

1 NOTE: This is a draft for review by the IEEE Security in Storage Working Group

2 To: EncryptionModes@nist.gov

3 From: Matthew V. Ball, Sun Microsystems, Chair of IEEE Security in Storage Working Group (P1619)

4 Subject: Follow-up to NIST's Consideration of XTS-AES as standardized by IEEE Std 1619-2007

5 Date: April 12, 2009 (Draft 3)

Table of Contents

Introduction.....	3
The XTS algorithm itself.....	3
Comments From Moses Liskov and Kazuhiko Minematsu.....	3
Comments From Vijay Bharadwaj and Neils Ferguson of Microsoft.....	3
Comment From Phillip Rogaway of Univ. of CA, Davis.....	5
Comment From Boaz Shahar of Entropic Communications.....	6
Comments From Michael Willett of Seagate Technology.....	7
The depth of support in the storage industry for which it was designed.....	8
Current Industry Support.....	8
Hard Disk Vendors.....	8
AMCC (from Paul Von Stamwitz).....	8
Bloomberg Technologies (from Harry Xiang).....	8
Brocade (from Scott Kipp).....	8
CipherMax (from Steven Tan).....	9
EMC (from David Black).....	9
Emulex (Larry Hofer).....	9
Freescale Semiconductor (from Michael Torla).....	9
GED-i (from Rony Shapiro).....	9
Helion (from Colin Sinclair).....	10
Hifn (from Doug Whiting).....	10
Oxford Semiconductor (from Gary Calder).....	10
PMC-Sierra (from Gary Nichols and Peichen Chang).....	10
SafeNet (from Marcel van Loon).....	10
Thales (from Bob Lockhart).....	10
TrueCrypt.....	10
WinMagic (from Gary McCracken).....	10
Comment from David Clunie.....	10
Comments From Michael Willett of Seagate Technology.....	11
The appeal of XTS for wider applications.....	13
Comment From David Clunie.....	13
Comments From Michael Willett of Seagate Technology.....	13
The proposal for the approved specification to be available only by purchase from IEEE.....	14
Comment From Vijay Bharadwaj and Neils Ferguson of Microsoft.....	14
Comment from Rich Schroepel.....	14
Comment From David Clunie.....	14
Comments From Michael Willett of Seagate Technology.....	14
General Response.....	14
Concerns of intellectual property rights.....	15
Comments From Vijay Bharadwaj and Neils Ferguson of Microsoft.....	15
Comment From David Clunie.....	15
Comments From Michael Willett of Seagate Technology.....	15
General Response.....	16
Investigation NeoScale Systems' IP Claims.....	16
IP Claims by Rogaway.....	17
Conclusions and Recommendations.....	17
References.....	17

1 Introduction

2 This is a follow-up to comments previously submitted to NIST concerning the adoption of the XTS-AES mode of
3 operation under FIPS 140 (currently at 140-2). This response is organized in the categories recommended by
4 NIST in the *Request for Public Comments on XTS* (see [Request]):

- 5 • The XTS algorithm itself;
- 6 • The depth of support in the storage industry for which it was designed;
- 7 • The appeal of XTS for wider applications;
- 8 • The proposal for the approved specification to be available only by purchase from IEEE;
- 9 • Concerns of intellectual property rights.

10 The following sections are presented in the form of comment/response pairs. Some comments were omitted if
11 they did not appear to need a response.

12 The XTS algorithm itself

13 **Comments From Moses Liskov and Kazuhiko Minematsu**

14 See [Liskov] for the full comments from Liskov and Minematsu. For brevity, I've only included the introduction
15 here.

16 *Overall, we believe that the XTS-AES algorithm, closely based on Rogaway's XEX mode [5] plus ciphertext*
17 *stealing, is a good choice for the purpose of block-oriented data storage encryption, and the use of an algorithm of*
18 *this type is well supported by research publications. We have two main criticisms of the publication. First, while*
19 *XEX uses one key, the proposed XTS algorithm uses two keys; Key1 is used to encipher the whitened plaintext,*
20 *while Key2 is used to compute the pre- and post-whitening values. We feel that only one key should be used, to*
21 *serve both purposes. Second, the draft incompletely analyzes the security of XTS-AES; it needs correction and*
22 *expansion in a couple of areas.*

23 Liskov and Minematsu continue by showing the security of one-key XTS-AES and suggest a sketch for a proof
24 that shows the security of ciphertext stealing within XTS-AES. Their recommendation is to allow one-key XTS-
25 AES.

26 **Comments From Vijay Bharadwaj and Neils Ferguson of Microsoft**

27 (see [MiscRsp])

28 ***Unclear Security Goals***

29 *In our opinion, one serious shortcoming of the proposal is that it does not contain a clear statement of what*
30 *application-level security goals XTS aims to achieve. This makes it difficult to analyze the proposal, or to*
31 *determine how widely applicable XTS may be. Appendix D contains some hints, and the XEX paper referred to*
32 *contains some low-level goals. However, it is extremely difficult for a security practitioner or system implementer*
33 *to understand what assurances can be provided with this mode. This differs from other recent standards such as*
34 *SP800-38D, which contain clear explanations of this sort.*

35 **Response**

36 While I believe that IEEE std 1619-2007 does provide a good discussion on the security of XTS, it is true that we
37 could have added more details about all the applications that are suitable for XTS. This is a somewhat
38 subjective area, and it is ultimately up to the designers of the cryptographic system to determine whether this
39 mode is appropriate, given their specific design requirements.

