

## iSCSI Security Requirements and SRP-based ESP Keys

### Status of this Memo

This document is an Internet-Draft and is subject to all provisions of [Section 10 of RFC2026](#). Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at  
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at  
<http://www.ietf.org/shadow.html>

Discussion and suggestions for improvement are requested. This draft will expire before February 2002. Distribution of this draft is unlimited.

### Abstract

This draft describes security requirements, status of security specification, operating environment characteristics, and related considerations for iSCSI. It also outlines an SRP-based [[RFC-2945](#)] mechanism to derive ESP [[RFC-2406](#)] keying material and associated rekeying mechanisms. These topics are combined in this draft for convenience; the security requirements, status and operating environment are generic to iSCSI, whereas the use of SRP to obtain ESP keys is an independent proposal and is not the only way to meet iSCSI's security requirements. The keying description focuses on the overall approach and structure, further details will be added if this proposal is pursued.

This draft is unlikely to become an RFC in its current form; its primary purpose is to capture a set of ideas as a basis for further discussions. Portions of this draft may be incorporated into other drafts that are intended to become RFCs. Caution should be exercised in drawing inferences from the fact that the author of

this draft is one of the chairs of the IP Storage WG. This draft is an individual submission that the IP Storage WG is free to adopt, modify, reject, fold, spindle, and/or mutilate as it sees fit.

---

## iSCSI Security Requirements and SRP-based ESP Keys Aug 2001

### Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC-2119](#)].

### Table of Contents

<a href="#">1. Context and Purposes.....</a>	<a href="#">2</a>
<a href="#">2. iSCSI Overview.....</a>	<a href="#">3</a>
<a href="#">3. Approach to Security Specification.....</a>	<a href="#">3</a>
<a href="#">4. Security Functionality Requirements.....</a>	<a href="#">4</a>
<a href="#">4.1 Negotiation of Security Functionality.....</a>	<a href="#">4</a>
<a href="#">4.2 Authentication.....</a>	<a href="#">4</a>
<a href="#">4.2.1 Multiple Authentication Considerations.....</a>	<a href="#">5</a>
<a href="#">4.3 Cryptographic Integrity and Data Origin Authentication.....</a>	<a href="#">5</a>
<a href="#">4.4 Confidentiality.....</a>	<a href="#">6</a>
<a href="#">4.5 Rekeying.....</a>	<a href="#">6</a>
<a href="#">5. iSCSI Implementation Characterization.....</a>	<a href="#">6</a>
<a href="#">5.1 Implementation Classes and Resource Availability.....</a>	<a href="#">6</a>
<a href="#">5.2 Implementation Scaling.....</a>	<a href="#">7</a>
<a href="#">5.3 Other Implementation Environment Concerns.....</a>	<a href="#">8</a>
<a href="#">6. SRP Keying of ESP.....</a>	<a href="#">9</a>
<a href="#">6.1 Overall structure.....</a>	<a href="#">9</a>
<a href="#">6.2 Negotiation.....</a>	<a href="#">10</a>
<a href="#">6.3 Key Derivation.....</a>	<a href="#">11</a>
<a href="#">6.4 ESP Startup.....</a>	<a href="#">11</a>
<a href="#">6.5 Rekeying.....</a>	<a href="#">13</a>
<a href="#">6.5.1 Rekeying Mechanism 1: No New Key Exchange.....</a>	<a href="#">14</a>
<a href="#">6.5.2 Rekeying Mechanism 2: New Key Exchange.....</a>	<a href="#">14</a>
<a href="#">7. Security Considerations.....</a>	<a href="#">14</a>
<a href="#">8. Changes from 00 Version.....</a>	<a href="#">14</a>
<a href="#">9. References.....</a>	<a href="#">15</a>
<a href="#">10. Acknowledgments.....</a>	<a href="#">16</a>
<a href="#">11. Author's Address.....</a>	<a href="#">16</a>

### [1. Context and Purposes](#)

As specification of iSCSI security mechanisms progresses, it seems useful to gather security requirements and reasonable security-

relevant assumptions about iSCSI operating environments in an accessible form; that is the purpose of Sections [2](#) through [5](#) of this draft. The IP Storage WG has recently completed work on iSCSI requirements [[ISCSI-REQTS](#)], which should be consulted for further information on iSCSI requirements in other areas. [Section 6](#) of this draft describes an SRP-based mechanism for keying ESP based on concepts discussed at the IP Storage WG's Nashua interim meeting (April 30 and May 1, 2001). [Section 6](#) is included in this draft primarily for convenience and to ensure that these ideas are not lost. While this mechanism appears to satisfy iSCSI's security requirements, it is not the only mechanism that satisfies them; IKE is another possibility, see [[ISCSI-IPSEC](#)].

