

Stateful Hash-Based Signatures

Public Comments on Draft SP 800-208
(February 28, 2020 deadline)

<i>Karsten Klein</i>	- 2 -
<i>AMD</i>	- 3 -
<i>Andreas Huelsing</i>	- 8 -
<i>Thales DIS</i>	- 10 -
<i>ETSI TC CYBER WG QSC</i>	- 14 -
<i>NSA's Center for Cybersecurity Standards</i>	- 29 -
<i>Crypto4A</i>	- 31 -
<i>The QuantumRISC Project</i>	- 40 -
<i>Stefan-Lukas Gazdag</i>	- 43 -
<i>Canadian Centre for Cyber Security</i>	- 46 -
<i>Panos Kampanakis</i>	- 48 -
<i>Google</i>	- 52 -

Karsten Klein

From: Karsten Klein

Date: Wednesday, January 15, 2020, at 3:19pm

Hi there,

Concerning - Draft NIST SP 800-208.

I just finished a first read (I've extracted some items for further follow up) and have a general and a minor comment:

Line 486 - With respect to how approved parameter sets are specified and footnote 3 in particular: In general, an overview of all schemes (approved by NIST and existing in the referenced materials) with an outcome (approved, approved with restrictions, not approved, not in scope) and a reasoning (inefficient, ineffective, less secure due to...) could be used to avoid confusion of which parameter sets are approved and which are not. This would also allow to unify the naming scheme and map the parameter set naming used in the referenced RFCs (as it appears to be not homogeneous). Eventually, this allows to omit the footnote. It really confused me, as it mixes scheme and parameter set level)

In short: please consider how not approved and approved parameter sets are represented to avoid confusion.

NIST Response: As noted in Section 5, the approved parameter sets are specified in Sections 5.1 – 5.4. We believe that providing a list of parameter sets that included both approved and non-approved parameter sets would be more likely to create confusion than to improve clarity. While only parameter sets in Section 5.1 came from RFC 8391, we tried to follow that RFC's naming scheme when assigning names to the parameter sets in Sections 5.2 – 5.4.

Line 502 - Numeric Identifier of XMSS-SHA2_20_256:

The Numeric Identifier for XMSS-SHA2_20_256 is 0x00000003 instead of 0x00000002; see RFC 8391 - Table 7.

NIST Response: Thank you. This was a typo and it has been corrected in the final version.

Best regards,
Karsten Klein

AMD

From: Don Matthews

Date: Thursday, January 23, 2020 at 4:39pm

NIST SP 800-208 Draft Review Comments

General Comments

This is a summary of requested modification found by reviewers from AMD.

Commonality of Parameter Sets for Both Algorithms

LMS and XMSS are similar algorithms that as defined in their respective RFCs contain different parameter sets. We believe that there is some benefit to unifying the LMS and XMSS parameter sets as much as possible. The specific parameters are: W value, tree height, and hierarchical definition (HSS and XMSSMT).

- For the W value, LMS support 1, 2, 4, and 8 where XMSS only supports 16 (equivalent to LMS w=4). We would expect that most devices would use 1, 2, or 4 trading off key size for performance, while 8 would be used for interface constrained devices with a subsequent drop in performance.
- For tree height, LMS supports 5, 10, 15, 20, and 25 while XMSS supports 10, 16, and 20. XMSS^{MT} has support for a height 5 tree along with heights of 10 and 20.
- For the hierarchical versions, XMSS^{MT} has a complete parameter set with a variety of options but HSS has no parameter set associated with it and therefore, leaves the use up to an individual creator's definition.

We expected that NIST would use the two IETF standards and create similar parameter sets for both. As an example, having both LMS and XMSS support tree heights of 5, 10, and 20.

By creating a common set of parameters, NIST can allow the implementors to choose an algorithm based on their analysis of implementation rather than choosing an algorithm that has parameters that best fits their requirements.

It is understood that any changes to the parameter sets may create additional work over what would be required to meet the IETF standard. Although, most implementations should be written such that new parameter sets should work with existing code for algorithm implementation. An LMS implementation should allow for heights of 5, 10, and 20. An XMSS + XMSSMT algorithm also should allow for heights of 5, 10, and 20. Therefore, a common parameter set that allowed for heights of 5, 10, 20 for both LMS and XMSS would be possible.

NIST Response: We do not believe there is a substantial benefit in defining new parameter sets for LMS/HSS or XMSS/XMSS^{MT} just for the sake of making the parameter sets of the two schemes more similar. The two schemes will need to have separate implementations, so harmonizing their parameter sets would provide minimal benefit.

Distributed Multi-Tree Hash-Based Signatures

We like this proposal that helps to alleviate concerns about system issues for solutions that require longevity of key signing capabilities.

Concerns

- Large number of cryptographic modules are required if each tree from level 0 and level 1 is a different module
 - The smallest LMS and XMSS^{MT} solution has a top-level tree of height 5
 - would require 33 cryptographic modules
 - The cryptographic module associated with the top level is a single point of failure.
 - Recommendations below make it possible for the signature system to still be functional even with the loss of the cryptographic module associated with the top level.

Recommendations

- Suggested changes to limit the number of cryptographic modules
 - Allow the top level to be based on LMS or XMSS with a height of 2, 3, or 4 and the lower levels to be based on 4, 8, or 16 (respectively) cryptographic.
 - Allow for the second layer to be implemented with LMS, HSS, XMSS, or XMSS^{MT} algorithms.
 - helps alleviate any concern about the total number of signatures that can be performed with a low height (height ≤ 5) top level tree.
 - Create a new system with associated parameters.
 - It would require parameters for the top level (LMS or XMSS with low height) plus parameters for the lower level (LMS, HSS, XMSS, XMSS^{MT}) and parameters for the OTS (LM-OTS, WOTS+).
 - Full parametrization of the complete system could get into long names.
- Suggested change to allow for loss of the cryptographic module associated with the top level
 - Have the signature of each lower level tree stored with the public key of the lower level tree.
 - Top level cryptographic module is only required for initial set up of sub level cryptographic modules allowing for the cryptographic module to be decommissioned.
 - Since the top-level module is only used at system initialization it prevents glitch attacks against the top level (as discussed in section 8).

NIST Response: The text in the final version of the document has been revised in order to make the idea behind distributed multi-tree hash-based signatures more clear.

- There is no need for all of the one-time keys of the top-level LMS or XMSS tree to be used. There is also no requirement for each bottom-level LMS or XMSS key to be created in a separate cryptographic module. For example, suppose the goal were to implement XMSS^{MT} and have two cryptographic modules each of which has 10 000 one-time keys available for signing ordinary messages. The parameter set XMSSMT-SHA2_20/2_256 could be used. A top-level XMSS tree would be created in one cryptographic module, and then 10 bottom-level XMSS trees could be created in that same cryptographic module. Another 10 bottom-levels XMSS trees could be created in a second cryptographic module. Only 20 of the 1024 top-level XMSS tree's one-time keys would be used. The resulting signatures will be slightly longer than necessary (since the top-level tree has a height of 10 instead of 5), but no additional cryptographic modules would be needed.
- While it is not described in Section 7, an HSS or XMSS^{MT} tree with more than two levels could be created. However, the goal is for the signature to be verifiable by an unmodified implementation. So, we would prefer avoiding the creation of new systems along with their associated parameters, as that would require signature verification software to be modified to be able to accept these new systems.
- The cryptographic module that holds the top-level tree is not a single point of failure in the manner described. The text in Section 7 assumes that there is an external device (e.g., laptop computer) that is sending commands to the cryptographic modules and receiving the responses. So, when a bottom-level tree is to be created, the external device sends the key generation command to the appropriate cryptographic module, receives the resulting public key (i.e., tree root), and then sends that public key to the cryptographic module holding the top-level tree for signature. This external device will then store the resulting signature for later use. So, the signature for each lower-level tree is not just stored in the cryptographic module holding the top-level tree. In fact, since the signatures can be made public, multiple copies of these signatures may be stored in multiple different locations.

FIPS Requirement

NIST has defined a process for algorithm validation (CAVP) and module validation (CMVP). SP 800-208 is defining an algorithm that should fall under CAVP but has a mandate that it only runs on a CMVP validated module.

In the past, FIPS has not posted a certification requirement for the solutions using NIST approved cryptographic algorithms. Many government contracts require FIPS certification, at different levels, but individual customers could determine if there was value in having a certified implementation. With the FIPS certification requirements as specified in line 741-743 in this draft, it mandates FIPS testing on the processing module to be able to implement these algorithms. This leads to two different concerns

- It is not possible for a company to implement 800-208, even for internal uses, without getting FIPS validation on their cryptographic module(s) or purchasing a module from an outside company
- It will lead to confusion over algorithm names. Referencing any of the algorithms (LMS, HSS, XMSS, XMSS^{MT}) doesn't indicate if it is compliant with 800-208 since it may only be compliant with the RFCs
 - AES, and other algorithms, are defined by NIST and is always NIST compliant no matter where used for cryptographic operations

- AES may or may not be CAVP tested

Some of the impetus for approval of Stateful Hash-Based Signatures was that companies may not be able to wait for the PQC algorithm selection process. Adding a FIPS 140 level 3 requirement for all implementations of 800-208 (see lines 741-743) could delay companies from using 800-208 as a solution. This requirement could be especially problematic for any company that has not been involved with a previous FIPS validation.

NIST Response: The “Authority” section near the beginning of the document explains that “This publication may be used by nongovernmental organizations on a voluntary basis....” So, companies are free to implement the schemes described in SP 800-208 (and RFCs 8391 and 8554) without getting their implementations validated under either CAVP or CMVP. (Federal government agencies, would not, however, be free to use implementations that had not been validated.)

The commenters may be correct that one could not claim that an unvalidated implementation is “compliant with SP 800-208,” but that is not unique to this document. As noted in the comment, one may implement AES without getting the implementation validated. However, FIPS 197 states that “Implementations of the algorithm that are tested by an accredited laboratory and validated will be considered as complying with this standard.” So, if the implementation were not validated, one could claim to have implemented AES, but could not claim that the implementation is “NIST compliant.” The same applies with the digital signature schemes in SP 800-208.

Federal government agencies cannot use cryptography that is not NIST-approved and must only use FIPS 140 validated cryptographic modules. So, for government agencies, the implementation would have to be validated whether SP 800-208 mentioned that requirement or not. While the document does impose certain additional requirements (e.g., level 3 physical security), we believe those requirements are appropriate given the risk that one-time key reuse poses.

Nongovernmental organizations that are not trying to sell to the Federal government do not need to follow any of the requirements in SP 800-208, although we hope that such organizations choosing to implement stateful hash-based signatures will review SP 800-208 and consider the benefits of following its guidelines.

Specific items

Line 358 (Figure 1) – Figure 1 is representing both the hash chain but also the signature and verification operation simultaneously. It may be easier for some to understand if this was broken into two different figures. The new Figure 1 would consist of (using shorthand):

$X \rightarrow \text{HASH} \rightarrow H(X) \rightarrow \text{HASH} \rightarrow H(H(X)) \rightarrow \text{HASH} \rightarrow H(H(H(X))) = \text{pub}$

The new figure 2 would consist of (using shorthand):

$X \rightarrow \text{HASH} \rightarrow H(X) = S$ $S \rightarrow \text{HASH} \rightarrow H(S) \rightarrow \text{HASH} \rightarrow H(H(S)) = \text{pub}$

|-----| |-----|

Signing Operation

Verification Operation

NIST Response: The final document has been modified based on this suggested change.

Lines 449-450 (related to HSS) – “Shall be used for every LMS tree at that level” –implies that one can have an HSS signature design that utilized a different LMS parameter set at each level. The only requirement is that they use the same hash algorithm. Is your intention to allow for that type of design?

NIST Response: Yes, it was intentional. A sentence has been added to this paragraph to make it clear that different parameter sets may be used at different levels.

Line 491 (and others) – As discussed in the general comment previously, for XMSS, it may make sense to allow for other W values (1, 2, 4, 8, 32) like what has been provided by LMS. This would allow for performance/signature size tradeoffs and allow for similar configurations between LMS and XMSS.

NIST Response: SP 800-208 is intended to profile the use of RFC 8391 and RFC 8554 rather than specifying yet another hash-based signature scheme. Section 3.1.1 of RFC 8391 specifies that w must be either 4 or 16, so we could not specify parameter sets for w values of 1, 2, 8, or 32, while remaining conformant with the RFC. While it would be possible to define parameter sets for $w = 4$ while still remaining conformant with RFC 8391, we believe it is better in this case to maintain consistency with the RFC by only defining parameter sets for $w = 16$.

LMS & HSS RNG requirements:

If a TRNG is available, should it not be possible to use the TRNG for all private keys? This capability is prevented in line 561-562 “shall be generated using the pseudorandom key generation method”

If TRNG is not allowed for private key generation (as currently written), then the content in parentheses should be removed from line 567 “and SEED (if using the pseudorandom key generation method)”

Line 577 – “generated using the pseudorandom key generation” – as with LMS comment above, if a TRNG is available, this specification prevents it from being used for private key generation.

NIST Response: The parenthetical on line 567 has been deleted. While any implementation of SP 800-208 would need access to a TRNG (or non-deterministic random bit generator (NRBG) in the terminology of SP 800-90), and it would be theoretically possible to use this NRBG to generate all of the random data for these schemes, doing so would not be practical, so there would be minimal benefit in allowing it as an option. In addition, allowing for all of the private elements to be generated using an NRBG rather than generating them deterministically from a seed would make testing with the Automated Cryptographic Validation Testing System (ACVTS) more complicated.

Andreas Huelsing

From: Andreas Huelsing

Date: Wednesday, January 29 at 10:58am

Dear NIST team,

Thanks for your work. I highly appreciate the current draft of SP 800-208. I only have a few remarks.