40 ***Temporal effects***

1 *The proposal appears to miss the effect of temporal effects on the security of XTS. It is possible for an attacker to*
2 *observe a disk for a period of time and thereby gain a significant advantage in cryptanalysis. For instance, on a*
3 *disk with 4 KB blocks where each block is a data unit, an attacker who observes approximately 4000 writes to a*
4 *given block will have obtained access to 2^{20} cipher blocks with the same tweak and key. Similarly, an attacker who*
5 *steals a 500 GB encrypted disk may get access to about 1 TB of ciphertext if they were also able to recover the*
6 *previous contents of each sector. Thus, it seems that the limits for key reuse given in Section 5 and Appendix D can*
7 *be reached fairly easily in practice. We believe that the proposal lacks the margin of safety that would be expected*
8 *of a mode which is to be used for 20-30 years.*

9 **Response**

10 *Traffic analysis (i.e., monitoring temporal effects) is a weakness of many encryption modes, and for XTS is*
11 *addressed in IEEE std 1619-2007, section D.2. However, this attack does not seem practical. To monitor*
12 *snapshots over time, the attacker would need a large local disk for storing the snapshots, or a high-bandwidth*
13 *connection for sending this data to the attacker's own disk array. An attacker with this level of access probably*
14 *has direct access to the plaintext and need not bother with traffic analysis. Ultimately, it is up to the system*
15 *designer to decide whether there is exposure to traffic analysis, and to choose the appropriate mode.*

16 ***An attack on large data units***

17 *The attack in D.4.2 can be extended. The attacker finds two positions i and j such that $P_i \oplus C_i = P_j \oplus C_j$. Let*
18 *us assume they are in the same data unit. With reasonable probability this is because $PP_i = PP_j$ and*

19 $CC_i = CC_j$. *This allows the attacker to compute*

20 $P_i \oplus P_j = PP_i \oplus T_i \oplus T_j \oplus PP_j = T_i \oplus T_j = (T \otimes \alpha^i) \oplus (T \otimes \alpha^j) = T \otimes (\alpha^i \oplus \alpha^j)$ *and a simple finite-field*
21 *division recovers the master tweak of the data unit T . This in turn leads to a knowledge of all tweaks in the data*
22 *unit. This makes the ciphertext manipulation attacks much easier. In particular, it allows the attacker to choose the*
23 *value for a subset of the bytes in each plaintext block being manipulated.*

24 *This shows that large data units significantly weaken the system. The standard should not allow data units larger*
25 *than the recommended 2^{20} blocks.*

26 **Response**

27 *Clearly, allowing large data units is dangerous. The working group discussed this limit at length and decided to*
28 *leave it as a 'should' instead of a 'shall' because it is up to the designers to determine the amount of risk they are*
29 *willing to accept for a collision and leak of the master tweak in the data unit.*

30 *That said, it is quite reasonable to limit FIPS-approved applications to 2^{20} blocks within a data unit.*

31 ***Ciphertext manipulation attacks***

32 *AES in XTS mode works with 16-byte blocks, and this allows for very fine-grained ciphertext manipulation attacks.*
33 *We believe this is a significant problem in practice.*

34 *As discussed in section D.2 and D.3 any transparent storage encryption scheme allows ciphertext manipulation*
35 *attacks. In general, the attacker can perform a number of manipulations to the ciphertext in order to influence the*
36 *decrypted plaintext, but is limited by the block size S . In particular, the attacker can randomize the plaintext of any*
37 *block by changing the ciphertext. The attacker's objective is to do this in such a way that some part of the plaintext*
38 *is changed to a value of his choosing. We look at this in two specific contexts: code modification and data*
39 *modification.*

40 ***Code modification***

41 *In a code modification attack the attacker randomizes a block of code and tries to corrupt the code in such a way*
42 *as to introduce a security hole in the system whilst keeping the system functional. Possibilities for the attacker*
43 *include corrupting one of the access control functions in the code such that it omits certain checks, and corrupting*
44 *an input validation routine such that it does not clean up untrusted input. Such a function might look like this:*

45 ***C prototype:***

46 *Bool IsOperationAllowed(...);*

1 *Assembler code:*
2 *IsOperationAllowed:*
3 <function prolog> (*pushes, setting up the stack frame, etc.*)
4 ...
5 *Mov eax,[...] // register eax is used during the computations*
6 ...
7 <set eax to return value>
8 <function epilog> (*pops, revert stack frame, return*)

9 *The attacker chooses a storage encryption block of bytes in the function implementation with the following*
10 *properties:*

- 11 • *When the CPU executes the first instruction in the block the eax register is nonzero (with high probability)*
- 12 • *The block does not contain the epilog code or other code that the system uses.*

13 *The attacker now replaces the first few bytes of the ciphertext block and hopes that the randomized plaintext*
14 *decrypts to a value such that the first instruction executed in the block is a jump instruction to the epilog code. On*
15 *the x86 a relative jump instruction is 2 bytes long and can jump up to 128 bytes forward. The probability of getting*
16 *the right plaintext is one in 2^{16} ; high enough for a practical attack. The small block size of XTS-AES makes this*
17 *attack rather easy. A larger block size makes it significantly harder. First of all, larger blocks have fewer starting*
18 *points, so it is harder to find a block inside this function where the eax register is nonzero when the first instruction*
19 *in the block is executed. Second, the larger block forces the attacker to randomize far more code. If the block size*
20 *is large enough, then the block extends far beyond the jump range of a 2-byte jump instruction so the attack ends*
21 *up jumping to some randomized plaintext. Thus, with wider blocks an attacker is much more likely to cause a crash*
22 *rather than a security compromise.*

23 **Data modification**

24 *A modern operating system has thousands of settings that are important for the security of the system. Each of*
25 *these settings is an attack target. The attacker tries to find a block of ciphertext that, when randomized, has a*
26 *reasonable chance of changing one such setting to an insecure value. To change a value in most data formats the*
27 *'header' has to be left unchanged (as the software will perform some consistency or sanity checks on the header)*
28 *but some of the contents has to be changed. This requires that there is a block boundary between the header and*
29 *the contents. Again, a larger block size significantly increases the security in two ways. There are fewer block*
30 *boundaries that can be exploited, and randomizing a block damages more of the overall data storage and*
31 *significantly increases the chance of triggering a software crash/failure, rather than creating a security hole.*

32 *Data modification can also be applied to data storage applications like a database. A small block size allows*
33 *targeted manipulations of just a few values in a single row of the database; a large block size severely restricts the*
34 *attacker and greatly increases the chances of damaging other (necessary) parts of the data structure.*

35 **Response**

36 IEEE Std 1619-2007 acknowledges in D.3 that the attacker has greater malleability with 16-byte narrow block
37 encryption modes like XTS, versus wide-block encryption modes like EME-2 or XCB (see IEEE P1619.2).
38 However, by making this trade-off, the encryption algorithm need only make one pass over the data unit instead
39 of the two passes required by wide-block encryption modes. It is important for the designers to understand this
40 trade-off and pick the appropriate mode. XTS is appropriate if the attacker has limited access in making
41 malicious attacks and if the programs have a high probability of detecting a randomized 16-byte block.