Black

[Page 2]

---

## iSCSI Security Requirements and SRP-based ESP Keys Aug 2001

As is the case for all Internet-Drafts, the contents of this draft are subject to discussion and revision. All material in this draft concerning status or content of specification efforts in the IP Storage WG is as of mid-August 2001 as understood by the author. It is hoped that this closely approximates reality, but no absolute assurances can be offered, and all WG decisions are potentially subject to change as work on the iSCSI specification(s) moves forward.

### [2.](#) iSCSI Overview

iSCSI [[ISCSI](#)] is a TCP/IP encapsulation of the SCSI protocols used to access disk, tape and other devices. iSCSI is a client-server protocol in which clients (Initiators) open connections to servers (Targets). This draft uses the SCSI terms Initiator and Target for clarity and to avoid the common assumption that clients have considerably less computational and memory resources than servers; the reverse is often the case for SCSI, as Targets are commonly dedicated devices of some form.

iSCSI Initiators and Targets are layer 5 entities that are independent of TCP ports and IP addresses. An Initiator or Target may use multiple communication endpoints (<TCP port, IP address> pair), and such endpoints may be shared by multiple Initiators or Targets. The common case for sharing will be that the sharing entities are all of the same type (i.e., all Initiators or all Targets). iSCSI entities have names that are independent of communication endpoints, and iSCSI defines its own naming syntax for such entities (i.e., Initiators and Targets), see [[ISCSI-NAME-DISC](#)].

The iSCSI protocol has a text based negotiation mechanism as part of

its initial (Login) procedure. The mechanism is extensible in what can be negotiated (new text keys and values can be defined) and also in the number of negotiation rounds (e.g., to accommodate functionality such as challenge-response authentication).

### 3. Approach to Security Specification

The IP Storage WG is currently pursuing an approach of selecting subsets of IETF security specifications for iSCSI. This requires that good engineering judgement be used and that the result be sound from a security standpoint. Motivations for this subset selection approach include:

- Avoiding specification of requirements that have little current justification, (e.g., AH).
- Matching the chosen subset more closely to iSCSI's security requirements (e.g., iSCSI does not require both transport and tunnel mode IPsec encapsulation).
- Matching implementation requirements more closely to iSCSI's expected operating environments, particularly resource-limited embedded systems.

Black

[Page 3]

---

iSCSI Security Requirements and SRP-based ESP Keys Aug 2001

In addition to the above examples, this approach may be applicable to areas such as IKE/ISAKMP [RFC-2407, [RFC-2408](#), [RFC-2409](#)] (e.g., subset of IKE modes and ISAKMP-negotiable items/values) and ESP algorithms (e.g., believe it or not, DES is still REQUIRED by ESP in [Section 5 of \[RFC-2406\]](#)).

### 4. Security Functionality Requirements

This section summarizes iSCSI security requirements and status.

#### 4.1 Negotiation of Security Functionality

The current status is that iSCSI's negotiation mechanism is capable of negotiating away any and all security functionality in environments where it is not needed (as determined by local policy). For example, a security administrator could determine that a physically isolated and physically secure network used only for iSCSI requires no additional iSCSI security beyond authentication. Another example is that a security administrator could determine that operation of iSCSI over a secure VPN requires no security beyond authentication because the VPN provides adequate defense against the relevant classes of threats.

Two caveats apply to the previous paragraph:

- Both of the examples above use authentication. Use of authentication should be RECOMMENDED by the iSCSI specification.
- The acronym SAN has sometimes been used to refer to environments in which iSCSI requires less security than is needed on a public network. Such usage is misleading, as it incorrectly implies that SANs have little to no need for communication protocol security.

iSCSI's inband negotiation mechanism has no intrinsic security because it is performed in the clear as far as iSCSI is concerned. If such a negotiation is conducted over an otherwise unsecured connection, man-in-the-middle attacks on security negotiation have to be considered. The usual admonition to not offer security mechanisms that are weaker than acceptable applies.

