



Key Management Interoperability Protocol Usage Guide 1.0

Committee Draft **03**

21 October 2009

Specification URIs:
This Version:

- <http://docs.oasis-open.org/kmip/ug/v1.0/cd03/kmip-ug-1.0-draft-03.html>
- <http://docs.oasis-open.org/kmip/ug/v1.0/cd03/kmip-ug-1.0-draft-03.doc>, (Authoritative)
- <http://docs.oasis-open.org/kmip/ug/v1.0/cd03/kmip-ug-1.0-draft-03.pdf>

Previous Version:

- <http://docs.oasis-open.org/kmip/ug/v1.0/cd02/kmip-ug-1.0-draft-02.html>
- <http://docs.oasis-open.org/kmip/ug/v1.0/cd02/kmip-ug-1.0-draft-02.doc>, (Authoritative)
- <http://docs.oasis-open.org/kmip/ug/v1.0/cd02/kmip-ug-1.0-draft-02.pdf>

Latest Version:

- <http://docs.oasis-open.org/kmip/ug/v1.0/kmip-ug-1.0.html>
- <http://docs.oasis-open.org/kmip/ug/v1.0/kmip-ug-1.0.doc>
- <http://docs.oasis-open.org/kmip/ug/v1.0/kmip-ug-1.0.pdf>

Technical Committee:

[OASIS Key Management Interoperability Protocol \(KMIP\) TC](#)

Chair(s):

Robert Griffin
Subhash Sankuratripati

Editor(s):

Indra Fitzgerald

Related work:

This specification replaces or supersedes:

- None

This specification is related to:

- [Key Management Interoperability Protocol Specification v1.0](http://docs.oasis-open.org/kmip/spec/v1.0/), <http://docs.oasis-open.org/kmip/spec/v1.0/>
- [Key Management Interoperability Protocol Profiles v1.0](http://docs.oasis-open.org/kmip/profiles/v1.0/), <http://docs.oasis-open.org/kmip/profiles/v1.0/>
- [Key Management Interoperability Protocol Use Cases v1.0](http://docs.oasis-open.org/kmip/usecases/v1.0/), <http://docs.oasis-open.org/kmip/usecases/v1.0/>

Declared XML Namespace(s):

None

Abstract:

This document is intended to complement the Key Management Interoperability Protocol Specification by providing guidance on how to implement the Key Management Interoperability Protocol (KMIP) most effectively to ensure interoperability.

Deleted: 2

Deleted: 14

Deleted: kmip-ug-1.0-draft-02.html

Deleted: kmip-ug-1.0-draft-02.doc

Deleted: kmip-ug-1.0-draft-02.pdf

Deleted: kmip-ug-1.0-draft-01.html

Field Code Changed

Field Code Changed

Deleted: kmip-ug-1.0-draft-01.doc

Field Code Changed

Deleted: kmip-ug-1.0-draft-01.pdf

Formatted: Font color: Auto

Deleted: kmip-ug-1.0-draft-02.html

Field Code Changed

Formatted: Font color: Auto

Field Code Changed

Formatted: Font color: Auto

Formatted: Font color: Auto

Deleted: kmip-ug-1.0-draft-02.doc

Field Code Changed

Formatted: Font color: Auto

Formatted: Font color: Auto

Deleted: kmip-ug-1.0-draft-02.pdf

Field Code Changed

Formatted: Bullets and Numbering

Field Code Changed

Formatted: Font color: Auto

Formatted: Bullets and Numbering

Deleted: TBD

Formatted: Contributor, None

Deleted: 2

Deleted: 14

Status:

This document was last revised or approved by the Key Management Interoperability Protocol TC on the above date. The level of approval is also listed above. Check the "Latest Version" or "Latest Approved Version" location noted above for possible later revisions of this document.

Technical Committee members should send comments on this specification to the Technical Committee's email list. Others should send comments to the Technical Committee by using the "Send A Comment" button on the Technical Committee's web page at <http://www.oasis-open.org/committees/kmip/>.

For information on whether any patents have been disclosed that may be essential to implementing this specification, and any offers of patent licensing terms, please refer to the Intellectual Property Rights section of the Technical Committee web page (<http://www.oasis-open.org/committees/kmip/ipr.php>).

The non-normative errata page for this specification is located at <http://www.oasis-open.org/committees/kmip/>.

Notices

Copyright © OASIS® 2009. All Rights Reserved.

All capitalized terms in the following text have the meanings assigned to them in the OASIS Intellectual Property Rights Policy (the "OASIS IPR Policy"). The full Policy may be found at the OASIS website.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published, and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this section are included on all such copies and derivative works. However, this document itself may not be modified in any way, including by removing the copyright notice or references to OASIS, except as needed for the purpose of developing any document or deliverable produced by an OASIS Technical Committee (in which case the rules applicable to copyrights, as set forth in the OASIS IPR Policy, must be followed) or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by OASIS or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and OASIS DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY OWNERSHIP RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

OASIS requests that any OASIS Party or any other party that believes it has patent claims that would necessarily be infringed by implementations of this OASIS Committee Specification or OASIS Standard, to notify OASIS TC Administrator and provide an indication of its willingness to grant patent licenses to such patent claims in a manner consistent with the IPR Mode of the OASIS Technical Committee that produced this specification.

OASIS invites any party to contact the OASIS TC Administrator if it is aware of a claim of ownership of any patent claims that would necessarily be infringed by implementations of this specification by a patent holder that is not willing to provide a license to such patent claims in a manner consistent with the IPR Mode of the OASIS Technical Committee that produced this specification. OASIS may include such claims on its website, but disclaims any obligation to do so.

OASIS takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on OASIS' procedures with respect to rights in any document or deliverable produced by an OASIS Technical Committee can be found on the OASIS website. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this OASIS Committee Specification or OASIS Standard, can be obtained from the OASIS TC Administrator. OASIS makes no representation that any information or list of intellectual property rights will at any time be complete, or that any claims in such list are, in fact, Essential Claims.

The names "OASIS", [insert specific trademarked names and abbreviations here] are trademarks of OASIS, the owner and developer of this specification, and should be used only to refer to the organization and its official outputs. OASIS welcomes reference to, and implementation and use of, specifications, while reserving the right to enforce its marks against misleading uses. Please see <http://www.oasis-open.org/who/trademark.php> for above guidance.

Deleted: 2

Deleted: 14

Table of Contents

1 Introduction	6	
1.1 Normative References	6	
1.2 Non-normative References	8	
2 Assumptions	8	
2.1 Island of Trust	9	Deleted: 8
2.2 Message Security	9	
2.3 State-less Server	9	
2.4 Extensible Protocol	9	
2.5 Support for Cryptographic Objects	9	
2.6 Client-Server Message-based Model	9	
2.7 Synchronous and Asynchronous Messages	10	Deleted: 9
2.8 Support for “Intelligent Clients” and “Key Using Devices”	10	
2.9 Batched Requests and Responses	10	
2.10 Reliable Message Delivery	10	
2.11 Large Responses	10	
2.12 Key Life-cycle and Key State	10	
3 Usage Guidelines	10	
3.1 Authentication	10	
3.2 Authorization for Revoke, Recover, Destroy and Archive Operations	11	
3.3 Using Notify and Put Operations	11	
3.4 Usage Allocation	12	Deleted: 11
3.5 Key State and Times	12	
3.6 Template	13	
3.7 Archive Operations	13	
3.8 Message Extensions	13	
3.9 Unique Identifiers	13	
3.10 Result Message Text	14	Deleted: 13
3.11 Query	14	Deleted: 13
3.12 Canceling Asynchronous Operations	14	
3.13 Multi-instance Hash	14	
3.14 Returning Related Objects	14	
3.15 Reducing Multiple Requests through Use of Batch	14	
3.16 Maximum Message Size	15	Deleted: 14
3.17 Using Offset in Re-key and Re-certify Operations	15	Deleted: 14
3.18 Locate Queries	15	
3.19 ID Placeholder	16	
3.20 Key Block	17	
3.21 Using Wrapped Keys with KMIP	18	Deleted: 17
3.21.1 Encrypt-only Example with a Symmetric Key for a Get Request and Response	18	
3.21.2 Encrypt-only Example with a Symmetric Key for a Register Request and Response	19	
3.21.3 Encrypt-only Example with an Asymmetric Key for a Get Request and Response	19	
3.21.4 MAC-only Example with an HMAC Key for a Get Request and Response	20	Deleted: 20
3.21.5 Registering a Wrapped Key as an Opaque Cryptographic Object	21	Deleted: 2
		Deleted: 14

3.22 Object Group	21	Deleted: 20
3.23 Certify and Re-certify	21	
3.24 Specifying Attributes during Create Key Pair	21	
3.24.1 Example of Specifying Attributes during Create Key Pair	22	Deleted: 21
3.25 Registering a Key Pair	23	
3.26 Non-Cryptographic Objects	24	Deleted: 23
3.27 Asymmetric Concepts with Symmetric Keys	24	
3.28 Application Specific Information	25	Deleted: 25
3.29 Mutating Attributes	26	Deleted: 25
3.30 Interoperable Key Naming for Tape	26	
3.30.1 Native Tape Encryption by a KMIP Client	26	
3.30.1.1 Method Overview	26	
3.30.1.2 Definitions	27	Deleted: 26
3.30.1.3 Algorithm 1. Numeric to text direction (tape format's KAD to KMIP ASI)	27	
3.30.1.4 Algorithm 2. Text to numeric direction (KMIP ASI to tape format's KAD)	28	Deleted: 27
3.30.1.5 Example Output	28	
3.30.1.6 Backward-compatibility assessment	30	
3.31 Revocation Reason Codes	30	
3.32 Certificate Renewal, Update, and Re-key	31	Deleted: 30
3.33 Key Encoding	31	
3.33.1 AES Key Encoding	31	
3.33.2 Triple-DES Key Encoding	31	
4 Deferred KMIP Functionality	32	Deleted: 31
A. Acronyms	34	
B. Acknowledgements	36	
C. Revision History	39	Deleted: 37
Tables		
Table 1: ID Placeholder Input and Output for KMIP Operations	17	
Table 2: Cryptographic Usage Masks Pairs	25	Deleted: 24

Deleted: 2
Deleted: 14

1 Introduction

This Key Management Interoperability Protocol Usage Guide is intended to complement the Key Management Interoperability Protocol Specification [KMIP-Spec] by providing guidance on how to implement the Key Management Interoperability Protocol (KMIP) most effectively to ensure interoperability. In particular, it includes the following guidance:

- Clarification of assumptions and requirements that drive or influence the design of KMIP and the implementation of KMIP-compliant key management.
- Specific recommendations for implementation of particular KMIP functionality.
- Clarification of mandatory and optional capabilities for conformant implementations.
- Functionality considered for inclusion in KMIP V1.0, but deferred to subsequent versions of the standard.

A selected set of conformance profiles and authentication suites are defined in the KMIP Profiles specification [KMIP-Prof].

Further assistance for implementing KMIP is provided by the KMIP Use Cases for Proof of Concept Testing document [KMIP-UC] that describes a set of recommended test cases and provides the TTLV (Tag/Type/Length/Value) format for the message exchanges defined by those use cases.

1.1 Terminology

For a list of terminologies refer to [KMIP-Spec].

1.2 Normative References

- | | |
|--------------|---|
| [FIPS186-3] | Digital Signature Standard (DSS), FIPS PUB 186-3, June 2009, http://csrc.nist.gov/publications/fips/fips186-3/fips_186-3.pdf |
| [FIPS197] | Advanced Encryption Standard (AES), FIPS PUB 197, November 26, 2001, http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf |
| [FIPS198-1] | The Keyed-Hash Message Authentication Code (HMAC), FIPS PUB 198-1, July 2008, http://csrc.nist.gov/publications/fips/fips198-1/FIPS-198-1_final.pdf |
| [IEEE1003-1] | IEEE Std 1003.1, <i>Standard for information technology - portable operating system interface (POSIX). Shell and utilities</i> , 2004. |
| [ISO16609] | ISO, <i>Banking -- Requirements for message authentication using symmetric techniques</i> , ISO 16609, 1991. |
| [ISO9797-1] | ISO/IEC, <i>Information technology -- Security techniques -- Message Authentication Codes (MACs) -- Part 1: Mechanisms using a block cipher</i> , ISO/IEC 9797-1, 1999. |
| [KMIP-Spec] | OASIS Draft, <i>Key Management Interoperability Protocol Specification v1,0</i> , Committee Draft, October 2009. |
| [KMIP-Prof] | OASIS Draft, <i>Key Management Interoperability Protocol Profiles v1,0</i> , Committee Draft, October 2009. |
| [PKCS#1] | RSA Laboratories, <i>PKCS #1 v2.1: RSA Cryptography Standard</i> , June 14, 2002, http://www.rsa.com/rsalabs/node.asp?id=2125 |
| [PKCS#5] | RSA Laboratories, <i>PKCS #5 v2.1: Password-Based Cryptography Standard</i> , October 5, 2006, http://www.rsa.com/rsalabs/node.asp?id=2127 |
| [PKCS#7] | RSA Laboratories, <i>PKCS#7 v1.5: Cryptographic Message Syntax Standard</i> , November 1, 1993, http://www.rsa.com/rsalabs/node.asp?id=2129 |
| [PKCS#8] | RSA Laboratories, <i>PKCS#8 v1.2: Private-Key Information Syntax Standard</i> , November 1, 1993, http://www.rsa.com/rsalabs/node.asp?id=2130 |

Deleted: S

Formatted: Bullets and Numbering

Formatted: Normal, No bullets or numbering, Hyphenate, Tabs: Not at 0"

Formatted: Bullets and Numbering

Deleted: .

