To: INCITS T10 Committee
From: David L. Black, EMC
Date: 7 November 2006
Document: T10/06-388r3
Subject: SPC-4: Security Goals and Threat Model

Revision History

r0: Original Version
r1: Editorial corrections and cleanup
r2: More corrections and cleanup, add Authorization
r3: Final cleanup in 7 November 2006 CAP Meeting

In order to provide appropriate security services to protect SCSI communications and functionality, it is important to describe the goals of security and the threats against which protection is appropriate. This sort of threat description is generally called a threat model. The purpose of this document is to describe security goals and a threat model for SCSI communication. It is heavily based on Internet security goals and the Internet threat model in RFC 3552.

The text that follows is intended for incorporation into SPC-4. An early application of this text will likely be SSC-3 protection of transfer of encryption keys (for encryption resident on a tape drive and related SCSI devices).

<SPC-4: Add the following IETF Reference to Section 2.5:>

RFC 3552, Guidelines for Writing RFC Text on Security Considerations

<SPC-4: Add the following glossary entry>

Information Unit (IU): A delimited and sequenced set of information in a format appropriate for transport by the service delivery subsystem (e.g., a CDB for a specific SCSI command).

<SPC-4: Add the following text as Section 5.13.1 and subsections, renumbering the current Section 5.13.1 to 5.13.2 - note that if 06-369 is adopted, further renumbering of the subsections of 5.13 will ensue.>

5.13.1 Security Goals and Threat Model

5.13.1.1 Overview

In many cases, the security goals and threat model used for the Internet are applicable to SCSI commands. The Internet security goals and threat model found in RFC 3552 as they apply to SCSI are summarized in 5.13.1. Terms, concepts, and classes of security techniques that are defined in RFC 3552 are discussed based on their RFC 3552 definitions without modification in this standard.

The security goals and threat model described in 5.13.1 are valid for all SCSI device types. SCSI command standards may modify this model to deal with threats appropriate to specific device types.

5.13.1.2 Security Goals

The overall goals of security may be divided into two broad categories:

a) Communications security (i.e., protecting communications); and
b) System security.

These goals interact as a result of communications being carried out by systems, with access to those systems provided through communications channels. It is possible to provide security services that independently meet these goals. A common methodology is to secure the communications first and then provide secure access to systems over the secured communication channels.

Communication security is subdivided into the following primary areas of protection:

a) Confidentiality: Preventing unintended entities from seeing the data.
b) Cryptographic Data Integrity: Ensuring that the data that arrives is identical to the data that was sent.
c) Peer Entity Authentication: Ensuring that the communicating endpoints are the intended peer entities.

Data Origin Authentication (i.e., ensuring that the received data was sent by the authenticated peer) is the combination of Peer Entity Authentication and Cryptographic Data Integrity.

Non Repudiation enhances Data Origin Authentication with the ability to prove to a third party that the sender sent the data that the receiver received.

Cryptographic Data Integrity is called Data Integrity in
RFC 3552. The term Cryptographic is added in this standard to
distinguish the class of integrity protection required to counter
malicious adversaries from the class of integrity protection required
to deal with random data corruption (e.g., caused by cosmic rays or
electrical noise). Mechanisms used to deal with random data corruption
(e.g., parity bits and CRCs) have minimal value against malicious
adversaries that are able to modify integrity checks to conceal
their modifications to the data. Cryptographic Data Integrity
requires knowledge of a secret key in order to modify an integrity
check without that modification being detectable. A well designed
system provides a high level of assurance that an attacker is unable
to learn, guess, discover, or otherwise obtain the required secret key.

In addition to the primary areas, there is another area of control:

a) Authorization: Controlling what an entity is allowed to do. For communications
security this is control of the entities with which an entity is allowed
to communicate.

Access Control (i.e., controlling what an entity is allowed to access)
is a form of Authorization.

Systems security consists of protecting systems
from unauthorized usage, inappropriate usage, and denial of service.