## [4.2](#) Authentication

Bi-directional authentication (Initiator to Target) and vice-versa is REQUIRED. This authentication is logically between the iSCSI Initiator and the iSCSI Target (as opposed to between the TCP/IP communication endpoints). There is no requirement that the identities used in authentication be kept confidential (e.g., from a passive eavesdropper). The intent of the iSCSI design is that the Initiator and Target represent the systems (e.g., host and disk array or tape system) participating in the communication, as opposed to network communication interfaces or endpoints.

The current status is that the Secure Remote Password protocol

Black

[Page 4]

---

iSCSI Security Requirements and SRP-based ESP Keys Aug 2001

(SRP) [[RFC-2945](#)] will be REQUIRED of all implementations based on iSCSI's text-based multi-round negotiation mechanism noted above. A number of additional authentication protocols have also been specified in the current iSCSI draft. Negotiation between Initiator and Target is used to determine which authentication algorithm to use (or whether to use one at all); the connection closes if either side requires authentication and no mutually acceptable algorithm can be agreed upon.

### [4.2.1](#) Multiple Authentication Considerations

If a secure communication channel (e.g., an IPsec Security Association) is used below iSCSI, authentication may occur both in setting up that channel and as part of iSCSI login. If these

authentications employ different identities, the security properties of the channel may not be constrained to the iSCSI authenticated identity (e.g., it may be possible for multiple iSCSI sessions with different authentication identities to use the same IPsec Security Association). This is particularly true when a machine identity is used for IPsec authentication, but a user identity is used for iSCSI authentication (the user identity may be that of some application rather than an actual human user). An analogous situation arises in the interaction of PPP and IPsec authentication when IPsec provides a secure communication channel for PPP encapsulated in L2TP, and most of the discussion of in see [Section 5.1.1](#) of [L2TP-SECURITY] is applicable to this iSCSI situation, as it involves a mix of user authentication (PPP) and machine authentication (IPsec).

Any multiple authentication approach for iSCSI MUST include some means of checking that the identities used in the two authentications for each connection correspond in some acceptable fashion, and these checks MUST be automatic because they may occur as part of a system boot procedure. A possible means of supporting such checks is to include both identities and communication endpoints as part of iSCSI access control and discovery information provided that the mechanisms for distributing and/or configuring such information are suitably secured.

#### [4.3](#) Cryptographic Integrity and Data Origin Authentication

Cryptographic integrity of all iSCSI communications after the login phase must be implemented, and this integrity mechanism MUST be keyed to provide data origin authentication for all received data. The key MUST be derived from or otherwise securely linked to the appropriate authentication; among the purposes of this is to prevent a TCP hijack attack from succeeding because the hijacker does not know the key required to generate the correct cryptographic integrity checks. This is a significantly stronger integrity requirement than the usual data integrity requirements for storage traffic; the integrity check MUST resist malicious tampering, and MUST be authenticated in a fashion that prevents an adversary from regenerating it. For example, an XOR-keyed CRC (CRC xor key) is not

a sufficient check because it has inadequate strength to resist tampering. An adversary who does not know the key can easily determine which bits to flip in a keyed CRC to offset bits flipped in the data covered by the CRC. Use of a cryptographic integrity mechanism in iSCSI is negotiable - see [Section 4.1](#).

The current status is that ESP [[RFC-2406](#)] with NULL encryption has been chosen as the implementation approach to meet this requirement, but the Authentication Algorithm (MAC, e.g., HMAC-SHA1) has not been selected. While iSCSI is clearly a "host" rather than a "security gateway" (see [Section 3 of \[RFC-2401\]](#)), ESP is likely to be used only in tunnel mode for reasons that include:

- (1) Only one IPsec encapsulation mode (transport or tunnel) is needed by a protocol like iSCSI.
- (2) Tunnel mode better accommodates NATs because the encapsulated IP and TCP checksums are correct in tunnel mode, whereas they are incorrect in transport mode. See [[NAT-REQTS](#)] for details.
- (3) Tunnel mode may allow the use of external IPsec gateways and related implementation structures, transport mode does not.

This is an example of the subset approach to security specification discussed in [Section 3](#).

#### [4.4](#) Confidentiality

A confidentiality mechanism MUST be implemented, but any such mechanism is OPTIONAL to use. The current status is also that ESP has been chosen as the implementation approach, but no preferred ESP transform (i.e., encryption algorithm) has been selected.