Deleted: 03

Deleted: ,, URI (TBD)

Deleted: Specification

Deleted: 01

Deleted: ,, URI (TBD)

Deleted: .

Deleted: .

Deleted: .

Deleted: .

Deleted: 2

Deleted: 14

- 45 [PKCS#10] RSA Laboratories, *PKCS #10 v1.7: Certification Request Syntax Standard*, May
 46 26, 2000, <http://www.rsa.com/rsalabs/node.asp?id=2132>
- 47 [RFC1319] B. Kaliski, *The MD2 Message-Digest Algorithm*, IETF RFC 1319, Apr 1992,
 48 <http://www.ietf.org/rfc/rfc1319.txt>
- 49 [RFC1320] R. Rivest, *The MD4 Message-Digest Algorithm*, IETF RFC 1320, Apr 1992,
 50 <http://www.ietf.org/rfc/rfc1320.txt>
- 51 [RFC1321] R. Rivest, *The MD5 Message-Digest Algorithm*, IETF RFC 1321, Apr 1992,
 52 <http://www.ietf.org/rfc/rfc1321.txt>
- 53 [RFC1421] J. Linn, *Privacy Enhancement for Internet Electronic Mail: Part I: Message
 54 Encryption and Authentication Procedures*, IETF RFC 1421, Feb 1993,
 55 <http://www.ietf.org/rfc/rfc1421.txt>
- 56 [RFC1424] B. Kaliski, *Privacy Enhancement for Internet Electronic Mail: Part IV: Key
 57 Certification and Related Services*, IETF RFC 1424, February 1993,
 58 <http://www.ietf.org/rfc/rfc1424.txt>
- 59 [RFC2104] H. Krawczyk, M. Bellare, R. Canetti, *HMAC: Keyed-Hashing for Message
 60 Authentication*, IETF RFC 2104, Feb 1007, <http://www.ietf.org/rfc/rfc2104.txt>
- 61 [RFC2119] S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*,
 62 <http://www.ietf.org/rfc/rfc2119.txt>, IETF RFC 2119, March 1997.
- 63 [RFC2898] B. Kaliski, *PKCS #5: Password-Based Cryptography Specification Version 2.0*,
 64 IETF RFC 2898, Sep 2000, <http://www.ietf.org/rfc/rfc2898.txt>
- 65 [RFC3394] J. Schaad, R. Housley, *Advanced Encryption Standard (AES) Key Wrap
 66 Algorithm*, IETF RFC 3394, Sep 2002, <http://www.ietf.org/rfc/rfc3394.txt>
- 67 [RFC3447] J. Jonsson, B. Kaliski, *Public-Key Cryptography Standards (PKCS) #1: RSA
 68 Cryptography Specifications Version 2.1*, IETF RFC 3447 Feb 2003,
 69 <http://www.ietf.org/rfc/rfc3447.txt>
- 70 [RFC3629] F. Yergeau, *UTF-8, a transformation format of ISO 10646*, IETF RFC 3629, Nov
 71 2003, <http://www.ietf.org/rfc/rfc3629.txt>
- 72 [RFC3647] S. Chokhani, W. Ford, R. Sabett, C. Merrill, and S. Wu, *RFC3647: Internet
 73 X.509 Public Key Infrastructure Certificate Policy and Certification Practices
 74 Framework*, November 2003, <http://www.ietf.org/rfc/rfc3647.txt>
- 75 [RFC4210] C. Adams, S. Farrell, T. Kause and T. Mononen, *RFC2510: Internet X.509
 76 Public Key Infrastructure Certificate Management Protocol (CMP)*, September
 77 2005, <http://www.ietf.org/rfc/rfc4210.txt>
- 78 [RFC4211] J. Schaad, *RFC 4211: Internet X.509 Public Key Infrastructure Certificate
 79 Request Message Format (CRMF)*, September 2005, <http://www.ietf.org/rfc/rfc4211.txt>
- 80 [RFC4868] S. Kelly, S. Frankel, *Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-
 81 512 with IPsec*, IETF RFC 4868, May 2007, <http://www.ietf.org/rfc/rfc4868.txt>
- 82 [RFC4949] R. Shirey, *RFC4949: Internet Security Glossary, Version 2*, August 2007,
 83 <http://www.ietf.org/rfc/rfc4949.txt>
- 84 [RFC5272] J. Schaad and M. Meyers, *RFC5272: Certificate Management over CMS (CMC)*,
 85 June 2008, <http://www.ietf.org/rfc/rfc5272.txt>
- 86 [RFC5280] D. Cooper, S. Santesson, S. Farrell, S. Boeyen, R. Housley and W. Polk, *RFC
 87 5280: Internet X.509 Public Key Infrastructure Certificate and Certificate
 88 Revocation List (CRL) Profile*, May 2008, <http://www.ietf.org/rfc/rfc5280.txt>
- 89 [RFC5649] R. Housley, *Advanced Encryption Standard (AES) Key Wrap with Padding
 90 Algorithm*, IETF RFC 5649, Aug 2009, <http://www.ietf.org/rfc/rfc5649.txt>
- 91 [SP800-38A] M. Dworkin, *Recommendation for Block Cipher Modes of Operation – Methods
 92 and Techniques*, NIST Special Publication 800-38A, Dec 2001,
 93 <http://csrc.nist.gov/publications/nistpubs/800-38a/sp800-38a.pdf>
- 94 [SP800-38B] M. Dworkin, *Recommendation for Block Cipher Modes of Operation: The CMAC
 95 Mode for Authentication*, NIST Special Publication 800-38B, May 2005,
 96 http://csrc.nist.gov/publications/nistpubs/800-38B/SP_800-38B.pdf
- 97

Deleted: .

Deleted: [RFC1424]. B. Kaliski, RFC1424: Privacy Enhancement for Internet Electronic Mail: Part IV: Key Certification and Related Services, February 1993. <http://www.ietf.org/rfc/rfc1424.txt>

Deleted: .

Deleted:

Deleted: .

Deleted: .

Deleted: .

Formatted: English (U.S.)

Deleted: .

Deleted: .

Deleted: .

Deleted: 2

Deleted: 14

- 98 [SP800-38C] M. Dworkin, *Recommendation for Block Cipher Modes of Operation: the CCM*
 99 *Mode for Authentication and Confidentiality*, NIST Special Publication 800-38C,
 100 May 2004, [http://csrc.nist.gov/publications/nistpubs/800-38C/SP800-](http://csrc.nist.gov/publications/nistpubs/800-38C/SP800-38C_updated-July20_2007.pdf)
 101 [38C_updated-July20_2007.pdf](http://csrc.nist.gov/publications/nistpubs/800-38C/SP800-38C_updated-July20_2007.pdf)
- 102 [SP800-38D] M. Dworkin, *Recommendation for Block Cipher Modes of Operation:*
 103 *Galois/Counter Mode (GCM) and GMAC*, NIST Special Publication 800-38D, Nov
 104 2007, <http://csrc.nist.gov/publications/nistpubs/800-38D/SP-800-38D.pdf>
- 105 [SP800-38E] M. Dworkin, *Recommendation for Block Cipher Modes of Operation: The XTS-*
 106 *AES Mode for Confidentiality on Block-Oriented Storage Devices*, NIST Special
 107 Publication 800-38E, Aug 2009 (draft), [http://csrc.nist.gov/publications/drafts/800-](http://csrc.nist.gov/publications/drafts/800-38E/draft-sp800-38E.pdf)
 108 [38E/draft-sp800-38E.pdf](http://csrc.nist.gov/publications/drafts/800-38E/draft-sp800-38E.pdf)
- 109 [SP800-57-1] E. Barker, W. Barker, W. Burr, W. Polk, and M. Smid, *Recommendations for Key*
 110 *Management - Part 1: General (Revised)*, NIST Special Publication 800-57 part
 111 1, March 2007, [http://csrc.nist.gov/publications/nistpubs/800-57/sp800-57-Part1-](http://csrc.nist.gov/publications/nistpubs/800-57/sp800-57-Part1-revised2_Mar08-2007.pdf)
 112 [revised2_Mar08-2007.pdf](http://csrc.nist.gov/publications/nistpubs/800-57/sp800-57-Part1-revised2_Mar08-2007.pdf)
- 113 [SP800-67] W. Barker, *Recommendation for the Triple Data Encryption Algorithm (TDEA)*
 114 *Block Cipher*, NIST Special Publication 800-67, Version 1.1, Revised 19 May
 115 2008, <http://csrc.nist.gov/publications/nistpubs/800-67/SP800-67.pdf>
- 116 [SP800-108] L. Chen, *Recommendation for Key Derivation Using Pseudorandom Functions*
 117 *(Revised)*, NIST Special Publication 800-108, October 2009,
 118 <http://csrc.nist.gov/publications/nistpubs/800-108/sp800-108.pdf>
- 119 [X.509] International Telecommunication Union (ITU)-T, X.509: Information technology
 120 – Open systems interconnection – The Directory: Public-key and attribute
 121 certificate frameworks, August 2005, [http://www.itu.int/rec/T-REC-X.509-200508-](http://www.itu.int/rec/T-REC-X.509-200508-l/en)
 122 [l/en](http://www.itu.int/rec/T-REC-X.509-200508-l/en)
- 123 [X9.24-1] ANSI, *X9.24 - Retail Financial Services Symmetric Key Management - Part 1:*
 124 *Using Symmetric Techniques*, 2004.
- 125 [X9.26] ANSI, *X9.26 - Financial Institution Sign-On Authentication for Wholesale*
 126 *Financial Transaction*, 1996.
- 127 [X9.31] ANSI, *X9.31-1992: Public Key Cryptography Using Reversible Algorithms for the*
 128 *Financial Services Industry: Part 2: The MDC-2 Hash Algorithm*, June 1993.
- 129 [X9.42] ANSI, *X9-42: Public Key Cryptography for the Financial Services Industry:*
 130 *Agreement of Symmetric Keys Using Discrete Logarithm Cryptography*, 2003.
- 131 [X9-57] ANSI, *X9-57: Public Key Cryptography for the Financial Services Industry:*
 132 *Certificate Management*, 1997.
- 133 [X9.62] ANSI, *X9-62: Public Key Cryptography for the Financial Services Industry, The*
 134 *Elliptic Curve Digital Signature Algorithm (ECDSA)*, 2005.
- 135 [X9-63] ANSI, *X9-63: Public Key Cryptography for the Financial Services Industry, Key*
 136 *Agreement and Key Transport Using Elliptic Curve Cryptography*, 2001.
- 137 [X9-102] ANSI, *X9-102: Symmetric Key Cryptography for the Financial Services Industry -*
 138 *Wrapping of Keys and Associated Data*, 2008.
- 139 [X9 TR-31] ANSI, *X9 TR-31: Interoperable Secure Key Exchange Key Block Specification for*
 140 *Symmetric Algorithms*, 2005.

Deleted: .
Formatted: Font: 10 pt

Deleted: [X9.24-1] . ANSI, X9.24 - Retail Financial Services Symmetric Key Management - Part 1: Using Symmetric Techniques, 2004.¶

Formatted: Bullets and Numbering

1.3 Non-normative References

- 141 [KMIP-UC] OASIS Draft, *Key Management Interoperability Protocol Use Cases v1,0*,
 142 Committee Draft, October 2009.

Deleted: 02
Deleted: , URI (TBD)

2 Assumptions

145 The section describes assumptions that underlie the KMIP protocol and the implementation of clients and
 146 servers that utilize the protocol.
 147

Deleted: 2
Deleted: 14

148 **2.1 Island of Trust**

149 Clients are provided key material by the server, but they only use that keying material for the purposes
150 explicitly listed in the delivery payload. Clients that ignore these instructions and use the keys in ways not
151 explicitly allowed by the server are non-compliant. There is no requirement for the key management
152 system, however, to enforce this behavior.

153 **2.2 Message Security**

154 KMIP relies on the chosen authentication suite as specified in **[KMIP-Prof]** to authenticate the client and
155 on the underlying transport protocol to provide confidentiality, integrity, message authentication and
156 protection against replay attack. KMIP offers a wrapping mechanism for the Key Value that does not rely
157 on the transport mechanism used for the messages; the wrapping mechanism is intended for importing or
158 exporting managed objects.

159 **2.3 State-less Server**

160 The protocol operates on the assumption that the server is state-less, which means that there is no
161 concept of “sessions” inherent in the protocol. State-less server operation is much more reliable and
162 easier to implement than stateful operation, and is consistent with possible implementation scenarios,
163 such as web-services-based servers. This does not mean that the server itself maintains no state, only
164 that the protocol does not require this.

165 **2.4 Extensible Protocol**

166 The protocol provides for “private” or vendor-specific extensions, which allow for differentiation among
167 vendor implementations. However, any objects, attributes and operations included in an implementation
168 are always implemented as specified, regardless of whether they are optional or mandatory.

169 **2.5 Support for Cryptographic Objects**

170 The protocol supports all reasonable key management system related cryptographic objects. This list
171 currently includes:

- 172 • Symmetric Keys
- 173 • Split (multi-part) Keys
- 174 • Asymmetric Key Pairs and their components
- 175 • Digital Certificates
- 176 • Derived Keys
- 177 • Secret Data
- 178 • Opaque (non-interpretable) cryptographic objects

179 **2.6 Client-Server Message-based Model**

180 The protocol operates primarily in a client-server, message-based model (the exceptions are the Put and
181 Notify operations). This means that most protocol exchanges are initiated by a client sending a request
182 message to a server, which then sends a response to the client. The protocol also provides optional
183 mechanisms to allow for unsolicited notification of events to clients, and unsolicited delivery of
184 cryptographic objects to clients, that is, the protocol allows a “push” model. These latter features are
185 optionally supported by servers and clients. Clients may register in order to receive such
186 events/notifications. Registration is implementation specific and not described in the specification.

