What can we learn from Bluffdale and Oak Ridge?

Stuart Wray

20 March 2012
www.stuartwray.net

This essay is an attempt to understand the implications of some claims made by James Bamford in an article for Wired magazine: “The NSA is building the country's biggest spy center” (http://www.wired.com/threatlevel/2012/03/ff_nsadatacenter/all/1).

James Bamford is an “NSA-watcher” with a long track-record. He appears to have maintained, over several decades, a number of contacts who are current or former employees of the NSA. (He has published several books on the NSA, starting with “The Puzzle Palace” in 1982.) The most interesting quotes in Bamford's recent article are from anonymous sources, and we have to wonder to what extent Bamford is merely retelling stories that the NSA in fact wants the world to hear. However, some of the story comes from on-the-record interviews with ex-NSA cryptographer William Binney and tallies with other reporting, for example the article “The Secret Sharer” by Jane Mayer (New Yorker, 23 May 2011).

In outline, Bamford's recent article tells the story of how the NSA has, over the past decade, built a vast technical apparatus for spying on US citizens. The NSA's illegal wiretapping programme, at first denied and then retroactively legalised in 2008, has continued to expand. William Binney states that at the start, around 2002, the NSA recorded around 320 million domestic phone calls a day, which at the time was about 75% of the total volume of their worldwide intercepts. Binney claims that this interception continues at perhaps 10 or 20 telecom central switching nodes, where deep-packet-inspectors now also copy all internet traffic of interest. If you are one of the million-or-so people on the NSA watch-list or you say something suspicious online, your communication is automatically recorded and filed away.

Bamford chooses not to explore what this domestic surveillance is meant to achieve, though he does point out that it has failed to prevent incidents such as the feeble Times Square “bombing” of 2010. Instead, the article highlights the implications of developments at two NSA sites, one in Bluffdale, Utah and the other at Oak Ridge, Tennessee. Although the technical details are at times rather garbled, some interesting points shine through and beg for further explanation. Let's see what we can work out.

Exhibit “A” is the data-centre at Bluffdale, Utah. Construction started at this brand-new site in January 2011 and the data-centre is expected to be operational in September 2013. With 100,000 square-feet of server-rooms and 900,000 square feet of office space the facility will cost around $2 billion. The site has a 65 megawatt electric power supply plus on-site backup generators with fuel for three days. Bamford puts this facility at the centre of the NSA's new “cloud” architecture, the warehouse for all those recorded communications.

The numbers at first sound impressive, but when we compare them with commercial data centres they are not quite so unusual. For example, Facebook recently spent $210 million on its Prineville, Oregon data-centre which has an area of 300,000 square-feet and a 28 megawatt power supply. No, the interesting thing is not the size, but the rationale for building the Bluffdale data centre at all.

Bamford quotes a senior intelligence manager involved with the planning at Bluffdale: “Why were we building this NSA facility? And boy, they rolled out all the old guys — the crypto guys.” These “crypto guys” then apparently told Director of National Intelligence Dennis Blair, “You've got to build this thing because we just don't have the capability of doing the code breaking.” What does that mean exactly?

Bamford gives a somewhat garbled explanation. The crypto guys were admitting that they couldn't currently break crypto-systems like AES, but according to Bamford, “the more messages from a given target, the more likely it is for the computers to detect telltale patterns, and Bluffdale will be able to hold a great many messages.” But that sounds a bit feeble. All that effort just for traffic analysis? And why do you have to store the messages themselves if you really can't read them, and never expect to read them? They are just so much random noise.

But Bamford claims that Bluffdale is not just about traffic analysis. He claims that breaking AES is one of the key reasons for building the Bluffdale data-centre. He says: “That kind of cryptanalysis requires two major ingredients: super-fast computers to conduct brute-force attacks on encrypted messages and a massive number of those messages for the computers to analyze.” Which sounds to me rather odd and implausible, but does hint at a real purpose for Bluffdale, in line with the NSA crypto guys' stated reason: that although they can't read the messages today, they anticipate somehow being able to read them in the future. How?

