From: Liu, Yi-Kai (Fed)
To: Moody, Dustin (Fed); Perlner, Ray A. (Fed)
Subject: Re: FW: Proposed edits to security strengths section
Date: Tuesday, October 18, 2016 12:28:11 PM

Hi Dustin and Ray,

Sorry for not replying sooner, I need a bit more time to think about this. I'll be at NIST this afternoon, if you'd like to chat...

Ray, I agree with your points: (1) we don't yet know what we NEED regarding quantum security, and (2) we want to give credit for having more quantum security than AES happens to have.

Really, it sounds like we want to set some very flexible guidelines now, so that people will propose cryptosystems with a variety of quantum/classical security levels, and then over the next 3-4 years, we will decide what kind of tradeoff we want between quantum security, classical security, and performance.

=> So, in order to give us more flexibility, I'd like to propose another strategy:

We list the different levels of classical and quantum security SEPARATELY. That is, we ask for 128, 192 and 256 bits of classical security, and we ask for 64, 80, 96 and 128 bits of quantum security.

To keep things simple, we say that people are allowed to submit 3 or at most 4 parameter settings. And we let people decide for themselves which combinations of classical and quantum security make the most sense for their PQC scheme.

Some people may propose parameter settings that provide security similar to AES, some people may propose parameter settings that provide security similar to SHA3, and some people may do something completely different.

Ray, I think this strategy may accomplish what you are aiming for? (I.e., asking people to hit a sufficient subset of our security levels, without making them feel like they have to hit all 5 of them.) This strategy is basically asking for the same thing, but in a different format.

Also, I like this strategy because it gives us a clearer picture of the classical and quantum security of the PQC schemes, without making them pretend to be AES or SHA3.

What do you think?

Cheers,

--Yi-Kai

From: Moody, Dustin (Fed)
Sent: Monday, October 17, 2016 11:41:10 AM
To: Liu, Yi-Kai (Fed); Perlner, Ray (Fed)
Subject: FW: FW: Proposed edits to security strengths section

Yi-Kai,
I just talked with Ray, (b) (6) He's going to make some edits in regards to the conversion table idea. In talking with him, he's not completely opposed to just three security levels, but he stated some of his concerns in the email below. One of his points is that levels 2 and 4 give algorithm designers who have more quantum security a way to get credit for that. I can understand that. I think some of the confusion from people was their understanding they needed to submit something for every security level. We need to make it much more clear they don't need to do this. I think his other point is in reaction to your summary: "here is a summary: For PQC, we should decide how much classical and quantum security we NEED, rather than aiming for the classical and quantum security levels that AES and SHA3 happen to ACHIEVE". He thinks we don't yet know what we NEED.

Can you re-open the dialogue with Ray so we can try and come up with text we can all agree with? Alternatively, we could express both viewpoints on the pqc-forum, and ask for comments from the mailing list.

Dustin

From: Ray Perlner [(b) (6)
Sent: Tuesday, October 11, 2016 10:52 PM
To: Liu, Yi-Kai (Fed) <yi-kai.liu@nist.gov>
Cc: Daniel Smith (b) (6) ; Perlner, Ray (Fed) <ray.perlner@nist.gov>; Alperin-Sheriff, Jacob (Fed) <jacob.alperin-sheriff@nist.gov>; Peralta, Rene (Fed) <rene.peralta@nist.gov>; Moody, Dustin (Fed) <dustin.moody@nist.gov>; Bassham, Lawrence E (Fed) <lawrence.bassham@nist.gov>; Chen, Lily (Fed) <lily.chen@nist.gov>; Jordan, Stephen P (Fed) <stephen.jordan@nist.gov>; Miller, Carl A. (Fed) <carl.miller@nist.gov>; Smith-Tone, Daniel (Fed) <daniel.smith@nist.gov>
Subject: Re: FW: Proposed edits to security strengths section

I think the fundamental problem with trying to specify how much quantum security we need is that we do not know how much quantum security we need. We simply do not know how rapidly quantum computers will develop relative to classical computers. It is certainly the case that the best quantum attack finding collisions on a good hash function is simply the classical attack (see Bernstein's paper.) It may also be the case that, in practice, the best attack on block ciphers is the classical attack and always will be. If so, then security strengths 1 and 2 are equivalent, as are 3 and 4. But then, it may also become the case that the best attack on block ciphers is a parallel version of Grover's algorithm. I think we need to plan for both possibilities. FWIW, for the vast majority of postquantum cryptosystems, I expect relative quantum/classical attack complexity to be intermediate between a hash function and a block cipher.

