them. If you kept on playing you discovered that Bacon was seldom more than two hops away from any actor, however removed in time and movie style. In a single-industry town like Hollywood, links like these might make intuitive sense. More surprising, if you did not spend much time around logarithms, was the distance traveled by one or two hops through the vastly larger NSA data set. Academic research suggested that an average of three hops—the same number Inglis mentioned—could trace a path between any two Americans.

Contact chaining on a scale as grand as a whole nation's phone records was a prodigious computational task, even for Mainway. It called for mapping dots and clusters of calls as dense as a star field, each linked to others by webs of intricate lines. Mainway's analytic engine traced hidden paths across the map, looking for relationships that human analysts could not detect. Mainway had to produce that map on demand, under pressure of time, whenever its operators asked for a new contact chain. No one could predict the name or telephone number of the next Tsarnaev. From a data scientist's point of view, the logical remedy was clear. If anyone could become an intelligence target, Mainway should try to get a head start on everyone.

"You have to establish all those relationships, tag them, so that when you do launch the query you can quickly get them," Rick Ledgett, the former NSA deputy director, told me years later. "Otherwise you're taking like a month to scan through a gazillion-line phone bill." And that, right there, was where *precomputation* came in. Mainway chained through its database continuously—"operating on a 7x24 basis," according to the classified project summary. You might compare its work, on the most basic level, to indexing a book—albeit a book with hundreds of millions of topics (phone numbers) and trillions of entries (phone calls). One flaw in this comparison is that it sounds like a job that will be finished

eventually. Mainway's job never ended. It was trying to index a book in progress, forever incomplete. The FBI brought the NSA more than a billion new records a day from the telephone companies. Mainway had to purge another billion a day to comply with the FISA Court's five-year limit on retention. Every change cascaded through the social graph, redrawing the map and obliging Mainway to update ceaselessly.

Mainway's purpose, in other words, was neither storage nor preparation of a simple list. Constant, complex, and demanding operations fed another database called the Graph-in-Memory.

When the Boston marathon bombs exploded in April 2013, the Graph-in-Memory was ready. Absent unlucky data gaps, it already held a summary map of the contacts revealed by the Tsarnaev brothers' calls. The underlying details—dates, times, durations, busy signals, missed calls, and “call waiting events”—were easily retrieved on demand. Mainway had already processed them. With the first hop precomputed, the Graph-in-Memory could make much quicker work of the second and the third.

To keep a Tsarnaev graph at the ready, Mainway also had to precompute a graph for everyone else. And if Mainway had your phone records, it also held a rough and ready diagram of your business and personal life.

As I parsed the documents and interviewed sources in the fall of 2013, the implications finally sank in. The NSA had built a live, ever-updating social graph of the US.

Our phone records were not in cold storage. They did not sit untouched. They were arranged in a one-hop contact chain of each to all. All kinds of secrets—social, medical, political, professional—were precomputed, 24/7. Ledgett told me he saw no cause for concern because “the links are unassembled until you launch a query.” I saw a database that was preconfigured to map anyone's life at the touch of a button.

I AM WELL aware that a person could take this line of thinking too far. Maybe I have. The US is not East Germany. As I pieced this picture together, I had no reason to believe the NSA made corrupt use of its real-time map of American life. The rules imposed some

restrictions on use of US telephone records, even after Bush's attorney general, Michael Mukasey, blew a hole in them. Only 22 top officials, according to the Privacy and Civil Liberties Oversight Board, had authority to order a contact chain to be built from data in Mainway's FISA partitions.

But history has not been kind to the belief that government conduct always follows rules or that the rules will never change in dangerous ways. Rules can be bypassed or rewritten—with or without notice, with or without malignant intent, by a few degrees at a time or more than a few. Government might decide one day to look in Mainway or a comparable system for evidence of a violent crime, or any crime, or any suspicion. Governments have slid down that slope before. Within living memory, Richard Nixon had ordered wiretaps of his political enemies. The FBI, judging Martin Luther King Jr. a “dangerous and effective Negro,” used secret surveillance to record his sexual liaisons. A top lieutenant of J. Edgar Hoover invited King to kill himself or face exposure.

Meaningful abuse of surveillance had come much more recently. The FBI illegally planted hundreds of GPS tracking devices without warrants. New York police spied systemically on mosques. Governments at all levels used the power of the state most heavy-handedly, sometimes illegally, to monitor communities disadvantaged by poverty, race, religion, ethnicity, and immigration status. As a presidential candidate, Donald Trump threatened explicitly to put his opposing candidate in jail. Once in office, he asserted the absolute right to control any government agency. He placed intense pressure on the Justice Department, publicly and privately, to launch criminal investigations of his critics.

The Graph-in-Memory knew nothing of such things. It had no awareness of law or norms or the nature of abuse. It computed the chains and made diagrams of our hidden relationships on a vast, ever-updating map. It obeyed its instructions, embedded in code, whatever those instructions said or might ever say.

*Adapted from **Dark Mirror: Edward Snowden and the American Surveillance State** by **Barton Gellman**. Copyright © 2020 by **Barton Gellman**. Published by*