This brings us on to exhibit “B”: the Multiprogram Research Facility (“Building 5300”) at the Oak Ridge National Laboratory in Tennessee. This 200,000 square-foot, five-story lab was completed in 2006 and houses around 300 NSA scientists and cryptanalysts. Bamford describes this lab as being the home of a top-secret programme running in parallel with the very public efforts, at the Department of Energy's nearby Leadership Computing Facility, to build the world's fastest super-computer. The successful public programme started in 2004 and built one of the fastest super-computers in the world (currently surpassed by only two machines, one in China, the other in Japan). Work is currently in progress to build a somewhat faster machine and reclaim the record.

As Bamford describes the NSA Oak Ridge lab, they were attempting to build an even faster super-computer than the DoE machine, with the intention of using it for brute-force cryptanalysis. Which is on the face of it ridiculous. As Bamford himself points out, to crack even AES-128 by brute force would take around 10^{36} attempts. Even if we were to equate AES attempts with floating-point operations, a machine comparable to the public DoE super-computer would only manage around 10^{15} attempts per second. It would be completely and utterly impractical to break AES-128 this way.

And yet, according to another former senior intelligence official, “They made a big breakthrough.” What sort of breakthrough? The former official said that it was “enormous” and disclosed to “only the chairman and vice chairman and the two staff directors of each intelligence committee.” Why so coy? “They were thinking that this computing breakthrough was going to give them the ability to crack the current public encryption.” An odd turn of phrase, and one which both Bamford and the former official take to mean the ability to crack AES-128. The official goes on to say, “Remember, a lot of foreign government stuff we've never been able to break is 128 or less. Break all that and you'll find out a lot more of what you didn't know — stuff we've already stored — so there's an enormous amount of information still in there.”

Bamford reports that whatever the nature of the “enormous” breakthrough, preparations are in train to exploit it. A 260,000 square-foot single-story extension — the Multiprogram Computational Data Center --- is proposed for the NSA Oak Ridge lab. This would have an electric power supply of 200 megawatts and be due for completion in 2018. But so what? What could it do? And what, if anything, can we deduce from all of the above claims?

First of all, we can discount any possibility of breaking AES by brute-force using any kind of conventional super-computer. Even one a million times faster than the world record holder would be impossibly feeble for that task. But what if the NSA was building a non-conventional super-computer? Say a quantum computer? Would it be possible to break AES that way? Research into quantum computers has still not produced significant results in the commercial or academic world, but the NSA has a track record of being, in some fields, a decade-or-so ahead of outside developments. Perhaps that's the case here. Oak Ridge would certainly be a good place to unobtrusively assemble a team of physicists and engineers to work on the problem.

But is a quantum cryptanalytic algorithm for AES possible, even if you could build a suitable quantum computer to run it on? Maybe, but maybe that's the wrong question to ask. Maybe by following Bamford and the former official we are barking up the wrong tree. What if the target was not AES after all?

There's a clue in the odd phrase used by the former official: “They were thinking that this computing breakthrough was going to give them the ability to crack the current public encryption.” Now a senior intelligence official is unlikely to be an expert in cryptography, and may have only a hazy understanding of the difference between symmetric and asymmetric algorithms. What if the former official was misquoting something he had been told? What if the phrase was not “the current public encryption” (which no “crypto guy” would ever say), but rather “the current public-key encryption”?

Now, that would indeed be an “enormous” breakthrough. And we know that quantum algorithms exist to crack at least some public-key encryption systems, if only we had quantum computers on which to actually run them. If we look again at the Bamford article and engage in textual analysis worthy of a bible scholar we can see another faint hint of this idea. Bamford says “The NSA believes it's on the verge of breaking a key encryption algorithm,” by which he understands that they are about to break a “key” encryption algorithm i.e. AES. But perhaps this is another phrase repeated without understanding? What if instead they are about to break a “key-encryption” algorithm e.g. RSA?

RSA is used for key-encryption in TLS, the protocol used for secure websites. Breaking

RSA would therefore enable the NSA to read and to *fake* all internet traffic transiting the USA. Although RSA is not the only public-key algorithm, it is unclear if any public-key algorithm could withstand attack from a quantum computer. Probably if one of them can be broken, they all can be.

Of course, by now we have run far ahead of the real evidence, and I'm probably drawing much more on my own biases than on the facts in front of us. This is just speculation. On the other hand, it is speculation which makes a great deal more sense than the conclusions drawn in Bamford's article itself.