5C Digital Transmission Content Protection White Paper

Hitachi, Ltd.

Intel Corporation

Matsushita Electric Industrial, Co., Ltd.

Sony Corporation

Toshiba Corporation

Revision 1.0

July 14, 1998

Introduction

As members of the Copy Protection Technical Working Group (CPTWG), Hitachi, Intel, Matsushita (MEI), Sony and Toshiba have jointly produced the Five Company (5C) Digital Transmission Content Protection (DTCP) specification, providing manufacturers with a simple and inexpensive implementation, while maintaining a high degree of protection.

The DTCP specification defines a cryptographic protocol for protecting audio/video entertainment content from illegal copying, intercepting and tampering as it traverses high performance digital buses, such as the IEEE 1394 standard. Only legitimate entertainment content delivered to a source device via another approved copy protection system (such as the DVD Content Scrambling System) will be protected by this copy protection system.

The DTCP specification relies on strong cryptographic technologies to provide flexible and robust copy protection across digital buses. These cryptographic techniques have evolved over the past 20 years to serve critical military, governmental, and commercial applications. They have been thoroughly evaluated by hackers and by legitimate cryptographic experts and have proven their ability to withstand attack. The cryptographic stability of the system is derived from the proven strength of the underlying technologies, rather than merely how well a certain algorithm can be kept secret.

A number of emerging technologies will take advantage of the IEEE 1394 high speed digital interface including desktop computers, DVD players, digital televisions and digital set-top-box receivers. The transparent DTCP framework allows consumers to enjoy high-quality digital pictures and sound without any noticeable performance or quality impact.

1394 Content Protection Architecture

Copy Protection Layers

The new content protection system addresses four fundamental layers of copy protection:

- Authentication and key exchange
- Content encryption
- Copy control information
- System renewability

Authentication and Key Exchange (AKE)

Before sharing valuable information, a connected device must first verify that another connected device is authentic. In an effort to balance the protection requirements of the film and recording industries with the real-world requirements of PC and CE users, the specification includes two authentication levels - full and restricted.

- **Full authentication** can be used with all content protected by the system, and must be used for copy-never content.
- **Restricted authentication** enables the protection of copy-one-generation and no-more-copies content. If a device handles either copy-one-generation or no-more-copies protection schemes, the device must support restricted authentication. Copying devices such as DV recorders or D-VHS recorders and devices communicating with them employ this kind of authentication and key exchange. No authentication is required for copy-freely content.

Both kinds of authentication involve the calculation of three encryption keys:

- an **authentication key**, established during authentication is used to encrypt the exchange key
- an **exchange key** used to set up and manage the security of copyrighted content streams
- a **content key**, used to encrypt the content being exchanged

When executing AKE, various information should be exchanged using 1394 asynchronous packets between source and sink devices. This mechanism of exchange using asynchronous 1394 packets is based upon the IEC-61883 specification and the AV/C Digital Interface Command Set. The necessary extensions to these specifications are described in detail in the DTCP specification.

Content Encryption

The content cipher, that is, the algorithm used to encrypt the digital content itself, must be robust enough to protect the content yet efficient enough to implement in PCs and CE devices. To ensure interoperability, all devices must support the specific cipher specified as the baseline cipher. The channel cipher subsystem can also support additional ciphers, the use of which is negotiated during authentication. All ciphers are used in the converted cipher block chaining mode. Converted cipher block chaining provides greater security than ordinary cipher block chaining.

The DTCP specification requires Hitachi's M6 as the baseline cipher. The M6 cipher is a common-key block cipher algorithm based on permutation-substitution. This rotation-based algorithm works the same way as encryption algorithms currently used in Japanese digital satellite broadcasting systems.

Optional, additional ciphers include the Modified Blowfish cipher and the Data Encryption Standard (DES) cipher.

The Content Cipher Subsystem must be able to support the bandwidth of an MPEG-2 compressed video stream. For PCs, this cipher subsystem may be implemented in software. Software M6 encryption/decryption of a 64 KB block of date on a 266-MHz Pentium® II Processor, had an approximate bandwidth of 200 Mbps.

For CE devices, the M6 channel cipher will typically be implemented in hardware. About 6K gates are estimated to be required for a 10-round M6 with C-CBC hardware implementation. This implementation is capable of encryption or decryption at 32 Mbps with a 25-MHz clock.

