From: Perlner, Ray (Fed)
To: Liu, Yi-Kai (Fed); Moody, Dustin (Fed); (b) (6) Chen, Lily (Fed); Alperin-Sheriff, Jacob (Fed)
Subject: RE: FAQ update
Date: Friday, October 21, 2016 3:59:09 PM
Attachments: new FAQ-1 Ray.docx

Not sure what happened. Try now.

-----Original Message-----
From: Liu, Yi-Kai (Fed)
Sent: Friday, October 21, 2016 3:57 PM
To: Perlner, Ray (Fed) <ray.perlner@nist.gov>; Moody, Dustin (Fed) <dustin.moody@nist.gov>; Daniel Smith (b) (6) ; Chen, Lily (Fed) <lily.chen@nist.gov>; Alperin-Sheriff, Jacob (Fed) <jacob.alperin-sheriff@nist.gov>
Subject: Re: FAQ update

Hi Ray,

Did you attach the latest version of the file? It doesn't seem to contain your new FAQ's.

Hi Daniel,

1. About your FAQ ("Is this a competition?"), I know we always struggle with this question, if it is not a competition, then what is it? I think I finally found a name for this thing we are doing... it is an "analysis of alternatives." This is a DoD term, but I rather like it.

"An analysis of alternatives (AoA) is an analytical comparison of the operational effectiveness, cost, and risks of proposed materiel solutions to gaps and shortfalls in operational capability. AoAs document the rationale for identifying and recommending a preferred solution or solutions to the identified shortfall(s) [1]."

Source: https://www.mitre.org/publications/systems-engineering-guide/acquisition-systems-engineering/acquisition-program-planning/performing-analyses-of-alternatives

What do you think?

2. Also, do you think we can delete or soften your last paragraph? Especially this part, it seems awfully pessimistic?

"It will likely be the case that NIST's selection is less universally agreed upon, and as such will likely be less universally judged as a fair selection of the best option(s). We cannot, therefore, promise the universal perception of fairness which is naturally implied by a competition; rather, the best we can hope for is to offer selections that most experts can agree are good options, since there will likely be no consensus of what constitutes a best option."

Cheers,

--Yi-Kai

From: Perlner, Ray (Fed)
Sent: Friday, October 21, 2016 2:26:19 PM
To: Moody, Dustin (Fed); Daniel Smith; Liu, Yi-Kai (Fed); Chen, Lily (Fed); Alperin-Sheriff, Jacob (Fed)
Subject: RE: FAQ update

I broke up the material previously in section 4.A.6 among 4 questions added to the Q&A. We may need some additional work on formatting, but. Please let me know if this approach seems good.

Thanks,
Ray

From: Moody, Dustin (Fed)
Sent: Thursday, October 20, 2016 8:10 AM
To: Daniel Smith (b) (6) ; Perlner, Ray (Fed) <ray.perlner@nist.gov>; Liu, Yi-Kai (Fed) <yi-kai.liu@nist.gov>; Chen, Lily (Fed) <lily.chen@nist.gov>; Alperin-Sheriff, Jacob (Fed) <jacob.alperin-sheriff@nist.gov>
Subject: Re: FAQ update

Daniel added a FAQ on the differences with a competition. Reading NISTIR 7977, we certainly share a lot of commonalities with the process described for competitions. I think it's good that we explain what's different.

I've added some revisions/comments. Let me know what you think.

From: Daniel Smith (b) (6)
Sent: Wednesday, October 19, 2016 10:45:18 AM
To: Moody, Dustin (Fed)
Subject: FAQ update

Hi, Dustin,

Here is a draft of a FAQ on the IP issues. This is a contentious point, and we'll probably want to talk about and revise this.

Cheers,
Daniel

Q: Does the requirement for ANSI C source code preclude the use of assembly language optimizations?

A: The optimized code required as part of the submission package should be ANSI C with no assembly (this includes inline assembly). This code is meant to be portable. If significant optimizations can be made with assembly, then it can be included as an additional implementation and discussed in the performance analysis.

Q: Will NIST consider platforms other than the “NIST PQC Reference Platform” when evaluating submissions?

A: The reference platform was defined in order to provide a common and ubiquitous platform to verify the execution of the code provided in the submissions. NIST will include performance metrics from a variety of platforms in our evaluation, including: 64-bit “desktop/server class”, 32-bit “mobile class”, microcontrollers (32-, 16-, and where possible, 8-bit), as well as hardware platforms (e.g., FPGA). Submitters are encouraged to provide additional implementations for these platforms if possible.