#### [4.5](#) Rekeying

A manual keying mechanism uses a pre-configured key without key exchanges or rekeying. This is insufficient because iSCSI will support long-lived connections that exchange enough traffic to cause ESP's 32-bit sequence number to rollover, exposing the connection to replay attacks. Other considerations may dictate or recommend rekeying considerably earlier than required solely to avoid rollover (e.g., to limit the amount of data encrypted or signed with the same session key). For these reasons, an interoperable automatic rekeying mechanism MUST be specified as part of iSCSI and will be REQUIRED of all conforming implementations.

### [5](#). iSCSI Implementation Characterization

#### [5.1](#) Implementation Classes and Resource Availability

iSCSI will be implemented on a variety of systems ranging from large servers running general purpose operating systems to embedded host bus adapters (HBAs). Host Bus Adapter is a generic term for a SCSI interface to other device(s); it's roughly analogous to the term

Network Interface Card (NIC) for a TCP/IP network interface, except that HBAs generally have on-board SCSI implementations, whereas most NICs do not implement TCP, UDP, or IP. In general, a host bus adapter is the most constrained iSCSI implementation environment, although an HBA may draw upon the resources of the system to which it is attached in some cases (e.g., authentication computations required for connection setup). More resources should be available to iSCSI implementations for embedded and general purpose operating systems. The following guidelines indicate the approximate level of resources that authentication, keying, and rekeying functionality can reasonably expect to draw upon:

- Low power processors with small word size are generally not used, as power is usually not a constraining factor, with the possible exception of HBAs, which can draw upon the computational resources of the system into which they are inserted). Computational horsepower should be available to perform a reasonable amount of exponentiation as part of authentication and key derivation for connection setup. The same is true of rekeying, although the ability to avoid exponentiation for rekeying may be desirable (but is not an absolute requirement).
- RAM and/or flash resources tend to be constrained in embedded implementations. 8-10 MB of code and data for authentication, keying, and rekeying is clearly excessive, 800-1000 KB is clearly larger than desirable, but tolerable if there is no other alternative and 80-100 KB should be acceptable. These sizes are intended as rough order of magnitude guidance, and should not be taken as hard targets or limits (e.g., smaller code sizes are always better). Software implementations for general purpose operating systems may have more leeway.

In summary, the primary resource concern for implementation of authentication and keying mechanisms is code size, as iSCSI assumes that the computational horsepower to do exponentiations (e.g., those required by SRP) will be available (SRP implementation is currently REQUIRED by iSCSI).

## [5.2](#) Implementation Scaling

There is no dominant iSCSI usage scenario - the scenarios range from a single connection constrained only by media bandwidth to hundreds of Initiator connections to a single Target or communication endpoint. SCSI sessions and hence the connections they use tend to be relatively long lived; for disk storage, a host typically opens a SCSI connection on boot and closes it on shutdown. Tape session length tends to be measured in hours or fractions thereof (i.e.,



rapid fire sharing of the same tape device among different Initiators is unusual), although tape robot control sessions can be short when the robot is shared among tape drives. On the other hand, tape will not see a large number of Initiator connections to a single Target or communication endpoint, as each tape drive is

## iSCSI Security Requirements and SRP-based ESP Keys Aug 2001

dedicated to a single use at a single time, and a dozen tape drives is a large tape device.

Given current networking technology, security solutions must work well at 1 Gbit/sec in terms of CPU overhead and/or availability of suitable hardware implementations. Solutions that scale up to 10 Gbits/sec are desirable but are not an absolute requirement as there are issues with a number of technologies at 10 Gbits/sec. This is of particular concern for HMAC [[RFC-2104](#)], which is the basis for the only MACs (cryptographic data integrity and data origin authentication) currently standardized for use with ESP. HMAC implementations are commercially available at 1 Gbit/sec (often based on hardware assistance of some form), but HMAC's computational structure and overhead raise serious concerns about scaling to 10 Gbits/sec.

### [5.3](#) Other Implementation Environment Concerns

This section gathers up several questions and answers on other iSCSI implementation environment topics.

Q: Will IPsec generally be present on systems supporting iSCSI due to other traffic requiring IPsec security?

A: No, especially for Targets, as they may have no important functionality other than iSCSI.

Q: What are the persistence requirements for security state across power off or loss of TCP connections?

A: Essentially none; most SCSI state does not survive power loss or system crash events with a few exceptions such as persistent reservations. Security state for open TCP connections need not survive the loss of those connections; any new connection will have to re-authenticate.