Copy Control Information (CCI)

Content owners need a way to specify whether their content can be duplicated. The content protection system must therefore support transmission of encrypted data between devices, utilizing **Copy Control Information (CCI)**. If source and sink devices have conflicting capabilities, they should follow the most restrictive CCI method(s) available, which is determined by the source device. Two methods can be used:

- The **Encryption Mode Indicator (EMI)** provides easily accessible yet secure transmission of CCI via the most significant two bits of the synch field of the isochronous packet header. The encoding used for the EMI bits distinguishes the content encryption/decryption mode: copy-freely, copy-never, copy-one-generation, or no-more-copies.
  - No authentication or encryption is required to protect content that can be copied freely.
  - Content that is never to be copied (e.g. content from prerecorded media with a Copy Generation Management System (CGMS) value of 11, such as a DVD Movie) can only be exchanged between devices that have successfully completed full authentication.
  - Content that can be copied one generation (e.g. content from prerecorded media with a CGMS value of 10, such as a pay TV program) can be exchanged between devices using either full or restricted authentication.
  - For content marked no-more-copies, future exchanges are marked to indicate that a single-generation copy has already been made. This content can be exchanged between devices using either full or restricted authentication.

  By locating the EMI in an easy-to-access location, devices can immediately determine CCI without needing to extract embedded CCI (e.g. In the MPEG transport stream). This ability is critical for bitstream recording devices (such as a digital VCR) that do not recognize and cannot decode specific

content formats. When multiple mechanisms are available, the most restrictive should be used. The EMI indicates the mode of encryption applied to a stream:

  - Source devices will choose the right encryption mode based on embedded CCI and set the EMI accordingly.
  - Sink devices will choose the right decryption mode by examining the EMI.

  If the EMI bits are tampered with, the encryption and decryption modes will not match, resulting in erroneous decryption of the content.

- Embedded CCI is carried as part of the content stream. Many content formats including MPEG have fields allocated for carrying the CCI associated with the stream. The integrity of embedded CCI is ensured since tampering with the content stream results in erroneous decryption of the content. Only devices capable of processing the content itself can process this form of CCI.

System Renewability

Devices that support full authentication can receive and process **System Renewability Messages (SRMs)**. These SRMs are generated by the Digital Transmission Licensing Administrator (DTLA) and delivered via content and new devices. System renewability ensures the long term integrity of the system and provides the capability for revoking unauthorized devices.

- Prerecorded content source devices such as DVD players should be able to update an SRM from prerecorded content media (such as a DVD disc). In addition, prerecorded content should carry a system renewability message current as of the time the content is mastered. They should also be able to update an SRM from another compliant device with a newer SRM.
- Devices such as a digital set-top box (STB) serving as a digital cable receiver or DBS digital broadcast satellite receivers are a real-time delivery source of copyrighted content. They should be able to update a SRM from content stream or from another compliant device with a newer SRM.
- Devices such as digital televisions are a receiver of copyrighted content. These devices should be able to update a SRM from another compliant device with a newer SRM.
- Devices such as DV recorders are a format-cognizant recording and playback device. Other recording devices such as D-VHS are a format-non-cognizant recording and playback device. SRM support by these devices is only necessary if they support prerecorded copyrighted content marked copy-never. Thus, full authentication is used. If SRM support is required, both types of devices should be able to update a SRM from another compliant device with a newer SRM.

Typical Device Components

Figure 1 shows the components typically required for a device to be compliant with digital transmission content protection, as applicable to the IEEE 1394 interface.