- a) As you do define a key generation mechanism, it might be worthwhile to define a forward-secure one (as for example in the original XMSS paper). XMSS and LMS with forward-secure key generation lead a forward-secure signature scheme. Forward-security can add a strong guarantee for old signatures in case of key-compromise and essentially comes for free in this setting.

NIST Response: While the cost of switching from the currently specified method of key generation to one that offered forward-security would be relatively small, it would add some complexity to the implementation, and the benefits are not entirely clear.

In theory, if a forward-secure key generation method is used, and the cryptographic module holding the private keying material becomes compromised, relying parties could be notified to distrust some of the one-time keys while continuing to trust others. However, this may not be practical to implement in practice. It would require having a secure means for distributing the revocation information, which would likely mean a separate signing key that has not been compromised. It would also require the issuer of the revocation information to be able to determine when the compromise occurred in order to ensure that all potentially compromised one-time keys are revoked.

Safely making use of the forward-secure property would also require sufficient knowledge about the implementation of the cryptographic module in use. Even if NIST mandated the use of the forward-secure key generation method specified in the original XMSS paper, that would not ensure that every implementation securely deleted older state information after each signature generation. Without knowing exactly when (or if) older state information was deleted, the only safe option for users of cryptographic modules would be to not take advantage of the forward security property. Given that the users of the cryptographic modules would not, in general, be able to safely make use of the forward security property, it may be better not to suggest in SP 800-208 that the property is available when it may not be in every implementation.

- b) I understand that there is no decision made about the NIST post-quantum standardization project. However, if NIST is even considering to keep SPHINCS+ it might be helpful to synchronize the addressing schemes of XMSS & SPHINCS+ as well as considering the tree-less WOTS-PK compression. This would allow to treat XMSS as a sub-step of SPHINCS+, requiring the same code-base. Especially, the code for XMSS signature verification would be almost a full sub-set of the SPHINCS+ verification code.

NIST Response: We believe that the benefits of referencing an existing standard (RFC 8391) outweigh the potential benefits of developing yet another stateful hash-based signature scheme that happens to be more closely aligned with SPHINCS+. For the intended use cases specified in Section 1.1 of the document, it would not be practical to implement the scheme in SP 800-208 and then later add an implementation of SPHINCS+, so increasing the amount of commonality between them would be of limited benefit.

- c) While I essentially do agree with your assessment of the security proofs there are a few nits:
- Line 1459: Should say second-preimage resistance.
 - On the whole paragraph starting at 1457: Following the analysis in our recent publication "The SPHINCS+ Signature Framework" we additionally require h_k to be post-quantum, multi-function, multi-target decisional second preimage resistant. Alternatively, one needs a statistical assumption about h_k which does not hold for random functions (see the discussion about the tight security proof for SPHINCS+ on the PQC mailing list).
 - You could mention that LMS and XMSS are (non-tightly) secure in the standard model if we are willing to assume collision resistance of the used hash function. In this case, all the bitmasks and the prepended values can be arbitrary bit strings.

NIST Response: Line 1459 has been changed to "decisional second-preimage-resistant" and a reference to "The SPHINCS+ Signature Framework" has been added. We decided against mentioning the security in the standard model.

- d) Regarding parameters: While I do think that limiting the choice of w and the used hash function is a good idea, I do not see any benefit in limiting the number of options for the total tree height and the number of layers. All implementations that I have seen are generic with regard to these values. This allows users to adapt the schemes to their constraints. What would be necessary in this case is defining upper bounds on both values.

NIST Response: While it may be relatively easy for implementations to support arbitrary tree heights, all implementations need to be tested under the Cryptographic Algorithm Validation Program (CAVP), and allowing implementations to use any tree height would make that testing more complicated. SP 800-208 does not impose any limits on the number of layers other than the limits that are imposed by RFC 8391 and RFC 8554.

Best wishes,

Andreas

Thales DIS

From: Aline Gouet

Date: Tuesday, February 18, 2020 at 8:59am

Hello

Please find below comments on NIST draft 800-208 as a contribution of Thales DIS.

Best regards

Aline

Comment 1: In sub-section 1.1, the statement from line 273 to 275 discourages the use of stateful HBS: “Stateful HBS schemes are only suitable for particular uses, as they require careful state management. The recommendations are summarized in section 1.2 and described in detail in [8]”. We believe that stateful HBS schemes can be efficient, secure and useful when some implementation conditions are met. As the document itself is meant to be a general recommendation, we would suggest to rephrase the sentence in a more assertive manner, e.g. highlighting the importance of securely manage the state/counter in the implementation whatever or independently from the use-case.

NIST Response: While it is certainly true that proper state management is important regardless of the use case, Section 1.1 is only intended to discuss intended applications. Section 1.2 already highlights the importance of proper maintenance of state.

Comment 2: In sub-section 1.1, lines 276 to 279, recommendation 2) “the implementation will have a long lifetime” seems to be different compared with and maybe contradict in some extent with initial answer from NIST on Gemalto comments: “we are keen to discourage the use of stateful hash-based signatures except in scenarios where signing is infrequent” (from <https://csrc.nist.gov/CSRC/media/Projects/Stateful-Hash-Based-Signatures/documents/stateful-HBS-misuse-resistance-public-comments-April2019.pdf>).

We believe that stateful HBS are suitable for long term, frequent usage, as long as the security recommendations are taken care of. Could you please clarify NIST’s position on this point? It would also make sense to add a fourth recommendation: “4) the implementation relies on hardware cryptographic modules, as described in section 8.1.”

NIST Response: We do not view these recommendations as contradictory. For example, if a manufacturer is selling a device that will be used for a very long time, there will be a need to be able to securely distribute firmware updates to that device as long as the device continues to be supported. However, firmware updates may only be issued a few times a year, so the firmware signing key may similarly be only used a few times a year.

We agree that stateful HBS is just as secure as other cryptographic algorithms as long as one-time keys are never reused and that organizations should generally be able to avoid one-time key reuse by following appropriate security recommendations. However, there is a risk that not all

organizations that implement or use stateful HBS will carefully follow all of the appropriate security recommendations. For this reason, we believe that stateful HBS is primarily intended for applications in which there is no practical alternative.

The suggested fourth recommendation would not be consistent with this section as it describes an implementation requirement rather than a type of application.

Comment 3: Section 3 on General Discussion describes mainly similarities of LMS and XMSS (in subsections 3.1, 3.2, 3.3) and only few differences between LMS and XMSS, i.e. mainly in subsection 3.4 on bitmasks and prefixes. It would be useful for the developer to describe in a similar way differences in final hashing of the Winternitz scheme and signature structure.

NIST Response: Section 3 is intended to provide a high-level description of hash-based signature schemes. For those who are not implementing the schemes this description may be sufficient. For those who will be implementing the schemes, reading this high-level description first may aid in understanding the detailed descriptions in RFCs 8391 and 8554. For those who are only interested in a high-level description, we believe the description of the final hashing is too detailed. For those who intend to implement one or both of these schemes, we believe the descriptions in the RFCs are sufficient.

Comment 4: In section 3 on General discussion, there is no guidance on how to select one signature scheme or the other one based on different criteria, such as for example the total number of hashing (including the hashes used for bitmask generation) for comparable parameter sets.

NIST Response: While this is true, it is no different from other NIST publications. For example, FIPS 186-4 describes RSA, DSA, and ECDSA, but does not provide any guidance on how to select one of these signature schemes.

Comment 5: In sections 4 and 5, the parameter sets of LMS and XMSS are described using original notation from RFC 8554 and RFC 8391. Since the naming for both schemes are not unified, that would be helpful to inform the reader in Section 3 and highlight some equivalences or differences in Notation. For example, it might be worth mentioning that p in LMS description is equal to len parameter in XMSS. Another example is w that has different meanings in XMSS and in LMS, w in LMS corresponds to logarithm of w in XMSS.

NIST Response: A footnote has been added to Section 3.1 explaining the different meanings of w in the two RFCs. For the other variables used in the descriptions of the two schemes, definitions have been added to Section 2.3.

Comment 6: The approved parameter sets for both LMS and XMSS are described in Section 4 and Section 5. For some parameter sets of XMSS, there are no equivalent parameters for LMS and vice versa. For example, there are no XMSS parameters for 32 signatures while it is possible for LMS. We believe that it is important to maintain the possibility to sign 32 messages which is suitable for implementation on constrained secured elements and it would be good to either provide similar parameter sets for both schemes or to explain the rationale of not having similar parameters for both schemes.

NIST Response: We do not believe the Special Publication itself is an appropriate place to provide rationale for design decisions. However, the parameters sets were mainly derived from those specified in RFCs 8391 and 8554, and so the two schemes do not have similar parameters in some cases since similar parameters were not defined in the RFCs.

An LMS tree with a height of 5 seems mainly useful as part of a multilevel HSS tree. As the parameter sets for XMSS only apply to single-level trees, the usefulness of a 5-level tree seems less clear. As described in Forward Secure Signatures on Smart Cards, even smart cards can support XMSS keys with trees of height 10 or more.

Comment 7: The four approved hash functions are defined in the beginning of Section 4 and Section 5. Since the most time consuming part of the signature is the OTS computation, it might be beneficial to have the possibility to use a function based on a block-cipher for this part, e.g. based on NIST SP 108 with PRF = CMAC-AES-256.

NIST Response: While it may be possible to specify a one-time signature scheme that uses CMAC-AES-256 rather than a hash function, doing so would not be straightforward. One could not simply use the output of one CMAC operation as the key for the next one, as the CMAC output would be at most 128 bits, whereas the key would need to be 256 bits. Also, even if a workable scheme were specified, it is not clear that the security proofs that have been developed for the OTS schemes would apply to a version in which the hash function was replaced by a CMAC operation. Finally, we believe that the OTS schemes using hash functions are already sufficiently fast. The key generation times in Table 3 of RFC 8554 suggest that generating each OTS key takes only about 0.16 milliseconds, and creating a signature should take only about half that time.

Comment 8: In Section 4, Tables 1, 3, 5 and 7, the parameter ls is indicated. Is there any reason for mentioning this parameter? We believe that it is used only for specific implementation described in original proposal, but it is not mandatory for different implementation, therefore it might be confusing placing it among the structure influencing parameters.

NIST Response: ls was included in these tables since the tables were somewhat modeled on Table 1 in RFC 8554. However, the ls column has been removed, since its value can be computed from n and w using the formulas in Appendix B of RFC 8554.

Comment 9: In Section 3, Figure 5, the symbol \otimes is used for XOR operation. Maybe, it would be better to use classical symbol \oplus instead.

NIST Response: Thank you. The symbol in Figure 5 has been changed.

Comment 10: In section 8, subsection 8.1, it is mentioned that “*The cryptographic module shall update the state of the private key in non-volatile storage before exporting a signature value or accepting another request to sign a message*”. Could you please clarify whether this requirement also enable the possibility to use external memory to store the encrypted private key. (The private key would be encrypted by a key of cryptographic module.)

NIST Response: The final sentence of the first paragraph in Section 8.1 states that “The cryptographic module **shall not** allow for the export of private keying material.” This prohibition

includes exporting private keying material in encrypted form, even if no copies of the decryption key are available outside of the cryptographic module.

For most cryptographic algorithms storing private keying material outside of the cryptographic module in encrypted form would not pose a security vulnerability. However, in a stateful HBS scheme allowing private keying material to be stored outside of the module could lead to problems with state management that could result in one-time key reuse. For example, when the key is loaded into the cryptographic module for use, the module may have no way to verify that the key being loaded represents the current state of the key rather than being an older copy of the key that is being loaded. The only safe way for the cryptographic module to ensure that state information is being properly maintained is to store this information internally.

Comment 11: General comment: beyond the algorithm standardization, there is a need to address the need for a standardized key parameter encoding. This applies not only to state full HBS schemes, but any new HBS scheme in general. A general recommendation is that implementations should rely on standardized key encoding techniques, which should be referenced.

NIST Response: As with other NIST standards, SP 800-208 will not specify encoding methods (other than the XDR that is included). Other standards organizations, such as the IETF, are in a better position to specify encoding methods.

ETSI TC CYBER WG QSC

From: ETSI CyberSupport

Date: Wednesday, February 19, 3:39am

LIAISON STATEMENT

Title: Responses to NIST's call for comments on Draft SP 800-208: Recommendation for Stateful Hash-Based Signature Schemes

Date:

From (source): TC CYBER WG QSC

Contact(s): cybersupport@etsi.org

To: NIST

Copy to:

Response to: [NIST's call for comments on Draft SP 800-208](#)
(if applicable)

Attachments:
(if applicable)

TC CYBER WG QSC – Responses to NIST's call for comments on Draft SP 800-208: Recommendation for Stateful Hash-Based Signature Schemes

This document contains a non-exhaustive collection of comments from ETSI TC CYBER WG QSC on NIST's draft Special Publication 800-208: Recommendation for Stateful Hash-Based Signatures Schemes.

The draft specifies approved profiles for the LMS/HSS and XMSS/XMSS^{MT} stateful hash-based signature schemes. This means that it lists parameter sets for the schemes, but it relies on RFC 8391 and RFC 8554 for detailed descriptions of the algorithms. This is problematic for several reasons:

- Although LMS and XMSS are very similar, the two RFCs use different and sometimes conflicting notation. The NIST draft keeps the same notation as the RFCs, which will inevitably cause confusion for readers who are not already familiar with the schemes. Harmonising the notation is preferable, but not straightforward. One possible, but imperfect, solution would be to add a section that defines the mappings between notations. An alternative may be to consider producing two separate documents, one profiling LMS/HSS, and the other profiling XMSS/XMSS^{MT}.

