I don't think that we need any language in the report stating that we may elevate an alternate to the status of finalist. I think that any circumstance that would make us go that route will likely be a big enough event that it would be no surprise to anyone that we may need to pause and regroup--- re-organize the project. I don't think that we should get hung up on low probability events that muddy the description of the project.

Cheers,
Daniel

On Tue, Jun 30, 2020 at 8:08 AM Dang, Quynh H. (Fed) <quynh.dang@nist.gov> wrote:

Hi all,

Frodo is standardized at the end of the 3rd round if:

1) all structured lattice KEMs (NTRU, Saber, Kyber and NTRUprime) and all "cyclic" code-based KEMs (Bike and HQC) are attacked in ways that make us not comfortable to standardize any of them at the end of the third round,

2) The industry pushes us to pick a KEM, and

3) SIKE's performance does not improve much by the end of the 3rd round and/or we don't have a solid confidence in its security at that time yet.

If nothing changes the current picture by the end of the 3rd round but we decide to standardize Frodo with the reason that we want a safe option. Standardizing both Frodo and a structured lattice KEM at the same time at the end of the 3rd round would not sound too good to me because standardizing Frodo implies that we don't have confidence in the structured lattice KEM but we still standardize it.

So, if all of the 3 conditions above are not met by the end of the 3rd round, I think we should wait on Frodo.

SPHINCS+ is standardized at the end of the 3rd round if:

1) Falcon and Dilithium are attacked and we don't have a strong confidence in their security.

2) The industry pushes us for a pq signature algorithm for widely used protocols.

PQ signatures are not urgent as KEMs in protocols because the signatures for authentication in the protocols do not need to be secure in the future (PQ world). The community might push us because it would take years for a new crypto to be widely deployed.

The case would become stronger for SPHINCS+ if Rainbow and GeMMSS's security are threatened by some new attack because SPHINCS+ and PICNIC would be the only ones standing strong.

But, regardless, the conditions 1 and 2 above are the most important factors because SPHINCS+ would work in widely used protocols even though it would have significant performance impacts on them generally and SPHINCS+ would have better performance than Rainbow and GeMMSS in those protocols generally. (Sure where keys are not needed to be sent over the wire regularly, Rainbow and GeMMSS may work much better than SPHINCS+).

Quynh.

From: Perlner, Ray A. (Fed) <ray.perlner@nist.gov>
Sent: Monday, June 29, 2020 7:08 PM
To: Alagic, Gorjan (Assoc) <gorjan.alagic@nist.gov>; Peralta, Rene C. (Fed) <rene.peralta@nist.gov>; Moody, Dustin (Fed) <dustin.moody@nist.gov>; Chen, Lily (Fed) <lily.chen@nist.gov>; Kelsey, John M. (Fed) <john.kelsey@nist.gov>; internal-pqc <internal-pqc@nist.gov>
Subject: RE: PQC

I second the thought that anything not listed as a finalist should only be standardized prior to a 4th round in the case of emergency (i.e. if a lot of the finalists are getting broken, or at least severely undermined.) If we think we might want to standardize Frodo or SPHINCS after round 3 in a non-emergency situation, I'd rather just list them as finalists.

Ray

From: Alagic, Gorjan (Assoc) <gorjan.alagic@nist.gov>
Sent: Monday, June 29, 2020 7:01 PM
To: Peralta, Rene C. (Fed) <rene.peralta@nist.gov>; Moody, Dustin (Fed) <dustin.moody@nist.gov>; Chen, Lily (Fed) <lily.chen@nist.gov>; Kelsey, John M. (Fed) <john.kelsey@nist.gov>; internal-pqc <internal-pqc@nist.gov>
Subject: Re: PQC

I guess count me in as one of the confused then. I was under the impression that, while it is true that each alternate is in that pile for a different reason, one thing they all share is that it

is *very unlikely* that they will be standardized in the third round *unless* the facts as we understand them now change significantly (e.g., a finalist gets seriously attacked, or we realize we asap need things with a very different security/performance profile than we thought.)

But if I'm understanding John's e-mail right (please correct me, John), one "could" course of action that doesn't require facts to change significantly is "at the end of the third round, that we're standardizing, say, Saber, Classic McEliece, and Frodo + Falcon, Rainbow, and SPHINCS+." (I think I initially missed the "third round" part of that.) That involves standardizing two alternates. If that's really a reasonably plausible scenario at the end of the third round, then I also don't fully understand the distinction between finalist and alternate.