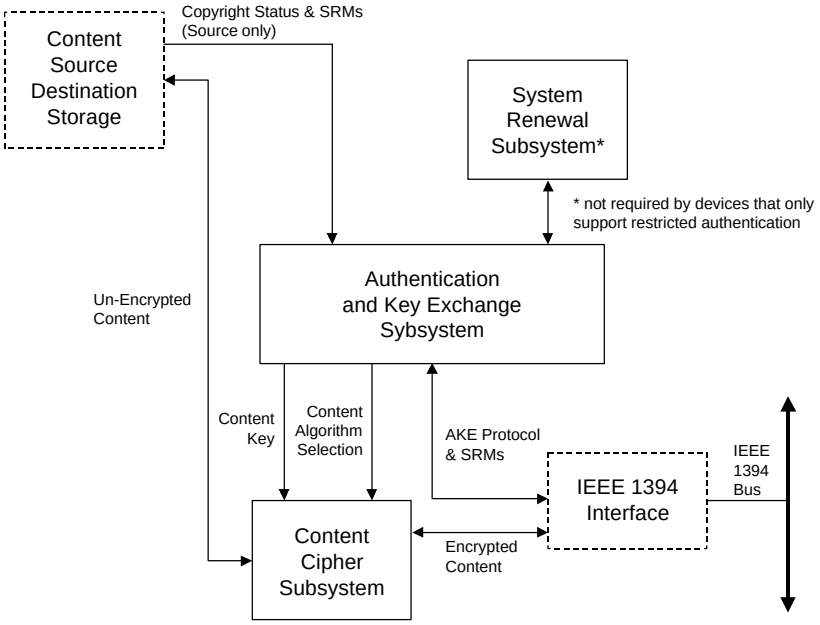

Figure 1. Typical Components of a Compliant Device

Subsystems in boxes with solid outlines are required for compliance. Boxes with dashed outlines are subsystems common to compliant and non-compliant devices. Depending on the device class, it will interact with a content source, a content destination, or content storage. For example, source devices receive content source, display devices send content to a destination, and recording and playback devices store content on media such as tape.

Components include:

- An **Authentication and Key Exchange (AKE) Subsystem** for performing full or restricted authentication
- A **Content Cipher Subsystem** for handling encryption/decryption of copyrighted content after authentication
- **System Renewal Subsystem** for supporting the system renewability mechanism associated with full authentication. The newest version of the SRM is stored here.

A robust **Random Number Generator (RNG)** is required for use as needed during authentication. For CE devices, the authentication and key-exchange mechanisms can be implemented in software running on an embedded micro-controller. To increase CE device performance, cryptographic acceleration hardware can be added. Currently, it is anticipated that the channel ciphers would be implemented in hardware. On a PC, the system can be implemented entirely in software. All implementations of this content protection system must be tamper-resistant.

1394 Content Protection Protocol

Figure 2 gives an overview of the content protection protocol flow. In this overview, the source device has been instructed to transmit a copy protection stream of content. In this and subsequent diagrams, a source device is one that can send a stream of content. A sink device is one that can receive a stream of content. Multifunction devices such as PCs and record/playback devices such as digital VCRs can be both source and sink devices.