NIST Response: Section 2.3 was extended to include definitions of the variables from RFCs 8391 and 8554 that are referenced in the SP, specifically noting instances in which a single variable is used in both RFCs, but with different meanings. NIST does not believe there is a need to have separate documents for LMS/HSS and XMSS/XMSS^{MT}. While it is true that RFC 8391 and RFC 8554 sometimes assign different meanings to the same identifier, we believe that implementers understand that the definition of an identifier in one context does not apply to the use of that identifier in a different context. For example, in FIPS 186-4 the identifiers p and q are used to identify prime numbers in both RSA and DSA, but in RSA they represent the private factors of the modulus n , whereas in DSA they represent public domain parameters. The identifiers n and d are used in the descriptions of both RSA and DSA, but with different meanings in each. Readers of FIPS 186-4 understand that the definitions of identifiers such as p , q , n and d in the description of RSA do not apply to DSA and vice versa.

- The RFCs were only intended to describe the schemes “with enough specificity to ensure interoperability between implementations”. Neither RFC gives a full description of signature generation. Indeed, RFC 8391 provides example pseudocode for computing the authentication path for XMSS, but strongly recommends that a different method is used. Further, both the RFCs, as well as the draft SP, omit discussion on tree management strategies; RFC 8391 mentions it briefly, but general discussion is omitted. While algorithms such as the Buchmann-Dahmen-Schneider (BDS) algorithm are not required for interoperability, some mention of them may be beneficial for prospective implementors.

NIST Response: Noted. As is the case with other NIST algorithm publications, SP 800-208 specifies the requirements that implementations must meet, not how to meet them. Those who are trying to implement these algorithms will read the RFCs. While RFC 8391 does not describe tree management strategies itself, it does refer readers to papers that provide this information. We strongly encourage implementers to review the existing literature for guidance on how to most efficiently and safely implement these schemes.

- There are some places where the RFCs are ambiguous. For example: when RFC 8554 refers to the LM-OTS or LMS public key it is not always clear whether it means the full public key including the typecodes and identifiers, or just the final hash values; RFC 8391 does not describe what should happen when idx_sig , which is incremented with each signature, exceeds the number of available one-time signatures.

NIST Response: We believe that if there are details in these RFCs that are ambiguous, then it would be better for them to be addressed in the RFCs (either via errata or updates). Implementers may also refer to the reference implementations, when necessary. If NIST were to attempt to address potential ambiguities in SP 800-208, then there would be the risk of NIST unintentionally deviating from the RFCs. Note that details about such issues as how to format public keys and signatures are out of scope for NIST algorithm specifications, except in the case of data objects that are to be digitally signed (e.g., the root of a lower-level tree). Other encoding details are better addressed elsewhere (e.g., [RFC 8708](#)).

Consequently, without further guidance it would be difficult for a non-expert to implement the signature schemes correctly and efficiently from the NIST draft and the RFCs.

NIST Response: As noted above, issues related to encoding and efficient implementation techniques are out of scope for SP 800-208. Implementers are strongly encouraged to review the

existing literature on hash-based signatures for information about how to implement these schemes efficiently.

More detailed comments follow:

Line 131: *“NIST would like feedback on whether there would be a benefit in reducing the number of parameter sets...”*

There are currently 80 LMS parameter sets, 12 XMSS parameter sets, and 32 XMSS^{MT} parameter sets. This seems excessive. Fewer choices of parameters generally increases interoperability of implementations, especially as there are now different choices of hash functions. In general, the choice of which parameter sets to eliminate and which to include is not straight-forward: parameter choices require different trade-offs, and those trade-offs may be compounded by other implementation choices, such as tree management strategies. However, certain parameter sets are impractical and can easily be eliminated. For example, RFC 8554 allows for up to 8 layers in an HSS hierarchy, and each layer can be of height at most 25, giving a maximum total tree height of 200. Time and compute resources for such a parameter set may not be readily available, and the benefits of using such large constructions are not clear. Conversely, it seems unlikely that the improved verification times are worth the increased signature sizes for the LMS parameters where $w = 1$ or $w = 2$. Therefore, we recommend NIST reduce the approved parameter sets to those that are practical or feasible to use.

There is an issue of redundancy in parameter sets: there exist multiple parameter sets that offer the same signature size but require a varying number of hash function invocations. Such parameter sets could be pruned down to the most performant options while the rest are discarded, perhaps based on the number of signatures required, or for obtaining specific trade-offs. Unfortunately, no closed-form formula currently exists that would exclude non-optimal parameter sets.

NIST Response: While we agree that an instance of HSS with 8 layers of trees each having a height of 25 would not be practical, we do not believe there is a practical way to prohibit its use, as parameter sets are only defined for LMS, not HSS. As noted in Section 7.1 of SP 800-208, a cryptographic module that implements LMS signing has no control over how LMS instances are combined to create HSS instances. On the verification side, we expect verifiers to process each LMS public key and signature in a chain individually, so that the overall height of the multi-tree will not be relevant.

We also believe that the primary use case for stateful hash-based signatures will be scenarios such as firmware signing, in which the same entity is in control of both the signing process and the verification software, so that interoperability between arbitrary signing cryptographic modules and arbitrary verification cryptographic modules will not be necessary.

Line 143: *“NIST would like feedback on whether there is a need to be able to create one-level XMSS or LMS keys in which the one-time keys are not all created or stored on same cryptographic module...”*

Resilience can already be provided by distributing a two-level HSS or XMSS^{MT} instance over different cryptographic modules. Distributing a single-level LMS or XMSS tree would

likely require more significant changes to the interfaces for key generation, but only saves the cost of an intermediate one-time signature.

NIST Response: Thank you for the feedback. Based on the feedback provided, we have decided not to include such an option.

Line 273: *“Stateful HBS schemes are not suitable for general use because they require careful state management that is often difficult to assure...”*

Another feature of HBS schemes that makes them less suitable for general use is that a given key pair can only sign a limited number of messages, and once that limit has been reached the long-term signing key is no longer useable.

NIST Response: Noted. However, the goal of this section is to highlight the potential security risk associated with the use of stateful hash-based signatures. While the number of signatures is limited, this can be addressed by choosing a parameter set that has enough one-time keys to ensure that the limit will never be reached.

Line 276: *“Instead, stateful HBS schemes are primarily intended for applications with the following characteristics...”*

It is also necessary to estimate the maximum number of messages that will need to be signed over the lifetime of the implementation, as this determines which parameter set should be used. This may be straightforward for some applications, but difficult for others; of course, it may be possible to be conservative and use a significant overestimate, but at the cost of reduced performance and increased signature sizes.

Further, there is also the notion of signature “loss” over the lifetime of the long-term key pair, depending on how state is managed. For example, an implementation may partition the state and advance it in distinct, non-overlapping blocks, accepting the risk that a system restart would lose the number of signatures in a single block. Over time, with large enough blocks, or with enough reboots, a significant portion of the total signatures may be lost.

It should be noted that the longer a hardware cryptographic module is in use, the greater the probability of device failure becomes. In such a case, existing signatures can still be verified, but no new signatures can be created under that same long-term key pair. As key back-up and recovery is restricted by the draft, the eventuality of no longer being able to generate signatures under a long-term key pair should be considered before deployment.

NIST Response: NIST is not aware of applications that would have the characteristics listed in Section 1.1 that would require frequent signing. So, for the intended applications it should not be overly difficult to choose a parameter set that provides enough one-time keys for the intended lifetime of the HBS scheme without significantly increasing signature sizes or reducing performance. In general, adding another tree level will add significantly to the signature size and to signing and verification time, but increasing the height of a tree by 5 (which multiplies the number of available OTS keys by 32) will add a relatively small amount to the signature size and signing/verification time.

As signing is not expected to be frequent, there should be no reason for HSMs to advance the state in blocks rather than a single one-time key at a time. It will be necessary for the

parameter set to have more one-time keys than the number of signatures needed in order to account for the possible loss (due to malfunction) of cryptographic modules, but it should not be problematic to account for this.

While there is a risk that a cryptographic module may malfunction (or cease working) during the intended lifetime of the HBS scheme, spreading the one-time keys across multiple cryptographic modules, as suggested in Section 7, should help to ensure that the ability to generate signatures is not lost.

Footnote 2: *“HSS allows for up to eight levels of trees and XMSS^{MT} allows for up to 12 levels of trees.”*

This restriction on the number of layers is important enough that it should be included in the main body of the text, as it could easily be missed. Implementors will select parameter sets from the tables within the SP, therefore parameter set restrictions should be explicit.

NIST Response: Noted. This footnote is merely repeating information from the base standard. While SP 800-208 does not restrict the number of levels beyond the restrictions imposed by the base standards, NIST does not envision use cases in which it would be necessary to use so many levels of trees. While two levels of trees are useful in order to easily spread one-time keys across multiple cryptographic modules as described in Section 7, the need for more than two levels of trees seems unclear.

Line 427: *“...which uniquely identifies where a particular hash invocation occurs within the scheme.”*

As per the comment below regarding Line 576, the addressing scheme used in RFC 8391 does not uniquely identify where every hash invocation occurs within the scheme.

NIST Response: Noted. Section 3.4 is referring to the generation of hash chains and Merkle trees, not to the generation of secret keying material. The suggested key generation mechanism in Section 3.1.7 of RFC 8391 uses neither prefixes nor bitmasks.

Line 428: *“This address is then hashed along with a unique identifier for the long-term public key (SEED) to create the prefix.”*

There is an unhelpful (and potentially dangerous) conflict of notation between the use of SEED in XMSS, where it is a public identifier, and in LMS, where it is a private value used to derive the one-time private keys (see line 563).

NIST Response: Noted. While it may be felt that it would have been better if RFC 8391 and RFC 8554 had not both used the identifier SEED, with one using it to refer to a public value and one using it to refer to a private key, NIST cannot change these RFCs. Text has been added to Section 2.3 to specifically note that SEED has different meanings in each of the RFCs, and the body of the text in SP 800-208 is written in such a way to ensure it is clear when the text is discussing LMS or XMSS.

Line 436: *Figure 5*

The diagram should use the symbol \oplus to denote exclusive or instead of \otimes .

NIST Response: Corrected. Thank you.

Line 438: *“This Special Publication approves the use of LMS and HSS...”*

The use of the word “and” implies that NIST approves stand-alone LMS implementations that are not themselves HSS with L=1. Section 6 of RFC 8554 states that “Since HSS with L=1 has very little overhead compared to LMS, all implementations MUST support HSS in order to maximize interoperability”; the somewhat ambiguous language “all implementations” is taken to mean “all implementations of LMS”. NIST should make it explicit if they wish to allow non-HSS implementations of LMS. However, as *LMS* is often used interchangeably with *HSS* (which could lead to undue confusion) it is recommended that NIST only allow HSS, where single-layer LMS is explicitly HSS.

NIST Response: Noted. The only difference between LMS and HSS with L=1 is the addition of some parameter information at the beginning of the encoding for public keys and signatures in HSS. As this encoding is outside the scope of SP 800-208, the difference is not relevant to this specification. Also, in order to implement the scheme suggested in Section 7 of SP 800-208, each cryptographic module would implement LMS and an external (non-cryptographic) device would use the outputs of these modules to form HSS public keys and signatures. So, while relying parties would receive HSS public keys and signatures, the signing cryptographic modules would implement LMS, and so the Cryptographic Algorithm Validation Program (CAVP) would be validating implementations of LMS key and signature generation.

Line 444: *“... the hash function used for the LMS system **shall** be the same as the hash function used in the LM-OTS keys.”*

RFC 8554 allows the use of different hash functions in LM-OTS and the LMS tree. If this restriction is intended to be enforced by verifiers, then Section 8.2 needs to mandate an explicit check of the typecodes in the public key, with the public key being rejected if they do not correspond to the same hash function.

NIST Response: The referenced text applies to “When generating a key pair for an LMS instance.” There is no intention to require verifiers to enforce this restriction.

Line 447: *“If the HSS instance has more than one level, then the hash function used for the tree at level 0 **shall** be used for every LMS tree at every other level.”*

As expressed in Section 6.1 of RFC 8554, the HSS public key only includes the typecodes for the LMS and LM-OTS signatures at level 0. The general HSS process described in RFC 8554 specifically allows the use of different parameter sets, and hence different hash functions, at different levels. If this restriction is to be enforced by verifiers, then Section 8.2 of the draft needs to mandate an explicit check of the typecodes contained in each signature; signatures are to be rejected if the typecodes do not correspond to the hash function specified in the HSS public key.

Because the long-term public key only includes the typecodes for the LMS and LM-OTS signatures at level 0, the signer could change the parameters used at other levels over time; that is, different signatures could use different parameters. Although Section 6 of RFC 8554 makes the explicit requirement “...the signer **MUST NOT** change the parameter sets for a specific level”, there is no way to detect or forbid this from the perspective of a verifier,

without storing extra state. Therefore, complete parameter sets (for all levels) should also be included in, or derivable from the public key.

Considering the above comment regarding Line 438, if an LMS instance is defined as an HSS instance with $L=1$, and if parameter sets are validated, there may be additional difficulty with signature verification if using the distributed method described in Section 7.1 of the draft, as each distinct module will use “ $L=1$ ”, although the “virtual hierarchy” is larger.

Section 5.3 of RFC 8554 (LMS public key) does not set explicit requirements for the LMS public key format. The language used is “*the LMS public key can be represented as the byte string $u32str(type) || u32str(otstype) || I || T[I]$* ”. In addition to the comments given above, NIST could make the public key formats explicit requirements. Similarly, there is a lack of requirements expressed for the LM-OTS or HSS public key formats.

There is also an unhelpful (and potentially dangerous) conflict of indexing conventions between HSS, where level 0 corresponds to the “root” tree used to compute the HSS public key, and XMSS^{MT}, where level 0 corresponds to the “leaf” trees used to sign messages.