1 **Comment From Phillip Rogaway of Univ. of CA, Davis**

2 (see [MiscRsp])

3 *XTS is the two-key version of my XEX construction (Asiacrypt 2004), but operating in ECB mode and with*
4 *ciphertext stealing for any short final block. While I have no serious objection to NIST approving XTS by reference,*
5 *I would like to make a couple of points.*

6 *First, it is unfortunate that there is nowhere described a cryptographic definition for what security property XTS is*
7 *supposed to deliver. When the data unit (sector) is a multiple of 128 bits, each 128-bit block within the data unit*
8 *should be separately enciphered as though by independent, uniformly random permutations. That part is clear. But*
9 *what security property does one expect for partial final blocks? One might hope, for example, that the final 128+b*
10 *bits ($b < 128$) would likewise be enciphered as if by a strong PRP. That would seem to be the cleanest natural*
11 *notion, and it's not too hard to achieve. But [XTS] does not achieve such an aim (because, for example, C_m*
12 *does not depend on P_m). One is left to wonder if ciphertext stealing actually works to buy you any strong*
13 *security property for the final 128+b bits.*

14 **Response:**

15 Concerning the security of ciphertext stealing, please see the security sketch offered by Liskov and Minematsu
16 in [Liskov]. While this is not an absolute proof, it does suggest that ciphertext stealing is sufficiently secure for
17 applications that use XTS. Also note that NIST has proposed incorporating a similar ciphertext stealing mode for
18 CBC (see [CBC-CS]). Despite lacking a formal security proof, ciphertext stealing still has general approval in the
19 cryptographic community.

20 *Second, I would like to express the opinion that the nominally "correct" solution for (length-preserving)*
21 *enciphering of disk sectors and the like is to apply a tweakable, strong PRP (aka wide-blocksize encryption) to the*
22 *(entire) data unit. That notion is strong, well-studied, easy to understand, and readily achievable. There are now*
23 *some 15+ proposed schemes in the literature for solving this problem. If NIST approves the "lite" enciphering*
24 *mode that is XTS, this should not be understood to diminish the utility of standardizing a (wide-blocksize) strong*
25 *PRP. In the end, because of its much weaker security properties, I expect that XTS is an appropriate mechanism*
26 *choice only in the case that one simply cannot afford the computation or latency associated to computing a strong*
27 *PRP.*

28 **Response:**

29 I agree with this sentiment, and I have heard interest in later submitting modes such as EME-2 or XCB from
30 IEEE P1619.2. However, there are many applications, like internal hard disk encryption, that cannot afford the
31 extra hardware or latency required by a two-pass wide-block encryption mode. These applications are still
32 secure with XTS because the attacker has limited access to the ciphertext.

33 **Comment From Boaz Shahr of Entropic Communications**

34 (see [MiscRsp])

35 *The XTS algorithm is using a multiplication by ALFA of $GF(2^{128})$ on each round. This multiplication over $GF(2^{128})$*
36 *is defined in little endian notation (See section 5.2 in the proposed standard). I see two difficulties with this:*

- 37 1. *This impose some extra complexity on implementation, as it is much easier to implement multiplication over GF*
38 *with Big Endian notation, where the multiplication reduces to simple left shift and xor; in some cases; And:*
- 39 2. *Usually NIST is using Big endian. For instance, SP800-38B about CMAC mode for authentication which is using*
40 *exactly the same operation is using Big Endian notation.*

41 **Response:**

42 The P1619 task group chose little-Endian order because this is optimal for software written on little endian
43 systems, which at this time is the most popular byte-ordering scheme, being on consumer grade x86 systems.
44 The performance boost is significant on these systems because there is no need to spend cycles to convert the
45 byte-ordering to the native format.

- 1 In hardware, the performance is the same, although the Verilog or VHDL may be less intuitive.
- 2 See threads on P1619 e-mail reflector for more details on the specifics behind this design choice.

3 **Comments From Michael Willett of Seagate Technology**

4 (see [Seagate1])

5 *Weaknesses of XTS*

- 6 1. *Unchanged data layout: unnecessary restriction. Data encryption SW, External encryption modules, Encrypting*
7 *controllers, and Encrypting devices can all:*
 - 8 1. *reserve space on the medium*
 - 9 2. *relocate blocks*
 - 10 3. *dissect/merge blocks, etc.*
- 11 2. *Much available metadata is left unutilized:*
 - 12 1. *in SW: file name, access date, file type, attributes, etc.*
 - 13 2. *in SCSI: protection info (checksum, data owner...)*
 - 14 3. *in disk drives: age of data, drive ID, track geometry, checksum...*
 - 15 4. *in disk arrays: disk ID, array name*
- 16 3. *Encryption during sequential (DMA) transfer is not optimized (speed, power consumption)*
- 17 4. *Encrypted data in transit needs further protection: Address is in the clear, data is susceptible to traffic analysis*
- 18 5. *Tweaked ECB mode: the same data in the same place is still recognizable, leaks at repeated peeks*
- 19 6. *Small block encryption:*
 - 20 1. *malleability of documents (especially the ones containing different versions)*
 - 21 2. *SW: jump address patch, short data can be changed to desired value with non-negligible probability, with little*
22 *collateral damage*
 - 23 3. *undetected data destruction*
 - 24 4. *DRM info, file fingerprint manipulation*
 - 25 5. *File system manipulation*
 - 26 6. *Hiding data from virus scanners, etc., by just moving it*
- 27 7. *Export control: the encryption module is a high speed, mass encryption device*
- 28 8. *XTS has extra complexity compared to CBC (2nd key, Galois multiplier, two XOR ops). What does it buy?*
 - 29 ○ *Only some protection against copy-and-paste attacks.*
- 30 9. *The XTS standard Annex D.4.3. says that there is a “strong security guarantee as long as the same key is not used*
31 *to encrypt much more than a terabyte of data”. This is not the capacity of the storage device, but the observed*
32 *amount of cipher-text by an adversary. If he can record the encrypted traffic or make repeated snapshots of the*
33 *stored data, this limit can be reached easily at smaller storage devices, too. Therefore, the user must periodically*
34 *re-encrypt large amounts of stored data, or partition the storage space into many bands, each encrypted with a*
35 *different key. It is very expensive in processing time, in power consumption, or in the complexity of securely storing*
36 *and managing many keys, etc.*