In addition, my recollection, from the public comments was that the reverse objection from yours was more common (the comments suggest we may be overengineering for classical security). I therefore see the current approach (taking into account all models of computation which may plausibly be relevant to future cryptanalytic technology) as somewhat of a compromise in the direction of public opinion. Where I see the current security section as pushing against current practice is that it does not consider purely serial models of quantum computation as being terribly plausible. I think it may be reasonable to simply classify anything whose best quantum attack is a variant of Grover's algorithm, based on classical security, but I don't think such a proposal would be terribly popular.

On Tuesday, October 11, 2016, Liu, Yi-Kai (Fed) <yi-kai.liu@nist.gov<mailto:yi-kai.liu@nist.gov>> wrote:
Hi everyone + Ray's vacation address,

I'm sorry for not replying earlier! After thinking about it some more, I think there are two separate issues here:

1. Should security be measured in AES operations, or elementary gates, or some other unit? It seems like there is a disconnect between people doing classical cryptanalysis, who talk about AES or arithmetic operations, and people doing quantum stuff, who talk about elementary gates and circuit depth. I like Jacob and Ray's idea of providing a translation table, that seems like a good compromise.

2. How much quantum security do we need for PQC? I agree that everyone who uses post-quantum public key crypto will also be using AES and SHA3, and therefore needs to offer similar levels of security. But I was getting at a different issue: it is possible that AES and SHA3 are over-engineered with respect to quantum security, and we shouldn't insist that PQC do the same thing.

For instance, it seems to me that SHA3, in order to reach its classical security targets, ends up having so much quantum security that it is almost certain that it will be broken classically before it is broken quantumly. In this sense, SHA3 has more quantum security than it needs. This is unavoidable for a hash function, because of the way that quantum security is related to classical security.

But for a PQC scheme, we're not necessarily forced into that situation; we might be able to tune the parameters to have a better balance of classical and quantum security. So, I think we should decide what is that balance, and ask for that.

TL;DR -- here is a summary: For PQC, we should decide how much classical and quantum security we NEED, rather than aiming for the classical and quantum security levels that AES and SHA3 happen to ACHIEVE, because some of those security levels are unnecessary and superfluous.

Cheers,

--Yi-Kai

From: Daniel Smith (b) (6)
Sent: Friday, October 7, 2016 10:37 PM
To: Perlner, Ray (Fed)
Cc: Alperin-Sheriff, Jacob (Fed); Peralta, Rene (Fed); Liu, Yi-Kai (Fed); Moody, Dustin (Fed); Bassham, Lawrence E (Fed); Chen, Lily (Fed); Jordan, Stephen P (Fed); Miller, Carl A. (Fed); Smith-Tone, Daniel (Fed)
Subject: Re: FW: Proposed edits to security strengths section

17th of September? Where are you going? Proxima Centauri b? What's your method of propulsion?

On Friday, October 7, 2016, Perlner, Ray (Fed) <ray.perlner@nist.gov<javascript:;><mailto:ray.perlner@nist.gov>>> wrote:
One additional note. I'm going to be on vacation throughout next week. I'm not planning on trying to VPN in, since that would require updating antivirus software on a machine I no longer use. If you need my input in the next week, I can be reached at rperlner@gmail.com<javascript:;>. Otherwise, see you all on the 17th of September.

From: Perlner, Ray (Fed)
Sent: Friday, October 07, 2016 5:16 PM
To: Alperin-Sheriff, Jacob (Fed) <jacob.alperin-sheriff@nist.gov<javascript:;>>; Peralta, Rene (Fed) <rene.peralta@nist.gov<javascript:;>>; Liu, Yi-Kai (Fed) <yi-kai.liu@nist.gov<javascript:;>>; Moody, Dustin (Fed) <dustin.moody@nist.gov<javascript:;>>; Bassham, Lawrence E (Fed) <lawrence.bassham@nist.gov<javascript:;>>; Chen, Lily (Fed) <lily.chen@nist.gov<javascript:;>>; 'Daniel Smith-Tone' <daniel-c.smith@louisville.edu<javascript:;>>; Jordan, Stephen P (Fed) <stephen.jordan@nist.gov<javascript:;>>; Miller, Carl A. (Fed) <carl.miller@nist.gov<javascript:;>>; Smith-Tone, Daniel (Fed) <daniel.smith@nist.gov<javascript:;>>
Subject: RE: Proposed edits to security strengths section

Jacob suggested to me that what I currently have written may be ok if it's supplemented with a reference to translate between units of AES operations and basic bit operations.