Deleted: 2

Deleted: 14

187 **2.7 Synchronous and Asynchronous Messages**

188 The protocol allows two modes of operation. Synchronous (mandatory) operations are those in which a
189 client sends a request and waits for a response from the server. Polled Asynchronous operations
190 (optional) are those in which the client sends a request, the server responds with a “pending” status, and
191 the client polls the server for the completed response and completion status. Server implementations may
192 choose not to support the Polled Asynchronous feature of the protocol.

193 **2.8 Support for “Intelligent Clients” and “Key Using Devices”**

194 The protocol supports intelligent clients, such as end-user workstations, which are capable of requesting
195 all of the functions of KMIP. It also allows subsets of the protocol, and possible alternate message
196 representations, in order to support less-capable devices, which only need a subset of the features of
197 KMIP.

198 **2.9 Batched Requests and Responses**

199 The protocol contains a mechanism for sending batched requests and receiving the corresponding
200 batched responses, to allow for higher throughput on operations that deal with a large number of entities,
201 e. g., requesting dozens or hundreds of keys from a server at one time, and performing operations in a
202 group. An option is provided to indicate whether to continue processing requests after an earlier one fails
203 or to stop processing the remaining requests in the batch. Note that there is no option to treat an entire
204 batch as atomic, that is, if a request in the batch fails, then preceding requests in the batch are undone or
205 rolled back. A special ID Placeholder (see Section 3.19) is provided in KMIP to allow related requests in
206 a batch to be pipelined.

207 **2.10 Reliable Message Delivery**

208 The reliable message delivery function is relegated to the transport protocol, and is not part of the key
209 management protocol itself.

210 **2.11 Large Responses**

211 For requests that are capable of large responses, a mechanism in the protocol allows a client to specify in
212 a request the maximum allowed size of a response. The server indicates in a response to such a request
213 that the response would have been too large and, therefore, is not returned.

214 **2.12 Key Life-cycle and Key State**

215 **[KMIP-Spec]** describes the key life-cycle model, based on the NIST SP 800-57 key state definitions
216 **[SP800-57-1]**, supported by the KMIP protocol. Particular implications of the key life-cycle model in terms
217 of defining time-related attributes of objects are discussed in Section 3.5 below.

218

219 **3 Usage Guidelines**

220 This section provides guidance on using the functionality described in the Key Management
221 Interoperability Protocol Specification.

222 **3.1 Authentication**

223 As discussed in **[KMIP-Spec]**, a conforming KMIP implementation establishes and maintains channel
224 confidentiality and integrity, and proves server authenticity for KMIP messaging. Client authentication is
225 performed according to the chosen KMIP authentication suite as specified in **[KMIP-Prof]**. Other
226 mechanisms for client and server authentication are possible and optional for KMIP implementations.

Deleted: 2
Deleted: 14

227 KMIP implementations that use other vendor-specific mechanisms for authentication may use the
228 Credential attribute to include additional identification information. Depending on the server's
229 configuration, the server may interpret the identity of the requestor from the Credential object if it is not
230 provided during the channel level authentication. For example, in addition to performing mutual
231 authentication during SSL/TLS, the client passes the Credential object (e.g. username and password) in
232 the request. If the requestor's username is not specified inside the client certificate and is instead
233 specified in the Credential object, the server interprets the identity of the requestor from the Credential
234 object. This supports use cases where channel level authentication authenticates a machine or service
235 that is used by multiple users of the KMIP server. If the client provides the username of the requestor in
236 both the client certificate and the Credential object, the server verifies that the usernames are the same. If
237 they differ, the authentication fails and the server returns an error. If no Credential object is included in the
238 request, the username of the requestor is expected to be provided inside the certificate.

Deleted: However, if this authentication option is not part of the chosen KMIP authentication suite, it should not be used to assert that an identity has been authenticated or be used as an alternative to the chosen KMIP authentication suite.

239 If it is possible to return an "authentication not successful" error, it should be returned in preference to any
240 other result status. This prevents status code probing by a client that is not able to authenticate.

241 Server decisions regarding which operations to reject if there is insufficiently strong authentication of the
242 client are not specified in the protocol. However, see Section 3.2 for particular operations for which
243 authentication and authorization are particularly important.

244 3.2 Authorization for Revoke, Recover, Destroy and Archive Operations

245 Neither authentication nor authorization is handled by the KMIP protocol directly. In particular, the
246 Credential attribute is not guaranteed to be an authenticated identity of the requesting client. However,
247 the authentication suite, as specified in **[KMIP-Prof]**, describes how the client identity is established for
248 KMIP-compliant implementations. This authentication is performed for all KMIP operations, with the single
249 exception of the Query operation.

250 Certain operations that may be requested by a client via KMIP, particularly Revoke, Recover, Destroy and
251 Archive, may have a significant impact on the availability of a key, on server performance and on key
252 security. When a server receives a request for one of these operations, it should ensure that the client
253 has authenticated its identity (see the Authentication Suites section in **[KMIP-Prof]**). The server should
254 also ensure that the client requesting the operation is an object creator, security officer or other identity
255 authorized to issue the request. It may also require additional authentication to ensure that the object
256 owner or a security officer has issued that request. Even with such authentication and authorization,
257 requests for these operations should be considered only a "hint" to the key management system, which
258 may or may not choose to act upon this request.

259 3.3 Using Notify and Put Operations

260 The Notify and Put operations are the only operations in the KMIP protocol that are initiated by the server,
261 rather than the client. As client-initiated requests are able to perform these functions (e.g., by polling to
262 request notification), these operations are optional for conforming KMIP implementations. However, they
263 provide a mechanism for optimized communication between KMIP servers and clients and have,
264 therefore, been included in **[KMIP-Spec]**.

Deleted: As the functionality provided by these operations are able to be accomplished through client-initiated requests (using a polling model from the client to request notification, for example)

265 In using Notify and Put, the following constraints and guidelines should be observed:

- 266 • The client registers with the server, so that the server knows how to locate the client to which a
267 Notify or Put is being sent and which events for the Notify are supported. However, such
268 registration is outside the scope of the KMIP protocol. Registration also includes a specification of
269 whether a given client supports Put and Notify, and what attributes may be included in a Put for a
270 particular client.
- 271 • Communication between the client and the server is properly authenticated to forestall man-in-
272 the-middle attacks in which the client receives Notify or Put operations from an unauthenticated
273 server. Authentication for a particular client/server implementation is at a minimum accomplished
274 using one of the mandatory authentication mechanisms (see **[KMIP-Prof]**). Further strengthening
275 of the client/server communications integrity by means of signed message content and/or
276 wrapped keys is recommended. Attribute values other than "Last Change Date" should not be
277 included in a Notify to minimize risk of exposure of attribute information.

Deleted: 2

Deleted: 14

- 278
- 279
- 280
- 281
- 282
- 283
- 284
- 285
- 286
- In order to minimize possible divergence of key or state information between client and server as a result of server-initiated communication, any client receiving Notify or Put messages returns acknowledgements of these messages to the server. This acknowledgement may be at communication layers below the KMIP layer, such as by using transport-level acknowledgement provided in TCP/IP.
 - For client devices that are incapable of responding to messages from the server, communication with the server happens via a proxy entity that communicates with the server, using KMIP, on behalf of the client. It is possible to secure communication between a proxy entity and the client using other, potentially proprietary mechanisms.

287 3.4 Usage Allocation

288 Usage should be allocated and handled carefully, since power outages or other types of client failures
289 (crashes) may render allocated usage lost. For example, in the case of a key being used for the
290 encryption of tapes, such a loss of the usage allocation information following a client failure during
291 encryption may result in the necessity for the entire tape backup session to be re-encrypted using a
292 different key, if the server is not able to allocate more usage. It is possible to address this through such
293 approaches as caching usage allocation information on stable storage at the client, and/or having
294 conservative allocation policies at the server (e.g., by keeping the maximum possible usage allocation per
295 client request moderate). In general, usage allocations should be as small as possible; it is preferable to
296 use multiple smaller allocation requests rather than a single larger request to minimize the likelihood of
297 unused allocation.

298 3.5 Key State and Times

299 **[KMIP-Spec]** provides a number of time-related attributes, including the following:

- 300
- 301
- 302
- 303
- 304
- 305
- 306
- 307
- 308
- 309
- 310
- 311
- 312
- 313
- 314
- 315
- 316
- Initial Date: The date and time when the managed cryptographic object was first created or registered at the server
 - Activation Date: The date and time when the managed cryptographic object may begin to be used for applying cryptographic protection to data
 - Process Start Date: The date and time when a managed symmetric key object may begin to be used for processing cryptographically protected data
 - Protect Stop Date: The date and time when a managed symmetric key object may no longer be used for applying cryptographic protection to data
 - Deactivation Date: The date and time when the managed cryptographic object may no longer be used for any purpose, except for decryption, signature verification, or unwrapping, but only under extraordinary circumstances and when special permission is granted
 - Destroy Date: The date and time when the managed cryptographic object was destroyed
 - Compromise Occurrence Date: The date and time when the managed cryptographic object was first believed to be compromised
 - Compromise Date: The date and time when the managed cryptographic object is entered into the compromised state
 - Archive Date: The date and time when the managed object was placed in Off-Line storage

317 These attributes apply to all cryptographic objects (symmetric keys, asymmetric keys, etc) with exceptions
318 as noted in **[KMIP-Spec]**. However, certain of these attributes (such as the Initial Date) are not specified
319 in template-related objects and are implicitly set by the server.

320 In using these attributes, the following guidelines should be observed:

- 321
- 322
- 323
- 324
- 325
- As discussed for each of these attributes in Section 3 of **[KMIP-Spec]**, a number of these times are set once and it is not possible for the client or server to modify these. However, several of the time attributes (particularly the Activation Date, Protect Start Date, Process Stop Date and Deactivation Date) may be set by the server and/or requested by the client. Coordination of time-related attributes between client and server, therefore, is primarily the responsibility of the server,

Deleted: 2

Deleted: 14

326 as it manages the cryptographic object and its state. However, special conditions related to time-
327 related attributes, governing when the server accepts client modifications to time-related
328 attributes, may be negotiated by policy exchange between the client and server, outside the Key
329 Management Interoperability Protocol.

330
331 In general, state transitions occur as a result of operational requests. However, clients may need
332 to specify times in the future for such things as Activation Date, Deactivation Date, Process Start
333 Date, and Protect Stop Date.

334
335 KMIP allows clients to specify times in the past for such attributes as Activation Date and
336 Deactivation Date. This is intended primarily for clients that were disconnected from the server at
337 the time that the client performed that operation on a given key.

- 338 • It is valid to have a Deactivation Date when there is no Activation Date. This means, however,
339 that the key is not yet active, even though its Deactivation Date has been specified. A valid
340 Deactivation Date is greater than or equal to the Activation Date.
- 341 • The Protect Stop Date may be equal to, but may not be later than the Deactivation Date.
342 Similarly, Process Start Date may be equal to, but may not precede, the Activation Date. KMIP
343 implementations should consider specifying both these attributes, particularly for symmetric keys,
344 as a key may be needed for decryption (process) long after it is no longer appropriate to use it for
345 encryption of new objects (protect).
- 346 • If a Destroy operation is performed, resulting in the Destroy Date being set, and the object has
347 not already been deactivated, the deactivation of the object is also performed prior to the Destroy
348 operation, so that Destroy Date is greater than or equal to the Deactivation Date. If other related
349 attributes (e.g., Protect Stop Date) have not already been set, the server should set these to the
350 deactivation date.

Deleted: then they should be set

- 351 • After a cryptographic object is destroyed, a key management server may retain certain
352 information about the object, such as the Unique Identifier.

353 KMIP allows the specification of attributes on a per-client basis, such that a server could maintain or
354 present different sets of attributes for different clients. This flexibility may be necessary in some cases,
355 such as when a server maintains the availability of a key for some clients even after a key moved to
356 inactive state for most clients. However, such an approach might result in significant inconsistencies
357 regarding the object state from the point of view of all participating clients and should, therefore, be
358 avoided. A server should maintain a consistent state for each object, across all clients that have or are
359 able to request that object.

360 3.6 Template

361 It is possible for a server to maintain different policy templates for different clients. As in the state
362 transitions described above, however, this practice is discouraged.

363 3.7 Archive Operations

364 When the Archive operation is performed, it is recommended that an object identifier and a minimal set of
365 attributes be retained within the server for operational efficiency. In such a case, the retained attributes
366 may include Unique Identifier and State.

367 3.8 Message Extensions

368 Any number of vendor-specific extensions may be included in the Message Extension optional structure.
369 This allows KMIP implementations to create multiple extensions to the protocol.

370 3.9 Unique Identifiers

371 For clients that require unique identifiers in a special form, out-of-band registration/configuration may be
372 used to communicate this requirement to the server.

Deleted: 2

Deleted: 14

373 **3.10 Result Message Text**

374 KMIP specifies the Result Status, the Result Reason and the Result Message as normative message
375 contents. For the Result Status and Result Reason, the enumerations provided in **[KMIP-Spec]** are the
376 normative values. The values for the Result Message text, on the other hand, are implementation-
377 specific. In consideration of internationalization, it is recommended that any vendor implementation of
378 KMIP provide appropriate language support for the Return Message. How a client specifies the language
379 for Result Messages is outside the scope of the KMIP.

380 **3.11 Query**

381 Query does not explicitly support client requests to determine what operations require authentication. To
382 determine whether an operation requires authentication, a client should request that operation.