37 **Response:**

38 Please see responses to correspondingly number sections below:

- 39 1. I do not see a restriction on data layout imposed by XTS. In IEEE Std 1619-2007, the definition of a
40 data unit states that “a data unit does not necessarily correspond to a physical or logical block on the
41 storage device.” It is possible to provide whatever necessary logical mapping is needed to assign
42 unique tweak values to each data unit.
- 43 2. If there is enough extra metadata for an authentication tag, then the system should use approved
44 authenticated encryption modes, such as CCM, GCM, XTS-HMAC, or CBC-HMAC (see IEEE std
45 1619.1-2007). If there is not enough extra metadata for an authentication tag, then XTS provides a
46 suitable solution for many systems.
- 47 3. I do not see how DMA transfer is not optimized under XTS. XTS can be processed in sequence, just like
48 CBC, CTR, ECB, etc. I could understand this argument for a wide-block mode that requires two passes
49 across the data (increasing buffer requirements and latency), but not for narrow-block encryption.

- 1 4. XTS is not designed to protect encrypted data in transit. XTS is for storage encryption. There are many
2 other good protocols, like ESP, TLS, or FC-SP, that protect data in transit.
- 3 5. This is addressed in the previous discussion on traffic analysis
- 4 6. See previous response to a similar comment by Bharadwaj and Ferguson. Narrow-block encryption
5 does have greater malleability than wide-block encryption, but there is a trade-off between processing
6 time and performance.
- 7 7. U.S. export control for mass encryption devices is under the Bureau of Industry and Security (BIS), not
8 by NIST, and as such I do not see the relevance in this context.
- 9 8. The extra complexity of XTS compared to CBC is not as large as is stated here. An implementation
10 does not need a Galois Multiplier, but can instead use a low-cost linear feedback shift register (LFSR).
11 CBC already has an XOR operation of its own, so XTS adds one XOR operation, not two. XTS currently
12 requires two keys, but I will recommend that NIST allow a one-key version so that the key storage costs
13 are equivalent to CBC. Overall, compared to the cost of an AES cipher, these additions are minimal,
14 and afford many advantages, like parallel operation and decreased malleability by the attacker.
- 15 9. When Annex D.4.3 states that the same key should not encrypt more than 2^{40} blocks (about 16
16 Terabytes – not 1 Terabyte), this is with the extremely conservative probability of a distinguishing attack
17 succeeding with a probability of 1 in 2^{53} , which is extraordinarily conservative. D.4.3 continues by
18 showing that encrypting a petabyte limits the probability of success of an attack to no more than 2^{37} , and
19 encrypting an exabyte limits the probability of success to no more than 2^{17} . This is the guarantee based
20 on the security proof, and the actual probability of a successful distinguishing attack could be lower if
21 further research tightens the security bounds. Also, the consequence of a success in this attack is being
22 able to distinguish XTS from an ideal tweakable block cipher, not to recover the encryption key or
23 plaintext. Essentially, this is the point where the security is roughly equivalent to that of ECB mode.

24 **The depth of support in the storage industry for which it was** 25 **designed**

26 ***Current Industry Support***

27 The industry support for XTS-AES is substantial, although this wasn't reflected in the public comments received
28 by NIST. After talking to representatives of these companies, many thought that sending a letter of support to
29 NIST was unnecessary because the support was already overwhelming.

30 As of March 2009, here is a list of companies and products that support or plan to support XTS-AES encryption:

31 ***Hard Disk Vendors***

32 I have been contacted by at least two major hard disk vendors who stated that a future hard disk drive will
33 support XTS-AES encryption. Other hard disk vendors are planning to support XTS, but are waiting on a
34 decision from NIST before making announcements.

35 ***AMCC (from Paul Von Stamwitz)***

36 *AMCC has a product, the PPC 460SX, that supports XTS. See:*

37 https://www.amcc.com/MyAMCC/jsp/public/productDetail/product_detail.jsp?productID=PPC460SX

38 ***Bloombase Technologies (from Harry Xiang)***

39 *Bloombase Spitfire storage encryption server supports XTS-AES for SAN, DAS, disk and tape encryption as usual*
40 *we do for other regional standard bodies in support of Camellia, SCB2, SSF33, etc.*

1 **Brocade (from Scott Kipp)**

2 *Brocade uses XTS-AES in our new line of encryption products including the Brocade Encryption Switch - a 32 port*
3 *switch, and the FS8-18 blade for the Brocade DCX Backbone. We use XTS-AES to encrypt disk-drive based*
4 *products. These are being sold through several companies such as EMC and NetApp with more to come. To find*
5 *out more about our products, you can go here:*

6 <http://www.brocade.com/products-solutions/solutions/security/products.page>

7 **CipherMax (from Steven Tan)**

8 *CipherMax newest family of CM200D, CM250 and CM500 FC security switch products uses XTS-AES mode for*
9 *disk encryption.*

10 *CM200D is a 1U 16port FC switch with support for disk encryption. The CM250 is a mid-range enterprise class*
11 *FC switch supporting up to 4 x 16 port ELC1624 line cards for disk encryption (XTS-AES) or tape encryption*
12 *(XTS-CCM) that can scale from 16 to 64ports. Similarly, the CM500 is an enterprise class FC switch that uses the*
13 *same ELC1624 line cards for disk or tape encryption supporting up to 16 x ELC1624 cards that can scale from 16*
14 *to 256ports.*

15 **EMC (from David Black)**

16 *EMC PowerPath Encryption with RSA supports AES XTS encryption:*

17 <http://www.emc.com/products/detail/software/powerpath-encryption-rsa.htm>

18 <http://www.emc.com/collateral/software/white-papers/rkmpp-sb-0808-lowres.pdf>