Q: What about load-balancing or fail-over middleboxes?

A: Discussions of iSCSI-aware middleboxes have usually assumed that such boxes serve as an endpoint for the iSCSI sessions on both sides of them. For example, this is why iSCSI specifies separate

CRCs for its header and data. The author has not seen any major objections in the IP Storage WG to the fashion in which IPsec can interact with such boxes (e.g., TCP header is encrypted, making it impossible to classify traffic based on TCP port number).

Q: What about NATs?

A: The IP Storage WG charter indicates that the ability to operate through NATs is important, but not an absolute requirement. Work is underway in the ipsec WG to specify transparent operation through NATs via UDP encapsulation.

Black

[Page 8]

---

iSCSI Security Requirements and SRP-based ESP Keys Aug 2001

## [6. SRP Keying of ESP](#)

This section describes a mechanism to provide keying material for ESP based on SRP inband authentication for iSCSI. This is a single-layer authentication approach that places all other security mechanism (i.e., ESP) below TCP. ESP is linked to the SRP authentication by deriving ESP's keys from the SRP authentication. ESP's position below TCP causes TCP retransmission to be invoked whenever ESP discards a packet due to a failed security check. If the failed check is due to data corruption, the result is TCP retransmission rather than an iSCSI or SCSI retry. This increases the integrity assurance of data delivered by TCP based on the more powerful integrity check in ESP MACs by comparison to TCP/IP checksums.

This approach to ESP keying requires an interface between iSCSI and the ESP implementation to transfer keying material from SRP to ESP; such an interface may make use of external IPsec gateways with iSCSI more difficult or impossible (e.g., software development may be required to securely transfer keys and related configuration information from iSCSI to the external IPsec gateway).

This section does not contain complete details; the major goal is to show how this mechanism works to enable it to be evaluated for use in iSCSI. The mechanism is currently described for SRP only, but appears to be extensible to any iSCSI inband authentication approach that provides good keying material.

### [6.1 Overall structure](#)

The overall structure of this keying mechanism consists of the following components:

- SRP inband iSCSI authentication (see [[ISCSI](#)]).
- iSCSI Negotiation of ESP SA parameters ([Section 6.2](#))
- ESP session key derivation from SRP results ([Section 6.3](#))
- ESP startup for an iSCSI TCP connection ([Section 6.4](#))
- Rekeying ([Section 6.5](#))

In summary, SRP is used as inband authentication for iSCSI, and iSCSI negotiates additional parameters for ESP Security Associations. Keys for these associations are derived from the results of the SRP authentication and used to initiate ESP security processing beneath an existing iSCSI TCP connection. The SAs used for such processing are tightly bound to the corresponding TCP connection and cannot be reused for other connections or traffic. Rekeying is handled by one of two methods described in [Section 6.5](#).

## [6.2](#) Negotiation

In addition to the SRP authentication conducted as part of iSCSI negotiation, the following additional parameters need to be negotiated by iSCSI:

- ESP authentication algorithm (MAC, e.g., HMAC-SHA-1) or "none"
- ESP encryption algorithm (e.g., AES) or "none"
- SPI value sets for both Initiator and Target

The scope of negotiation is a single TCP connection; every iSCSI TCP connection that uses this keying mechanism MUST perform SRP authentication and negotiate these parameters. The result will be different ESP keys for each TCP connection.

iSCSI should limit the supported ESP encryption and authentication algorithms to a small number that all implementations are REQUIRED to support, except that implementations may choose not to support any encryption algorithms. As indicated previously, these algorithms have yet to be chosen as of the date this draft was

written.

The Initiator and Target independently announce SPI value ranges via iSCSI negotiation. The Initiator tells the Target what SPI values to apply to traffic sent to the Initiator, and the Target likewise informs the Initiator what SPI values the Initiator to apply to traffic sent to the Target. To avoid having to renegotiate SPI values when rekeying, ranges of SPI values are announced. The two Least Significant Bits of the announced SPI values MUST be zero; the four SPI values obtained from all possible values for the two LSBs. The initial SPI values used when ESP security processing is initiated MUST have their two LSBs set to zero. Rekeying will cycle through each set of four SPI values (see [Section 6.5](#)). A range of four values is negotiated because two is too small for effective rekey signaling (see [section 6.5](#)) and eight seems excessive.