383 **3.12 Canceling Asynchronous Operations**

384 If an asynchronous operation is cancelled by the client, no information is returned by the server in the
385 result code regarding any operations that may have been partially completed. Identification and
386 remediation of partially completed operations is the responsibility of the server.

387 It is the responsibility of the server to determine when to discard the status of asynchronous operations.
388 The determination of how long a server should retain the status of an asynchronous operation is
389 implementation-dependent and not defined by KMIP.

390 Once a client has received the status on an asynchronous operation other than “pending”, any
391 subsequent request for status of that operation may return either the same status as in a previous polling
392 request or an “unavailable” response.

393 **3.13 Multi-instance Hash**

394 The Digest attribute contains the output of hashing a managed object, such as a key or a certificate. The
395 server always generates the SHA-256 hash value when the object is created or generated. KMIP allows
396 multiple instances of the digest attribute to be associated with the same managed object. For example, it
397 is common practice for public trusted CAs to publish two digests (often referred to as the fingerprint or the
398 thumbprint) of their certificate one calculated using the SHA-1 algorithm and another using the MD-5
399 algorithm. In this case, each digest would be calculated by the server using a different hash algorithm.

400 **3.14 Returning Related Objects**

401 The key block is intended to return a single object, with associated attributes and other data. For those
402 cases in which multiple related objects are needed by a client, such as the private key and the related
403 certificate specified by RACF and JKS, the client should issue multiple Get requests to obtain these
404 related objects.

405 **3.15 Reducing Multiple Requests through Use of Batch**

406 KMIP supports batch operations in order to reduce the number of calls between the client and server for
407 related operations. For example, Locate and Get are likely to be commonly accomplished within a single
408 batch request.

409 KMIP does not ensure that batch operations are atomic on the server side. If servers implement such
410 atomicity, the client is able to use the optional “undo” mode to request roll-back for batch operations
411 implemented as atomic transactions. However, support for “undo” mode is optional in the protocol, and
412 there is no guarantee that a server that supports “undo” mode has effectively implemented atomic
413 batches. The use of “undo”, therefore, should be restricted to those cases in which it is possible to assure
414 the client, through mechanisms outside of KMIP, of the server effectively supporting atomicity for batch
415 operations.

Deleted: 2
Deleted: 14

416 **3.16 Maximum Message Size**

417 When a server is processing requests in a batch, it should compare the response size after each request
418 with the specified Maximum Response Size. If the message is too large, it should prepare a maximum
419 message size response at that point, rather than continuing with operations in the batch. This increases
420 the client's ability to understand what operations have and have not been completed.

421 When processing individual requests within the batch, the server that has encountered a Maximum
422 Response Size error should not return attribute values or other information as part of the error response.

423 **3.17 Using Offset in Re-key and Re-certify Operations**

424 Both the Re-key and the Re-certify operations allow the specification of an offset interval.

425 The Re-key operation allows the client to specify an offset interval for activation of the key. This offset
426 specifies the duration of time between the time the request is made and when the activation of the key
427 occurs. If an offset is specified, all other times for the new key is determined from the new Activation
428 Date, based on the intervals used by the original key, i.e., from the Activation Date to the Process
429 Start Date, Protect Stop Date, etc.

430 The Re-certify operation allows the client to specify an offset interval that indicates the difference between
431 the Initial Date of the new certificate and the Activation Date of the new certificate. As with the Re-key
432 operation, all other times for the certificate are determined using the intervals used for the original
433 certificate.

434 **3.18 Locate Queries**

435 It is possible to formulate Locate queries to address any of the following conditions:

- 436 • Exact match of a transition to a given state. Locate the key(s) with a transition to a certain state at
437 a specified time (t).
- 438 • Range match of a transition to a given state. Locate the key(s) with a transition to a certain state
439 at any time at or between two specified times (t and t').
- 440 • Exact match of a state at a specified time. Locate the key(s) that are in a certain state at a
441 specified time (t).
- 442 • Match of a state during an entire time range. Locate the key(s) that are in a certain state during
443 an entire time specified with times (t and t'). Note that the Activation Date could occur at or before
444 t and that the Deactivation Date could occur at or after t'+1.
- 445 • Match of a state at some point during a time range. Locate the key(s) that are in a certain state at
446 some time at or between two specified times (t and t'). In this case, the transition to that state
447 could be before the start of the specified time range.

448 This is accomplished by allowing any date/time attribute to be present either once (for an exact match) or
449 at most twice (for a range match).

450 For instance, if the state we are interested in is Active, the Locate queries would be the following
451 (corresponding to the bulleted list above):

- 452 • Exact match of a transition to a given state: Locate (ActivationDate(t)). Locate keys with an
453 Activation Date of t.
- 454 • Range match of a transition to a given state: Locate (ActivationDate(t), ActivationDate(t')). Locate
455 keys with an Activation Date at or between t and t'.
- 456 • Exact match of a state at a specified time: Locate (ActivationDate(0), ActivationDate(t),
457 DeactivationDate(t+1), DeactivationDate(MAX_INT), CompromiseDate(t+1),
458 CompromiseDate(MAX_INT)). Locate keys in the Active state at time t, by looking for keys with a
459 transition to Active before or until t, and a transition to Deactivated or Compromised after t
460 (because we don't want the keys that have a transition to Deactivated or Compromised before t).
461 The server assumes that keys without a DeactivationDate or CompromiseDate is equivalent to
462 MAX_INT (i.e., infinite).

Deleted: 2

Deleted: 14

- 463 • Match of a state during an entire time range: Locate (ActivationDate(0), ActivationDate(t),
464 DeactivationDate(t'+1), DeactivationDate(MAX_INT), CompromiseDate(t'+1),
465 CompromiseDate(MAX_INT)). Locate keys in the Active state during the entire time from t to t'.
- 466 • Match of a state at some point during a time range: Locate (ActivationDate(0), ActivationDate(t'-
467 1), DeactivationDate(t+1), DeactivationDate(MAX_INT), CompromiseDate(t+1),
468 CompromiseDate(MAX_INT)). Locate keys in the Active state at some time from t to t', by looking
469 for keys with a transition to Active between 0 and t'-1 and exit out of Active on or after t+1.

470 The queries would be similar for Initial Date, Deactivation Date, Compromise Date and Destroy Date.

471 In the case of the Destroyed-Compromise state, there are two dates recorded: the Destroy Date and the
472 Compromise Date. For this state, the Locate operation would be expressed as follows:

- 473 • Exact match of a transition to a given state: Locate (CompromiseDate(t), State(Destroyed-
474 Compromised)) and Locate (DestroyDate(t), State(Destroyed-Compromised)). KMIP does not
475 support the OR in the Locate request, so two requests should be issued. Locate keys that were
476 Destroyed and transitioned to the Destroyed-Compromised state at time t, and locate keys that
477 were Compromised and transitioned to the Destroyed-Compromised state at time t.
- 478 • Range match of a transition to a given state: Locate (CompromiseDate(t), CompromiseDate(t'),
479 State(Destroyed-Compromised)) and Locate (DestroyDate(t), DestroyDate(t'), State(Destroyed-
480 Compromised)). Locate keys that are Destroyed-Compromised and were Compromised or
481 Destroyed at or between t and t'.
- 482 • Exact match of a state at a specified time: Locate (CompromiseDate(0), CompromiseDate(t),
483 DestroyDate(0), DestroyDate(t)) nothing else needed since there is no exit transition. Locate keys
484 with a Compromise Date at or before t, and with a Destroy Date at or before t. These keys are
485 therefore in the Destroyed-Compromised state at time t.
- 486 • Match of a state during an entire time range: Locate (CompromiseDate(0), CompromiseDate(t),
487 DestroyDate(0), DestroyDate(t)). Same as above. As there is no exit transition from the
488 Destroyed-Compromised state, the end of the range (t') is irrelevant.
- 489 • Match of a state at some point during a time range: Locate (CompromiseDate(0),
490 CompromiseDate(t'-1), DestroyDate(0), DestroyDate(t'-1)). Locate keys with a Compromise Date
491 at or before t'-1, and with a Destroy Date at or before t'-1. As there is no exit transition from the
492 Destroyed-Compromised state, the start of the range (t) is irrelevant.

493 3.19 ID Placeholder

494 A number of operations are affected by a mechanism referred to as the ID Placeholder. This is a
495 temporary variable consisting of a single Unique Identifier that is stored inside the server for the duration
496 of executing the batch of operations. The ID Placeholder is obtained from the Unique Identifier returned
497 by certain operations; the applicable operations are identified in [Table 1](#), along with a list of operations
498 that accept the ID Placeholder as input.

Deleted: Table 1

	ID Placeholder input	ID Placeholder output (in case of operation failure, a batch using ID Placeholder stops)
Create	-	new Object
Create Key Pair	-	new Private Key (the new Public Key may be obtained in the batched via a Locate)
Register	-	new Object
Derive Key	- (because there may be more than one object)	new Symmetric Key

Deleted: 2

Deleted: 14

Locate	-	Object
Get	Object	no change
Request Object	Object	no change
Validate	-	-
Get Attributes List/Modify/Add/Delete	Object	no change
Activate	Object	no change
Revoke	Object	no change
Destroy	Object	no change
Archive/Recover	Object	no change
Certify	Public Key	new Certificate
Re-certify	Certificate	new Certificate
Re-key	Symmetric Key	new Symmetric Key
Obtain Lease	Object	no change
Get Usage Allocation	Keys	no change

Table 1: ID Placeholder Input and Output for KMIP Operations

499

500 3.20 Key Block

501 The protocol uses the Key Block structure to transport a key to the client or server. This Key Block
502 consists of the Key Value Type, the Key Value, and the Key Wrapping Data. The Key Value Type
503 identifies the format of the Key Material, e.g., Raw format or Transparent Key structure. The Key Value
504 consists of the Key Material and optional attributes. The Key Wrapping Data provides information about
505 the wrapping key and the wrapping mechanism, and is returned only if the client requests the Key Value
506 to be wrapped by specifying the Key Wrapping Specification inside the Get Request Payload. The Key
507 Wrapping Data may also be included inside the Key Block if the client registers a wrapped key.

508 The protocol allows any attribute to be included inside the Key Value and allows these attributes to be
509 cryptographically bound to the Key Material (i.e., by signing, MACing, encrypting, or both encrypting and
510 signing/MACing the Key Value). Some of the attributes that may be included include the following:

- 511 • Unique Identifier – uniquely identifies the key
- 512 • Cryptographic Algorithm (e.g., AES, 3DES, RSA) – this attribute is either specified inside the Key
513 Block structure or the Key Value structure.
- 514 • Cryptographic Length (e.g., 128, 256, 2048) – this attribute is either specified inside the Key
515 Block structure or the Key Value structure
- 516 • Cryptographic Usage Mask– identifies the cryptographic usage of the key (e.g., Encrypt, Wrap
517 Key, Export)
- 518 • Cryptographic Parameters – provides additional parameters for determining how the key may be
519 used
 - 520 – Block Cipher Mode (e.g., CBC, NISTKeyWrap, GCM) – this parameter identifies the mode of
521 operation, including block cipher based MACs or wrapping mechanisms
 - 522 – Padding Method (e.g., OAEP, X9.31, PSS) – identifies the padding method and if applicable
523 the signature or encryption scheme.

Deleted: 2

Deleted: 14

- 524 – Hashing Algorithm (e.g., SHA-256) – identifies the hash algorithm to be used with the
- 525 signature/encryption mechanism or Mask Generation Function; note that the different HMACs
- 526 are defined individually as algorithms and do not require the Hashing Algorithm parameter to
- 527 be set
- 528 – Role Type – Identifies the financial key role (e.g., DEK, KEK)
- 529 • State (e.g., Active)
- 530 • Dates (e.g., Activation Date, Process Start Date, Protect Stop Date)
- 531 • Custom Attribute – allows vendors and clients to define vendor-specific attributes; may also be
- 532 used to prevent replay attacks by setting a nonce

533 3.21 Using Wrapped Keys with KMIP

534 KMIP provides the option to register and get keys in wrapped format. Clients request the server to return
 535 a wrapped key by including the Key Wrapping Specification in the Get Request Payload. Similarly, clients
 536 register a wrapped key by including the Key Wrapping Data in the Register Request Payload. The
 537 Wrapping Method identifies the type of mechanism used to wrap the key, but does not identify the
 538 algorithm or block cipher mode. It is possible to determine these from the attributes set for the specified
 539 Encryption Key or MAC/Signing Key. If a key has multiple Cryptographic Parameters set, clients may
 540 include the applicable parameters in Key Wrapping Specification. If omitted, the server chooses the
 541 Cryptographic Parameter attribute with the lowest index.

542 The Key Value includes both the Key Material and, optionally, attributes of the key; these may be
 543 provided by the client during in the Register Request Payload; the server only includes attribute when
 544 requested in the Key Wrapping Specification of the Get Request Payload. The Key Value may be
 545 encrypted, signed/MACed, or both encrypted and signed/MACed (and vice versa). In addition, clients
 546 have the option to request or import a wrapped Key Block according to standards, such as ANSI TR-31,
 547 or vendor-specific key wrapping methods.

548 It is important to note that if the Key Wrapping Specification is included in the Get Request Payload, the
 549 Key Value may not necessarily be encrypted. If the Wrapping Method is MAC/sign, the returned Key
 550 Value is in plaintext and the Key Wrapping Data includes the MAC or Signature of the Key Value.