19 *See p.8 of the latter whitepaper.*

20 **Emulex (Larry Hofer)**

21 *Emulex supports XTS-AES for inclusion in future products. We suggest that any proposed additional*
22 *options/changes from P1619 be designated as "allowed".*

23 **Freescale Semiconductor (from Michael Torla)**

24 *XTS can be found in the following shipping Freescale products:*

25 ** PowerQuicc devices:*

26 *o MPC8314E*

27 *o MPC8315E*

28 *o MPC8536E*

29 *and XTS is planned for the following announced but not-yet-released Freescale products:*

30 ** QorIQ devices:*

31 *o P1010*

32 *o P1020*

33 *o P2020*

34 *o P4080*

35 ** DSP devices:*

36 *o MSC8156*

37 *All future QorIQ devices will include XTS.*

1 **GED-i (from Rony Shapiro)**

2 *GED-i uses XTS-AES in its GSA-2000 storage disk encryption product line. GSA-2000 devices provide storage disk*
3 *data encryption and security for iSCSI and FC SAN servers. Our products are certified as interoperable with Sun*
4 *and IBM storage solutions.*

5 *For more information, see <http://www.ged-i.com>*

6 **Helion (from Colin Sinclair)**

7 *Helion have FPGA and ASIC IP Cores available for AES-XTS mode, also providing CBC mode for legacy*
8 *compatibility. We support both key sizes, variants for encrypt-only or full encrypt-decrypt, with optional ciphertext-*
9 *stealing support. Our product pages can be found here:*

10 *<http://www.heliontech.com/storage.htm>*

11 *http://www.heliontech.com/aes_xex.htm*

12 **Hifn (from Doug Whiting)**

13 *Hifn has a chip that is taped out and coming back very soon that includes AES-XTS mode. It's a generic*
14 *crypto/compression lookaside (PCIe) processor.*

15 **Oxford Semiconductor (from Gary Calder)**

16 *Oxford Semiconductor supports XTS-AES for inclusion in future products.*

17 **PMC-Sierra (from Gary Nichols and Peichen Chang)**

18 *PMC-Sierra has both a FibreChannel Protocol Controller with XTS and a SAS Protocol Controller with XTS.*

19 *PMC is currently shipping the Tachyon QE8e+ Quad Port Fibre Channel protocol controller with XTS-AES*
20 *encryption. The primary use model for the encryption function is to protect data written to disks. Tachyon*
21 *products power approximately 70% of the world's disk array systems. Moving forward, PMC plans to continue*
22 *offering XTS-AES functionality in our Fibre Channel and SAS/SATA protocol controllers*

23 **SafeNet (from Marcel van Loon)**

24 *SafeNet provides a family of Silicon IP cores supporting XTS-AES in a range of performances. This IP gets*
25 *included in our customer's products that require support for XTS-AES. We've seen considerable interest in these*
26 *cores from partners.*

27 *A description of the SafeXcel EIP-38 family of cores can be found here:*

28 *http://www.safenet-inc.com/products/ip/safeXcel_IP_AESGCMXTS_Acc.asp*

29 **Thales (from Bob Lockhart)**

30 *Thales plans to implement XTS if it is accepted by NIST as a certified mode of operation.*

31 **TrueCrypt**

32 *TrueCrypt, the leading open source disk encryption software, supports XTS as the default mode.*

33 **WinMagic (from Gary McCracken)**

34 *WinMagic will implement XTS-AES when it becomes a FIPS 140 approved mode of operation.*

1 WinMagic produces full disk encryption software for desktops and laptops. FIPS 140 certification is required by
2 many of our customers.

3 **Comment from David Clunie**

4 (see [MiscRsp])

5 *For XTS to be useful for interchange (removable) media protection, standardization of XTS or a similar mechanism*
6 *as an encryption mode of operation of AES is desirable, but insufficient; there would also need to be an*
7 *appropriate means of exchange of keys in a standard and inter-operable manner defined. It is understood that this*
8 *is outside the scope of the current request, but some indication of whether or not there is such an effort in progress*
9 *would be helpful.*

10 *Or to put this another way, it may be premature to standardize the mode of operation until it is well established*
11 *that there will be support in the operating system industry for its usage in the context of removable media device*
12 *drivers (for removable optical media like CD and DVD as well as USB media (flash drives)).*

13 **Response:**

14 NIST has standardized many previous modes of operation (e.g., ECB, CTR, CBC, GCM, CCM) without
15 standardizing the corresponding key management of the keys used in these modes. SP 800-56A describes a
16 framework for establishing secrets using Diffie-Hellman and ECC, and SP 800-56B describes using RSA for
17 establishing a shared secret. SP 800-57 provides recommendations for key management, but does not go into
18 specific protocols or recommendations for interoperability. In general, NIST has kept its requirements general
19 and left it to other organizations, like IETF or IEEE, for defining the specifics.

20 For current key management efforts that are standardizing key interchange, see IETF KEYPROV, OASIS EKMI,
21 IEEE P1619.3, or other related efforts.

22 **Comments From Michael Willett of Seagate Technology**

23 (see [Seagate1])

24 *In sealed secure storage the following are possible and desirable, but not considered with XTS:*

- 25 • *access control*
- 26 • *user authentication*
- 27 • *data authentication (integrity check)*
- 28 • *tamper detection.*

29 *For example, with the (mandatory) access control, copy-and-paste attacks are automatically thwarted, but their*
30 *prevention is the only distinguishing security feature of XTS.*

31 *For secure removable media storage application (like tape) the following are desirable, but missing from XTS:*

- 32 • *data authentication*
- 33 • *large block encryption*
- 34 • *independent encryption from the (possibly changing or unknown) data location*
- 35 • *unpredictable tweak (user supplied, stored IV)*

36 **Response:**

37 IEEE Std 1619-2007 does consider access control and authentication in in D.2. Integrity checks and IVs are
38 out-of-scope for XTS because the constraint of the mode requires that no extra data be added. If you have room
39 for an integrity check, then CCM, GCM, CBC-HMAC, XTS-HMAC (e.g., modes standardized in IEEE std
40 1619.1-2007) are more appropriate.

