To: INCITS T10 Committee
From: David L. Black, EMC
Date: 25 August 2006
Document: T10/06-388r0
Subject: SPC-4: Security Goals and Threat Model

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 a 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 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

SCSI interactions between an application client and a device server over a networked transport are an example of network communications, an area in which significant security analysis has been performed. In particular, the security goals and threat model used for the Internet are generally applicable to SCSI. This section summarizes the Internet security goals and threat model found in [RFC 3552]; the reader is strongly encouraged to consult the more comprehensive discussion in Sections 2-3 of [RFC 3552], as well as the discussion of classes of security techniques in Section 4 of [RFC 3552].

The Security Goals and Threat Model described here 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.

5.13.1.1 Security Goals

The overall goals of security can be divided into two broad categories, as stated in Section 2 of [RFC 3552]:

> We can loosely divide security goals into those related to protecting communications (COMMUNICATION SECURITY, also known as COMSEC) and those relating to protecting systems (ADMINISTRATIVE SECURITY or SYSTEM SECURITY). Since communications are carried out by systems and access to systems is through communications channels, these goals

obviously interlock, but they can also be independently provided.

Communication security can be subdivided into three areas of protecting communicated data:

- Confidentiality: Preventing unintended listeners from seeing the data.
- Cryptographic Data Integrity: Ensuring that the data that arrives is identical to the data that was sent.
- Peer Entity Authentication: Ensuring that the other endpoint of the communication in the intended peer entity.

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.

Note: Cryptographic Data Integrity is called Data Integrity in RFC 3552; the word "Cryptographic" is added here to distinguish the class of integrity protection required to counter a malicious adversary from the class of integrity protection required to deal with random data corruption (e.g., caused by cosmic rays, electrical noise, etc.). Mechanisms used for the latter purpose (e.g., parity bits and CRCs) have minimal, if any, value against a malicious adversary who can modify integrity checks to cover her tracks. Cryptographic Data Integrity generally requires knowledge of a secret key in order to successfully modify an integrity check; for a properly-designed system, an attacker will not know the required key.

Systems security also consists of protecting systems, often those involved in communication from Unauthorized Usage, Inappropriate Usage and Denial of Service.

5.13.1.2 Threat Model

The following explanation of the contents and purpose of a threat model is reproduced from RFC 3552:

> A Threat Model describes the capabilities that an attacker is assumed to be able to deploy against a resource. It should contain such information 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. First, we wish to identify the threats we are concerned with. Second, we wish to rule some threats explicitly out of scope. Nearly every security system is vulnerable to a sufficiently dedicated and resourceful attacker.

The last sentence above has important implications. 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. Design of security measures that resist attackers with essentially unlimited capabilities (e.g., certain nation-states) is out of scope, as being well beyond the expertise that a standards organization can reasonably bring to bear on this subject. Further, 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 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 over FCIP or iFCP). The full threat model can be found in Section 3 of RFC 3552, but its basic assumptions can be summarized as:

- Assumption: End systems engaging in communication are not under the control of the attacker.
- Assumption: The attacker can read any communicated PDU (Protocol Data Unit) and undetectably remove, change, or inject forged PDUs, including injection of PDUs that appear to be from a known and/or trusted system.

The computational capability of an attacker is treated as a variable in Internet communication security designs, as that capability is inherently a moving target as more powerful processors appear, and it directly influences design aspects such as key length. Good security designs are agile in that they can operate not only with different key lenghts, but also with different cryptographic algorithms, including operating modes for encryption ciphers.

5.13.1.3 Types of Attacks

It is useful to distinguish attacks that only require reading PDUs (Passive Attacks) from those that require the attacker to change communication and/or engage in communication herself (Active Attacks). More in-depth discussion of all of these attack types can be found in RFC 3552.

Simple passive attacks involve reading communicated data that the attacker was not intended to see (e.g., a password or credit card number sent in the clear). More complex passive attacks involve post-processing the communicated data, for example 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, including, spoofing, replay, insertion, deletion, and modification of communications, as well as a particularly pernicious class of attacks called Man-in-the-Middle that involve the attacker inserting herself in the middle of a communication, enabling her to intercept all communications without the knowledge of the communicating parties for the purpose of insertion, deletion, and/or modification of the communication.

5.13.1.4 SCSI Security Considerations

A discussion of communication security techniques for Internet protocols can be found in Section 4 of RFC 3552; the application of these and other communication security techniques to SCSI is a matter for individual command set standards. This section covers specific design considerations in applying RFC 3552's Internet Threat Model to SCSI.

SCSI environments of even moderate size tend not to be fully connected because mechanisms such as physical and logical connectivity restrictions (e.g., in SCSI to SCSI gateways across different transports), LUN mapping and masking, and transport zoning (e.g., Fibre Channel) place restrictions on which SCSI device servers a given SCSI application client can send commands to. The resulting connectivity is more limited than the usual Internet security assumption that an off-path host can generally transmit to an arbitrary victim (see Section 3.5 of [RFC 3552]).

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 (cf. SAM-4) that can be realized over a number of different SCSI Transport Protocols and Interconnects. Security functionality can be provided at the SCSI level (i.e., specified as part of a Command Set) or at the SCSI transport level. The standards for at least two SCSI Transport Protocols, FCP and iSCSI, specify security functionality that provides Confidentiality, Cryptographic Data Integrity, and Peer Entity Authentication for all communicated data. However, there are environments 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); these are two classes of situations in which SCSI level security protection can be appropriate.