551 Prior to wrapping or unwrapping a key, the server should verify that the wrapping key is allowed to be
 552 used for the specified purpose. For example, if a symmetric key is used for key encryption in response to
 553 a Get request, the symmetric key should have the “Wrap Key” bit set in Cryptographic Usage Mask.
 554 Similarly, if the client registers a signed key, the server should verify that the Signature Key, as specified
 555 by the client inside Key Wrapper Data, has the “Verify” bit set in Cryptographic Usage Mask. If the
 556 wrapping key is not **permitted** to be used for the requested purpose (e.g., when the Cryptographic Usage
 557 Mask is not set), **the server should return an error.**

Deleted: able

Deleted: then

558 3.21.1 Encrypt-only Example with a Symmetric Key for a Get Request and 559 Response

560 The client sends a Get request to obtain a key that is stored on the server. When the client sends a Get
 561 request to the server, a Key Wrapping Specification may be included. If a Key Wrapping Specification is
 562 included in the Get request, and a client wants the requested key and its Cryptographic Usage Mask
 563 attribute to be wrapped using AES key wrap, clients include the following information in the Key Wrapping
 564 Specification:

- 565 • Wrapping Method: Encrypt
- 566 • Encryption Key Information
 - 567 – Unique Key ID: Key ID of the AES wrapping key
 - 568 – Cryptographic Parameters: The Block Cipher Mode is NISTKeyWrap (not necessary if default
 - 569 block cipher mode for wrapping key is NISTKeyWrap)
- 570 • Attribute Name: Cryptographic Usage Mask

Deleted: 2

Deleted: 14

571 The server uses the Unique Key ID specified by the client to determine the attributes set for the key. For
572 example, the algorithm of the wrapping key is not explicitly specified inside the Key Wrapping
573 Specification; the server determines the algorithm to be used for wrapping the key by identifying the
574 Algorithm attribute set for the specified Encryption Key.

575 The Cryptographic Parameters attribute should be specified by the client if multiple instances of the
576 Cryptographic Parameters exist, and the lowest index does not correspond to the NIST key wrap mode of
577 operation. The server should verify that the AES wrapping key has NISTKeyWrap set as an allowable
578 Block Cipher Mode, and that the "Wrap Key" bit is set in the Cryptographic Usage Mask.

579 If the correct data was provided to the server, and no conflicts exist, the server wraps the Key Value for
580 the requested key using the AES key wrap algorithm and wrapping key specified in the Encryption Key
581 Information; the Key Value contains both the Key Material and the Cryptographic Usage Mask attribute,
582 and return the encrypted result (octet string) as the Key Value in the Key Block of the server's response.

583 The Key Wrapping Data of the Key Block in the Get Response Payload includes the same data as
584 specified in the Key Wrapping Specification of the Get Request Payload except for the Attribute Name.

585 3.21.2 Encrypt-only Example with a Symmetric Key for a Register Request and 586 Response

587 The client sends a Register request to the server and includes the wrapped key and the unique ID of the
588 wrapping key inside the Request Payload. The wrapped key is provided to the server inside the Key
589 Block. The Key Block includes the Key Value Type, the Key Value, and the Key Wrapping Data. The Key
590 Value Type identifies the format of the Key Material, the Key Value consists of the Key Material and
591 optional attributes that may be included to cryptographically bind the attributes to the Key Material, and
592 the Key Wrapping Data identifies the encryption key used to wrap the object and the wrapping
593 mechanism.

594 Similar to example 3.21.1 the key is wrapped using the AES key wrap. The Key Value includes four
595 attributes: Cryptographic Algorithm, Cryptographic Length, Cryptographic Parameters, and Cryptographic
596 Usage Mask.

597 The Key Wrapping Data includes the following information:

- 598 • Wrapping Method: Encrypt
- 599 • Encryption Key Information
 - 600 – Unique Key ID: Key ID of the AES wrapping key
 - 601 – Cryptographic Parameters: The Block Cipher Mode is NISTKeyWrap (not necessary if default
 - 602 block cipher mode for wrapping key is NISTKeyWrap)

603 Attributes do not need to be specified in Key Wrapping Data. When registering a wrapped Key Value with
604 attributes, clients may include these attributes inside the Key Value without specifying them inside the
605 Template-Attribute.

606 Prior to unwrapping the key, the server determines the wrapping algorithm from the Algorithm attribute set
607 for the specified Unique ID in Encryption Key Information. The server verifies that the wrapping key may
608 be used for the specified purpose. In particular, if the client includes the Cryptographic Parameters in
609 Encryption Key Information, the server verifies that the specified Block Cipher Mode is set for the
610 wrapping key. The server also verifies that the wrapping key has the "Unwrap Key" bit set in the
611 Cryptographic Usage Mask.

612 The Register Response Payload includes the Unique ID of the newly registered key and an optional list of
613 attributes that were implicitly set by the server.

614 3.21.3 Encrypt-only Example with an Asymmetric Key for a Get Request and 615 Response

616 The client sends a Get request to obtain a key (either symmetric or asymmetric) that is stored on the
617 server. When the client sends a Get request to the server, a Key Wrapping Specification may be
618 included. If a Key Wrapping Specification is included, and the key is to be wrapped with an RSA public

Deleted: 2

Deleted: 14

619 key using the OAEP encryption scheme, the client includes the following information in the Key Wrapping
620 Specification. Note that for this example, attributes for the requested key are not requested.

- 621 • Wrapping Method: Encrypt
- 622 • Encryption Key Information
 - 623 – Unique Key ID: Key ID of the RSA public key
 - 624 – Cryptographic Parameters:
 - 625 Padding Method: OAEP
 - 626 Hashing Algorithm: SHA-256

627 The Cryptographic Parameters attribute is specified by the client if multiple instances of Cryptographic
628 Parameters exist for the wrapping key, and the lowest index does not correspond to the associated
629 padding method. The server should verify that the specified Cryptographic Parameters in the Key
630 Wrapping Specification and the “Wrap Key” bit in Cryptographic Usage Mask are set for the
631 corresponding wrapping key.

632 The Key Wrapping Data returned by the server in the Key Block of the Get Response Payload includes
633 the same data as specified in the Key Wrapping Specification of the Get Request Payload.

634 For both OAEP and PSS, KMIP currently assumes that the Hashing Algorithm specified in the
635 Cryptographic Parameters of the Get request is used for both the Mask Generation Function (MGF) and
636 hashing data. The example above requires the server to use SHA-256 for both purposes.

637 3.21.4 MAC-only Example with an HMAC Key for a Get Request and Response

638 The client sends a Get request to obtain a key that is stored on the server. When the client sends a Get
639 request to the server, a Key Wrapping Specification may be included. If a key and Custom Attribute (i.e.,
640 x-Nonce) is to be MACed with HMAC SHA-256, the following Key Wrapping Specification is specified:

- 641 • Wrapping Method: MAC/sign
- 642 • MAC/Signature Key Information
 - 643 – Unique Key ID: Key ID of the MACing key (note that the algorithm associated with this key
644 would be HMAC-256)
- 645 • Attribute Name: x-Nonce

646 For HMAC, no Cryptographic Parameters need to be specified, since the algorithm, including the hash
647 function, may be determined from the Algorithm attribute set for the specified MAC Key. The server
648 should verify that the HMAC key has the “MAC Generate” bit set in Cryptographic Usage Mask. Note that
649 an HMAC key does not require the “Wrap Key” bit to be set in the Cryptographic Usage Mask.

650 The server creates an HMAC value over the Key Value if the specified MACing key may be used for the
651 specified purpose and no conflicts exist. The Key Value is returned in plaintext and the Key Block
652 includes the following Key Wrapping Data:

- 653 • Wrapping Method: MAC/sign
- 654 • MAC/Signature Key Information
- 655 • Unique Key ID: Key ID of the MACing key
- 656 • MAC/Signature: HMAC result of the Key Value

657 In the example, the custom attribute x-Nonce was included to help clients, who are relying on the proxy
658 model, to detect replay attacks. End-clients, who communicate with the key management server, may not
659 support SSL/TLS and may not be able to rely on the message protection mechanisms provided by a
660 security protocol. A custom attribute may be created to hold a random number, counter, nonce, date, or
661 time. The custom attribute needs to be created before requesting the server to return a wrapped key and
662 is recommended to be set if clients frequently wrap/sign the same key with the same wrapping/signing
663 key.

Deleted: 2

Deleted: 14

664 **3.21.5 Registering a Wrapped Key as an Opaque Cryptographic Object**

665 Clients may want to register and store a wrapped key on the server without the server being able to
666 unwrap the key (i.e., the wrapping key is not known to the server). Instead of storing the wrapped key as
667 an opaque object, clients have the option to store the wrapped key inside the Key Block as an opaque
668 cryptographic object, i.e., the wrapped key is registered as a managed cryptographic object, but the
669 encoding of the key is unknown to the server. Registering an opaque cryptographic object allows clients
670 to set all the applicable attributes that apply to cryptographic objects (e.g., Cryptographic Algorithm and
671 Cryptographic Length),

672 Opaque cryptographic objects are set by specifying the following inside the Key Block structure:

- 673 • Key Format Type: Opaque
674 • Key Material: Wrapped key as a Byte String

675 The Key Wrapping Data does not need to be specified.

676 **3.22 Object Group**

677 The key management system may specify rules for the valid group names which may be created by the
678 client. Clients are informed of such rules by a mechanism that is not specified by **[KMIP-Spec]**. In the
679 protocol, the group names themselves are character strings of no specified format. Specific key
680 management system implementations may choose to support hierarchical naming schemes or other
681 syntax restrictions on the names. Groups may be used to associate objects for a variety of purposes. A
682 set of keys used for a common purpose, but for different time intervals, may be linked by a common
683 Object Group. Servers may create predefined groups and add objects to them independently of client
684 requests.

685 **3.23 Certify and Re-certify**

686 The key management system may contain multiple embedded CAs or may have access to multiple
687 external CAs. How the server routes a certificate request to a CA is vendor-specific and outside the scope
688 of KMIP. If the server requires and supports the capability for clients to specify the CA to be used for
689 signing a Certificate Request, then this information may be provided by including the Certificate Issuer
690 attribute in the Certify or Re-certify request.

691 **[KMIP-Spec]** supports multiple options for submitting a certificate request to the key management server
692 within a Certify or Re-Certify operation. It is vendor decision as to whether the key management server
693 offers certification authority (CA) functionality or proxy the certificate request on to a separate CA for
694 processing. It is also a vendor decision as to the type of certificate request formats it supports and this
695 may in part be based upon the request formats supported by any CA to which it proxies the certificate
696 requests.

697 All certificate request formats for requesting X.509 certificates specified in **[KMIP-Spec]** (i.e., PKCS#10,
698 PEM and CRMF) provide a means for allowing the CA to verify that the client that created the certificate
699 request possess the private key corresponding to the public key in the certificate request. This is referred
700 to as Proof-of-Possession (POP). However, it should be noted that in the case of the CRMF format that
701 some CAs may not support the CRMF POP option, but instead rely upon the underlying certificate
702 management protocols (i.e., CMP and CMC) to provide POP. In the case where the CA (including CA
703 functionality within the key management server) does not support POP via the CRMF format an
704 alternative certificate request format (i.e., PKCS#10, PEM) would need to be used if POP needs to be
705 verified.

706 **3.24 Specifying Attributes during Create Key Pair**

707 The Create Key Pair operation allows clients to specify attributes using the Common Template-Attribute,
708 Private Key Template-Attribute, and Public Key Template-Attribute. The Common Template-Attribute
709 object includes a list of attributes that apply to both the public and private key. Attributes that are not
710 common to both keys may be specified using the Private Key Template-Attribute or Public Key Template-

Deleted: 2

Deleted: 14

711 Attribute. If a single-instance attribute is specified in multiple Template-Attribute objects, the server obeys
712 the following order of precedence:

- 713 1. Attributes specified explicitly in the Private and Public Key Template-Attribute, then
- 714 2. Attributes specified via templates in the Private and Public Key Template-Attribute, then
- 715 3. Attributes specified explicitly in the Common Template-Attribute, then
- 716 4. Attributes specified via templates in the Common Template-Attribute

717 **3.24.1 Example of Specifying Attributes during Create Key Pair**

718 A client specifies several attributes in the Create Key Pair Request Payload. The Common Template-
719 Attribute includes the template name RSACom and other explicitly specified common attributes:

720 Common Template-Attribute

- 721 • RSACom Template
 - 722 – Cryptographic Algorithm: RSA
 - 723 – Cryptographic Length: 2048
 - 724 – Cryptographic Parameters: Padding Method OAEP
 - 725 – Custom Attribute: x-Serial 1234
 - 726 – Object Group: Key encryption group 1
- 727 • Attribute
 - 728 – Cryptographic Length: 4096
 - 729 – Cryptographic Parameters: Padding Method PKCS1 v1.5
 - 730 – Custom Attribute: x-ID 56789

731 The Private Key Template-Attribute includes the template name RSAPriv and other explicitly specified
732 private key attributes:

733 Private Key Template-Attribute

- 734 • RSAPriv Template
 - 735 – Object Group: Key encryption group 2
- 736 • Attribute
 - 737 – Cryptographic Usage Mask: Unwrap Key
 - 738 – Name: PrivateKey1

739 The Public Key Template Attribute includes explicitly specified public key attributes:

740 Public Key Template-Attribute

- 741 • Attribute
 - 742 – Cryptographic Usage Mask: Wrap Key
 - 743 – Name: PublicKey1

744 Following the attribute precedence rule, the server creates a 4096-bit RSA key. The following client-
745 specified attributes are set:

747 Private Key

- 748 • Cryptographic Algorithm: RSA
- 749 • Cryptographic Length: 4096
- 750 • Cryptographic Parameters: OAEP
- 751 • Cryptographic Parameters: PKCS1 v1.5
- 752 • Cryptographic Usage Mask: Unwrap Key

Deleted: 2

Deleted: 14

- 753 • Custom Attribute: x-Serial 1234
- 754 • Custom Attribute: x-ID 56789
- 755 • Object Group: Key encryption group 1
- 756 • Object Group: Key encryption group 2
- 757 • Name: PrivateKey1

758 Public Key

- 759 • Cryptographic Algorithm: RSA
- 760 • Cryptographic Length: 4096
- 761 • Cryptographic Parameters: OAEP
- 762 • Cryptographic Parameters: PKCS1 v1.5
- 763 • Cryptographic Usage Mask: Wrap Key
- 764 • Custom Attribute: x-Serial 1234
- 765 • Custom Attribute: x-ID 56789
- 766 • Object Group: Key encryption group 1
- 767 • Name: PublicKey1

768 **3.25 Registering a Key Pair**

769 During a Create Key Pair operation, a Link Attribute is automatically created by the server for each object
 770 (i.e., a link is created from the private key to the public key and vice versa). Certain attributes are the
 771 same for both objects and are set by the server while creating the key pair. The KMIP protocol does not
 772 support an equivalent operation for registering a key pair. Clients are able to register the objects
 773 independently and manually set the Link attributes to make the server aware that these keys are
 774 associated with each other. When the Link attribute is set for both objects, the server should verify that
 775 the registered objects indeed correspond to each other and apply similar restrictions as if the key pair was
 776 created on the server.