1 *Is XTS a good tradeoff between security and complexity (costs)? NO:*

- 2 • *There are asymmetric large block encryption modes (e.g. BitLocker of Windows Vista) with similar*
- 3 *complexity but better security for storage systems*
- 4 • *There are faster, simpler encryption modes used together with access control (e.g. CBC)*

5 **Response:**

6 Large block encryption modes need two passes over the data. In applications that require streaming reads and
7 writes, the delay or buffer requirements of the second pass can be prohibitive. For example, a hard disk
8 implementing a wide-block encryption mode would likely need to buffer 512 bytes, which is not trivial. That said,
9 if a particular implementation has extra buffer space, computing power, and latency tolerance, then wide-block
10 encryption provides better protection than XTS.

11 CBC is not universally faster and simpler than XTS. XTS can operate in parallel, which makes it faster when
12 multiple AES execution units are available. Even when creating a virtually parallel CBC implementation by
13 interleaving several CBC chains, the implementation still needs an extra AES operation to start each CBC chain,
14 and the format has to be interleaved ahead of time. For XTS, it's possible to process the same format either
15 serially, or in parallel, without changing the format.

16 *What security problem XTS addresses? (Threat model, attack scenarios)*

- 17 • *Lost laptop scenario: CBC with encrypted address as IV is adequate, faster, simpler, cheaper*
- 18 • *Janitor access of locked PC: If the key is properly destroyed in memory, it is similar to the lost laptop case*
- 19 • *Disk theft from datacenter: ~ lost laptop*
- 20 • *When repeated access to the ciphertext is assumed (removable media, external encryption module): Only*
- 21 *copy-and-paste attacks are mitigated by XTS, all other small block problems remain.*

22 **Response:**

23 XTS provides a solution for storage encryption systems where the applications using the plaintext are able to
24 detect when 16-byte blocks are randomized, and where using wide-block encryption or authenticated encryption
25 is not feasible due to limited processing power or space limitations. XTS is less malleable than CBC because an
26 attacker can flip particular bits within a block, but with XTS can only randomize the whole block. Also, XTS is a
27 parallel design and CBC is serial, making XTS faster in environments that have multiple AES execution units.

28 *Ciphertext stealing of XTS is defined with a block swap. There is no reason to destroy the straight block order,*
29 *causing unnecessary implementation difficulties, confusion, incompatibilities with existing modules.*

- 30 • *Without the swap of the last two blocks, the data buffer could be accessed sequentially backward or*
- 31 *forward.*
- 32 • *More recent definitions of ciphertext stealing do not swap blocks (see Henk, Tillborg: Encyclopedia of*
- 33 *Cryptography and Security, Springer 2005, p387.)*

34 **Response:**

35 There is no reason that you could not add a special 'formatter' that swaps the last two blocks before writing it to
36 disk. I will recommend that NIST clarify this. The ordering of the blocks in ciphertext stealing is not a security
37 issue, and there are situations where either order is preferable.

38 *Danger of standardizing a "disk" encryption mode, which is not optimal for storage encryption:*

- 39 • *Customers demand it in place of more secure solutions (as "the" secure storage standard)*
- 40 • *Computer manufacturers relay the demand of their customers to disk manufacturers, SW houses,*
- 41 *controller designers; forcing wide spread implementation of a less secure, more expensive encryption*
- 42 *mode*
- 43 • *In the most common applications there are unnecessary costs in key generation/storage, algorithm*
- 44 *complexity, speed, power consumption*

- *Poor security systems can be sold as standards conformant, thereby reducing the overall security in storage*

[Therefore, if the XTS mode gets approved by NIST, it must not be labeled as “storage/disk encryption”, but something like as an “encryption mode, resistant to moving blocks”.]

Response

This last comment seems subjective and is only an issue if XTS truly is inferior. However, as has been explained multiple times, XTS has particular strengths that make it suitable for applications that have fixed-size sectors and lack the processing power to do two passes over the data for a wide-block encryption. NIST offers a toolbox of encryption modes that all have their strengths and weaknesses in particular contexts, and XTS is no different in this regard.

The appeal of XTS for wider applications

Comment From David Clunie

(see [MiscRsp])

In digital medical image exchange (and digital medical record exchange in general), though online transactions are envisaged, the reality of the state of the art is that most images are exchanged on removable media due to their bulk (hence slow speed of network transfer) and the lack of an authentication infrastructure between loosely coupled providers and users. Current exchange is generally unprotected (no encryption at all), since the threat is presumed to be low and any barrier to reception risks patient safety. However, it is anticipated that encryption may be required in the long term, and it is desirable that a standard, inter-operable, cross-platform storage encryption infrastructure be available in consumer operating systems so that it can be leveraged by medical applications.

Comments From Michael Willett of Seagate Technology

(see [Seagate1])

Why standardize XTS at all? No good reason:

- 1. Interoperability: dumb disks with encrypted data can be connected to different controllers. BUT most new disk drives have internal encryption, and their platters cannot be mixed and matched.*
- 2. Archived data recovery in the distant future: Small audience. Only specialized archival devices benefit from a standard, if their encryption module is separate from the medium (tape, CD-ROM...). However, these devices are better served by higher security encryption*
- 3. By employing a standard there is no need to perform security analysis for the encryption mode (as for a proprietary solution). BUT the encryption mode is the least of the worries. Some other harder issues to be considered:*
 - the applicability of XTS needs to be shown*
 - the implementation has to be validated*
 - implementation errors have to be avoided: audits, reviews, tests (buffer overrun, ill formed messages, weak parameters...)*
 - robust protocols (message replay, man-in-the-middle attacks, non-random nonces, weak session keys...)*
 - firmware code, crypto keys, sensitive memory content need protection*
 - authenticated secure firmware (prevent ROM swap/patch, safe update)*
 - protection of wires, pins, RAM chips (freezing attacks)*

- *protection of debug ports*

Response:

Standardizing XTS provides a parallel encryption mode that has roughly the same computation complexity as CBC, but provides less malleable ciphertext (i.e., in XTS the adversary cannot randomize with a granularity of less than 16-bytes, whereas with CBC, the adversary can modify specific bits within a 16-byte block – with the cost of randomizing the previous 16-bytes). Clearly there are other issues beside the encryption mode needed to make a system secure. But this is true for all the SP 800-38 series encryption modes.