NIST Response: The referenced text is in a paragraph that applies “When generating a key pair for an HSS instance.” There is no intention to require verifiers to enforce this restriction. It is also understood that using the scheme described in Section 7.1, there is no way for the signing cryptographic modules to enforce this restriction.

While Section 5.3 of RFC 8554 does not set explicit requirements, Section 3.3 notes that “The signature and public key formats are formally defined in XDR to provide an unambiguous, machine-readable definition [RFC4506].” The text quoted above from Section 5.3 of RFC 8554 is consistent with the XDR for public keys that is specified in Section 3.3.

Line 449: “For each level, the same LMS and LM-OTS parameters set **shall** be used for every LMS tree at that level.”

For clarity, it may be worth explicitly stating that different levels may use different LMS and LM-OTS parameters; e.g., they are allowed to have different tree heights. However, as mentioned above, the verifier cannot check whether this statement has been adhered to.

NIST Response: A sentence has been added explicitly noting that different tree levels may use different LMS and LM-OTS parameters as long as the same hash function is used at every level.

Line 452: “The parameters n , w , ls , m , and h specified in the tables are defined in Sections 4.1 and 5.1 of [2].”

There is an unhelpful (and potentially dangerous) conflict of notation between the use of w in XMSS, where the Winternitz chains have length w , and in LMS, where they have length 2^w . If the parameters are not explained, then there should at least be a warning that they represent different things for the two schemes. Similarly, there is a conflict of notation between the use of h in HSS, where it is the height of the trees in a single level, and in XMSS^{MT}, where it is the total height of the hypertree.

NIST Response: Section 2.3 has been extended to include definitions of these variables, highlighting instances in which a single variable has one meaning in RFC 8391 and a different meaning in RFC 8554.

Line 459: *Table 1*

Although the signature lengths for LM-OTS are taken directly from RFC 8554, they are rather misleading when taken in isolation, as the one-time signature scheme will never be used by itself. It would be useful to have a separate table listing the public key and signature sizes for the different LMS parameters.

NIST Response: The signature length column has been removed from Tables 1, 3, 5, and 7.

Line 516: *“For the parameter sets in this section, the functions F , H , H_{msg} , and PRF are defined as follows.”*

In RFC 8391, the SHA-256 and SHA-512 parameter sets pad the key so that it completely fills a SHA-2 message block for F , H and PRF , or two SHA-2 message blocks for H_{msg} . If the same approach is used for the truncated SHA-256/192 parameter sets, then the functions should be defined as:

$$\begin{aligned} F(KEY, M) &= T_{192}(\text{SHA-256}(\text{toByte}(0, 40) \parallel KEY \parallel M)) \\ H(KEY, M) &= T_{192}(\text{SHA-256}(\text{toByte}(1, 40) \parallel KEY \parallel M)) \\ H_{msg}(KEY, M) &= T_{192}(\text{SHA-256}(\text{toByte}(2, 56) \parallel KEY \parallel M)) \\ PRF(KEY, M) &= T_{192}(\text{SHA-256}(\text{toByte}(3, 40) \parallel KEY \parallel M)) \end{aligned}$$

In the current draft the text reads “*toByte(i, 4)*”, representing integer i only in 4 bytes.

NIST Response: We are aware that this is how the padding was specified for SHA-256 and SHA-512 in RFC 8391. However, we chose not to follow that convention for the 192-bit hash functions. Using the smaller amount of padding allows $F(KEY, M)$ to be computed with just one iteration of the compression function, resulting in faster computation.

Line 545: *“For the parameter sets in this section, the functions F , H , H_{msg} , and PRF are defined as follows.”*

In RFC 8391 it is explained that although a shorter identifier could be used with SHA3, n bytes are used for consistency with the SHA2 implementations. The draft appears to stick with this convention in the case where $n = 32$, in lines 533 to 536, so it is recommended that the functions defined in lines 547 to 550 pad their identifiers to 24 bytes.

NIST Response: As noted in the previous comment, the padding length was set in RFC 8391 so that the padded prefix and the key completely fills a message block. It is just coincidence that the amount of padding needed to do this is n bytes. As was done in RFC 8391, the padding for SHAKE256/192 was set to be consistent with the amount of padding specified for SHA256/192.

Line 566: *“If more than one LMS instance is being created (e.g., for an HSS instance), then a separate key pair identifier I , and $SEED$ (if using the pseudorandom key generation method) **shall** be generated for each LMS instance.”*

The previous paragraph of Section 6.1 mandates the use of the pseudorandom key generation method.

NIST Response: This was an editing mistake. The parenthetical, “(if using the pseudorandom key generation method)” should have been deleted.

Line 569: “When generating a signature, the n -byte randomizer C (see Section 4.5 of [2]) **shall** be generated...”

The LM-OTS signatures are not deterministic because of the randomizer C . Therefore, if a leaf node on a higher level signs a root node on a lower level more than once, the resulting signatures will be different, which could allow an attacker to forge signatures. Section 6 of RFC 8554 implicitly addresses this issue by stating that “It is expected that the above arrays are maintained for the course of the HSS key.” NIST should make storage of these arrays a requirement, or propose an alternative, deterministic, signing method.

NIST Response: Noted. The final version of the document has been modified to specify that each one-time key may only be used to sign a message once, even in the case where multiple levels of trees are implemented in a single cryptographic module and the OTS key is being used to sign the root of a lower-level tree.

Line 576: “The private n -byte strings in the WOTS+ private keys ($sk[i]$ in Section 3.1.3 of [1]) **shall** be generated using the pseudorandom key generation method specified in Section 3.1.7 of [1].”

There is a serious flaw in the pseudorandom key generation process described in RFC 8391 and mandated in the NIST draft. The private key value $sk_{i,j}$ for one-time signature instance j is derived from the private seed $seed_j$ via

$$sk_{i,j} = PRF\left(seed_j, toByte(i, 32)\right)$$

The private key index i acts as the address for the PRF, but this address does not depend on the index j of the one-time signature. Consequently, after observing r one-time signatures there is a multi-target attack that recovers a private seed with around $2^{8n+4}/r$ calls to the PRF. This reduces the classical security of XMSS and XMSS^{MT} with tree height h by around $h - 4$ bits.

Expanding on the above, suppose we observe a WOTS+ signature $S = (s_0, s_1, \dots, s_{l-1})$ on the message $M = (m_0, m_1, \dots, m_{l-1})$, where we implicitly include the checksum. For most values of i the message word m_i can be viewed as a uniformly random element of $\{0, 1, \dots, w - 1\}$, with the obvious exception being the most significant word of the checksum. The probability that $m_i = 0$, and so the probability that s_i reveals the private value sk_i , will be $1/w$.

Now suppose that we observe r WOTS+ signatures S_1, S_2, \dots, S_r on the messages M_1, M_2, \dots, M_r . For a fixed i we expect r/w of the messages to have $m_{i,j} = 0$, so we expect the r signatures to reveal r/w private values; that is, we expect there to be r/w values where

$s_{i,j} = sk_{i,j}$. In general, we can choose the index i that reveals the most private values, which will be higher than r/w , but not significantly so.

Because the r/w private values all have the form $sk_{i,j} = H(seed_j || i)$ for private one-time seeds $seed_j$ and a fixed index i , we can try to guess a seed by choosing a putative n -byte value $seed'$, computing $sk' = H(seed' || i)$, and then comparing sk' with our r/w target values $sk_{i,j}$. The probability that our guess will match one of the targets is $r/w2^{8n}$, so we would expect to recover one of the seeds after $w2^{8n}/r$ guesses.

For XMSS, the Winternitz parameter is always chosen to be $w = 2^4$. Given a tree of height h , the maximum number of one-time signatures that can be observed is $r = 2^h$. Consequently, the attack requires 2^{8n+4-h} guesses.

A possible fix to this attack is to adopt the addressing method used for XMSS^{MT} in the NIST PQC Round 2 SPHINCS+ submission.

NIST Response: Thank you for pointing out this weakness in the pseudorandom key generation offered as an example in RFC 8391. Just as all other aspects of XMSS and LMS were designed to protect against multi-target attacks, we agree that the key generation mechanism should protect against such attacks.

As suggested, we reviewed the key generation method specified in the NIST PQC Round 2 SPHINCS+ submission, which uses $sk[i] = \text{PRF}(\text{SK.seed}, \text{ADRS})$. This construction does provide better multi-target protection, since ADRS has a different value for each chain in a XMSS or XMSS^{MT} key. We note, however, that ADRS does not identify the XMSS or XMSS^{MT} key itself. As a result, this construction does not fully protect against multi-target attacks, if the attacker has access to signatures created using many different XMSS or XMSS^{MT} keys.

In order to fully protect against multi-target attacks, the final version of SP 800-208 specifies that the private strings shall be generated using $sk[i,j] = \text{PRF}_{\text{keygen}}(S_XMSS, \text{SEED} || \text{ADRS})$. A new function, $\text{PRF}_{\text{keygen}}$, was defined since RFC 8391 defines PRF as a function that “takes as input an n -byte key and a 32-byte index.”

Line 586: *“Distributed Multi-Tree Hash-Based Signatures”*

The methods described in this section of the draft effectively describe “virtual hypertree” schemes, distributed across multiple hardware cryptographic modules, where no keying material is exported from any module. To use this approach in practice will require a significant amount of supporting software to facilitate communication between hardware modules, keep track of which trees belong to which device, prevent malicious re-routing of requests to inauthentic modules, and other operational requirements.

Consequently, such techniques will be difficult to deploy or use practically. With that in mind, NIST may want to consider relaxing the constraints on exporting private data. Below are some options NIST may consider that would allow for secure key backup and recovery:

- Backup and recovery should happen between two distinct machines that share the same code (e.g., both are HSMs).

- This communication should be supported by a KEM, where the shared secret is ephemeral and securely deleted after one use; this prevents redeployment.
- The state must be deleted from source machine after it has been exported to the other device. This prevents redeployment as well.

NIST Response: Section 7.2.2 provides a high-level description of the functionality that supporting software would need to implement for XMSS^{MT}. This does not seem to be a significant amount of supporting software. In a typical implementation, there *may* be two cryptographic modules (one for the top-level tree and one for a bottom-level tree) connected to the external device (desktop computer) during key generation. As described in Section 7.2.2 the amount of “communication” between the two modules would be very limited – submitting the public key from one module to the other module to be signed. During normal operations, only a single cryptographic module (holding a bottom-level tree) would be connected. There is no need to keep track of which trees belong to which device, as the signature itself indicates from which tree it was generated. It is unclear why the commenters believe that the supporting software would need to “prevent malicious re-routing of requests to inauthentic modules.”

Line 620: *“Distributing the implementation of an XMSS^{MT} instance across multiple cryptographic modules requires each cryptographic module to implement slightly modified versions of the XMSS key and signature generation algorithms provided in [1].”*

Distributing HSS across multiple cryptographic modules is reasonably straightforward, as each intermediate signature is an independent instance of LMS. However, in XMSS^{MT} the intermediate signatures are instances of a reduced variant of XMSS, which are all implicitly viewed as being part of the same hypertree of total height h ; e.g., the hash function addresses are given in terms of their locations in this hypertree.

The method of distributing XMSS^{MT} across multiple cryptographic modules suggested in Section 7.2 preserves interoperability with RFC 8391 by modifying the standard XMSS key generation and signing algorithms but is significantly more complicated to implement and use. Further, if the process for provisioning a bottom-level cryptographic module fails for some reason (see line 719) then this wastes a valuable signature from the top-level module.

A simpler approach would be to adapt the approach from HSS and use independent instances of (full) XMSS for the intermediate signatures. The disadvantages of doing this are that it would increase the length of the signatures, and the scheme would not be interoperable with XMSS^{MT} as specified by RFC 8391.

Given that NIST is allowing additional parameter sets and hash functions for both HSS and XMSS^{MT}, RFC-compliant implementations may not be able to verify all NIST-compliant signatures. This raises the question of how much interoperability should be preserved? NIST may want to break away from the RFCs entirely and set their own, distinct, requirements.

NIST Response: Noted. As indicated later, line 719 was added as an extra check that the operation worked as expected. The need to try signature generation a second time should be rare. While the failure would result in the “loss” of a one-time key in the top-level module, this should not be a significant problem, as these keys are only used to sign the roots of bottom-level keys. One would expect only a handful of the top-level tree’s one-time keys to be used, and the top-level tree will have at least 32 one-time keys (and will more likely have at least 1024 one-time keys).

Given that there are already two competing standards for stateful hash-based signatures, LMS and XMSS, NIST does not believe that it would be beneficial to create yet another standard in this same space.

Line 641: *“7.2.1 Modified XMSS Key Generation and Signature Algorithms”*

The LMS and XMSS RFCs both contain explicit return statements in their pseudocode, which improves clarity, but the pseudocode in the NIST draft does not. This is particularly confusing in, for example, lines 708 and 712 where assignments are made to public key values using information returned from calls to `XMSS' _keygen`.

It may be worth stating explicitly that Algorithm 10' is a modified version of Algorithm 10 in RFC 8391; the same applies to Algorithm 12'. Similarly, it may be worth stating explicitly that `XMSS^MT external device keygen` replaces Algorithm 15, and that `XMSS^MT external device sign` replaces Algorithm 16.

There is a lack of clarity about where the structure `SigPK` lives in relation to the provisioned cryptographic modules, and whether it needs to be protected.

NIST Response: While the outputs of Algorithms 10' and 12' were already specified on lines 647 and 686, explicit return statements have been added.

While it does not specifically state “Algorithm 10,” the text at the beginning of Section 7.2 states that the algorithms in Section 7.2.1 are slightly modified versions of the XMSS key and signature generation algorithms provided in RFC 8391. The XMSS key and signature generation algorithms are provided in RFC 8391 as Algorithm 10 and Algorithm 12.