To be clear: to whatever extent there is already a consensus on all this, I am not arguing that this consensus should change. I'm just saying that I'm getting confused about what that consensus actually is.

David: I of course agree that our actions should not shock anyone. I meant that the scenario "Frodo > Saber/Kyber because structured lattices are broken" as being shocking, not that our reasonable reactions to such a scenario would be shocking. I really wish I hadn't used that word now though.

Gorjan

From: Peralta, Rene C. (Fed) <rene.peralta@nist.gov>
Sent: Monday, June 29, 2020 5:03 PM
To: Moody, Dustin (Fed) <dustin.moody@nist.gov>; Chen, Lily (Fed) <lily.chen@nist.gov>; Alagic, Gorjan (Assoc) <gorjan.alagic@nist.gov>; Kelsey, John M. (Fed) <john.kelsey@nist.gov>; internal-pqc <internal-pqc@nist.gov>
Subject: Re: PQC

OK, my two cents as well.

We get confused every time somebody proposes a short description of what is an alternate, or a single plan for alternates. I don't see that we need to do that. The text that we do have about each alternate does explain why that proposal is not a finalist and why we would like to keep that proposal around.

Rene.

Like David said, if our group isn't clear on this, then it is likely other people will be confused. We need to agree on what our position is, and then make sure we clearly communicate that position.

Andy suggested we be a little more up front with the possible changes that we could conceivably do during the third round. These are unlikely events, but it will go over better with the crypto community if we have explained what we're thinking in the report. I think this idea makes sense.

The text suggested was:

"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."

We currently have in the report that it is unlikely we would standardize an alternate at the end of the third round. We already agreed on that. But we kept our options open - we left in the possibility of standardizing an alternate.

Here are a few questions to clarify:

- Should we be more clear and explain the reasons this could possibly happen? In the write-up for SPHINCS+, we have text saying if one of the signature finalists were to be attacked, then it could be ready. Not quite as strongly, we say something similar for Frodo in regards to KEMs. Is there any other circumstance (besides losing confidence in a finalist) so that we would consider an alternate for standardization at the end of the third round? The only other explanation I can think of right now is that we would want to include a very conservative option for some reason (other than a finalist being broken). In our previous discussions, we'd said that if that was our reason we could likely wait until after we get our high priority KEM and signatures standardized.
- When we think about the situation we would be in if we were considering to standardize an alternate, how would we let the community know? I think Andy's suggestion of explicitly telling the community that we're changing the alternate to a finalist is fine. It makes it clear.

    - In this case, the suggested text explains the circumstances in which we would consider doing this. If we have other possible circumstances, we could list them. Or

alternatively, we could not list the circumstances. I think it's better to provide an explanation.

- ○ We can keep the report text which says it is "not likely" we would consider standardizing an alternate at the end of the third round, or we could change that to say we won't standardize an alternate at the end of the third round (because anything we'd be considering we'd bump up to be a finalist).   I think being more definite is easier to understand.

Thoughts?

Dustin

From: Chen, Lily (Fed) <lily.chen@nist.gov>
Sent: Monday, June 29, 2020 3:55 PM
To: Alagic, Gorjan (Assoc) <gorjan.alagic@nist.gov>; Kelsey, John M. (Fed) <john.kelsey@nist.gov>; internal-pqc <internal-pqc@nist.gov>
Subject: Re: PQC

I agree that the reasoning behind why a scheme is an alternate depends a lot on the scheme. Lily

---

From: "Alagic, Gorjan (Assoc)" <gorjan.alagic@nist.gov>
Date: Monday, June 29, 2020 at 3:49 PM
To: John Kelsey <john.kelsey@nist.gov>, internal-pqc <internal-pqc@nist.gov>
Subject: Re: PQC

My earlier comment using the term "shocking" was specifically about the scenario "we end up standardizing Frodo as a replacement for structured lattices" and I think that's basically true. But this discussion is about a lot more than that, I guess.

I agree that standardizing *both* Saber and Frodo (the latter for the paranoid) would not be shocking, or even very surprising.

I also agree that the reasoning behind why a scheme is an alternate depends a lot on the scheme, and that we shouldn't discuss them in the report in a way that would confuse this.

Gorjan

From: Kelsey, John M. (Fed) <john.kelsey@nist.gov>
Sent: Monday, June 29, 2020 1:38 PM
To: internal-pqc <internal-pqc@nist.gov>
Subject: Re: PQC

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

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