The proposal for the approved specification to be available only by purchase from IEEE

Comment From Vijay Bharadwaj and Neils Ferguson of Microsoft

(see [MiscRsp])

We believe that it is highly undesirable to standardize an algorithm whose specification, unlike those of other FIPS approved algorithms, is not freely available.

Comment from Rich Schroepfel

(see [MiscRsp])

It seems inappropriate for the government to be lending its imprimatur to a standard that will cost \$105 per copy, and be sold by a private organization. Standards should be available for free for the asking, from a government website. Of course the IEEE is free to do whatever they want, but NIST shouldn't be endorsing it. Especially for an encryption standard, where wide and perpetual public review is needed to assure sufficient security scrutiny. Moreover, if NIST is going to include an external reference in a standard, they should include a suitable hash code to make sure that IEEE doesn't make alterations in the standard.

Comment From David Clunie

(see [MiscRsp])

The IEEE and the participants in P1619 presumably expect significant commercial benefits to flow from the inclusion of XTS as an FIPS, and should be satisfied in that regard rather than expect a flow of revenue from purchase of the P1619 XTS standard itself. Alternatively, there are other means by which IEEE may be remunerated by the federal government that do not involve purchase fees for the standard.

Response:

The IEEE-SA (IEEE Standards Association) and the members of the IEEE P1619 task group are financially independent entities. The IEEE-SA oversees the standards creation process and pays for editors and other publication costs, and recuperates its costs by selling the standard. The IEEE P1619 task group members include representatives of companies who create products with XTS-AES.

IEEE-SA does offer a way for the standards developers to pay for IEEE's publication costs, and have the standard offered to the public for free (see, for example, the IEEE 802 series). However, the group members were unable (and in some cases unwilling) to pay this fee.

Comments From Michael Willett of Seagate Technology

(see [Seagate1])

1 *As stated in the NIST e-mail, the IEEE has provided a “free” copy of the XTS algorithm description on a*
2 *temporary basis for the purpose of supporting this Call for Comment. But, after 90 days, the IEEE will charge for*
3 *a copy of the algorithm (\$85/\$105). As far as we know, this is without precedent for NIST: to certify an algorithm*
4 *for which the NIST public has to pay. Seagate is against being charged. But, Seagate is against NIST certification*
5 *of this algorithm in the first place.*

6 **General Response**

7 Most of the comments were highly critical of requiring that implementors purchase the IEEE 1619 standard. This
8 is a natural response. We have come to expect free standards from NIST and this goes against that
9 expectation. NIST itself prefers free standards, as do the members of the IEEE P1619 working group.

10 However, all standards cost money. It's just a matter of who pays. In this case, the members of IEEE P1619 did
11 not pay anything fees to participate, nor did NIST pay to develop this standard using U.S. taxpayer money. This
12 structure allowed the participation of top cryptographers to develop the details of the XTS-AES mode, free of
13 cost to them.

14 Unfortunately, if NIST doesn't pay, and the IEEE P1619 members don't pay, then the publication costs get
15 passed on to the user. For corporations, this is generally not an issue because there is an expectation to pay for
16 standards as part of product development. Additionally, many companies already have company-wide access to
17 the IEEE library, allowing employees to download IEEE 1619 at no additional cost.

18 Also note that NIST has an existing precedent to reference standards that cost money. Here is a partial list of
19 standards that NIST references and that cost money now, or used to cost money:

- 20 1. ANS X9.31-1998, Digital Signatures Using Reversible Public Key Cryptography for the Financial
21 Services Industry (rDSA). (as referenced in FIPS 186-3 and FIPS 140-2, Annex C)
- 22 2. ANS X9.62-2005, Public Key Cryptography for the Financial Services Industry: The Elliptic Curve Digital
23 Signature Algorithm (ECDSA). (as referenced in FIPS 186-3 and FIPS 140-2 Annex C)
- 24 3. ANS X9.80, Prime Number Generation, Primality Testing and Primality Certificates. (as referenced in
25 FIPS 186-3)
- 26 4. IEEE 802.11i (as reference in FIPS 140-2 Implementation Guidance)

27 Overall, referencing standards that cost money is a common industry practice, and in this case is a good way for
28 NIST to save money for tax payers, although it unfortunately does cause some inconvenience to some
29 implementors.

30 **Concerns of intellectual property rights**

31 **Comments From Vijay Bharadwaj and Neils Ferguson of Microsoft**

32 (see [MiscRsp])

33 *As stated in the call for comments, the current situation of IP rights with respect to XTS-AES is unclear. We believe*
34 *that standardizing an algorithm which is so encumbered is undesirable, and that IP issues could inhibit adoption of*
35 *such a standard.*

36 **Comment From David Clunie**

37 (see [MiscRsp])

38 *This proposal is regarded as unacceptable.*

39 *FIPS standards should be available to the public for download without charge.*

40 *Any charge, no matter how apparently insignificant, limits the accessibility, opportunity for scientific review and*
41 *likelihood of adoption of the standard.*

1 The IEEE and the participants in P1619 presumably expect significant commercial benefits to flow from the
2 inclusion of XTS as an FIPS, and should be satisfied in that regard rather than expect a flow of revenue from
3 purchase of the P1619 XTS standard itself. Alternatively, there are other means by which IEEE may be
4 remunerated by the federal government that do not involve purchase fees for the standard.

5 **Comments From Michael Willett of Seagate Technology**

6 (see [Seagate1])

7 As stated in the NIST e-mail and web site:

8 “The chair of the (Security in Storage) SISWG informed NIST that he is unaware of any patent claims on XTS, but
9 that NeoScale Systems, subsequently acquired by nCipher, submitted a Letter of Assurance of Essential Patents to
10 the IEEE, without elaborating on what aspect of IEEE 1619 was patented.”

11 This information suggests that the XTS is encumbered with possible patent claims of unknown quantity and
12 ownership, aided by the voluntary nature of the IEEE IP declaration process. NIST should not certify an algorithm
13 under such circumstances, since the NIST certification would encourage mandatory adoption of XTS in the storage
14 industry, with unknown royalty consequences on the implementers. AES candidates were solicited with royalty-free
15 pre-conditions and the granting of any patent rights to NIST. We expect no less for a certified mode for AES.