This negotiation is completely in the clear, and hence is vulnerable to man-in-the-middle tampering, but over a small number of possibilities. A single negotiation may not be particularly vulnerable to this tampering provided that care is taken in configuring the minimum levels of security protection offered and accepted (e.g., an Initiator who offers "none" for ESP authentication and encryption algorithms deserves the absence of security that may result). If this is insufficient, additional measures may be needed, such as requiring the exchange of security parameter offers and negotiated results (and/or hashes thereof) over the secure channel after ESP security processing is initiated.

Negotiation retries with weaker security than initially offered are dangerous and MUST NOT be performed. Consider an adversary who changes the ESP authentication algorithm parameter to "none" on iSCSI negotiation messages when neither party offered "none". If

both parties are prepared to offer and accept "none" on a negotiation retry, the man-in-the-middle can cause "none" to be the agreed value, even though it was not included in the initial negotiation.

### [6.3](#) Key Derivation

The same session keys are used for traffic in both directions on a single TCP connection used by iSCSI. In IPsec terminology this means that the two Security Associations which carry the connection's TCP traffic (one in each direction) share session keys.

SRP produces 320 bits of session keying material (K) as part of its authentication. These 320 bits of material are the result of a secure hash after the Diffie-Hellman exchange embedded in SRP. The ESP keying material is then derived from K as follows:

- (1) Set the 160 most significant bits of K aside as K\_SAVE. This is for use in rekeying.
- (2) Apply the SHA\_Interleave function from [Section 3.1 of RFC 2945](#) to the 160 least significant bits of K to yield K2.
- (3) The ESP encryption keying material is the 160 most significant bits of K2, and the ESP authentication keying material is the 160 least significant bits of K2. Use of this keying material is defined by the specifications for the encryption and authentication algorithms; any algorithm requiring less than 160 bits of keying material MUST use the most significant bits of the appropriate keying material.

It's easy enough to generate different keying material for each SA by following IKE's approach of incorporating the SPI value for each direction into the hashes (see the definition of KEYMAT in [section 5.5 of \[RFC-2407\]](#)). At (2) above, the SPI for the SA would be appended to the 160 least significant bits of K2 before applying SHA\_Interleave. The definition of SHA\_Interleave generalizes to arbitrary length inputs. The most obvious benefit of this is lengthening rekeying intervals when rekeying is based on the amount of data for which a key is used. The SPI values are publicly known, but this is also the case for IKE and does not appear to cause a security issue there.

#### [6.4](#) ESP Startup

At the completion of security negotiation, all communication MUST switch to using ESP with the negotiated algorithms, SPI numbers, and associated keys. This means that all traffic sent by an Initiator or Target after the iSCSI PDU containing SecurityContextComplete=Yes must be processed by ESP. This switch MUST occur at the end of the PDU containing this security completion indication. iSCSI requires that both sides send the security completion indication before

switching into full feature mode; access to SCSI data can only occur in full feature mode.

This approach of starting ESP processing during of an active connection may raise security monitoring/enforcement and implementation concerns. There is no meaningful security for traffic prior to the point at which ESP processing is initiated as described in the previous paragraph. Firewalls and network traffic monitoring systems may have difficulty verifying that iSCSI switches to use of ESP prior to entering full feature mode. Implementation concerns have also been raised in that some IPsec implementations may have difficulty in switching security processing from "off" to "on" underneath an existing TCP connection. All of these concerns are in need of further discussion, exploration and explanation.

There are at least three alternatives for initiating ESP security processing of iSCSI traffic after negotiation:

- (1) Just start ESP security processing under TCP as described in the first paragraph of this section. Traffic on the wire changes from "TCP in IP" to "TCP in IP in ESP in IP". The "TCP in IP" portion may not be visible in the latter format if encryption is in use.
- (2) Encapsulate all traffic prior to the start of security processing in tunnel mode ESP, but use an SPI value of zero and omit the authentication data prior to the start of ESP security processing.
- (3) Employ a specified universal SPI value, universal key and universal authentication algorithm for use prior to starting security processing. All traffic on any iSCSI connection will be authenticated with this key and algorithm prior to initiation of ESP security processing.

