To: INCITS T10 Committee
From: David L. Black, EMC
Date: 13 October 2006
Document: T10/06-388r2
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

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 elements 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

SCSI interactions between an application client and a device server over a service delivery subsystem are an example of network communications, an area in which significant security analysis has been performed. The security goals and threat model used for the Internet are generally applicable to SCSI. 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 redefining them in this standard.

The security goals and threat model described in 5.13.1 are useful

for all SCSI device types. SCSI command standards may elaborate, specialize and/or adapt 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 because communications are carried out by systems and access to systems is through communications channels, but 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 three primary areas of protecting communicated data, plus an additional area of control:

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 other endpoint of the communication in the intended peer entity.
d) Authorization: Controlling what an entity can do; for communications security, this is control of what an entity can communicate with.

The combination of Peer Entity Authentication and Cryptographic Data Integrity is also known as Data Origin Authentication, namely ensuring that the received data was sent by the authenticated peer.

Access Control (controlling what an entity can access) is a form of Authorization.

Note: Cryptographic Data Integrity is called Data Integrity in RFC 3552; the word "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 for the latter purpose (e.g., parity bits and CRCs) have minimal, if any, value against malicious adversaries who can modify integrity checks to cover their tracks. Cryptographic Data Integrity requires knowledge of a secret key in order to successfully modify an integrity check. A properly-designed system provides a high level of assurance that an attacker is unable to learn, guess, discover, or otherwise obtain the required secret key.

Systems security consists of protecting systems (e.g., communications systems) from unauthorized usage, inappropriate usage and denial of service.

5.13.1.3 Threat Model

Almost every security system is vulnerable to a sufficiently dedicated

and resourceful attacker. 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 against a resource.

The threat model contains information such as the resources available to an attacker in terms of information, computing capability, and control of the system. The purpose of a threat model is twofold:

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 they raise the difficulty of successfully 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
than can be overcome with 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 more powerful processors are built. The computational capability of an attacker influences design aspects (e.g., key length). Good security designs are agile in that they can 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 can 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 cannot 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

It is useful to distinguish passive attacks (i.e., attacks that only require reading IUs) from active attacks (i.e., attacks that require the attacker to change communication and/or engage in communication). More in-depth discussion of all 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). A particularly pernicious class of active attacks, called man-in-the-middle, 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.

5.13.1.4 SCSI Security Considerations

The application of communication security techniques (see RFC 3552) is defined in individual command standards. This subclause covers specific design considerations in applying the threat model (see 5.13.1.2) to all SCSI device types.

SCSI environments of even moderate size tend not to be fully connected because mechanisms such as the following place restrictions on which SCSI device servers a specific SCSI application client can send commands to:

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 usual Internet security assumption (see RFC 3552) that an off-path attacker is able to transmit to an arbitrary victim.

SCSI security designs are also influenced by the fact that SCSI is not a protocol in and of itself. Rather SCSI is 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 [FC-SP] 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, out of choice or necessity, 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.