16 **General Response**

17 Many commenters were uneasy with the ambiguity of the intellectual property (IP) rights situation, in regards to
18 the Letter of Assurance (LoA) submitted by NeoScale systems. In truth, this ambiguity exists for all technology.
19 There are many potentially unenforceable patents that are continually submitted to the patent office, and which
20 are thrown out or substantially changed after years of prosecution. This is just a case where we were made
21 aware of a claim that *might* be relevant.

22 **Investigation NeoScale Systems' IP Claims**

23 Subsequent to submitting the LoA to IEEE, NeoScale Systems declared bankruptcy, and the bulk of the assets
24 were purchased by nCipher (now owned by Thales). However, after nCipher investigated the matter further,
25 they discovered that they did not purchase the IP in question (i.e., that covers IEEE Std 1619). This IP is still the
26 property of the bank that acquired NeoScale's assets as a result of the bankruptcy.

27 From searching for patents owned by NeoScale Systems, the following patents applications are current as of
28 Jan 2009. Please note that that I am not a lawyer and that these comments are only my opinion (commentary in
29 *italics*):

- 30 • **20090013016** (Noll, LeBlanc), SYSTEM AND METHOD FOR PROCESSING DATA FOR DATA
31 SECURITY, Filed July 2007: *This application does not mention XTS or IEEE 1619, although it does*
32 *patent some general storage encryption methods.*
- 33 • **20080273706** (Noll), System and Method for Controlled Access Key Management, Filed May 2007:
34 *Covers key management and does not mention XTS or disk encryption.*
- 35 • **20060277387** (Rhoades) System and method for hardware allocation of memory resources, filed April
36 2006: *Does not apply to cryptography.*
- 37 • **20050080761** (Sundararajan, et al) Data path media security system and method in a storage area
38 network , filed Oct 2003: *Does not discuss XTS or related encryption techniques.*
- 39 • **20050041812** (Sundararajan, et al) Method and system for stateful storage processing in storage area
40 networks, filed Oct 2003: *Does not discuss XTS or related encryption techniques.*
- 41 • **20050033988** (Chandrashekar, et al) Method and system for transparent encryption and authentication
42 of file data protocols over internet protocol, Filed Oct 2003: *Mentions AES encryption, but does not*
43 *mention XTS or XEX. The part to watch for is a claim on a NAS server where the “[encryption or*

1 *decryption] is provided by a NIST approved process”.*

2 I could not find any granted U.S. patents that are assigned to NeoScale Systems.

3 NeoScale Systems submitted a Letter of Assurance to IEEE in March 2007 (see [NLoA]). It is unclear which
4 specific patent application this was intended to cover, but none of the published applications specifically cover
5 the XTS or XEX mode of encryption.

6 **IP Claims by Rogaway**

7 The creator of XEX, Philip Rogaway, does not have any IP claims on XEX (and by extension, XTS) (see
8 [RogIP]). Here is the statement from Rogaway to Shai Halevi:

9 *I know of no IP claims that would in [fact] touch upon XEX itself; I myself never claimed anything in the*
10 *direction of a tweakable conventional blockcipher.*

11 **Conclusions and Recommendations**

12 Overall, I recommend that NIST accept XTS-AES (as defined in IEEE Std 1619-2007) as an Approved Mode of
13 Operation with the following changes:

- 14 • Allow 'one-key' variants of XTS-AES, where the same 128-bit or 256-bit key is used within both AES
15 ciphers.
- 16 • Prohibit data unit sizes larger than 2^{20} blocks.
- 17 • Clarify that the ordering of the blocks in ciphertext stealing is a logical ordering, and can be mapped to a
18 different physical ordering (i.e., it is possible to swap the order of the ciphertext stealing blocks)

19 **References**

- [Ball] Ball, Matthew, “NIST’s Consideration of XTS-AES as standardized by IEEE Std 1619-2007.” Sept 2008. See http://csrc.nist.gov/groups/ST/toolkit/BCM/documents/comments/XTS/XTS_comments-Ball.pdf
- [CBC-CS] National Institute of Standards and Technology (NIST), *Proposal To Extend CBC Mode By “Ciphertext Stealing”*, May 2007. See <http://csrc.nist.gov/groups/ST/toolkit/BCM/documents/ciphertext%20stealing%20proposal.pdf>
- [Liskov] Liskov, Moses and Kazuhiko Minematsu, “Comments on XTS-AES.” Sept 2008. See http://csrc.nist.gov/groups/ST/toolkit/BCM/documents/comments/XTS/XTS_comments-Liskov_Minematsu.pdf
- [MiscRsp] Bharadwaj, Clunie, Ferguson, Rogaway, Shahar, and Shroepfel, “Public Comments on the XTS-AES Mode.” (Collected XTS Comments) Sept 2008. See http://csrc.nist.gov/groups/ST/toolkit/BCM/documents/comments/XTS/collected_XTS_comments.pdf
- [NLoA] NeoScale’s Letter of Assurance to IEEE on IEEE Std 1619, March 2007, see <http://standards.ieee.org/db/patents/loa-1619-neoscale-08Mar2007.pdf>
- [Request] NIST, Request for Public Comment on XTS, 2008, see http://csrc.nist.gov/groups/ST/documents/Request-for-Public-Comment-on_XTS.pdf
- [RogIP] Rogaway, Phillip, “Quick check on IP situation of XEX”, Sept 2006, See <http://grouper.ieee.org/groups/1619/email/msg01309.html>

[Seagate1] Willet, Michael, Seagate Technology, "Comments provided to NIST in response to: Request for Public Comment on XTS/AES." Sept 2008. See http://csrc.nist.gov/groups/ST/toolkit/BCM/documents/comments/XTS/revise_XTS_comments-Seagate.pdf

[Seagate2] Hars, Laszlo, Seagate Technology, "Response from Laszlo Hars." Sept 2008. See http://csrc.nist.gov/groups/ST/toolkit/BCM/documents/comments/XTS/XTS_comments-Hars.pdf