Re: A few more PQC comments

Alperin-Sheriff, Jacob (Fed)

Mon 9/26/2016 12:57 PM

To: Moody, Dustin (Fed) <dustin.moody@nist.gov>; Chen, Lily (Fed) <lily.chen@nist.gov>; Perlner, Ray (Fed) <ray.perlner@nist.gov>; Peralta, Rene (Fed) <rene.peralta@nist.gov>; Liu, Yi-Kai (Fed) <yi-kai.liu@nist.gov>; Jordan, Stephen P (Fed) <stephen.jordan@nist.gov>; Miller, Carl A. (Fed) <carl.miller@nist.gov>; Daniel C Smith (daniel-c.smith@louisville.edu) (daniel-c.smith@louisville.edu) <daniel-c.smith@louisville.edu>;

I know the 2^{64} question was already asked by at least one person (I think Vadim).

But I don't think the "very long term" thing is relevant, unless there are any concrete uses for signatures that don't involve some sort of certificate with an expiration date.

Otherwise (if all concrete uses do involve a certificate), it should be much easier to upper bound the maximum number of possible chosen messages one could realistically expect by answering:

1. What is the NIST standard on lifecycle length for a certificate? Is it a year? Six months? Two years?
2. What are the maximum number of signatures any given entity (that is to say, holder(s) of a specific signing key) issues per second in today's world?

Then note that ~ 2^{25} seconds/year, and add some padding of 1000 or so, and end up with a bound that should hold.

From: "Moody, Dustin (Fed)" <dustin.moody@nist.gov>
Date: Monday, September 26, 2016 at 12:39 PM
To: "Chen, Lily (Fed)" <lily.chen@nist.gov>, "Perlner, Ray (Fed)" <ray.perlner@nist.gov>, "Peralta, Rene (Fed)" <rene.peralta@nist.gov>, "Liu, Yi-Kai (Fed)" <yi-kai.liu@nist.gov>, "Jordan, Stephen P (Fed)" <stephen.jordan@nist.gov>, "Alperin-Sheriff, Jacob M. (Fed)" <jacob.alperin-sheriff@nist.gov>, "Miller, Carl A. (Fed)" <carl.miller@nist.gov>, "Daniel C Smith (daniel-c.smith@louisville.edu) (daniel-c.smith@louisville.edu)" <daniel-c.smith@louisville.edu>
Subject: A few more PQC comments

While at the ETSI workshop last week, we received the following comment from Peter Campbell, of UK's CESG:

I think my main comment about the draft NIST call is that they need to be careful with their discussion following the target security strengths listed in section 4.A.4. It's important that they make it absolutely clear that this discussion is about parallelisation of *quantum* attacks and is not applicable to classical attacks. This is particularly true of the clause "... NIST recognizes that extremely serial or extremely parallel attacks (e.g., those that have a time depth or space complexity exceeding 2^{100}) may be of minimal practical importance". The danger is that someone reading this might incorrectly interpret it as a giving a memory bound for classical attacks. Given recent heated arguments around the NTRUPrime security analysis, this should be avoided if at all possible. For example, it might be better to remove the specific figure 2^{100} from the clause above. (On the other hand, if NIST do intend to give a memory bound then it should be applied consistently to *all* large-memory attacks.)

We also heard from the French agency ANSSI, who wondered if 2^{64} chosen messages might not be too low, since we should be having a very long-term perspective. He also questioned "replacing" our SP's, as we'll still be using our current SP's for a while.

Dustin