Q: In Sections 4.A.2 and 4.A.~~3~~4, NIST’s CFP sets the number of decryption (resp. signature) queries, that an attacker against a proposed encryption (resp. signature) scheme can make, to at most 2 to the 64. What is the rationale for not letting the adversary make essentially as many queries as the target security?

A~~)~~: Our reason for primarily considering attacks involving fewer than 2 to the 64 decryption/signature queries is that the number of queries is controlled by the amount of work the honest party is willing to do, which one would expect to be significantly less than the amount of work an attacker is willing to do. Any attack involving more queries than this looks more like a denial of service attack than an impersonation or key recovery attack. Furthermore, effectively protecting against online attacks requiring more than 2 to the 64 queries using NIST standards would require additional protections which are outside the scope of the present postquantum standardization effort, most notably the development of a block cipher with a block size larger than 128 bits. This may be something NIST pursues in the future, but we do not feel it is necessary for addressing the imminent threat of quantum computers. That said, as noted in the proposed call for algorithms, NIST is open to considering attacks involving more queries, and would certainly prefer algorithms that did not fail catastrophically if the attacker exceeds 2 to the 64 queries.

Q: Why does NIST’s CFP ask submitters to provide a classical security analysis, when the intent is to plan for a world with quantum computers?

A: For algorithms not subject to dramatic quantum attacks, such as those involving Shor’s algorithm, NIST believes that classical measures of security will continue to be highly relevant. Currently envisioned quantum computing technologies would be orders of magnitude slower and more energy intensive than today’s classical computing technology, when performing the same sorts of operations. When these considerations are combined with the poor parallelization of Grover’s algorithm, it becomes quite likely that variants of Grover’s algorithm will provide no advantage to an adversary wishing to perform a cryptanalytic attack that can be completed in a matter of years, or even decades. As most quantum attacks on proposed postquantum cryptosystems have involved some variant of Grover’s algorithm, it may be the case that the best attack in practice will simply be the classical attack.

Additionally, even if the classical attack does not prove to be the most practical attack in the future, the best classical attack will often be highly predictive of the complexity of the best quantum attack (provided there is no dramatic quantum speedup similar to Shor’s algorithm.) At present, the science involved in assessing classical security is better developed than that for assessing quantum security.

Q: How does NIST plan to measure the complexity of attacks involving quantum computers? How does NIST feel that existing measures of quantum security that appear in academic literature are inadequate?

A: Existing measures of quantum security will often ignore the very real challenges involved in implementing a quantum attack. Most notably they ignore the difficulties involved in parallelizing quantum algorithms, and the fact that quantum gates are expected to be much slower and more expensive than classical gates. This leads to claims, for example, that an algorithm with a 128 bits of classical security and no known quantum speedup, should be considered just as secure as a 256-bit block cipher like AES256, since both have “128 bits of quantum security”. NIST strongly disagrees with this analysis and can think of no realistic model of future computing technology where this is likely to be the case. The requirement that submitters provide classical security analysis, and that their parameter sets be categorized in a way that takes classical security into account, mitigates this problem somewhat, but it does not completely eliminate it. In addition to classical attacks, an algorithm may be susceptible to quantum attacks that either parallelize well, or can be made inexpensive by performing a significant portion of the necessary computation classically. NIST feels its categorization of the security such schemes should reflect these threats.

While NIST will continue throughout the evaluation process to accept input from the cryptographic community regarding the best way to measure quantum security, so that it is likely to reflect the real world cost of attacking a scheme, NIST currently favors an approach where quantum attacks are given a strict maximum bound on depth, and then measured based on circuit size. Such attacks may also have a cost metric that rates logical quantum gates as being several orders of magnitude more expensive than classical gates.

Plausible limitations on depth range from 2^{40} logical gates (the approximate number of gates that presently envisioned quantum computing architectures like quDOS are expected to serially perform in a year) through 2^{64} logical gates (the approximate number of gates that current classical computing architectures can perform serially in a decade), to 2^{96} logical gates (the approximate number of gates that atomic scale qubits with speed of light propagation times could perform in a millennium). Presently envisioned quantum computing architectures typically indicate that the cost per quantum gate could be billions or trillions of times the cost per classical gate. However, especially when considering algorithms claiming a high security strength (e.g. equivalent to AES256 or SHA384), it is likely prudent to consider the possibility that this disparity will narrow significantly or even be eliminated.

Commented [PR(1]: Someone please check my math on this paragraph. I’m pretty sure it’s right, but …

Commented [PR(2]: Cite: https://arxiv.org/pdf/1010.5022.pdf

Q: In section 4.A.5 of its CFP, NIST defines security strength categories for submitted algorithms in terms of the security of its existing standards in symmetric cryptography. Can NIST quantify the security of these standards?