777 Clients should perform the following steps when registering a key pair:

- 778 1. Register the public key and set all associated attributes:
 - 779 a. Cryptographic Algorithm
 - 780 b. Cryptographic Length
 - 781 c. Cryptographic Usage Mask
- 782 2. Register the private key and set all associated attributes
 - 783 a. Cryptographic Algorithm is the same for both public and private key
 - 784 b. Cryptographic Length is the same for both public and private key
 - 785 c. Cryptographic Parameters may be set; if set, the value is the same for both the public and
 786 private key
 - 787 d. Cryptographic Usage Mask is set, but does not contain the same value for both the public
 788 and private key
 - 789 e. Link is set with Link Type *Public Key Link* and the Linked Object Identifier of the
 790 corresponding Public Key
 - 791 f. Link is set for the Public Key with Link Type *Private Key Link* and the Linked Object Identifier
 792 of the corresponding Private Key

Deleted: 2
 Deleted: 14

793 **3.26 Non-Cryptographic Objects**

794 The KMIP protocol allows clients to register Secret Data objects. Secret Data objects may include
795 passwords or data that are used to derive keys.

796 KMIP defines Secret Data as cryptographic objects. Even if the object is not used for cryptographic
797 purposes, clients still set certain attributes, such as the Cryptographic Usage Mask, for this object unless
798 otherwise stated. Similarly, servers set certain attributes for this object, including the Digest, State, and
799 certain Date attributes, even if the attributes seem relevant only for cryptographic objects.

800 When registering a Secret Data object, the following attributes are set by the server:

- 801 • Unique Identifier
- 802 • Object Type
- 803 • Digest
- 804 • State
- 805 • Initial Date
- 806 • Last Change Date

807 When registering a Secret Data object for non-cryptographic purposes, the following attributes are set by
808 either client or server:

- 809 • Cryptographic Usage Mask

810 **3.27 Asymmetric Concepts with Symmetric Keys**

811 The Cryptographic Usage Mask attribute is intended to adequately support asymmetric concepts using
812 symmetric keys. This is fairly common practice in established crypto systems: the MAC is an example of
813 an operation where a single symmetric key is used at both ends, but policy dictates that one end may
814 only generate cryptographic tokens using this key (the MAC) and the other end may only verify tokens.
815 Security of the system fails if the verifying end is able to use the key to perform generate operations.

816 In these cases it is not sufficient to describe the usage policy on the keys in terms of cryptographic
817 primitives like "encrypt" vs. "decrypt" or "sign" vs. "verify". There are two reasons why this is the case.

- 818 • In some of these operations, such as MAC generate and verify, the same cryptographic primitive
819 is used in both of the complementary operations. MAC generation involves computing and
820 returning the MAC, while MAC verification involves computing that same MAC and comparing it
821 to a supplied value to determine if they are the same. Thus, both generation and verification use
822 the "encrypt" operation and the two usages are not able to be distinguished by considering only
823 "encrypt" vs. "decrypt".
- 824 • Some operations which require separate key types use the same fundamental cryptographic
825 primitives. For example, encryption of data, encryption of a key, and computation of a MAC all
826 use the fundamental operation "encrypt", but in many applications securely differentiated keys are
827 used for these three operations. Simply looking for an attribute that permits "encrypt" is not
828 sufficient.

829 Allowing use of these keys outside of their specialized purposes may compromise security. Instead,
830 specialized application-level permissions are necessary to control the use of these keys. KMIP provides
831 several pairs of such permissions in the Cryptographic Usage Mask (3.14), such as:

MAC GENERATE MAC VERIFY	For cryptographic MAC operations. Although it is possible to compose using a series of encrypt calls, the security of the MAC relies on the operation being atomic and specific.
GENERATE CRYPTOGRAM VALIDATE CRYPTOGRAM	For composite cryptogram operations such as financial CVC or ARQC. To specify exactly which cryptogram the key is used for it is also necessary to specify a <i>role</i> for the key (see Section 3.6

Deleted: 2

Deleted: 14

	"Cryptographic Parameters" in [KMIP-Spec] .
TRANSLATE ENCRYPT TRANSLATE DECRYPT TRANSLATE WRAP TRANSLATE UNWRAP	<p>To accommodate secure routing of traffic and data. In many areas that rely on symmetric techniques (notably but not exclusively financial networks), information is sent from place to place encrypted using shared symmetric keys. When encryption keys are changed it is desirable for the change to be an atomic operation, otherwise distinct unwrap-wrap or decrypt-encrypt steps risk leaking the plaintext data in the middle.</p> <p><i>TRANSLATE ENCRYPT/DECRYPT</i> is used for data encipherment.</p> <p><i>TRANSLATE WRAP/UNWRAP</i> is used for key wrapping.</p>

832 **Table 2: Cryptographic Usage Masks Pairs**

833 In order to support asymmetric concepts using symmetric keys in a KMIP system the server
 834 implementation needs to be able to differentiate between clients for generate operations and clients for
 835 verify operations. As indicated by Section 3 ("Attributes") of **[KMIP-Spec]** there is a single key object in
 836 the system to which all relevant clients refer, but when a client requests that key the server is able to
 837 choose which attributes (permissions) to send with it based on the identity and configured access rights of
 838 that specific client. There is thus no need to maintain and synchronize distinct copies of the symmetric
 839 key: just a need to define access policy for each client or group of clients.

840 The internal implementation of this feature at the server end is a matter of choice for the vendor: storing
 841 multiple key blocks with all necessary combinations of attributes or generating key blocks dynamically are
 842 both acceptable approaches.

843 3.28 Application Specific Information

844 The Application Specific Information attribute is used to store data which is specific to the application(s)
 845 using the object. Some examples of Application Name Space and Application Data pairs are given below.

- 846 • SMIME, 'someuser@company.com'
- 847 • SSL, 'some.domain.name'
- 848 • Volume Identification, '123343434'
- 849 • File Name, 'secret.doc'
- 850 • Client Generated Key ID, '450994003'

851 The following Application Name Spaces are recommended:

- 852 • SMIME
- 853 • SSL
- 854 • IPSEC
- 855 • HTTPS
- 856 • PGP
- 857 • Volume Identification
- 858 • File Name
- 859 • LTO4
- 860 • LIBRARY-LTO4

861 KMIP provides optional support for server-generated Application Data. Clients may request the server to
 862 generate the Application Data for the client by omitting Application Data while setting or modifying the

Deleted: 2
 Deleted: 14

863 Application Specific Information attribute. A server only generates the Application Data if Application Data
864 is completely omitted from the request and the client-specified Application Name Space is recognized and
865 supported by the server. An example for requesting the server to generate the Application Data is shown
866 below:

```
867     AddAttribute(UID, AppSpecInfo{AppNameSpace='LIBRARY-LTO4'});
```

868 If the server does not recognize the name space, the “Application Name Space Not Supported” error is
869 returned to the client.

870 If Application Data is set to null, as shown in the example below, and the Application Name Space is
871 recognized by the server, the server does not generate the Application Data for the client. The server
872 stores the Application Specific Information attribute with the Application Data value set to null.

```
873     AddAttribute(UID, AppSpecInfo{AppNameSpace='LIBRARY-LTO4', AppData=null});
```

874 3.29 Mutating Attributes

875 KMIP does not support server mutation of client-supplied attributes. If a server does not accept an
876 attribute value that is being specified inside the request by the client, the server returns an error and
877 specifies “Invalid Field” as Result Reason.

878 Attributes that are not set by the client, but are implicitly set by the server as a result of the operation, may
879 optionally be returned by the server in the operation response inside the Template–Attribute.

880 If a client sets a time-related attribute to the current date and time, but as a result of a clock skew the
881 specified date of the attribute is received after the set time, it is up to the server policy to decide whether
882 to accept the backdated attribute. KMIP does not require the server to fail a request if a backdated
883 attribute is set by the client.

884 If a server does not support backdated attributes and cryptographic objects are expected to change state
885 at the specified current date and time, clients are recommended to issue the operation that would
886 implicitly set the date for the client. For example, instead of explicitly setting the Activation Date, clients
887 could issue the Activate operation. This would require the server to set the Activation Date to the server’s
888 current date and time.

889 If it is not possible to set a date attribute via an operation and the server does not support backdated
890 attributes, clients need to take into account that potential clock skew issues may cause the server to
891 return an error even if a date attribute is set to the client’s current date and time.

892 For additional information, refer to the sections describing the State attribute and the Time Stamp field in
893 **[KMIP-Spec]**.

894 3.30 Interoperable Key Naming for Tape

895 This section describes methods for creating and storing key identifiers that are interoperable across multi-
896 vendor KMIP clients.

897 3.30.1 Native Tape Encryption by a KMIP Client

898 This method is primarily intended to promote interoperable key naming between tape library products
899 which already support non-KMIP key managers, where KMIP support is being added.

900 When those existing library products are KMIP clients, a common method for naming and storing keys
901 may be used to support moving tape cartridges between the libraries, and successfully retrieving keys,
902 assuming the clients have appropriate access privileges. The library clients may be from multiple vendors,
903 and may be served from a KMIP key manager from a different vendor.

904 3.30.1.1 Method Overview

905 • The method uses the KMIP Application Specific Information (ASI) attribute’s Application Data field
906 to store the key name. The ASI Application Name Space is used to identify the namespace (such
907 as LIBRARY-LTO4).

Deleted: 2

Deleted: 14

- 908 • The method also uses the tape format's Key Associated Data (KAD) fields to store the key name.
909 Tape formats may provide both authenticated and unauthenticated storage for the KAD data. This
910 method ensures optimum utilization of the authenticated KAD data, when the tape format
911 supports authentication.
- 912 • The method supports both client-generated and server-generated key names.
- 913 • The method, in many cases, is backward-compatible if tapes are returned to a non-KMIP key
914 manager environment.
- 915 • Key names stored in the KMIP server's ASI attribute are always text format. Key names stored on
916 the KMIP client's KAD fields are always numeric format, due to space limitations of the tape
917 format. The method basically consists of implementing a specific algorithm for converting
918 between text and numeric formats.
- 919 • The algorithm used by this conversion is reversible.

920 3.30.1.2 Definitions

- 921 • Key Associated Data (KAD). Part of the tape format. May be segmented into authenticated and
922 unauthenticated fields. KAD usage is detailed in the SCSI SSC-3 standard, from the T10
923 organization.
- 924 • Application Specific Information (ASI). A KMIP attribute.
- 925 • Hexadecimal numeric characters. Case sensitive printable single byte ASCII characters
926 representing the numbers 0 through 9 and uppercase alpha A through F. (US-ASCII characters
927 30h-39h and 41h-46h)
928 Hexadecimal numeric characters are always paired, each pair representing a single 8-bit numeric
929 value. A leading zero character is provided, if necessary, so that every byte in the tape's KAD is
930 represented by exactly 2 hexadecimal numeric characters.
- 931 • N(k). The number of bytes in the tape format's combined KAD fields (both authenticated and
932 unauthenticated).
- 933 • N(a), N(u). The number of bytes in the tape format's authenticated, and unauthenticated KAD
934 fields, respectively.
935

936 3.30.1.3 Algorithm 1. Numeric to text direction (tape format's KAD to KMIP ASI)

937 Description: All information contained in the tape format's KAD fields is converted to a null-terminated
938 ASCII string consisting of hexadecimal numeric character pairs. First the unauthenticated KAD data is
939 converted to text. Then the authenticated KAD data is converted and appended to the end of the string.
940 The string is then null-terminated.

941

942 Implementation Example:

- 943 1. Define an input buffer sized for N(k). For LTO4, N(k) is 44 bytes (12 bytes authenticated, 32
944 unauthenticated).
- 945 2. Define an output buffer sufficient to contain a null-terminated string with a maximum length of
946 $2*N(k)+1$.
- 947 3. Define the standard POSIX (also known as C) locale. Each character in the string is a single-byte US-
948 ASCII character.
- 949 4. Copy the tape format's KAD data, from the unauthenticated KAD field first, to the input buffer.
950 Effectively, the first byte (byte 0) of the input buffer is the first byte of unauthenticated KAD. Bytes
951 from the authenticated KAD are concatenated, after the unauthenticated bytes.
- 952 5. For each byte in the input buffer, convert to US-ASCII as follows:

Deleted: 2

Deleted: 14

- 953 a. Convert the byte's value to exactly 2 hexadecimal numeric characters, including a leading 0
 954 where necessary. Append these 2 numeric characters to the output buffer, with the high-nibble
 955 represented by the left-most hexadecimal numeric character.
- 956 b. After all byte values have been converted, null terminate the output buffer.
- 957 6. When storing the string to the KMIP server, use the object's ASI attribute's Application Data field.
 958 Store the namespace (such as LIBRARY-LTO4) in the ASI attribute's Application Name Space field.