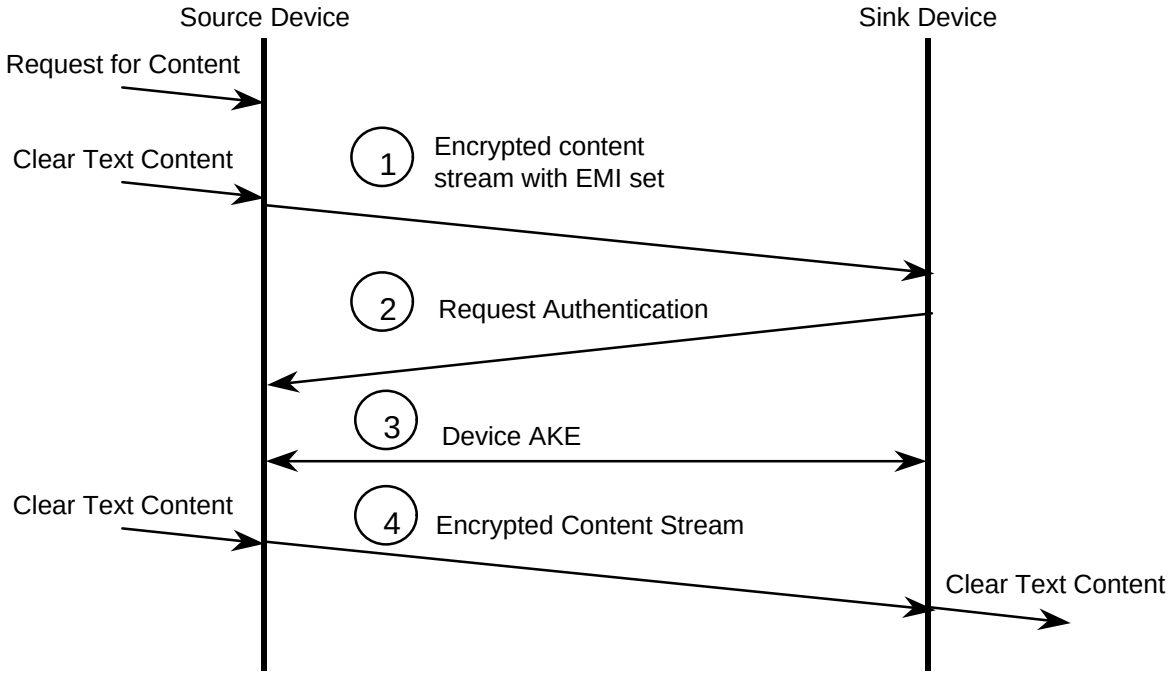

Figure 2 Content Protection Overview

1. The source device initiates the transmission of a stream of encrypted content marked with the appropriate copy protection status (e.g. copy-one-generation, copy-never, or no-more-copies) via the EMI bits.
2. Upon receiving the content stream, the sink device inspects the EMI bits to determine the copy protection status of the content. If the content is marked copy-never the sink device requests that the source device initiate Full AKE. If the content is marked copy-one-generation or no-more-copies the sink device will request Full AKE, if supported, or Restricted AKE. If the sink device has already performed the appropriate authentication, it can immediately proceed to Step 4.
3. When the source device receives the authentication request it proceeds with the type of authentication requested by the sink device. If the sink device requests Full AKE and the source device is only capable of Restricted AKE, the authentication performed will be Restricted Authentication .
4. Once the devices have completed the required AKE procedure, a content channel encryption key (content key) can be exchanged between them. This key is used to encrypt the content at the source device and decrypt the content at the sink.

Table 1 illustrates the authentication method performed, based on the source and sink device authentication capabilities:

Source	Sink	Authentication Performed
Full	Full	Full
Full	Full / Restricted	Full
Full / Restricted	Full	Full
Full / Restricted	Full / Restricted	Full
Full / Restricted	Restricted	Restricted
Restricted	Full / Restricted	Restricted
Restricted	Restricted	Restricted
Full	Restricted	None
Restricted	Full	None

Table 1. Authentication Method Matrix

Full Authentication

Full authentication can be used with all content protected by the system, and must be used for copy-never content. The full authentication protocol employs the public-key-based **Digital Signature Algorithm (DSA)** algorithm and the **Diffie-Hellman (DH)** key-exchange algorithm. Both the DSA and Diffie-Hellman implementations for the system employ **Elliptic Curve (EC)** cryptography. This technique offers superior performance compared to systems based on calculating discrete logarithms in a finite field.

- EC-DSA is a method for digitally signing and verifying the signatures of digital documents to verify the integrity of the data.
- EC-DH key exchange is used during full authentication to establish control channel symmetric cipher keys, allowing two or more parties to generate a shared key. Developed more than 20 years ago, the DH algorithm is considered secure when combined with digital signatures to prevent a so-called "man-in-the-middle" attack.

Figure 3 gives an overview of full authentication:

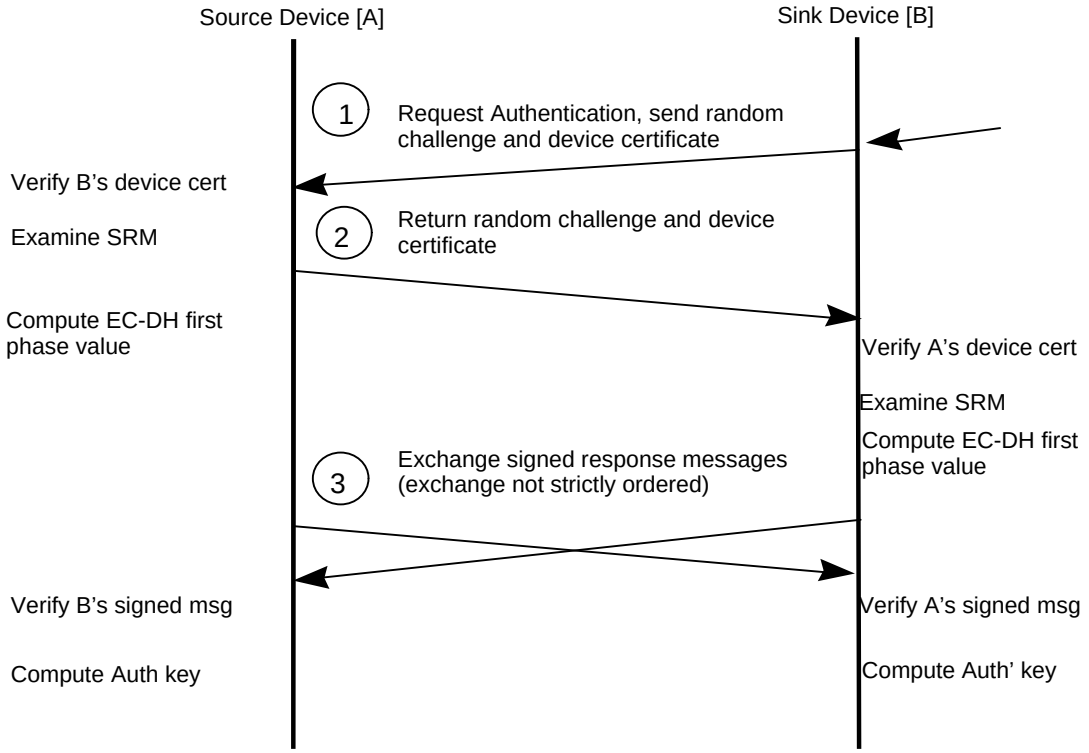

Figure 3. Full Authentication

During Full Authentication:

1. The sink device requests authentication by sending a random challenge and its device certificate. This can be the result of the sink device attempting to access a protected content stream (whose EMI is set to copy-never, no-more-copies, or copy-one-generation). The sink device may request authentication on a speculative basis, before attempting to access a content stream. If a sink device attempts speculative authentication, the request can be rejected by the source.
2. Device A then returns a random challenge and its device certificate. After the random challenge and device certificate exchange, each device verifies the integrity of the other device's certificate using EC-DSA. If the signature is determined to be valid, the devices examine the CRL embedded in their SRMs to verify that the other device has not been revoked. If the other device has not been revoked, each device calculates an EC-DH key exchange first-phase value .
3. The devices then exchange messages containing the EC-DH key exchange first-phase value, and the Renewability Message Version Number and Generation of the SRM stored by the device, and a message signature containing the other device's random challenge concatenated to the signed components. The devices verify the signed messages received by checking the message signature using EC-DSA with the other device's public key. This verifies that the message has not been tampered with. If the signature cannot be verified, the device refuses to continue. In addition, by comparing the exchanged version numbers and generations of the SRM, devices can invoke the system renewability at a later time. Each device calculates an authentication key by completing the EC-DH key exchange.

Restricted Authentication

Restricted authentication is an AKE method for devices with limited computing resources. This method is used by copying devices of any kind (such as DV recorders or D-VHS recorders) and devices communicating with them for authenticating copy-one-generation and no-more-copies contents.

The restricted authentication protocol employs asymmetric key management and common key cryptography and relies on the use of shared secrets and hash functions to respond to a random challenge. This method is based on a device being able to prove that it holds a secret shared with other devices. One device authenticates another by issuing a random challenge that is responded to by modifying it with the shared secrets and multiple hashings. Figure 4 gives an overview of restricted authentication:

