

From: [Kelsey, John M. \(Fed\)](#)
To: [internal-pqc](#)
Subject: Re: PQC
Date: Monday, June 29, 2020 1:38:56 PM

I think this makes our position ***way*** less clear.

Part of the problem here is that alternates fall into at least three categories, and we're sticking them into the same bin. There are algorithms where:

- a. We are pretty confident in the security, but they're not finalists because of lousy performance.
Frodo, SPHINCS+, probably HQC and GeMSS
- b. We still have questions about the security, but the performance is promising.
BIKE, maybe SIKE, NTRU Prime (sort-of)
- c. We think the design just isn't quite cooked yet
Picnic

The way it looks to me is that (a) are the things we might decide to standardize at the end of the third round. And we might do that ***either*** because we've got concerns about the better-performing options (maybe we want to wait another year or two before nailing down the parameters for the structured lattice schemes), ***or*** because we've decided we want to standardize a paranoid option.

That is, we could just decide, at the end of the third round, that we're standardizing, say, Saber, Classic McEliece, and Frodo + Falcon, Rainbow, and SPHINCS+, with Frodo and SPHINCS+ explicitly chosen as paranoid options for people who want postquantum security but also are concerned that these structured lattice algorithms aren't as well-understood as they should be. There would be nothing shocking about that, and it wouldn't require a shocking new sequence of cryptanalysis results.

My basic claim is that the schemes in (a) are about as solid in security terms as the finalists--if SPHINCS+ gave us signatures twice the size and half the speed of Dilithium, it would probably be a finalist. The reason it's not is because its signatures are 4x the size and like 1/50 as fast as Dilithium. That means that there is no strong reason why we ***couldn't*** standardize SPHINCS+ at the end of the third round, if we decided that was something we wanted to do. By contrast, something like BIKE probably just needs another round to get its security proofs and parameters nailed down, and maybe we'd like to see the field mature around SIKE and the implementations improve before we standardized it. And I think there is no possible world in which we are standardizing Picnic at the end of the third round.

--John

From: "Moody, Dustin (Fed)" <dustin.moody@nist.gov>

Date: Monday, June 29, 2020 at 12:51

To: "Cooper, David A. (Fed)" <david.cooper@nist.gov>, "Regenscheid, Andrew R. (Fed)" <andrew.regenscheid@nist.gov>, internal-pqc <internal-pqc@nist.gov>

Subject: Re: PQC

Good points.

I agree that with this new text, we can probably just simply say we won't standardize an alternate at the end of the third round (because we would make it a finalist first). The alternates we want to keep at the end of the third round would then get a 4th round.

From: David A. Cooper <david.cooper@nist.gov>

Sent: Monday, June 29, 2020 12:41 PM

To: Regenscheid, Andrew R. (Fed) <andrew.regenscheid@nist.gov>; internal-pqc <internal-pqc@nist.gov>

Subject: Re: PQC

I agree entirely. When we were talking about the two track approach, I thought there was to be clear distinction: decisions about finalists would be made in the third round and decisions about alternates would not be made until the fourth round. This would make it clear to those who would be looking at all of the candidates, e.g., groups developing hardware implementations, that it would be okay to just work on the finalists during round three, and that work on the alternates could wait until later.

Our current text isn't so clear. By merely saying that we are "unlikely" to standardize an alternate at the end of round three, that creates confusion. If my goal is to implement all algorithms that might be standardized before the standardization decision is made, can I implement just the seven finalists during the third round or do I need to implement all 15 remaining candidates since any of the alternates "could" be standardized at the end of the third round.

If we aren't going to impose a strict rule of "no selecting alternates at the end of the third round," then I think we should at least say that we won't select an alternate for standardization at the end of the third round unless we make an announcement about it at some point during the third round of evaluation. The amount of time between the announcement and the end of the third round needs to be long enough that people feel they have been given a fair chance to review the algorithm.

David

On 6/29/20 12:12 PM, Regenscheid, Andrew R. (Fed) wrote:

One of the main things you want in these processes is predictability. It's not enough to say we might do something- people have to expect it. We learned that one in SHA-3.

I've been somewhat concerned that we're sending mixed messages the

alternates. In general, we're saying we don't plan to standardize any of them right away (until after a 4th round) except that we want to carve out some leeway so that we could if we really wanted to. The main case for that would probably be SPHINCS+, which we allude to in the report. Perhaps you could imagine Frodo being another case for that.

I don't think we want there to be any surprise if we get to the end of round 3 and we decide we're going to standardize SPHINCS+, Frodo, or one of the other four examples John cited. I think we'd want to signal that clearly, and somewhat formally, in advance. That's where the idea of "elevating" an alternate to a finalist came in.

-Andy

From: Kelsey, John M. (Fed) <john.kelsey@nist.gov>

Sent: Monday, June 29, 2020 12:03 PM

To: Moody, Dustin (Fed) <dustin.moody@nist.gov>; internal-pqc <internal-pqc@nist.gov>

Subject: Re: PQC

It seems weird to phrase it that way. I think the point of Andy's sentence there is that we may decide to standardize one of the alternates at the end of the third round, right? But I don't think that would change the fact that we had already named some things as finalists and others as alternates. I mean, if all the structured lattice KEMs get broken or dented and we decide to standardize Frodo at the end of the third round, it wouldn't mean that Kyber and Saber and NTRU got demoted to being alternates—it would mean that we just decided to standardize one of our alternates instead of one of our finalists.

That's a plausible outcome, as far as I can tell, for five or six alternates: SPHINCS+, GeMSS, HQC, SIKE, Frodo, and maybe BIKE. For example, imagine that over the next 18 months, we get a bunch of results that make us uneasy about the parameter selection for structured lattice schemes, and at the same time, there's a very clear upper bound on error rate for BIKE that lets them get CCA security. It seems very plausible to me that we standardize Frodo and BIKE as KEMs in that world. Then maybe we standardize a structured lattice KEM in another couple years when we feel like we know how the parameters should be selected.

But I don't think that would change the fact that Frodo and BIKE were both alternates instead of finalists. I can't imagine that we'd want to, say, announce that we'd demoted Saber to an alternate and Frodo to a finalist, six months from now.

--John

From: "Moody, Dustin (Fed)" <dustin.moody@nist.gov>

Date: Monday, June 29, 2020 at 11:49

To: internal-pqc <internal-pqc@nist.gov>

Subject: PQC

Everyone,

I don't have any plans for a meeting tomorrow. Let me know if you think we need one. The reviews for the report are still on going, and I'll make changes to suggestions we get back. Here's one Andy recommended we add in:

"It is possible that new analysis could result in an alternate candidate being elevated to being a finalist, in the case that NIST's confidence in the security of any of the finalists is greatly reduced."

Seems reasonable to me. It doesn't tie our hands and keeps our options open in case of an unexpected advance that breaks a finalist.

Dustin