959 **3.30.1.4 Algorithm 2. Text to numeric direction (KMIP ASI to tape format's KAD)**

960 Description: Hexadecimal numeric character pairs in the null-terminated ASCII string are converted to
 961 single byte numeric values, and stored in the tape format's KAD fields. The authenticated KAD field is
 962 populated first, from a sub-string consisting of the last 2*N(a) characters in the full string. Any remaining
 963 characters in the string are converted and stored to the unauthenticated KAD field. The null termination
 964 byte is not converted.
 965

966 Implementation Example:

- 967 1. Obtain the key's name from the KMIP server's ASI attribute for that object. Copy the null terminated
 968 string to an input buffer of size 2*N(k) + 1 bytes. For LTO4, an 89 character string, including null
 969 termination, is sufficient for all possible key descriptors when names are directly referenced.
- 970 2. Define output buffers for unauthenticated KAD, and authenticated KAD, of size N(u) and N(a)
 971 respectively. For LTO4, this would be 32 bytes of unauthenticated data, and 12 bytes of authenticated
 972 data.
- 973 3. Define the standard POSIX (also known as C) locale. Each character in the string is a single-byte US-
 974 ASCII character.
- 975 4. First populate the authenticated KAD buffer, converting a sub-string consisting of the last 2*N(a)
 976 characters of the full string, not including the null termination byte.
- 977 5. When the authenticated KAD is filled, next populate the unauthenticated KAD buffer, by converting
 978 the remaining hexadecimal character pairs in the string.

979 **3.30.1.5 Example Output**

980 Following are examples illustrating some results of this method. In the following examples, the sizes of
 981 the KAD for LTO4 are used. Different tape formats may utilize different KAD sizes.

982
 983 Example 1. Full combined KAD

984
 985 This LTO4 tape's combined KAD contains the following data (represented in hexadecimal). For LTO4, the
 986 unauthenticated KAD contains 32 bytes, and the authenticated KAD contains 12 bytes.
 987

988 Example 1a. Hexadecimal numeric data from a tape's KAD.

989 Shaded data is authenticated by the tape drive.

990
 991 02 04 17 11 39 43 42 36 30 41 33 34 39 31 44 33
 992 41 41 43 36 32 42 07 F6 54 54 32 36 30 38 4C 34
 993 **30 30 30 39 30 35 32 38 30 34 31 32**
 994

995 The algorithm converts the numeric KAD data to the following 89 character null-terminated string, for
 996 storage in the Application Data field of a KMIP object's Application Specific Information attribute. The ASI
 997 Application Name Space contains "LIBRARY-LTO4".

Deleted: 2
 Deleted: 14

998
999 Example 1b. Text string from KMIP ASI Application Data.
1000 Shaded characters are derived from authenticated data. The null character is represented as
1001 <null>

1002
1003 0204171139434236304133343931443341414336324207F65454323630384C343030303930353
1004 23830343132<null>

1005
1006 Example 1c. The hexadecimal values of the 89 US-ASCII characters in string 1b, from the KMIP
1007 ASI Application Data. Note: these values are always in the range 30h-39h, or in the range 41h-
1008 46h, or the 0h null.

1009 30 32 30 34 31 37 31 31 33 39 34 33 34 32 33 36 33 30 34 31 33 33 33 34 33 39 33 31 34 34 33
1010 33 34 31 34 31 34 33 33 36 33 32 34 32 30 37 46 36 35 34 35 34 33 32 33 36 33 30 33 38 34 43
1011 33 34 33 30 33 30 33 30 33 39 33 30 33 35 33 32 33 38 33 30 33 34 33 31 33 32 00

1012
1013 For the reverse transformation, a client would retrieve the string in 1b from the server, derive the numeric
1014 values shown in 1a, and store them to the tape format's KAD data. First, the sub-string containing the
1015 right-most 24 characters of the full string 1b are used to derive the 12-byte authenticated KAD. The
1016 remaining characters are used to derive the 32-byte unauthenticated KAD.

1017
1018 Example 2. Authenticated KAD only

1019 This LTO4 tape's KAD contains the following data (represented in hexadecimal), all 12 bytes obtained
1020 from the authenticated KAD field. There is no unauthenticated KAD data.

1021
1022 Example 2a. Hexadecimal numeric data from a tape's KAD.
1023 Shaded data is authenticated.

1024
1025 17 48 33 C6 20 42 10 A7 E8 05 F8 C7

1026
1027 The algorithm converts the numeric KAD data to the following 24 character null-terminated string, for
1028 storage in the Application Data field of a KMIP object's Application Specific Information attribute.

1029
1030 Example 2b. Text string from KMIP ASI Application Data.
1031 Shaded characters are derived from authenticated data. The null character is represented as
1032 <null>

1033
1034 174833C6204210A7E805F8C7<null>

1035
1036 For the reverse transformation, a client would derive the numeric values in 2a, and store them to the tape
1037 format's KAD data. The right-most 24 characters of the string in 2b are used to derive the 12 byte
1038 authenticated KAD. In this example, there is no unauthenticated KAD data.

1039
1040 Example 3. Partially filled authenticated KAD originating from a non-KMIP method

1041 This LTO4 tape's KAD contains the following data (represented in hexadecimal). The unauthenticated
1042 KAD contains 10 bytes, and the authenticated KAD contains 8 bytes.

1043

Deleted: 2

Deleted: 14

1044 Since the authenticated KAD was not filled, but the unauthenticated data was populated, the method
1045 creating this key name is potentially not backward-compatible with the KMIP key naming method. See
1046 backward-compatibility assessment, below.

1047
1048 Example 3a. Hexadecimal numeric data from a non-KMIP tape's KAD.
1049 Shaded data is authenticated.

1050
1051 02 04 17 11 39 43 42 36 30 41 30 30 30 39 30 35
1052 32 38

1053
1054 The algorithm converts the numeric KAD data to the following 36 character null-terminated string, for
1055 storage in the Application Data field of a KMIP object's Application Specific Information attribute.

1056
1057 Example 3b. Text string from KMIP ASI Application Data.
1058 Shaded characters are derived from authenticated data. The null character is represented as
1059 <null>

1060
1061 02041711394342363041303030393035238<null>

1062
1063 For the reverse transformation, a client would derive the same numeric values shown in 3a, and store
1064 them to the tape's KAD. But their storage locations within the KAD now differs (see 3c). The right-most 24
1065 characters from the text string in 3b are used to derive the 12-byte authenticated KAD. The remaining
1066 characters are used to fill the 32-byte unauthenticated KAD.

1067
1068 Example 3c. Hexadecimal numeric data from a tape's KAD.
1069 Shaded data is authenticated.

1070
1071 02 04 17 11 39 43 42 36 30 41 30 30 30 39 30 35
1072 32 38

1073 3.30.1.6 Backward-compatibility assessment

1074 Where all the following conditions exist, a non-KMIP solution may encounter compatibility issues during
1075 the Read and Appended Write use cases.

Deleted: 1

- 1076 1. The tape format supports authenticated KAD, but the non-KMIP solution does not use, or only
1077 partially uses, the authenticated KAD field.
- 1078 2. The non-KMIP solution is sensitive to data position within the combined KAD.
- 1079 3. The media was written in a KMIP environment, using this method, then moved to the non-KMIP
1080 environment.

1081 3.31 Revocation Reason Codes

1082 The enumerations for the Revocation Reason attribute specified in KMIP (see table 9.1.3.2.17 in [KMIP-
1083 Spec]) are aligned with the Reason Code specified in X.509 and referenced in RFC 5280 with the
1084 following exceptions. The *certificateHold* and *removeFromCRL* reason codes have been excluded from
1085 [KMIP-Spec], since this version of KMIP does not support certificate suspension (putting a certificate
1086 hold) or unsuspension (removing a certificate from hold). The *aaCompromise* reason code has been
1087 excluded from [KMIP-Spec] since it only applies to attribute certificates and at this point of time attribute
1088 certificates are considered out-of-scope for [KMIP-Spec]. The *privilegeWithdrawn* reason code is

Deleted: 2

Deleted: 14

1089 included in **[KMIP-Spec]** since it may be used for either attribute or public key certificates. In the context
1090 of its use within KMIP it is assumed to only apply to public key certificates.

1091 3.32 Certificate Renewal, Update, and Re-key

1092 The process of generating a new certificate to replace an existing certificate may be referred to by
1093 multiple terms based upon what data within the certificate is changed when the new certificate is created.
1094 In all situations, the new certificate includes a new serial number and new validity dates. **[KMIP-Spec]**
1095 uses the following terminology which is aligned with the definitions found in IETF RFCs 3647 and 4949:

- 1096 • *Certificate Renewal*: The issuance of a new certificate to the subject without changing the subject
1097 public key or other information (except the serial number and certificate validity dates) in the
1098 certificate.
- 1099 • *Certificate Update*: The issuance of a new certificate due to changes in the information in the
1100 certificate other than the subject public key.
- 1101 • *Certificate Rekey*: The generation of a new key pair for the subject and the issuance of a new
1102 certificate that certifies the new public key.

1103 The current KMIP Specification supports certificate renewals using the Re-Certify operation and certificate
1104 updates using the Certify operation. Support for certificate rekey is not currently supported by KMIP, since
1105 certificate rekey requires the ability to rekey an asymmetric key pair a capability not currently supported
1106 by KMIP. Support for rekey of asymmetric key pairs along with certificate rekey may be considered for a
1107 future KMIP release.

1108 3.33 Key Encoding

1109 Two parties receiving the same key as a Key OCTET STRING make use of the key in exactly the same
1110 way in order to interoperate. To ensure that, it is necessary to define a correspondence between the
1111 abstract syntax of Key and the notation in the standard algorithm description that defines how the key is
1112 used. The next sections establish that correspondence for the algorithms AES **[FIPS197]** and 3DES
1113 **[SP800-67]**.

Formatted: Ref term, Font: Not Bold, English (U.S.)

Deleted: [SP800-67]

1114 3.33.1 AES Key Encoding

1115 **[FIPS197]** section 5.2, titled Key Expansion, uses the input key as an array of bytes indexed starting at 0.
1116 The first octet of Key becomes the key byte in AES labeled index 0 in **[FIPS197]** is the first octet of Key,
1117 and the other key bytes follow in index order.

1118 Proper parsing and key load of the contents of Key for AES is determined by using the following Key octet
1119 string to generate and match the key expansion test vectors in **[FIPS197]** Appendix A for AES Cipher
1120 Key: 2B 7E 15 16 28 AE D2 A6 AB F7 15 88 09 CF 4F 3C.

1121 3.33.2 Triple-DES Key Encoding

1122 A Triple-DES key consists of three keys for the cryptographic engine (Key1, Key2, and Key3) that are
1123 each 64 bits (even though only 56 are used); the three keys are also referred to as a key bundle (KEY)
1124 **[SP800-67]**. A key bundle may employ either two or three mutually independent keys. When only two are
1125 employed (called two-key Triple-DES), then Key1 = Key3.

Formatted: Ref term, Font: Not Bold, English (U.S.)

Deleted: [SP800-67]

Deleted: [SP800-67]

1126 Each key in a Triple-DES key bundle is expanded into a key schedule according to a procedure defined in
1127 **[SP800-67]** appendix A. That procedure numbers the bits in the key from 1 to 64, with number 1 being
1128 the left-most, or most significant bit. The first octet of Key is bits 1 through 8 of Key1 with bit 1 being the
1129 most significant bit. The second octet of Key is bits 9 through 16 of Key1, and so forth, so that the trailing
1130 octet of KEY is bits 57 through 64 of Key3 (or Key2 for two-key Triple-DES).

Formatted: Ref term, Font: Not Bold, English (U.S.)

Formatted: Ref term, Font: Not Bold, English (U.S.)