It should be clear from the context that variables used in Section 7.2.2 are stored on the external device, not the provisioned cryptographic modules, as the algorithms described are to be run on the external device.

Line 647: `“Output: XMSS public key PK”`

There may be scope for confusion here, as in RFC 8391 the output of Algorithm 10 is the XMSS public key and the XMSS private key.

NIST Response: As noted in Section 7 and in Section 8, private keying material cannot be exported. So, the private key is not an output. It is generated within the cryptographic module and remains there.

Line 651: `“wots_sk[i] = WOTS_genSK();”`

In RFC 8391, `WOTS_genSK()` (as described in Algorithm 3) sets each element of `wots_sk` to a uniformly random n -byte string, but the NIST draft mandates the use of the pseudorandom key generation method described in Section 3.1.7 of RFC 8391. This has the potential to cause confusion as the `WOTS_genSK()` function requires access to a uniformly random n -byte string S that should be stored as part of the private key.

NIST Response: Algorithm 10' has been modified to explicitly mention the generation of S_{XMSS} and to describe the method for pseudorandomly generating the one-time keys, rather than just referring to $WOTS_genSK()$.

Line 679: `“SK = L || t || idx || wots_sk || SK_PRF || root || SEED”`

No terminating semicolon. The same comment applies to lines 681, 683, 696, and 697.

This definition also conflicts with the use of “setter methods” in lines 657, 669, and 670.

NIST Response: The terminating semicolons have been added. It is unclear in what way the definition conflicts with the use of “setter methods.”

Line 683: `“PK = OID || root || SEED”`

The format of the OID is not defined in RFC 8391, and it is not entirely clear how it relates to the identifiers in Section 5 of the NIST draft. There may be some confusion between the identifiers for XMSS and XMSS^{MT} as they appear to overlap.

NIST Response: This line has been clarified by adding a comment that refers to the XDR syntax for `xmssmt_public_key` in Appendix C.3 of RFC 8391.

Line 719: `“if (getIdx(SigPK[t]) ≠ t) {”`

This should be a `while` loop rather than an `if` statement. This process probably deserves more detailed explanation in the surrounding text.

NIST Response: The “If” statement has been changed to a “while” statement. However, a single failure of this test should be very uncommon. If a second attempt fails, this is likely an indication of a serious problem with the top-level signing module. A footnote has been added to the comment preceding this line to provide even more explanation for the reason that this step may be needed.

Line 729: `“// Send XMSS'_sign() command to one of the bottom-level key pairs”`

In the example XMSS^{MT} signing algorithm described in RFC 8391, when one bottom-level key pair is exhausted a new key pair is generated automatically for the next signature. The method of external device operations presented in Section 7.2.2 suggest that the bottom-level cryptographic modules are provisioned first during key generation, and then one of the available modules is chosen for use during each signing call. In practice, there will likely need to be a mechanism for switching between modules and dynamically re-provisioning them when their key pairs have been exhausted.

NIST Response: Re-provisioning when key pairs have been exhausted is not recommended, as any newly generated key bottom-level key pair would have to be signed by the top-level key. If the module holding the top-level key has failed, this will be impossible.

In most cases, we expect each cryptographic module that holds bottom-level keys to be pre-provisioned with enough one-time keys to generate all signatures that may need to be created over the life of the HBS key pair. This could be a single bottom-level key with 2^{10} or 2^{20} one-time keys, or a relatively small number of bottom-level keys that each has 2^{10} or 2^{20} one-time keys. For example, in the case of signing firmware updates, it is highly likely that the total number of firmware images that would need to be signed using a single HBS key would be far less than 1000.

Line 815: *“The faulted signature remains a valid signature, so checking that the signature verifies is insufficient to detect or prevent this attack.”*

The faulted signature is highly likely to be valid, but it depends where the fault occurs. If it is during one of the hash function calls that needs to be recomputed for verification, then the signature will not be valid.

NIST Response: The section in the document discussing fault injection resistance has been removed.

Line 816: *“The only reliable way to prevent this attack is to compute each one-time signature once, cache the result, and output it whenever needed.”*

There are alternative mitigations. For example, one approach is to use redundancy: compute the full signature twice, compare the results and only release a signature if the results match; an attacker would need to induce two identical faults in order to obtain an exploitable signature.

NIST Response: The section in the document discussing fault injection resistance has been removed.

Line 841: *“The randomized hashing process does not, however, impact the ability for a signer to create a generic collision since the signer, knowing the private key, could choose the random value to prepend to the message.”*

It is not entirely clear why this discussion is included, since, as pointed out on line 851, this should not really be considered an attack on the signature scheme. Randomised hashing is intended to prevent someone other than the signer preparing a pair of colliding messages; see, for example, the discussion in NIST SP 800-106. This is only a threat if the values r in RFC 8391 and C in RFC 8554 are not sufficiently random.

NIST Response: In most cases it is not an issue if the private key holder is able to generate two messages that have the same signature, which is why preventing this is not required to achieve EUF-CMA security. However, there are some limited cases, such as commitment protocols, in which there is a need to ensure that there is a one-to-one relationship between signatures and messages. The verifier is presented with a signature, which is intended to represent the signer’s “commitment” to a particular message. If a signer could create two messages with the same signature, the signer could cheat in such a protocol.

Line 844: *“The 196-bit hash functions in this recommendation...”*

They are 192-bit hash functions.

NIST Response: Corrected. Thank you.

Line 898: “union lmots_signature switch”

The indenting of the case statements is inconsistent.

The same comment holds for case statements beginning on lines 947 and 982.

NIST Response: Corrected. Thank you.

Line 1452: *“However, in the current version of XMSS^{MT} [1], the security analysis differs somewhat. In the standard model, [17] shows that XMSS^{MT} is EUF-CMA. Further, [16] shows that XMSS^{MT} is post-quantum existentially unforgeable under adaptive chosen message attacks with respect to the QROM.”*

Appendix C.4 somewhat overstates the provable security results for XMSS^{MT}. The standard model result by Malkin et al in [17] holds for a general signature framework which covers both XMSS^{MT} and HSS. It shows that hierarchical signature schemes are secure provided that the underlying one-time signature schemes are secure, but with a significant tightness gap.

The tight QROM proof by Hülsing et al from [16] does not apply to XMSS^{MT} as described in [1]. Firstly, the result from [16] requires an assumption about the hash function family F that is almost certainly not satisfied by any NIST approved cryptographic hash function; a recent paper presented by Bernstein and Hülsing at ASIACRYPT 2019 replaces this with a brand-new security notion which they call *(multi-target) decisional second-preimage resistance* and which they believe should be difficult to attack. Secondly, the scheme analysed in [16] differs from the version of XMSS^{MT} described in [1] in a few important details; for example, the method for generating one-time private keys in [16] involves the address of the one-time signature, which prevents the attack described above.

NIST Response: The reference to “Mitigating Multi-Target Attacks in Hash-based Signatures” has been replaced by a reference to “The SPHINCS+ Framework.”

Line 1469: *“The main difference between these schemes’ security analyses comes down to the use (and the degree of use) of the random oracle model or quantum random oracle models.”*

It is also arguable that the complexity of the security reduction and the number of assumptions involved are also important. A simpler argument gives more confidence in the correctness of the result.

NIST Response: Noted. We agree that this is generally true, but do not feel the security reductions for either scheme are so complex that they would influence our confidence in either scheme’s security in this case.

NSA's Center for Cybersecurity Standards

From: Sharon Ehlers

Date: Monday, February 24, 2020, at 1:13pm

Comments for SP 800-208.

- The option of using SHA384 or SHA512 could be useful.

NIST Response: Using SHA-256 or SHAKE256 with a 256-bit output already provides security comparable to AES-256. We believe this provides a sufficient level of security for any use case.

- The parameter sets for LMS and XMSS use similar but different notation and this could cause some confusion. For example, w has two different meanings between the two schemes and SEED is a private value in LMS and a public value in XMSS. Consider making these differences clear.

NIST Response: Section 2.3 has been extended to provide definitions of the variables from LMS and XMSS that are referenced in SP 800-208, highlighting instances in which the same variable is used in both RFCs, but with different meanings.

- Section 7.1, page 20 line 618: Unable to find an Algorithm 9 in [2].

NIST Response: This was a typo and should have referred to Algorithms 7 and 8 rather than 8 and 9.

- Sections 7.2.1 and 7.2.2:
 - Calls to XMSS' sign need to know to which module it's being sent so layer/tree can be tracked in the external device keygen and external device sign.

NIST Response: Section 7.2 only describes an algorithm that will work with a two-layer tree. In this case, it is important to distinguish between the top-level key and the bottom-level keys, but it is not necessary to keep track of which bottom-level key is which. When a message is signed, the signature includes an index value, idx_sig . As noted in Section 7.2.2, the high-order bits of idx_sig identify which XMSS key (tree) was used to create the signature. Even in the case of the top-level key, the cryptographic module should reject attempts to sign messages that aren't exactly n bytes in length, which should reduce the likelihood of an ordinary message being accidentally signed using the top-level key.

If a user were to attempt to implement XMSS^{MT} with more than two layers, with each layer being implemented in a separate cryptographic module, this would impose somewhat more of a requirement on the user to keep track of which key is which.

- Lines 716-723: It is not clear what the purpose of this if statement is. Please Clarify.

NIST Response: A footnote has been added to provide additional explanation for the reason for this “if” statement (now a “while” statement).

- Line 732: The definition of t is misleading. In the RFC, it is $h-(h/d)$ most significant bits of idx_sig . Here, since $d=2$, $t=h/d$ most significant bits is correct, but using $t=h-(h/d)$ or $t=h/2$ most significant bits would be clearer. Furthermore, the definition from the RFC, $t=h-(h/d)$ most significant bits of idx_sig , is misleading as well.
If idx_sig has exactly h bits, this is fine, but idx_sig has $\text{ceil}(h/8)$ bytes, which is not always h bits. In that case, the definition of t might not be grabbing the intended bits of idx_sig . This definition comes up in the XMSS^{MT} sign and verify algorithms.

NIST Response: The text in line 732 has been modified to compute t in an unambiguous manner.

- p26: 196's should be 192's

NIST Response: This has been corrected. Thank you.

Crypto4A

From: Jim Goodman

Date: Monday, February 24, 2020 at 3:25pm

Crypto4A's Comments on NIST SP800-208 Draft Specification

Crypto4A's comments are provided in two distinct parts: first we provide editorial comments regarding the draft's proposed language, and then we provide comments regarding the concepts being proposed within the draft itself.

Editorial Comments

First, our editorial comments:

- **Line 266:** replace “some but not all of” with “some, but not all, of”

NIST Response: Draft SP 800-208 underwent an editorial review prior to being published, and the copy-editor requested removal of the commas that this comment proposes to insert.
- **Line 268:** consider adding references for SHA-256 and SHAKE256 (i.e., [3] and [5] respectively)
NIST Response: Accepted.
- **Line 280:** change “is firmware” to “is authenticating firmware”

NIST Response: Accepted.
- **Line 342:** consider changing “public keys.” to “public keys using a Merkle tree construction.”

NIST Response: Declined. The Merkle tree construction is what is being referred to by “a method.” Section 3.2 explains what that method is.
- **Line 348:** consider deleting “, as follows”

NIST Response: Declined. Without “as follows,” readers may be confused, believing that they are expected to already know how a digest is signed using a hash chain. Including “as follows” assures the reader that this will be explained later.
- **Line 358:** consider changing figure title to “A sample Winternitz chain for $b = 4$ ”

NIST Response: Accepted.
- **Line 376:** fix formatting to avoid CRLF's in $H^{*i}(x_j)$ elements in the figure

NIST Response: Noted. This error was introduced when the document was converted to PDF, and an updated version of the draft that only corrected this formatting problem was posted a few days after the initial version was posted.

- Line 385-386:** consider changing “value, which will” to “value at the root of the tree, which will”

NIST Response: Resolved by adding “the root of the tree” in parenthesis.
- Line 389:** consider changing “public keys.” to “public keys ($k_i, i \in [0, 7]$).”

NIST Response: Accepted.
- Line 390:** consider changing “the tree.” to “the tree ($h_j, j \in [0, 7]$).”

NIST Response: Accepted.
- Line 391:** consider changing “the tree.” to “the tree (i.e., h_{01}, h_{23}, h_{45} , and h_{67}).”

NIST Response: Accepted.
- Line 419:** change “different values” to “different prefix values”

NIST Response: Accepted.
- Line 436:** the symbol for XORing x_k and the bitmask looks an awful lot like some form of multiplication, perhaps there’s a more “XOR-like” symbol that could be used instead?

NIST Response: The XOR symbol has been changed to \oplus .
- Line 489:** change “functions is specified” to “functions are specified”

NIST Response: Accepted.
- Line 502:** change XMSS-SHA2_20_256 entry’s Numeric Identifier from “0x00000002” to “0x00000003”

NIST Response: Thank you. This was a typo and it has been corrected in the final version.
- Line 518:** change “toByte(0, 4)” to “toByte(0, 24)” (or perhaps you’d prefer to stay with 32?)
- Line 519:** change “toByte(1, 4)” to “toByte(1, 24)” (or perhaps you’d prefer to stay with 32?)
- Line 520:** change “toByte(2, 4)” to “toByte(2, 24)” (or perhaps you’d prefer to stay with 32?)
- Line 521:** change “toByte(3, 4)” to “toByte(3, 24)” (or perhaps you’d prefer to stay with 32?)
- Line 547:** change “toByte(0, 4)” to “toByte(0, 24)” (or perhaps you’d prefer to stay with 32?)
- Line 548:** change “toByte(1, 4)” to “toByte(1, 24)” (or perhaps you’d prefer to stay with 32?)
- Line 549:** change “toByte(2, 4)” to “toByte(2, 24)” (or perhaps you’d prefer to stay with 32?)
- Line 550:** change “toByte(3, 4)” to “toByte(3, 24)” (or perhaps you’d prefer to stay with 32?)