Alternative (1) is an honest reflection of what is actually happening. It is also the best fit with iSCSI's ability to negotiate "none" for both ESP security algorithms (in which case ESP is not used). On the other hand, it does not work through Network Address Port Translation (see [Section 4.1.2 of \[RFC 2663\]](#)) because adding ESP disables the port translation in the network without providing sufficient information to the destination to enable it to be reconstructed (i.e., the TCP port number will change as part of the change from  $\hat{\text{TCP in IP}}$  to  $\hat{\text{TCP in ESP in IP}}$ . Alternative (2) is not consistent with [\[RFC-2406\]](#) because it amounts to ESP with NULL authentication and NULL encryption prior to initiation of security processing. Alternative (3) provides a potentially false indication of security, as any serious adversary can be assumed to know the universal key, but it may provide some protection against less capable adversaries.

If this SRP/ESP keying mechanism is employed, one of the above choices must be selected and REQUIRED for all iSCSI TCP connections.

### 6.5 Rekeying

Rekeying only needs to provide new keys; it is not necessary to renegotiate any SA parameters in order to rekey. The Initiator or Target may initiate rekeying, and the results are applied to the h ESP SAs in both directions. This avoids negotiating key lifetimes (both initially and as part of rekeying), as either party to a connection can initiate rekeying when it determines that a key needs to be replaced. [Section 6.2](#) describes the means used to negotiate a range of four SPI values in each direction; rekeying causes the two least significant bits of the SPI value in each direction to be incremented, modulo 4 (i.e., incrementing 11 yields 00).

A range of four SPI values in each direction is negotiated to avoid confusion that could occur if only two values were negotiated. ESP may receive reordered IP packets, causing a packet with an unincremented SPI value to arrive after packets with incremented SPI values, even though the unincremented packet was sent first. If only two SPI values are available in each direction, it's not possible to determine a priori whether such a packet is a rekeying signal and needs to be processed with new keys versus a packet prior to the last rekeying event that needs to be processed with old keys. This is primarily a performance issue as there is only one correct set of keys for any packet. A range of four SPI values eliminates this confusion, and specifying sufficient minimum rekeying intervals should eliminate any need for larger ranges.

Rekeying MUST NOT be initiated until the initiating party receives confirmation that the other party has completed any previous rekeying event (i.e., packets arrive with the SPI value from that rekeying event and are verified to have been processed with the associated keys). Spurious rekeying is prevented by checking that an inbound IP packet is correctly processed with the corresponding keys before initiating rekeying in response to its reception. If ESP authentication is in use, the Authentication Data in the ESP trailer (see [\[RFC-2406\]](#)) MUST be verified before initiating rekeying. If only Encryption is in use, the Next Header, Pad Length and Padding fields in the ESP trailer (see [\[RFC-2406\]](#)) MUST be verified before initiating rekeying. Receipt of stale packets based on keys from a prior use of an SPI value will not pass these tests, and hence will not cause spurious rekeying.

Two different rekeying mechanisms are proposed depending on whether

a new key exchange is desired to decouple the new session keys from past keying material. Minimum time and/or amount of data transmitted between rekeying events SHOULD be specified to avoid excessive rekeying, and these times should be sufficient so that at the time of rekeying event N+1, no packets with keys from event N-1 are expected to be alive in the network. The initial SRP

## iSCSI Security Requirements and SRP-based ESP Keys Aug 2001

authentication on an iSCSI TCP connection is considered to be rekeying event 0, and hence the minimum rekeying specifications apply to the time and/or amount of data until the first rekeying event. If this keying mechanism is employed, further discussion will be needed to decide whether to specify one or both rekeying mechanisms. Specifying both mechanisms will entail also specifying a means of preventing the Initiator and Target from concurrently initiating both rekeying mechanisms. For this reason, it may be better to only specify one rekeying mechanism.

### [6.5.1](#) Rekeying Mechanism 1: No New Key Exchange

The new keying material is obtained by applying SHA\_Interleave to the keying material K\_SAVE saved by step (1) in [Section 6.3](#), and using the resulting 320 bit value as the input to steps (1) to (3) in [Section 6.3](#) to yield new keying material and a new K\_SAVE. Either communicating party may initiate rekeying by calculating new keys, incrementing the two least significant bits of the SPI value used to send data (increment of 11 yields 00), and immediately using the new keys to send data using the new SPI value.

A party who receives an incremented SPI value MUST process the data using the new keys (which it can calculate by itself) and MUST commence using the new keys and corresponding incremented SPI value to send data. This rekeying and transition to use of new keys and SPI value to send data MUST be completed before sending any iSCSI PDU that depends on the PDU whose arriving IP packet contained the incremented SPI value. Implementations MAY precalculate keys for future rekeying events in order to avoid delays caused by a need to perform rekeying calculations when a new SPI value is received.