1131 Proper parsing and key load of the contents of Key for Triple-DES is determined by using the following
1132 Key octet string to generate and match the key expansion test vectors in **[SP800-67]** appendix B for the
1133 key bundle:

Deleted: [SP800-67]

Deleted: 2

1134 Key1 = 0123456789ABCDEF

Deleted: 14

1135 Key2 = 23456789ABCDEF01
1136 Key3 = 456789ABCDEF0123
1137

1138 4 Deferred KMIP Functionality

1139 The KMIP Specification is currently missing items that have been judged candidates for future inclusion in
1140 the specification. These items currently include:

- 1141 • Registration of Clients. This would allow in-band registration and management of clients, which
1142 currently may only be registered and/or managed using off-line mechanisms.
- 1143 • Client-requested specification of additional clients allowed to use a key. This requires coordinated
1144 identities between the client and server, and as such, is deferred until registration of clients is
1145 addressed.
- 1146 • Registration of Notifications. This would allow clients to specify, using an in-band mechanism,
1147 information and events that they wish to be notified of, and what mechanisms should be used for
1148 such notifications, possibly including the configuration of pushed cryptographic material. This
1149 functionality would assume Registration of Clients as a prerequisite.
- 1150 • Key Migration. This would standardize migration of keys from one HSM to another, using
1151 mechanisms already in the protocol or ones added for this purpose.
- 1152 • Server to Server key management. This would extend the protocol to support communication
1153 between key management servers in different key management domains, for purposes of
1154 exporting and importing of cryptographic material and potentially policy information.
- 1155 • Multiple derived keys. This would allow creation of multiple derived keys from one or more input
1156 keys. Note, however, that the current version of KMIP provides the capability to derive multiple
1157 keys and initialization vectors by creating a Secret Data object and specifying a cryptographic
1158 length equal to the total length of the derived objects.
- 1159 • XML encoding. Expression of KMIP in XML rather than in tag/type/length/value may be
1160 considered for the future.
- 1161 • Specification of Mask Generation Function. KMIP does not currently allow clients to specify the
1162 Mask Generation Function and assumes that encryption or signature schemes, such as OAEP or
1163 PSS, use MGF1 with the hash function as specified in the Cryptographic Parameters attribute.
1164 Client specification of MGFs may be considered for the future.
- 1165 • Certificate creation without client-provided Certificate Request. This would allow clients to request
1166 the server to perform the Certify or Re-certify operation from the specified key pair IDs without
1167 providing a Certificate Request.
- 1168 • Server monitoring of client status. This would enable the transfer of information about the client
1169 and its cryptographic module to the server. This information would enable the server to generate
1170 alarms and/or disallow requests from a client running component versions with known
1171 vulnerabilities.
- 1172 • Symmetric key pairs. Only a subset of the cryptographic usage bits of the Cryptographic Usage
1173 Mask attribute may be permitted for keys distributed to a particular client. KMIP does not currently
1174 address how to securely assign and determine the applicable cryptographic usage for a client.
- 1175 • Hardware-protected attribute. This attribute would allow clients and servers to determine if a key
1176 may only be processed inside a secure cryptographic device such as an HSM. If this attribute is
1177 set, the key may only exist in cleartext from inside a secure hardware device, and all security-
1178 relevant attributes are bound to it in such a way that they may not be modified outside of such a
1179 secure device.
- 1180 • Alternative profiles for key establishment. Less capable end-clients may not be able to support
1181 TLS and should use a proxy to communicate with the key management system. The KMIP
1182 protocol does not currently support alternative profiles nor does it allow end-clients relying on the
1183 proxy model to securely establish a key with the server.

Deleted: 2

Deleted: 14

- 1184 • Attribute mutation. The possibility for the server to use attribute values different than requested by
1185 the client if these values are not suitable for the server, and return these values in the response,
1186 instead of failing the request.
- 1187 • Cryptographic Domain Parameters. KMIP allows a limited number of parameters to be specified
1188 during a Create Key Pair operation. Additional parameters may be considered for the future.
- 1189 • Re-key support for other cryptographic objects. The Re-key operation is currently restricted to
1190 symmetric keys. Applying Re-key to other cryptographic objects, such as asymmetric keys and
1191 certificates, may be considered for the future.
- 1192 • Certificate Suspension/Unsuspend. KMIP does not currently support certificate suspension
1193 (putting a certificate on hold) or unsuspension (removing a certificate from hold). Adding support
1194 for certificate suspension/unsuspension into KMIP may be considered for the future.
- 1195 • Namespace registration. Establishing a registry for namespaces may be considered for the
1196 future.
- 1197 • Registering extensions to KMIP enumerations. Establishing a registry for extensions to defined
1198 KMIP enumerations, such as in support of profiles specific to IEEE P1619.3 or other
1199 organizations, may be considered for the future.

1200 In addition to the functionality listed above, the KMIP TC is interested in establishing a C&A (certification
1201 and accreditation) process for independent validation of claims of KMIP conformance. Defining and
1202 establishing this process is a candidate for work by the KMIP TC after V1.0.

1203 A. Acronyms

1204 The following abbreviations and acronyms are used in this document:

- 1205 3DES - Triple Data Encryption Standard specified in ANSI X9.52
- 1206 AES - Advanced Encryption Standard specified in FIPS 197
- 1207 ANSI - American National Standards Institute
- 1208 ARQC - Authorization Request Cryptogram
- 1209 ASCII - American Standard Code for Information Interchange
- 1210 CA - Certification Authority
- 1211 CBC - Cipher Block Chaining specified in NIST SP 800-38A
- 1212 CMC - Certificate Management Messages over CMS specified in RFC 5275
- 1213 CMP - Certificate Management Protocol specified in RFC 4210
- 1214 CRL - Certificate Revocation List specified in RFC 5280
- 1215 CRMF - Certificate Request Message Format specified in RFC 4211
- 1216 CVC - Card Verification Code
- 1217 DES - Data Encryption Standard specified in FIPS 46-3
- 1218 DEK - Data Encryption Key
- 1219 DH - Diffie-Hellman specified in ANSI X9.42
- 1220 FIPS - Federal Information Processing Standard
- 1221 GCM - Galois/Counter Mode specified in NIST SP 800-38D
- 1222 HMAC - Keyed-Hash Message Authentication Code specified in FIPS 198-1
- 1223 HSM - Hardware Security Module
- 1224 HTTP - Hyper Text Transfer Protocol
- 1225 HTTP(S) - Hyper Text Transfer Protocol (Secure socket)
- 1226 ID - Identification
- 1227 IP - Internet Protocol
- 1228 IPsec - Internet Protocol Security
- 1229 JKS - Java Key Store
- 1230 KEK - Key Encryption Key
- 1231 KMIP - Key Management Interoperability Protocol
- 1232 LTO4 - Linear Tape-Open 4
- 1233 MAC - Message Authentication Code
- 1234 MD5 - Message Digest 5 Algorithm specified in RFC 1321
- 1235 MGF - Mask Generation Function
- 1236 NIST - National Institute of Standards and Technology
- 1237 OAEP - Optimal Asymmetric Encryption Padding specified in PKCS#1
- 1238 PEM - Privacy Enhanced Mail specified in RFC 1421
- 1239 PGP - Pretty Good Privacy specified in RFC 1991

Deleted: 2

Deleted: 14

- 1240 PKCS - Public-Key Cryptography Standards
- 1241 POP - Proof of Possession
- 1242 POSIX - Portable Operating System Interface
- 1243 PSS - Probabilistic Signature Scheme specified in PKCS#1
- 1244 RACF - Remote Access Control Facility
- 1245 RSA - Rivest, Shamir, Adelman (an algorithm)
- 1246 SHA - Secure Hash Algorithm specified in FIPS 180-2
- 1247 SP - Special Publication
- 1248 SSL - Secure Sockets Layer
- 1249 S/MIME - Secure/Multipurpose Internet Mail Extensions
- 1250 TCP - Transport Control Protocol
- 1251 TLS - Transport Layer Security
- 1252 TTLV - Tag, Type, Length, Value
- 1253 URI - ~~Uniform~~ Resource Identifier
- 1254 X.509 - Public Key Certificate specified in RFC 5280
- 1255 XML - Extensible Markup Language

Deleted: que

Deleted: 2

Deleted: 14

1256

B. Acknowledgements

1257 The following individuals have participated in the creation of this specification and are gratefully
1258 acknowledged:

1259 Original Authors of the initial contribution:

- 1260 David Babcock, HP
- 1261 Steven Bade, IBM
- 1262 Paolo Bezoari, NetApp
- 1263 Mathias Björkqvist, IBM
- 1264 Bruce Brinson, EMC
- 1265 Christian Cachin, IBM
- 1266 Tony Crossman, Thales/nCipher
- 1267 Stan Feather, HP
- 1268 Indra Fitzgerald, HP
- 1269 Judy Furlong, EMC
- 1270 Jon Geater, Thales/nCipher
- 1271 Bob Griffin, EMC
- 1272 Robert Haas, IBM
- 1273 Timothy Hahn, IBM
- 1274 Jack Harwood, EMC
- 1275 Walt Hubis, LSI
- 1276 Glen Jaquette, IBM
- 1277 Jeff Kravitz, IBM
- 1278 Michael McIntosh, IBM
- 1279 Brian Metzger, HP
- 1280 Anthony Nadalin, IBM
- 1281 Elaine Palmer, IBM
- 1282 Joe Pato, HP
- 1283 René Pawlitzek, IBM
- 1284 Subhash Sankuratripati, NetApp
- 1285 Mark Schiller, HP
- 1286 Martin Skagen, Brocade
- 1287 Marcus Streets, Thales/nCipher
- 1288 John Tattan, EMC
- 1289 Karla Thomas, Brocade
- 1290 Marko Vukolić, IBM
- 1291 Steve Wierenga, HP

1292 Participants:

- 1293 [Gordon Arnold, IBM](#)
- 1294 [Todd Arnold, IBM](#)
- 1295 [Matthew Ball, Sun Microsystems](#)
- 1296 [Elaine Barker, NIST](#)
- 1297 [Peter Bartok, Venafi, Inc.](#)
- 1298 [Mathias Bjorkqvist, IBM](#)
- 1299 [Kevin Bocek, Thales e-Security](#)
- 1300 [Kelley Burgin, National Security Agency](#)
- 1301 [Jon Callas, PGP Corporation](#)
- 1302 [Tom Clifford, Symantec Corp.](#)
- 1303 [Graydon Dodson, Lexmark International Inc.](#)

← Formatted: Indent: First line: 0.5"

Deleted: 2

Deleted: 14

- 1304 [Chris Dunn, SafeNet, Inc.](#)
- 1305 [Paul Earsy, SafeNet, Inc.](#)
- 1306 [Stan Feather, HP](#)
- 1307 [Indra Fitzgerald, HP](#)
- 1308 [Alan Frindell, SafeNet, Inc.](#)
- 1309 [Judith Furlong, EMC Corporation](#)
- 1310 [Jonathan Geater, Thales e-Security](#)
- 1311 [Robert Griffin, EMC Corporation](#)
- 1312 [Robert Haas, IBM](#)
- 1313 [Thomas Hardjono, M.I.T.](#)
- 1314 [Marc Hocking, BeCrypt Ltd.](#)
- 1315 [Larry Hofer, Emulex Corporation](#)
- 1316 [Brandon Hoff, Emulex Corporation](#)
- 1317 [Walt Hubis, LSI Corporation](#)
- 1318 [Wyllys Ingersoll, Sun Microsystems](#)
- 1319 [Jay Jacobs, Target Corporation](#)
- 1320 [Glen Jaquette, IBM](#)
- 1321 [Scott Kipp, Brocade Communications Systems, Inc.](#)
- 1322 [David Lawson, Emulex Corporation](#)
- 1323 [Robert Lockhart, Thales e-Security](#)
- 1324 [Shyam Mankala, EMC Corporation](#)
- 1325 [Marc Massar, Individual](#)
- 1326 [Don McAlister, Cipheroptics](#)
- 1327 [Hyrum Mills, Mitre Corporation](#)
- 1328 [Landon Noll, Cisco Systems, Inc.](#)
- 1329 [Rene Pawlitzek, IBM](#)
- 1330 [Rob Philpott, EMC Corporation](#)
- 1331 [Bruce Rich, IBM](#)
- 1332 [Scott Rotondo, Sun Microsystems](#)
- 1333 [Anil Saldhana, Red Hat](#)
- 1334 [Subhash Sankuratripati, NetApp](#)
- 1335 [Mark Schiller, HP](#)
- 1336 [Jitendra Singh, Brocade Communications Systems, Inc.](#)
- 1337 [Servesh Singh, EMC Corporation](#)
- 1338 [Sandy Stewart, Sun Microsystems](#)
- 1339 [Marcus Streets, Thales e-Security](#)
- 1340 [Brett Thompson, SafeNet, Inc.](#)
- 1341 [Benjamin Tomhave, Individual](#)

Deleted: TBD

Deleted: 2

Deleted: 14

1342 | [Sean Turner, IECA, Inc.](#)
1343 | [Paul Turner, Venafi, Inc.](#)
1344 | [Marko Vukolic, IBM](#)
1345 | [Rod Wideman, Quantum Corporation](#)
1346 | [Steven Wierenga, HP](#)
1347 | [Peter Yee, EMC Corporation](#)
1348 | [Krishna Yellepeddy, IBM](#)
1349 | [Peter Zelechowski, Election Systems & Software](#)
1350 |
1351 |

Deleted: 2

Deleted: 14

C. Revision History

Revision	Date	Editor	Changes Made
ed-0.98	2009-04-29	Indra Fitzgerald	Initial conversion of input document to OASIS format.
ed-0.98	2009-07-28	Indra Fitzgerald	Added clarifications, examples, and deferred items.
ed-0.98	2009-09-08	Indra Fitzgerald	Added approved proposals and incorporated Elaine Barker's comments.
ed-0.98	2009-09-23	Indra Fitzgerald	Removed KMIP Profiles section and incorporated the Interoperable Key Naming for Tape proposal.
ed-0.98	2009-09-24	Indra Fitzgerald	Removed the Conformance section; added additional Certificate Request and POP text to Certify and Re-certify; added the Revocation Reason Codes section.
draft-01	2009-10-07	Indra Fitzgerald	Incorporated the Certificate Renewal, Update, Re-key proposal, the Key Encoding proposal; removed normative words "must", "shall", "required", "will", and "can"; added Create Key Pair example; updated the references and acronyms list; incorporated comments from RobertH and SubhashS; updated the Authentication section; added minor edits and clarifications.
draft-02	2009-10-09	Indra Fitzgerald	Incorporated Rod Wideman's comments on the language. Changed the heading indentation, paragraph style, and list styles according to the OASIS template guidelines. Added additional references. Replaced the TBDs. Added a use-case for registering a wrapped key as an opaque cryptographic object.
draft-03	2009-10-21	Indra Fitzgerald	Added the list of participants to Appendix B. Clarified the Authentication section (section 3.1) and added examples. Modified the title page. Performed minor editorial changes.

Deleted: 2

Deleted: 14