NIST Response: The choice of 4 rather than 24 or 32 was intentional, not an editorial mistake. When using SHA-256/192, padding the prefix to only 4 bytes allows for the function F() to be computed with one iteration of the compression function rather than 2, thus saving time. The amount of padding for SHAKE256/192 was set to be consistent with the amount of padding for SHA-256/192.

- **Line 587:** consider changing “of time and” to “of time, and”

NIST Response: Draft SP 800-208 underwent an editorial review prior to being published, and the copy-editor did not propose inserting a comma in this location.

- **Line 683:** consider adding additional line after 683 that states “return (PK)”

NIST Response: Accepted.

- **Line 685:** consider changing “Message M” to “Message M, XMSS private key SK”

NIST Response: Declined. A fundamental difference between Algorithms 10' and 12' in SP 800-208 and the algorithms in RFC 8391 is that the private key is generated and stored within the cryptographic module, and it can never be exported. So, the private key is never an input or an output, it is always stored internally.

- **Line 686:** consider changing “signature Sig” to “Updated SK, XMSS signature Sig”

NIST Response: Declined. See previous comment.

- **Line 703:** consider adding additional line after 703 that states “return (SK || Sig)”

NIST Response: Resolved by adding a line that states “return Sig”. “SK” was not included for the reason stated previously.

- **Line 907:** consider adding additional space at start of line for proper alignment
- **Line 915:** consider adding additional space at start of line for proper alignment
- **Line 952:** consider adding additional space at start of line for proper alignment
- **Line 958:** consider adding additional space at start of line for proper alignment
- **Line 961:** consider adding additional space at start of line for proper alignment
- **Line 964:** consider adding additional space at start of line for proper alignment
- **Line 995:** consider adding additional space at start of line for proper alignment
- **Line 1279:** consider adding two additional spaces at start of line for proper alignment

NIST Response: Accepted.

Qualitative Comments Regarding Concepts

In addition to the aforementioned editorial comments, we have identified several primary concerns with the document, as well as just some general comments regarding various sections of the document:

- There is no disaster recovery (DR) option given the manner NIST is proposing to generate HBS private keys, and the restrictions you're imposing in Section 8.1. On line 745 you clearly state that the cryptographic module **shall not** allow for the export of private keying material. While we don't expect NIST to have to provide guidance on DR, we also don't believe it should be explicitly precluding options by putting this sort of restriction on the cloning/exporting of HBS private keys. Yes, state management is difficult to do, but processes can be put in place to manage the activity (more on this later), and the benefits of being able to archive keys to avoid having the entire hierarchy come crashing down if the top level HSM were to fail. Your proposed solution attempts to mitigate this by distributing the private key generation across multiple devices such that the top level HSM signs public keys presented by other HSMs (more on this in a later comment) which have generated private keys for lower layers of the hierarchy. This approach is still dependent on the top level HSM being present and operational so that it can sign new public keys as they come online, which could be difficult for a long-lived keying hierarchy. One way to overcome that is to have all of the subordinate HSMs present and accounted for soon after the top level HSM has generated its HBS private key, so that they can all request their public keys get signed before the top level HSM fails. Unfortunately, you're just moving the problem around as now those subordinate HSMs need to survive long enough to carry out their roles as HBS signing authorities, and the amount of capital expenditure to finance the bulk purchase of HSM devices may prove prohibitive. Hence, we think it would be best for NIST to **not** preclude exporting private key materials, but rather focus on devising best practices related to managing the risks associated with that operation, so that operators can devise their own DR solutions.

NIST Response: If HBS private keys could be cloned/exported, then there would be a need to rely on procedural mechanisms to avoid one-time key reuse. In the comments that NIST received last year (<https://csrc.nist.gov/CSRC/media/Projects/Stateful-Hash-Based-Signatures/documents/stateful-HBS-misuse-resistance-public-comments-April2019.pdf>), Adam Langley noted instances in which root CAs issued multiple CA certificates with the same serial number. As noted in https://bugzilla.mozilla.org/show_bug.cgi?id=1405815, it seems that the problem was that the CAs were relying on procedural mechanisms to avoid reusing serial numbers. So, we have to accept that even in environments in which procedures should be carefully followed, mistakes can happen. While NIST acknowledges that some organizations will be able to develop procedures to avoid one-time key reuse and ensure that those procedures will always be carefully followed, there is no assurance that every user of stateful hash-based signatures would do the same.

Section 7 of SP 800-208 was written with the idea that all of the keys (the top-level key and all of the bottom-level keys) would be generated at the same time as part of a key generation ceremony. In most scenarios, we envision that only one of the cryptographic modules holding bottom-level keys would be used for signing, and the rest would be set aside for DR purposes. Modules that are merely sitting in storage should be unlikely to fail, so the total number of cryptographic modules needed shouldn't be too great.

- Over the past 25 years of handling DR principles around critical PKI root keys, we have evolved very strong procedures for the secure extraction and re-injection of critical root key material in

HSMs. This has provided us with a high guarantee of having preserved the integrity, confidentiality and availability of the keys by enforcing the tracking of private key material whether it's within an HSM or some form of secure external storage such as a safe or vault. This was possible as the RSA/ECC keys were complete objects with no additional state that needed to be maintained. Unfortunately, HBS introduces state to the management equation so attempting to distribute HBS private key material across multiple HSMs is tantamount to scattering the private key in both space and time. Hence, the proposed multi-HSM approach for implementing a distributed multitree HBS (Section 7) is concerning to us from a security perspective in its current form. What guarantees does the top level HSM have regarding the validity of the signing request it receives from parties looking to have the public key of the HSS private key they've generated on their HSM devices? Mechanically anyone could present a public key for signing, thereby introducing the possibility of rogue parties now being able to generate valid signatures. In a PKI CA world, they would manage this with revocation to punish the bad actors who managed to fool the CA into signing their illegitimate certificate. In the proposed 2-level HBS scheme there are no such revocation methods to save us after the fact, so we need to do everything we can to prevent this situation from happening. Hence, there needs to be some robust mechanism in place to validate requests BEFORE they are signed, which we have found to be a very difficult problem to solve unless very rigid procedures are put in place to eliminate the possibility (e.g., force the subordinate HSM to be brought into the room where the root HSM is so that the root HSM operators can witness the HSS key generation process and perform some sort of attestation that the HBS public key the subordinate HSM generates corresponds to a private key generated on that subordinate HSM). This is likely to prove to be a very onerous process akin to a full-on traditional root key generation ceremony in a conventional PKI, so this needs to be considered and addressed somehow (e.g., guidance on procedures, introduction of requirements to guarantee attestation of the authenticity of the signing request, etc.).

NIST Response: As noted in this comment and in a later comment, Crypto4A appears to have interpreted the procedures proposed in Section 7 of SP 800-208 as being comparable to establishing a root CA and then using that root CA's key to sign the keys of subordinate CAs. This was never the intention. The procedures in Section 7 are intended to have the same effect as cloning an RSA or (EC)DSA key, but without making copies of one-time keys. Just as the operator of a root CA would not provide a copy of its private RSA or (EC)DSA key to the operator of a subordinate CA, the top-level LMS or XMSS key in the scheme described in Section 7 should not be used to sign a key that is held by some other entity.

As suggested, the generation of the top-level key and of all of the bottom-level keys should be performed at the same time as part of the traditional key generation ceremony. When using RSA or (EC)DSA, copies of the key may be made onto additional HSMs as part of the key generation process, and there is no mechanism available that would allow for revoking one copy of this key without revoking the other copies. There is no more of a need for a special, but standardized, mechanism for revoking some of the one-time keys in an HBS scheme without revoking the entire HBS key.

- The existing hash-sigs github repository that provides a reference implementation for LMS-HSS includes functions to pseudo-randomly generate LMS subtree {I, SEED} values from a master seed value for a given LMS-HSS instance (i.e., `hss_generate_root_seed_I_value()` and `hss_generate_child_seed_I_value()` in `hss.c`), which allows the implementor to optimize the

private key data storage requirements by eliminating the need to store discrete pairs of {I, SEED} for each layer of the tree since we can just recompute them from a single master seed value. This method of pseudo-random value generation for I in particular was identified as an option in RFC 8554 Section 7.1, so we don't believe it represents a security compromise of any proposed solution. Lines 566-568 of Section 6.1 appears to preclude this sort of implementation option by forcing the implementer to generate a separate {I, SEED} pair for each LMS instance. However, this requirement is itself quite vague as you put no requirements on how those values are generated (i.e., can they be pseudo-random or do we need to generate using a random bit generator that supports at least $8n$ bits of security strength)? We would prefer to be able to continue using a pseudo-random method, but if that isn't acceptable then perhaps the language of the requirement can be made more precise to remove the aforementioned ambiguity.

NIST Response: It is unclear why the commenter believes that the requirement is quite vague, or that no requirements are placed on how the values are generated. The text clearly states that the values I and SEED “**shall** be generated using an approved random bit generator [6] where the instantiation of the random bit generator supports at least $8n$ bits of security strength.” This seems quite unambiguous. The specified reference, [6], does include definitions of both nondeterministic random bit generators and deterministic random bit generators, but that does not seem to be the source of the confusion.

NIST does not believe that a cryptographic module would need to store very many LMS keys at any given time, so storing a separate I and SEED value for each LMS key should not be a significant issue.

- Section 6.1 also enforces the requirement that the same SEED value shall be used to generate every private element in a single LMS instance (line 563). We feel this is overly restrictive, and an implementor should be able to use one or more values/SEEDs provided they are generated in a manner that meets the stated security criteria (i.e., using an approved random bit generator where the instantiation of the random bit generator supports at least $8n$ bits of security strength). Relaxing this constraint opens up the possibility of proposing novel DR-compatible solutions, one of which we describe below.

NIST Response: If all of the one-time keys are to be generated and stored on the same cryptographic module, then there does not seem to be any benefit in using multiple different SEED values. While allowing such an option may not introduce a security vulnerability, the key generation process needs to be fully specified so that it is testable by the Cryptographic Algorithm Validation Program (CAVP), and allowing for additional flexibility, such as proposed here, would make the compliance testing very difficult.

- Would NIST consider a mechanism whereby the top-level LMS instance (we're applying things to LMS-HSS in the interest of simplicity, but the comments should extend to XMSS/XMSS^{MT} as well) is sectorized into cryptographically-isolated segments, each of which shares the same I value but which has its own SEED value that was generated using a manner similar to the pseudo-random generation of LMS-OTS private keys (but using a unique format to ensure it doesn't collide with that pseudo-random process, or any of the processes used in hash-sigs to generate {I, SEED} pairs, and which can't be used to guess another sector's SEED value). Sectorization would segment the 2^h leaves of the top-level tree into 2^s groups (a.k.a., sectors), each containing 2^{h-s} leaves. Each sector's SEED value allows a device to generate signatures

from that sector's set of leaves and NOT any other sectors' leaves. Hence, you have cryptographically-enforced state reuse protection if you assign different sectors to different cryptographic modules (i.e., HSM_i can't generate valid signatures from the sector assigned to HSM_j). However, the sector generation process can ensure that all sectors share the top-level public key value, so all sectors are part of the same HBS signing authority. These sectors can then safely be exported from the top-level HSM and stored in a secure fashion using the same techniques and procedures that have been proven over the years to handle the secure extraction and handling of any regular private keys so that they can be loaded onto other HSMs (once and only once) when needed (e.g., the existing HSM(s) fail and we need to recover the signing capability for the given HBS public key, we use up all of the existing allocated sectors' signatures and need to load new sectors into the HSM many years down the road, or we want to load unique sectors into multiple HSMs in parallel to allow higher signing throughput). We believe this will yield a feasible means of providing DR for HBS on HSMs (albeit with potential over-allocation of the total tree size in order to accommodate the redundancies that facilitate DR). Note that this approach can be used to create a one-layer tree with OTS keys being created and stored on different HSMs as per the request made in the paragraph on lines 143-146 within the Note to Reviewers section. In that use case, each sector would be loaded into a different HSM, where the resulting unique SEED values would facilitate the generation of unique OTS keys on each device.

[NIST Response: The proposal in this comment involves generating one-time keys on an HSM and then exporting some of those keys to other HSMs. As noted previously, allowing the export of one-time keys creates a risk that one-time keys could be reused. It might be possible to develop a mechanism that avoids having multiple copies of a one-time key by always erasing the one-time key from the source module when the key is transferred to another module. However, in order for such a mechanism to be safe, in addition to ensuring that any one-time key exported by a cryptographic module could not later be used by that same module, there would need to be a way to ensure that any exported keys could not be imported into more than one module. \(As noted previously, we do not believe it is sufficient to rely on procedural means to ensure that copies of one-time keys are not copied onto multiple modules, after which they are used more than once.\) While such a protocol could be designed, it would not be possible for the exporting cryptographic module to ensure that the protocol was being followed correctly, which means that it would not be possible for the Cryptographic Module Validation Program \(CMVP\) to validate that a module's implementation prevents one-time key reuse. The only reliable way for the CMVP to ensure that one-time keys cannot be reused is to validate that they cannot be exported at all.](#)

- An additional note on revocation as per the proposed 2-level scheme described in Section 7. Our interpretation is that the subordinate cryptographic modules are generating a single certificate that verifies back to the primary cryptographic module's top-level public key. In a typical PKI the root CA would sign a subordinate CA's public key, generating a certificate for that subordinate CA public key that the user/application could validate. In the proposed approach we'd have the root CA (i.e., top-level CM) sign the subordinate CA's (i.e., subordinate CM) public key, but that result would just appear as part of any HSS/XMSS^{MT} signature the subordinate CA generates (i.e., the first LMS/XMSS signature component that precedes the subordinate CA's public key element, and LMS/XMSS signature on the message). Hence there is no discrete certificate that could be checked and revoked. Furthermore, if another subordinate CA has been stood up, and it hasn't been compromised, then will it be affected as a consequence of revoking the other subordinate CA given it shares the same root CA public key as all other

subordinate CA's in this stratified approach, and we don't have a discrete top-level certificate to use to achieve finer-grained revocation. We've kicked around ideas related to atypical revocation mechanisms based on longest prefix-matching against portions of the HBS and its components, but these are all custom hacks that don't lend themselves well to a standardization effort. How does NIST envision revocation working with the proposed 2-level scheme? Is it an all-or-nothing sort of thing?