### [6.5.2](#) Rekeying Mechanism 2: New Key Exchange

This rekeying mechanism repeats the SRP authentication in order to generate new keys based on the key exchange embedded in SRP. An iSCSI Initiator may initiate rekeying by using iSCSI Text PDUs to initiate a new SRP authentication. The keys derived from that



authentication are then used with incremented (mod 4) SPI values for ESP processing. An iSCSI Target may initiate rekeying by sending an iSCSI Asynchronous Message PDU to the Initiator to make this request; a new iSCSI Asynchronous Message code would need to be defined for this purpose.

## 7. Security Considerations

This entire draft is about security.

## 8. Changes from 00 Version

Added statement that there is no requirement to keep authentication identities confidential.

Black

[Page 14]

---

iSCSI Security Requirements and SRP-based ESP Keys Aug 2001

Removed discussion of IKE text from authentication requirements section to remove any implication that IKE is required. Added short discussion of multiple authentications and pointer to L2TP security draft.

Updated confidentiality section to indicate that confidentiality is now "MUST implement" but "OPTIONAL to use".

Added note that starting ESP on an active TCP connection will not work correctly when TCP port translation is in use.

Added statement that SAs used with the proposed key management are tightly bound to iSCSI TCP connections and hence not reusable to key management description.

## 9. References

[ISCSI] Satran, J., et.al., "iSCSI", [draft-ietf-ips-iscsi-07.txt](#), Work in Progress, July, 2001.

[ISCSI-IPSEC] Aboba, B., et.al., "Securing iSCSI with IPsec", [draft-aboba-ips-iscsi-security-00.txt](#), Work in Progress, August 2001.

[ISCSI-NAME-DISC] Bakke, M., et.al., "iSCSI Naming and Discovery", [draft-ietf-ips-iscsi-name-disc-02.txt](#), Work in Progress, August 2001.

[ISCSI-REQTS] Krueger, M., et.al., "iSCSI Requirements and Design

Considerations", [draft-ietf-ips-iscsi-reqmts-05.txt](#), Work in Progress, July 2001.

[L2TP-SECURITY] Patel, B., et.al., "Securing L2TP using IPsec", [draft-ietf-l2tpext-security-04.txt](#), Work in Progress, July 2001.

[NAT-REQTS] Aboba, B., "IPsec-NAT Compatibility Requirements", [draft-ietf-ipsec-nat-reqts-00.txt](#), Work in Progress, June 2001.

[RFC-2104] Krawczyk, H., M. Bellare, and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", [RFC 2104](#), February 1997.

[RFC-2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), March 1997.

[RFC-2401] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", [RFC 2401](#), November 1998.

[RFC-2406] Kent, S. and R. Atkinson, "IP Encapsulating Security Payload (ESP)", [RFC 2406](#), November 1998.

[RFC-2407] Piper, D., "The Internet IP Security Domain of Interpretation for ISAKMP", [RFC 2407](#), November 1998.

Black

[Page 15]

---

## iSCSI Security Requirements and SRP-based ESP Keys Aug 2001

[RFC-2408] Maughan, D., M. Schertler, M. Schneider, and J. Turner, "Internet Security Association and Key Management Protocol (ISAKMP)", [RFC 2408](#), November 1998.

[RFC-2409] Harkins, D. and D. Carrel, "The Internet Key Exchange (IKE)", [RFC 2409](#), November 1998.

[RFC-2663] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", [RFC 2663](#), August 1999.

[RFC-2945] Wu, T., "The SRP Authentication and Key Exchange System", [RFC 2945](#), September 2000.

## 10. Acknowledgments

This draft expands on topics originally discussed during the May 1st, 2001 interim meeting of the IP Storage WG. Ted Ts'o and Bernard Aboba have made significant contributions to these topics

since that meeting. A number of other people have contributed via discussions and email exchanges. All contributions, discussions and comments are gratefully appreciated and hereby acknowledged. This draft will expire before March 2002.

11. Author's Address

David L. Black  
EMC Corporation  
42 South Street  
Hopkinton, MA 01748  
Phone: +1 (508) 435-1000 x75140  
Email: black\_david@emc.com

Black

[Page 16]