I was able to find a decent reference for the complexity of Grover's algorithm on AES: https://arxiv.org/pdf/1512.04965v1.pdf<https://urldefense.proofpoint.com/v2/url?u=https-3A__arxiv.org_pdf_1512.04965v1.pdf&d=AwMGaQ&c=SgMrq23dbjbGX6e0ZsSHgEZX6A4IAf1SO3AJ2bNrHlk&r=dpZDgy9SDYABmPp7fYQVcgiPGbaUt1gqjrK5JyA1g-4&m=10MjlSsCtTJj01d_0dfh0fNluiESTCzAIUpSIGhmXn8&s=q50NqNfB0z7QwQ8yXr7G0biAqaRUzmb1k1PbzqrXXx4&e=>

This would suggest that an attack on AES with a maximum depth MAXDEPTH requires:
2^170/MAXDEPTH for AES128 (up to MAXDEPTH=2^81)
2^233/MAXDEPTH for AES192 (up to MAXDEPTH=2^113)
2^298/MAXDEPTH for AES256 (up to MAXDEPTH=2^145)
(all measured based on total number of gates in the Clifford-T gate set.)

I wasn't able to find a comparable classical gate count for SHA2 or SHA3 (and I'm not sure that's the classical security metric Yi-Kai wants) but I'd guess total gate count for either compression function is somewhere around 2^20 classical gates. This would suggest that, for any plausible value of MAXDEPTH, Van Oorschot-Wiener parallel collision search requires a total of

2^148 classical gates for SHA256/SHA3-256
2^212 classical gates for SHA384/SHA3-384
(2^276 classical gates for SHA512/SHA3-512)

As a side note, this would imply that, even ignoring the probability that quantum gates are more expensive than classical gates, security strengths 2 and 4 are less than security strengths 3 and 5 as long as MAXDEPTH<2^85. (This is very close to the limit of what's possible with atomic scale qubits and speed of light propagation time.)
From: Perlner, Ray (Fed)
Sent: Friday, October 07, 2016 11:25 AM
To: Alperin-Sheriff, Jacob (Fed) <jacob.alperin-sheriff@nist.gov<javascript:;>>; Peralta, Rene (Fed) <rene.peralta@nist.gov<javascript:;>>; Liu, Yi-Kai (Fed) <yi-kai.liu@nist.gov<javascript:;>>; Moody, Dustin (Fed) <dustin.moody@nist.gov<javascript:;>>; Bassham, Lawrence E (Fed) <lawrence.bassham@nist.gov<javascript:;>>; Chen, Lily (Fed) <lily.chen@nist.gov<javascript:;>>; Daniel Smith-Tone <daniel-c.smith@louisville.edu<javascript:;>>; Jordan, Stephen P (Fed) <stephen.jordan@nist.gov<javascript:;>>; Miller, Carl A. (Fed) <carl.miller@nist.gov<javascript:;>>; Smith-Tone, Daniel (Fed) <daniel.smith@nist.gov<javascript:;>>
Subject: RE: Proposed edits to security strengths section

True. I agree that we should have a "purely ephemeral only" KEM with IND-CPA security. Typically you still need symmetric crypto to get IND-CPA, though, so even there I think comparing security to that of symmetric primitives is unavoidable.

From: Alperin-Sheriff, Jacob (Fed)
Sent: Friday, October 07, 2016 11:19 AM
To: Perlner, Ray (Fed) <ray.perlner@nist.gov<javascript:;>>; Peralta, Rene (Fed) <rene.peralta@nist.gov<javascript:;>>; Liu, Yi-Kai (Fed) <yi-kai.liu@nist.gov<javascript:;>>; Moody, Dustin (Fed) <dustin.moody@nist.gov<javascript:;>>; Bassham, Lawrence E (Fed) <lawrence.bassham@nist.gov<javascript:;>>; Chen, Lily (Fed) <lily.chen@nist.gov<javascript:;>>; Daniel Smith-Tone <daniel-c.smith@louisville.edu<javascript:;>>; Jordan, Stephen P (Fed) <stephen.jordan@nist.gov<javascript:;>>; Miller, Carl A. (Fed) <carl.miller@nist.gov<javascript:;>>; Smith-Tone, Daniel (Fed) <daniel.smith@nist.gov<javascript:;>>
Subject: Re: Proposed edits to security strengths section

I'm still not entirely convinced that CCA security is required for fully ephemeral key exchange, though, as any information learned by a CCA-style attack on a single session will be useless for any other sessions. Since by other means (i.e the authentication aspect of the protocol), the CCA attack should only be mountable by Eve if Alice (or Bob) intended to start a session key exchange with Eve in the first place, a CCA-style attack is only going to potentially yield Eve information about Alice (or Bob)'s secret key information for that particular session, and it's of no concern if she learns that since she already knows everything that the secret key information is used for, namely deriving the session key.