NIST Response: As noted previously, the scheme described in Section 7 should not be compared to the establishment of a 2-level PKI. The top-level LMS or XMSS key should never be used to sign a key belonging to a different entity. The top-level and bottom-level keys should be thought of as one (HSS or XMSS^{MT}) key. If a root CA makes multiple copies of its RSA key for disaster recovery purposes, and one of the copies of the key becomes compromised, the "entire" RSA key must be revoked. There is no way to revoke just the compromised copy of the key. If a root CA were using the scheme in Section 7, then generating a bottom-level LMS or XMSS key should be thought of as being the equivalent of making a copy of the root CA's RSA private key, with the security consequences being the same if that key is compromised.

- A general comment regarding Figure 4, and the differences between HSS and XMSS^{MT}: Figure 4 shows the top-level tree being marked as level 0, with the level value increasing as we progress from top-to-bottom of the multi-level tree. This approach is fine for HSS, where a similar numbering convention is used, but in XMSS^{MT} we believe the standard numbers the top-level tree as level (d-1) and proceeds to decrease the level value as we progress from top-to-bottom. This may lead to confusion later on, and we think the difference merits some form of mention in the text.

NIST Response: In Figure 4 (now Figure 5), the 0's and 1's representing the levels were replaced by *a*'s and *b*'s in order to avoid using the notation of one scheme versus the other.

- The description/pseudocode for XMSS^{MT} external device key generation is confusing to us. Under what conditions would the IF statement in line 719 evaluate to true given the generation calls on Lines 712 and 715, thereby necessitating us to essentially repeat the generation calls using lines 721 and 722 respectively? Would the given code not just adjust the incorrect *t* value by at most 1 given the correction is not iterative, but just a one-off? This confusion is somewhat compounded by what seems to us to be under-specified inputs/outputs for Algorithms 10' and 12' in section 7.2.1. which are used extensively in Section 7.2.2.

NIST Response: It is certainly hoped that the "If" statement on line 719 would not evaluate to true. However, a tree index value, *t*, needs to be specified as an input to the key generation process on line 712, while the caller to the signing process on line 715 cannot be allowed to specify which one-time key to use to perform the signing. Since Algorithm 12' specifies that the signer should iterate through its one-time keys, the caller should be able to know which one-time key will be used for signing and thus specify that same value for the input on line 715. As noted in the comment on lines 716-718, line 719 is simply a check to ensure that the expected one-time key was used to sign the bottom-level key's root. If, for some unknown reason, that did not happen, then the key cannot be used, and so one needs to try the key generation again.

The "if" statement has been replaced with a "while" statement, and a footnote has been added to provide more explanation for why the key generation may need to be repeated.

- In Appendix A and Appendix B, the text indicates we're extending the XDR syntax for [2] and [1] respectively, but the subsequent descriptions in Lines 859-1002 and Lines 1007-1391 read like they are the entire XDR specifications. Would it make sense to add comments into the XDR elements to remind the reader that you're supposed to also include all existing XDR specification code into each definition? For example, for LMS-OTS algorithm type (`lmots_algorithm_type`), add a new line between Lines 861 and 862 that says something along the lines of `/* includes all existing lmots_algorithm_type values */` or some similar language to remind the reader that existing definitions are retained as well.

NIST Response: Adding the proposed comment to each data structure seems excessive. However, a sentence was added to the beginning of each appendix.

The QuantumRISC Project

From: Marc Stöttinger

Date: Friday, February 28, 2020, at 3:43am

Dear Author team of the document SP800-208,

as consortium members of the German nationally funded research project “QuantumRISC”, we would like to provide you feedback on the draft NIST Special Publication 800-208 (SP 800-208).

The QuantumRISC project is funded by the German Federal Ministry of Education and Research (BMBF) and brings together partners from both academia and industry. The project partners jointly develop and improve post-quantum secure cryptographic schemes for low-end devices with severe limitations on memory usage and power consumption while maintaining a high level of security. The practical implementation of such schemes highly depends on their operability on embedded devices. The main focus of the project is the development of quantum secure solutions for the automotive domain; however, research findings will be transferable to other domains and use cases. We investigate the interaction between existing vehicle systems and architectures as well as the integration of PQC into the vehicle while allowing a future exchange of cryptographic primitives (crypto agility).

The project consortium consists of the following partners: Continental AG, Elektrobit Automotive GmbH, Fraunhofer Institute for Secure Information Technology SIT, RheinMain University of Applied Sciences, MTG AG, Ruhr-University Bochum and Technical University of Darmstadt.

We have the following three feedback comments to the current draft version:

- 1) Past experience has shown that developers find it difficult to deploy cryptography if the specifications are distributed among different standards or if ambiguous representations exist (e.g. RSA parameters with explicit NULL or empty). In order to improve interoperability and to be able to use algorithms between different applications, object identifiers and standardized representations of public keys are necessary. Therefore, object identifiers (OID) should be specified for the two algorithms XMSS and LMS and for the signatures and public keys. Public keys should be uniquely represented in ASN.1 to make it possible to issue interoperable certificates that contain public XMSS or LMS keys.

For example, a public key could be represented as:

```
SubjectPublicKeyInfo ::= SEQUENCE {  
    algorithm      AlgorithmIdentifier,  
    subjectPublicKey BIT STRING }
```


The 'algorithm' field could specify an OID and an explicit statement regarding the parameters and the 'subjectPublicKey' field could provide a concrete specification of the encoded public key (e.g., a 1-to-1 mapping to the specifications of the RFC). With SP 800-208, there is a chance to specify OIDs and representations in one document to facilitate the use of XMSS and LMS.

NIST Response: NIST agrees that it is important to clearly specify encodings and identifiers (e.g., OIDs), but does not believe that SP 800-208 is the appropriate place to specify this information. Just as the underlying signature schemes are defined in IETF publications RFC 8391 and RFC 8554, work is underway in the IETF to specify how to use these signature schemes in various protocols (see RFC 8708, <https://tools.ietf.org/html/draft-ietf-cose-hash-sig>, <https://tools.ietf.org/html/draft-vangeest-x509-hash-sigs>, and <https://tools.ietf.org/html/draft-murciple-ssh-xmss>).

- 2) What is the reason that XMSS and LMS variants are not harmonized to provide parameter sets with the same tree heights? Different usage scenarios have different requirements and more flexibility for the maximum number of signatures should be provided. Hence, we would like to see similar parameter sets for XMSS and LMS with respect to the tree heights and ideally with a smaller step size in the tree height in order to choose a number of 2^5 , 2^8 , 2^{10} , 2^{15} , 2^{16} , 2^{20} , 2^{25} , 2^{32} , 2^{40} , ... signatures. Alternatively, the tree height could not be specified in the parameter set but freely chosen (in a certain range) for each key pair.

NIST Response: The differences in the parameter set options between LMS/HSS and XMSS/XMSS^{MT} in SP 800-208 are the result of their differences in RFC 8391 and RFC 8554. We do not believe there is a substantial benefit in defining new parameter sets for LMS/HSS or XMSS/XMSS^{MT} just for the sake of making the parameter sets of the two schemes more similar. Allowing for any tree height to be chosen would make testing more difficult, and there does not seem to be a compelling need. In practice, working with a tree of height 10 instead of one of height 8 would not result in substantially longer signatures, or signing and verification times. The additional tree height would also not significantly increase the amount of memory required to implement the schemes.

- 3) SP 800-208 references RFC 8391, which also provides a description of the XMSS algorithm. Alongside the RFC document, there is also a C reference implementation of XMSS. We note that each of these documents provides different algorithm definitions. For example, algorithm 10 in SP 800-208 and algorithm 10 in RFC 8391 both specify the XMSS key generation; yet they provide different implementations. Though the algorithms are semantically identical, a uniformly standardized basis of algorithms would likely prevent misunderstandings and implementation flaws. Similarly, the implementation of algorithms in the C reference implementation does not follow the pseudocode from RFC 8391. For example, algorithm 2 (WOTS Chaining) is defined recursively in the RFC but implemented iteratively in the reference code. Having a unified definition of algorithms throughout the provided documents would presumably ease understanding and implementation.

NIST Response: Algorithm 10' in Section 7.2.1 of SP 800-208 does not specify XMSS key generation, so it is not semantically identical to Algorithm 10 in RFC 8391. Algorithm 10' in

Section 7.2.1 of SP 800-208 is intended for use in implementing XMSS^{MT} in a distributed manner. Algorithm 10 in RFC 8391 could not be used for this purpose, as it only works for XMSS (i.e., it assumes that $L=0$, $t=0$, and $d=1$).

NIST had no involvement in the development of either RFC 8391 or the C reference implementation. However, RFC 8391 explicitly notes that the algorithms in that document are written for simplicity, not efficiency, and so recommends against using them in implementations.

Best regards,

Marc Stöttinger

Stefan-Lukas Gazdag

From: Stefan-Lukas Gazdag

Date: Friday, February 28, 2020, at 12:50pm

Hi,

thanks to NIST for all the great work regarding the PQC standardization process! Please find enclosed some comments on draft SP 800-208.

We (genua GmbH) provide hybrid signatures (ECDSA and XMSS) for our latest software updates. Both signatures have to be verified as valid, otherwise the update is rejected. Key generation and signing is done on a secure key server. Authorized build servers in a restricted development network may ask for a signature via an OpenSSH connection. First updates have been applied to machines in the field. We look forward to HBS being used more widely by others.

Open topic: OIDs

For the use in practice (explicitly taking a look at X.509 certificates) object identifiers (OIDs) are needed. This far there are no OIDs defined by any organization (neither by any agency, corporation, university or the IETF/IRTF). Without going into details about former discussions on who should publish OIDs I just want to raise awareness that this should be dealt with. Software using HBS so far uses "temporary" or private OIDs (that have somewhat been agreed on between some software projects) or use software specific identifiers.

[NIST Response: OIDs have been specified in RFC 8708 for LMS/HSS. While this is not sufficient, NIST believes that it would be more appropriate for other standards organizations, such as the IETF, to specify OIDs and encoding formats rather than NIST.](#)

Line 273-275:

Yet another peculiarity is that you should choose a proper parameter set suiting your specific use case (e. g. which signature size is still ok, while maintaining a specific security level). This also means how many signatures will be written as the key has a limited life-time. Whereas classical keys have an implicit life-time (forced by a validity date or due to the need of increasing the security level due to advances in supercomputing, cryptanalysis, ...), for HBS maybe a small key writing e.g. a million keys would be enough (or may be exchanged in time) for a specific use case while other scenarios would require a huge multi-level tree. All in all decisions that have to be made beforehand in a different way than with classical schemes.

[NIST Response: While the fixed limit in the number of signatures that can be created is a difference from traditional signature schemes, the concern in this section is mainly the security risk resulting from the need to maintain state. We believe that for applications that have the characteristics described in Section 1.1, the number of signatures that will need to be created will be relatively small, so it shouldn't be too difficult to choose a parameter set that can generate enough signatures.](#)

Line 278/279:

I'd argue that using HBS now is important in many other, probably most use cases of software updates and code signing. History shows that software runs for way longer in the field than often expected as users stick to their running systems. Thus old systems are likely to be found running pre-quantum update mechanisms once a large enough quantum computer exists. Therefore it is recommendable to apply HBS now to existing systems even it is "just" to ensure a proper transition to other quantum-safe signature schemes later on. Also implementing and distributing update mechanisms using hybrid signatures now might help having somewhat modular mechanisms where exchanging a single scheme might be easier.

NIST Response: The referenced text specifically notes that stateful HBS may be applicable if the implementation will have a long lifetime and it would not be practical to transition to a different digital signature scheme once the implementation has been deployed. In most cases of software update, we believe it would be practical to update the software update mechanism itself in order to be able to work with a newer digital signature scheme.

While it is true that not all systems are updated as quickly as they should be, this is becoming less of an issue as automatic updating becomes more prevalent. In addition, while we cannot predict when quantum computers that will be capable of breaking current signature algorithms will become available, it is likely that it will be possible to start deploying updated software update mechanisms with new (stateless) post-quantum signature algorithms years before the current signature algorithms become insecure.

Line 436:

Please use the \oplus symbol for exclusive-or

NIST Response: This has been corrected in the final version. Thank you.

Line 502:

The correct numeric identifier of XMSS-SHA2_20_256 is 0x00000003

NIST Response: This has been corrected in the final version. Thank you.

Line 587:

Not the most sophisticated solution but practicable: as the public keys for all the schemes are quite small, a specific device or software might be provided with several HBS public keys.

NIST Response: Thank you for noting this. The final version of SP 800-208 notes that this is an option.

Line 641 and following:

Some pseudo-code lines are missing semicolons. Also, sometimes setter / getter methods are used as in the RFC but sometimes they are omitted

NIST Response: This issue has been corrected in the final version of the document.