A: At the present time, NIST would give the following estimates for the classical and quantum gate counts for the optimal key recovery and collision attacks on AES and SHA3, respectively, where circuit depth is limited to MAXDEPTH.

AES128: 2^{170}/MAXDEPTH quantum gates or 2^{143} classical gates.
SHA3-256: 2^{146} classical gates.
AES192: 2^{233}/MAXDEPTH quantum gates or 2^{207} classical gates.
SHA3-384: 2^{210} classical gates.
AES256: 2^{298}/MAXDEPTH quantum gates or 2^{272} classical gates.

SHA3-512: 2^{274} classical gates.

(Quantum circuit sizes are based on https://arxiv.org/abs/1512.04965 which was also published at http://link.springer.com/chapter/10.1007%2F978-3-319-29360-8_3. NIST believes the above estimates are accurate for the majority of values of MAXDEPTH that are relevant to its security analysis, but the above estimates may understate the security of SHA for very small values of MAXDEPTH, and may understate the quantum security of AES for very large values of MAXDEPTH.)

Q: In section 4.A.5, it is stated that NIST will assume that its 5 security categories are correctly ordered (i.e. that a collision attack on SHA256 (resp. SHA384) will be harder to perform than a key search attack on AES192 (resp. AES 256.)) How realistic is this assumption?

A: Even assuming no disparity in the cost of quantum and classical gates, the assumption holds as long as the adversary is depth limited to fewer than about 2^{87} logical quantum gates. This is quite near the limit of what NIST considers to be a plausible technology for the foreseeable future.

Formatted: Superscript

Q: Is the NIST PQC Standardization Process~~In Section 2.D, the section on Intellectual Property Statement, there is no explicit requirement of royalty-free licensing. Doesn't NISTIR 7977 specify that a NIST competition will require submitters to relinquish intellectual property rights?~~ a competition?

Commented [Office3]: The way you framed the question, it is addressing Microsoft's comment. I think we should make the Q&A more general. I don't think we need to mention IPR in the question, as we don't even mention it in the answer.

A~~)~~ : ~~NISTIR 7977 specifies the rules for NIST competitions. NISTIR 7977 does not specify rules for this process which is NOT a competition.~~ This process shares many features with NIST competitions, and is modelled after the successes we have had with competitions in the past. There are, however, some important requirements that the current research climate demands we require for this process which constitute significant distinctions between this process and a competition.

First, our handling of the applicants does not coincide with a competition as specified in NISTIR 7977. There will not be a single "winner". Our intention is to select a couple of options for more immediate standardization~~, (in competition lingo, one might call these "winners")~~ as well as to eliminate some submissions as unsuitable. ~~(again, in competition lingo, one might call these "losers")~~ T~~but t~~here will likely be some submissions that we ~~cannot classify as either option~~do not select for standardization, but that we also do not eliminate and which may be ~~an~~ excellent options for a specific application that we're not ready or don't have the contemporaneous resources to standardize. In such a circumstance, we would communicate with the submitters to allow these to remain under a public license for study and practice and to remain under consideration for future standardization. There is no specification for the handling of such an applicant in a competition.

Second, the state of the science in the competitions of the past, i.e. for the ~~block cipher and hash function~~AES and SHA-3 competitions, was far more developed than for~~the~~ post-quantum cryptograph~~nom~~y. Though differences of opinion are inevitable~~,~~, the selection of the past winners ~~it~~ should not have been to~~n~~o surprising~~surprise for anyone involved which options were selected as winners~~. Rijndael was obviously one of the best choices as the winner for the AES competition. Keccak was a leading performer, had solid theoretical security and offered more functionality and originality than ~~any~~ other competitors, and was therefore~~, hence~~ ~~, also obviously one of~~ the best possible selection~~s~~. The situation in post-quantum cryptography is less clear and opinions of required properties are less unanimous. It will likely be the case that NIST's~~our~~ selection is less universally agreed upon, and as such will likely be less universally judged as a fair selection of the best option(s). We cannot, therefore, promise the universal perception of fairness which is naturally implied by a competition; rather, the best we can hope for is to offer selections that most experts can agree are good options, since there will likely be no consensus of what constitutes a best option.

Commented [Office4]: Do we want to mention that the criteria/timeline could change as well, due to the uncertainties in the field?

Commented [Office5]: Should we address the issue that people will call it a competition anyway? Say we prefer the phrase "competition-like" or invent some new word like quasi-competition or something?