I suppose this situation could be more easily screwed up by cryptographically less-than-knowledgeable implementers, more so than if it had the additional CCA security, but I can't think of any other reason to require CCA for fully ephemeral exchange.

From: "Perlner, Ray (Fed)" <ray.perlner@nist.gov<javascript:;>>
Date: Friday, October 7, 2016 at 10:19 AM
To: "Alperin-Sheriff, Jacob M. (Fed)" <jacob.alperin-sheriff@nist.gov<javascript:;>>, "Peralta, Rene (Fed)" <rene.peralta@nist.gov<javascript:;>>, "Liu, Yi-Kai (Fed)" <yi-kai.liu@nist.gov<javascript:;>>, "Moody, Dustin (Fed)" <dustin.moody@nist.gov<javascript:;>>, "Bassham, Lawrence E (Fed)" <lawrence.bassham@nist.gov<javascript:;>>, "Chen, Lily (Fed)" <lily.chen@nist.gov<javascript:;>>, Daniel Smith-Tone <daniel-c.smith@louisville.edu<javascript:;>>, "Jordan, Stephen P (Fed)" <stephen.jordan@nist.gov<javascript:;>>, "Miller, Carl A. (Fed)" <carl.miller@nist.gov<javascript:;>>, "Smith-Tone, Daniel (Fed)" <daniel.smith@nist.gov<javascript:;>>
Subject: RE: Proposed edits to security strengths section

I strongly disagree that the definition comparing security to AES/SHA3 is harder to use than one that specifies concrete numbers in terms of a gate set

1) What we're standardizing are modes that provide CCA/CMA security. This almost certainly means that block ciphers or hash functions will be used in the padding and/or in the message representative. So, cryptanalysts cannot avoid thinking about the cost of attacking a block cipher or hash function.

2) The natural units for a lot of cryptanalytic attacks are not basic machine instructions or gates in the first place. Rather, it is very common to think in terms of field or matrix operations, which are not much easier to analyze than AES in terms of basic gates or real machines. And, unlike standard symmetric crypto primitives, you're less likely to find someone who went to the effort to create a genuinely optimized implementation, especially if the field/matrix size is not a round number.

3) Classical cryptanalysis needs to remain a big part of our evaluation, since, if the best theoretical attack is a variant of Grover's algorithm, there's a significant chance that the best attack in practice will simply be the classical attack. Specifying everything in terms of gate sets etc. will be incomprehensible to people used to doing classical cryptanalysis.

The literature in classical cryptanalysis almost exclusively gives estimates in bits of security. Granted, the unit of work isn't always explicitly an AES operation. Cryptanalysts are happy enough to consider an AES operation, a SHA compression function, a modular multiplication over thousand bit integers etc. to be "about the same amount of work" because they really don't care whether something has 80 bits of security, 75 or 87. If you're cutting your security margin close enough that a few bits of security matter, you're probably cutting it too close. I think this goes double for our process, since the best attacks on asymmetric primitives have a habit of moving significantly.

Regarding Yi-Kai's last suggesting about measuring the "real world" cost of classical and quantum attacks. I believe that is what my currently proposed text is doing. The 5 security categories are merely meant as landmarks along the way so that we can compare like with like when doing performance comparisons, and so that cryptanalysts can be free to choose their own security metric without making their attacks completely incommensurable with every other piece of research out there.

From: Alperin-Sheriff, Jacob (Fed)
Sent: Thursday, October 06, 2016 2:31 PM
To: Peralta, Rene (Fed) <rene.peralta@nist.gov<javascript:;>>; Liu, Yi-Kai (Fed) <yi-kai.liu@nist.gov<javascript:;>>; Perlner, Ray (Fed) <ray.perlner@nist.gov<javascript:;>>; Moody, Dustin (Fed) <dustin.moody@nist.gov<javascript:;>>; Bassham, Lawrence E (Fed) <lawrence.bassham@nist.gov<javascript:;>>; Chen, Lily (Fed) <lily.chen@nist.gov<javascript:;>>; Daniel Smith-Tone <daniel-c.smith@louisville.edu<javascript:;>>; Jordan, Stephen P (Fed) <stephen.jordan@nist.gov<javascript:;>>; Miller, Carl A. (Fed) <carl.miller@nist.gov<javascript:;>>; Smith-Tone, Daniel (Fed) <daniel.smith@nist.gov<javascript:;>>
Subject: Re: Proposed edits to security strengths section

I wasn't going to heavily rock the boat (although I did ask about it), but that was also the main weird issue I had when I first read the draft proposal myself when I started here …