Line 647:

Algorithm 10' / XMSS'_keyGen should also output the secret key SK

NIST Response: As noted in Sections 7 and 8, cryptographic modules conforming to SP 800-208 cannot export private keying material. So, the cryptographic module cannot output SK. The cryptographic module must generate and store SK, but must not output it.

Line 774 and following:

In some use cases performance might improve by the reservation approach described in [8], which we've tried in practice. Reserving an interval of OTS keys, meaning writing an updated secret key according to the interval chosen to non-volatile memory before signing alleviates performance issues in practice. In case of any interrupt, some OTS keys stay unused, which in most scenarios should not be a problem with somewhat stable cryptographic modules / key servers.

NIST Response: We believe that for the applications described in Section 1.1, signing will be infrequent, and so reserving multiple OTS keys at a time will not be beneficial.

Line 844:
s/196/192/

NIST Response: This has been corrected in the final version. Thank you.

Kind Regards,
Stefan-Lukas Gazdag

Canadian Centre for Cyber Security

From: David E. Smith

Date: Friday, February 28, 2020, at 3:58pm

Please find below our editorial and technical comments on the Draft SP 800-208 issued for comment in December 2019.

David Smith

Canadian Centre for Cyber Security

Line	Type	Comment
288, 775, 809	Technical	<p>Starting at line 288, "If an attacker were able to obtain digital signatures for two different messages created using the same OTS key, then it would become computationally feasible for that attacker to forge signatures on arbitrary messages".</p> <p>Similarly, starting at line 775 and line 809 "...this is acceptable since it just involves using an OTS key multiple times to sign the same message".</p> <p>Comment: Per Section 6.1, 9.3 and [2], it seems that in LMS the OTS generates a random prefix for every message to be signed (Algorithm 3 in Section 4.5 of [2]). In particular, a forgery would be possible given two distinct signatures even if they were for the same message. Also, it would not be acceptable to use an OTS key multiple times, even for the same message, unless the random prefix was forced to be the same. XMSS also generates a random prefix before signing, but it appears to be deterministically derived from the private key and signature index (Algorithm 12 of Section 4.1.9 of [1]), so signing the same message with the same OTS would result in the same signature.</p> <p>NIST Response: The final version of SP 800-208 has been modified to forbid using an OTS key more than once, even to sign the same message a second time (e.g., recomputing the signature on the root of a lower-level tree). This avoids the potential mistake described here and also protects against fault injection attacks.</p> <p>Replace "checksum is computed as $\sum_{k=0, n-1} (b-1-N_k)$" with "checksum is computed as $\sum_{k=0, n-1} (b-1-N_k)$, which requires $\text{ceil}(\log_b(n*(b-1)))$ digits".</p>
368	Editorial	<p>NIST Response: We believe this level of detail is not needed for the high-level description. The checksum formula itself may be unnecessary, but it was included in case it helps some people understand the sentence that follows: "The checksum is designed so that the value is non-negative and any increase in a digit in the message digest will result in the checksum becoming smaller."</p>
389	Editorial	<p>Replace "Figure 3 depicts a hash tree containing eight OTS public keys." with "Figure 3 depicts a hash tree containing eight OTS public keys k_0, \dots, k_7".</p>

NIST Response: Accepted.

Replace SHA2 with either SHA-256 (to match earlier in the draft) or SHA2-256 (to match [1]).

506 Editorial

NIST Response: Noted. The current text aligns with RFC 8391 [1], which says: “To implement the keyed hash functions, the following is used for SHA2 with $n = 32$ ”

Panos Kampanakis

From: Panos Kampanakis

Date: Friday, February 28, 2020 at 4:49pm

Dear Quynh, NIST,

I would like to provide some more feedback regarding the SP 800-208 Draft for HBS after discussing with some of our HSM peers implementing HBS. They pointed out to us some practical concerns:

- 1) Section 8.1 mandates that private keys should not be extractable. Today HSMs allow for extracting a classical private key using some Shamir sharing scheme so that key can be reconstructed and reused in case of an HSM failure. I don't think LMS is different. In a hierarchical scenario where a top level HSM signs subordinate LMS trees, the top HSM would need to survive for a long time (30 years for a traditional CA root) in order to be able to sign any new subordinate tree coming online. That may not always be practical. We should allow for the OTS private keys to be extractable using similar methods (Shamir secret sharing or so) so someone could reconstruct the top HBS tree and sign new messages in case of failure.

NIST Response: As noted in the response to Crypto4A, Section 7 of SP 800-208 was written under the assumption that all of the operational and spare cryptographic modules would be instantiated with keys at the beginning so that the cryptographic module holding the top-level key would not be a single point of failure.

While there may be many mechanisms available to protect traditional private keys (e.g., secret sharing) while allowing the keys to be copied, these mechanisms do not apply to stateful HBS, since they cannot address the state management issue. As long as copies of keys can be made, there is the risk that the same key (with the same state) will be loaded onto more than one cryptographic module and that this will result in one-time key reuse. While we understand that many organizations would be able to put procedural mechanisms in place to prevent this one-time key reuse, past experience has shown that we cannot rely on all organizations that may choose to use stateful HBS to be able to do this reliably.

- 2) Section 6.1 requires a separate I and SEED value for each LMS instance. If someone wanted to generate I with a PRF he should be able to, so that the subtrees of a hypertree can be generated by using a master value instead of storing separate (I, SEED) pairs for each tree in the hypertree. Generating I in a deterministic pseudorandom could point to SP800-90A.

NIST Response: Noted. The requirement for both I and SEED is that they “**shall** be generated using an approved random bit generator (see the SP 800-90 series of publications [6]) where the instantiation of the random bit generator supports at least 128 [8n] bits of security strength” This includes the option to use any of the deterministic random bit generators from SP 800-90A.

- 3) Section 6.1 requires one SEED per LMS tree. By allowing more SEED values, HSMs can use them to be able to generate non-overlapping sections of the tree in order to prevent

state reuse in a DR scenario. Using different SEEDs in some of the LM-OTS leaves does not compromise the security of LMS tree.

NIST Response: Noted. However, as long as all of the SEEDs are generated in the same cryptographic module and cannot be exported, there is no security benefit in allowing the use of more than one SEED. Allowing this option, though, would complicate testing key generation under the Cryptographic Module Validation Program (CAVP).

Panos
Cisco

From: Panos Kampanakis
Date: Friday, February 28, 2020 at 10:26pm

We would also like to propose for the SP to include the following parameters that are suitable for all our (Cisco and probably many more vendor) image signing usecases

~~~~~

- LMS\_SHA256\_M16\_H5 with LMOTS\_SHA256\_N16\_W8
- LMS\_SHA256\_M24\_H5 with LMOTS\_SHA256\_N24\_W8
- LMS\_SHA256\_M32\_H5 with LMOTS\_SHA256\_N32\_W8
  
- LMS\_SHA256\_M16\_H10 with LMOTS\_SHA256\_N16\_W8
- LMS\_SHA256\_M24\_H10 with LMOTS\_SHA256\_N24\_W8
- LMS\_SHA256\_M32\_H10 with LMOTS\_SHA256\_N32\_W8
  
- LMS\_SHA256\_M16\_H15 with LMOTS\_SHA256\_N16\_W8
- LMS\_SHA256\_M24\_H15 with LMOTS\_SHA256\_N24\_W8
- LMS\_SHA256\_M32\_H15 with LMOTS\_SHA256\_N32\_W8
  
- LMS\_SHA256\_M16\_H20 with LMOTS\_SHA256\_N16\_W8
- LMS\_SHA256\_M24\_H20 with LMOTS\_SHA256\_N24\_W8
- LMS\_SHA256\_M32\_H20 with LMOTS\_SHA256\_N32\_W8
  
- HSS (with 2-4 levels) with any of the above LMS trees at any level.

~~~~~

For $N=M=16$ we realize that that would provide 64-bit PQ security, but given NIST's stance with AES-128 (Grover not being parallelizable and thus AES-128 is considered secure) we could use it when needing very small signatures at acceptable security.

Thank you,
Panos
Cisco Systems

NIST Response: We understand that the Post-Quantum Cryptography standardization effort is considering digital signature algorithms that offer Level 1 security, which is defined as security comparable to that offered by AES 128. As noted in the [Post-Quantum Cryptography FAQs](#), NIST believes it is likely that AES 128 (and other algorithms that provide a comparable level of security) will remain secure for decades to come. However, the FAQs also suggest that if quantum computers turn out to be much less expensive than anticipated, then AES 128 may not be sufficiently secure, whereas AES 192 and AES 256 would continue to be safe. In the unlikely event that this were to happen, NIST would issue guidance regarding any transitions of symmetric key algorithms and hash functions that may be needed.

As noted in Section 1.1 of SP 800-208, NIST believes that stateful hash-based signatures are primarily intended for applications that have long lifetimes and for which transitioning to a

different signature scheme would not be practical. For applications that have such properties, which may need to remain secure for decades and which would be unable to follow any potential future transition guidance, we believe it is appropriate to be more conservative in specifying acceptable security levels.

Google

From: Stefan Kölbl, Roy D'Souza

Date: Friday, February 28, 2020, at 6:18pm

Google's Comments on the NIST SP800-208 Draft Specification

Stefan Kölbl, Roy D'Souza

February 28, 2020

Google anticipates deployment of post-quantum hash-based signature schemes for verified boot, and over-the-air updates, for a range of hardware modules. These modules vary significantly in available power, computational capabilities and related resources.

When deciding between stateless and stateful schemes, for scenarios that are amenable to the larger signature sizes of stateless schemes we would leverage a NIST-recommended scheme, such as the anticipated SPHINCS+. Whereas for other contexts, where it is an imperative to limit signature sizes, we would deploy a NIST-recommended stateful scheme such as LMS/HSS.

Deployment Scenarios

The following three deployment scenarios would most likely be constrained to usage of a stateful scheme:

- **Google Security Chips:** All Chromebooks are deployed with an embedded Google Security Chip that is candidate for being a quantum-ready hardware root of trust. It would probably have computational abilities similar to an ARM Cortex M3, with limited memory and flash.
- **Battery Operated IoT Sensors:** These include sensor devices such as Nest Detect, the motion and perimeter sensors used by the Nest Guard secure alarm system. This class of devices has the resource constraints of the previous category, and also needs to operate on the equivalent of an AAA battery for over two years.
- **Powered IoT Devices and Chromebooks:** These are powered devices based on Intel/AMD and ARM chips, and these lower cost devices have space and other resource constraints that would benefit from compact signatures.

Our choice of stateful hash-based standardization candidates is LMS/HSS, and the following two categories of parameters would be important for addressing the resource constraints of the scenarios outlined above.

Variable (Sub-)Trees

It would be beneficial to have different parameters depending on the level of a multi-tree. The cryptographic modules at a lower level might be deployed in more constrained environments,

while a higher-level tree, perhaps belonging to a more trustworthy third party, could afford more expensive computations.

The cadence of firmware updates to devices, even within each category, could differ significantly. A Chromebook might be updated every six weeks, while some IoT devices might only be updated occasionally. Therefore it would be useful to have a choice of parameters for LMS/HSS:

- LMS_SHA256_M24_H5 with LMOTS_SHA256_N24_W8
- LMS_SHA256_M32_H5 with LMOTS_SHA256_N32_W8

- LMS_SHA256_M24_H10 with LMOTS_SHA256_N24_W8
- LMS_SHA256_M32_H10 with LMOTS_SHA256_N32_W8

- LMS_SHA256_M24_H15 with LMOTS_SHA256_N24_W8
- LMS_SHA256_M32_H15 with LMOTS_SHA256_N32_W8

- LMS_SHA256_M24_H20 with LMOTS_SHA256_N24_W8
- LMS_SHA256_M32_H20 with LMOTS_SHA256_N32_W8

- HSS (with 2-4 levels) with any of the above LMS trees at any level.

[NIST Response: Thank you for the feedback. All of the above parameter sets are included in SP 800-208. Given the limited feedback that was provided about parameter sets, none of the parameter sets listed in the draft version of SP 800-208 were removed.](#)

Security Targets

In the ongoing NIST post-quantum cryptography standardization process five security levels have been defined and the proposed schemes seem to fall into NIST security level 3 and 5, as they do not rely on the collision resistance of the underlying hash function.

In some of our scenarios it might be useful to have variants of LMS/XMSS that target NIST security level 1, as this would provide security comparable to ECDSA with P-256 or Ed25519, while still providing a buffer against quantum adversaries given the limitations of Grover's algorithm (e.g., limited parallelization or that the quantum circuit of the hash functions will be fairly large). Introducing new variants with $n = 16$ would reduce the signature size for the OTS by over 50%:

- LMOTS_SHA256_N16_W1: 2196 bytes
- LMOTS_SHA256_N16_W2: 1108 bytes
- LMOTS_SHA256_N16_W4: 580 bytes
- LMOTS_SHA256_N16_W8: 308 bytes

[NIST Response: We understand that the Post-Quantum Cryptography standardization effort is considering digital signature algorithms that offer Level 1 security, which is defined as security comparable to that offered by AES 128. As noted in the \[Post-Quantum Cryptography FAQs\]\(#\),](#)

NIST believes it is likely that AES 128 (and other algorithms that provide a comparable level of security) will remain secure for decades to come. However, the FAQs also suggest that if quantum computers turn out to be much less expensive than anticipated, then AES 128 may not be sufficiently secure, whereas AES 192 and AES 256 would continue to be safe. In the unlikely event that this were to happen, NIST would issue guidance regarding any transitions of symmetric key algorithms and hash functions that may be needed.

As noted in Section 1.1 of SP 800-208, NIST believes that stateful hash-based signatures are primarily intended for applications that have long lifetimes and for which transitioning to a different signature scheme would not be practical. For applications that have such properties, which may need to remain secure for decades and which would be unable to follow any potential future transition guidance, we believe it is appropriate to be more conservative in specifying acceptable security levels.