5.13.1.3 Threat Model

Most secured systems are vulnerable to an attacker equipped with
sufficient resources, time, and skills.
In order to make designing a security system
practical, a threat model is defined to describe the capabilities that
an attacker is assumed to be able to deploy (e.g., knowledge, computing
capability, and ability to control the system).

The main purposes of a threat model are as follows:

a) To identify the threats of concern; and
b) To rule some threats explicitly out of scope.

Most security
measures do not provide absolute assurance that an attack has not

occurred. Rather, security measures raise the difficulty of accomplishing the attack to well beyond the attacker's assumed capabilities and/or resources. Design of security measures that resist attackers with essentially unlimited capabilities (e.g., certain nation-states) is outside the scope of this standard. Security measures that are susceptible to a level of capability available to some attackers may still be useful for deterring attackers who lack that level of capability, especially when combined with non-technical security measures such as physical access controls.

The computational capability of an attacker is treated as a variable because that capability is inherently a moving target as a result of more powerful processors. The computational capability of an attacker influences design aspects (e.g., key length). Well designed security systems are agile in that they are able to operate not only with different key lengths, but also with different cryptographic algorithms.

The Internet threat model described in RFC 3552 is generally applicable to SCSI, and is specifically applicable when Internet Protocols are used by the SCSI transport (e.g., iSCSI, Fibre Channel via FCIP or iFCP). Its basic assumptions can be summarized as:

a) End systems engaging in communication are not under the control of the attacker.
b) The attacker is able to read any communicated IU and undetectably remove, change, or inject forged IUs, including injection of IUs that appear to be from a known and/or trusted system.

Communications security designs are based on an additional assumption that secrets (e.g., keys) used to secure the communications are protected so that an attacker is unable to learn, guess, discover, or otherwise obtain them. A consequence of this assumption is that attacks against secured communications are assumed to begin without with advance knowledge of the secrets used to secure the communications.

5.13.1.4 Types of Attacks

The following types of attacks are considered:
a) Passive attacks (i.e., attacks that only require reading IUs), and
b) Active attacks (i.e., attacks that require the attacker to change communication and/or engage in communication).
More information on attack types is available in RFC 3552.

Simple passive attacks involve reading communicated data that the attacker was not intended to see (e.g., password, credit card number). More complex passive attacks involve post-processing the communicated data (e.g., checking a challenge-response pair against a dictionary to see if a common word was used as a password).

There are a wide variety of active attacks (e.g., spoofing, replay, insertion, deletion, and modification of communications). Man-in-the-middle attacks are a sinister class of active attacks that involve the attacker inserting itself in the middle of communication, enabling it to intercept all communications without the knowledge of the communicating parties for the purpose of insertion, deletion, and/or modification of the communications.

5.13.1.5 SCSI Security Considerations

The application of communication security techniques (see RFC 3552) is defined by command standards. This subclause describes specific design considerations in applying the threat model (see 5.13.1.3) to all SCSI device types.

SCSI environments tend not to be fully connected (i.e., there are restrictions on the SCSI device servers with which a SCSI application client is able to communicate) due to the following mechanisms:

a) physical and logical connectivity restrictions (e.g., in SCSI to SCSI gateways across different transports)
b) LUN mapping and masking, and
c) transport zoning.

The resulting connectivity is more limited than the Internet security assumption that an off-path attacker is able to transmit to an arbitrary victim (see RFC 3552).

SCSI security designs are also influenced by SCSI being a client-server distributed service model (see SAM-4) that is realized over a number of different SCSI transport protocols and interconnects.

Security functionality may be defined as part of a command set or at the SCSI transport level. Some SCSI transport protocols (e.g., Fibre Channel and iSCSI)

define security functionality that provides confidentiality, cryptographic integrity, and peer entity authentication for all communicated data. However, there are situations in which some or all of those mechanisms are not used and there are SCSI communications
whose scope spans more than one SCSI Transport Protocol (e.g., via a gateway between iSCSI and FCP). Security that is defined by a command set is appropriate for such situations.