From: "Peralta, Rene (Fed)" <rene.peralta@nist.gov<javascript:;>>
Date: Thursday, October 6, 2016 at 2:29 PM
To: "Liu, Yi-Kai (Fed)" <yi-kai.liu@nist.gov<javascript:;>>, "Perlner, Ray (Fed)" <ray.perlner@nist.gov<javascript:;>>, "Moody, Dustin (Fed)" <dustin.moody@nist.gov<javascript:;>>, "Alperin-Sheriff, Jacob M. (Fed)" <jacob.alperin-sheriff@nist.gov<javascript:;>>, "Bassham, Lawrence E (Fed)" <lawrence.bassham@nist.gov<javascript:;>>, "Chen, Lily (Fed)" <lily.chen@nist.gov<javascript:;>>, Daniel Smith-Tone <daniel-c.smith@louisville.edu<javascript:;>>, "Jordan, Stephen P (Fed)" <stephen.jordan@nist.gov<javascript:;>>, "Miller, Carl A. (Fed)" <carl.miller@nist.gov<javascript:;>>, "Smith-Tone, Daniel (Fed)" <daniel.smith@nist.gov<javascript:;>>
Cc: "Peralta, Rene (Fed)" <rene.peralta@nist.gov<javascript:;>>
Subject: Re: Proposed edits to security strengths section

I tend to agree with Yi-Kai. If we are concerned about some sort of common complexity

measure with other NIST standards, maybe we can look into that as a separate item as

this project progresses.

Rene.

From: Liu, Yi-Kai (Fed)
Sent: Thursday, October 6, 2016 12:02 PM
To: Perlner, Ray (Fed); Moody, Dustin (Fed); Alperin-Sheriff, Jacob (Fed); Bassham, Lawrence E (Fed); Chen, Lily (Fed); Daniel Smith-Tone; Jordan, Stephen P (Fed); Miller, Carl A. (Fed); Peralta, Rene (Fed); Smith-Tone, Daniel (Fed)
Subject: Re: Proposed edits to security strengths section

Thanks Ray! I've been thinking about this, however, and I am becoming more and more convinced that it is a bad idea to define our security levels for PQC in terms of the resources needed to break block ciphers and hash functions. I have two objections:

1. This definition is hard for cryptanalysts to use, because it forces them to convert their results from natural units (e.g., number of arithmetic operations, number of gates, circuit depth) to unnatural ones (e.g., the amount of computational effort needed to break AES or SHA-3).

2. This definition is our way of ducking an important question: how much security do we want for PQC? A proper answer to that question would be something like "at least 2^128 basic operations on a classical PRAM machine with 64-bit registers, and at least 2^64 quantum gates over the basis {CNOT, Hadamard, Phase, Pi/8}." Instead, we are telling the community we want "as much security as certain instantiations of AES and SHA-3," which is unhelpful, because they then have to go figure that out themselves (which is not easy).

Instead, I would favor a different approach:

1. I would much rather we specify some concrete security levels, in terms of some explicit models of computation. This will remove a lot of uncertainty in evaluating the complexity of cryptanalytic attacks. After seeing the public comments, I think this is a significant concern, and it has the potential to turn into a huge source of confusion when we start to evaluate the security of the different cryptosystems.

2. I think it is fine if the PQC security levels don't match the security of AES or SHA-3, because that is comparing apples and oranges anyway. I see this as a less serious problem than having a community-wide meltdown during the PQC competition.

3. Finally, I would suggest that maybe AES and SHA-3 are *not* good guides for the amount of quantum security we want for PQC. Mathematically, there is no reason why we should force the quantum security of PQC to behave the same way as it does for a hash function or block cipher. Instead, perhaps we should think about the "real world" cost of quantum versus classical attacks, and use THAT as the guide for how much quantum security we want for PQC.

Cheers, and sorry for the long email!

--Yi-Kai

From: Perlner, Ray (Fed)
Sent: Wednesday, October 5, 2016 12:20:03 PM
To: Moody, Dustin (Fed); Alperin-Sheriff, Jacob (Fed); Bassham, Lawrence E (Fed); Chen, Lily (Fed); Daniel Smith-Tone; Jordan, Stephen P (Fed); Liu, Yi-Kai (Fed); Miller, Carl A. (Fed); Peralta, Rene (Fed); Smith-Tone, Daniel (Fed)
Subject: Proposed edits to security strengths section

Since I will be gone for the next PQC meeting, Dustin asked me to try rewriting the security strengths section (which I have now divided into 4.A.4 and 4.A.5) See attached
Thanks,
Ray