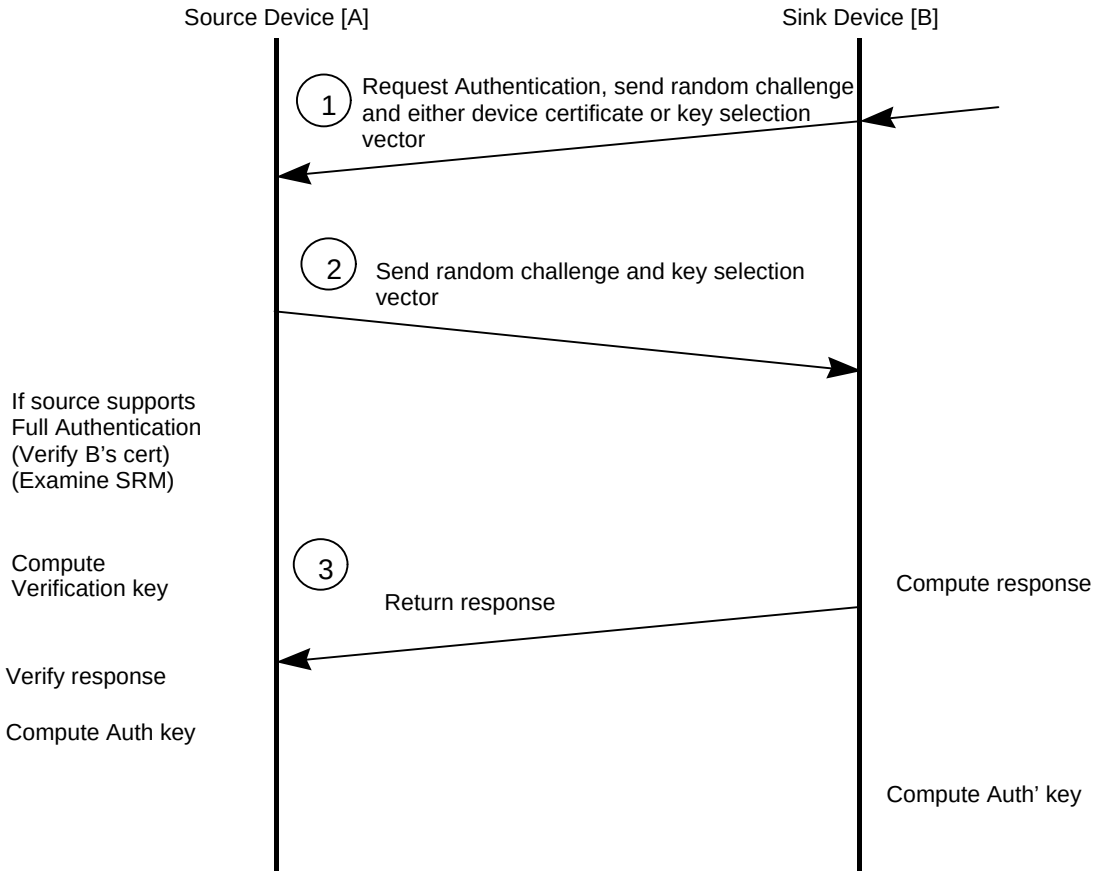

Figure 4. Restricted Authentication

1. The sink device initiates the authentication protocol by sending an asynchronous challenge request to the source device. This request contains the type of the authentication key to be shared by both the source device and the sink device corresponding to the service being requested, a random number generated by the sink. If the sink device knows that the source device does not have a capability for Full Authentication the sink sends its key selection vector to the source, otherwise the sink sends its Restricted Authentication device certificate.
2. The source device generates a random challenge and sends it to the sink device. If the source device supports Full Authentication, it extracts the device ID of the sink device from the certificate sent by the sink. It then checks 1) that the certificate sent by the sink is valid and 2) that the sink's device ID is not listed in the certification revocation list in the system renewability message stored in the source device. The source continues the protocol, only if both checks complete successfully. Then the source computes the verification key.
3. After receiving a random challenge back from the source device, the sink device computes a response using a verification key that it has computed and sends it to the source.
4. After the sink device returns a response, the source device compares this response with similar information generated at the source side using its verification key. If the comparison matches its own calculation, the sink device has been verified and authenticated. If the comparison does not match it, the source device shall reject the sink device. Finally, each device computes the authentication key.

Content Channel Management and Protection

Content channel management and protection mechanisms are used to establish and manage the encrypted channel through which protected content flows. Either full or restricted authentication (depending on the capability of the device) must be completed before establishing a content channel. Upon authentication of the devices, the source device sends an exchange key, encrypted with the authentication key, to the sink device.

Table 2 illustrates the allocation of the EMI and Odd/Even bits in the 1394 isochronous packet header. The EMI and Odd/Even bits exist only when the tag field is 01_2.

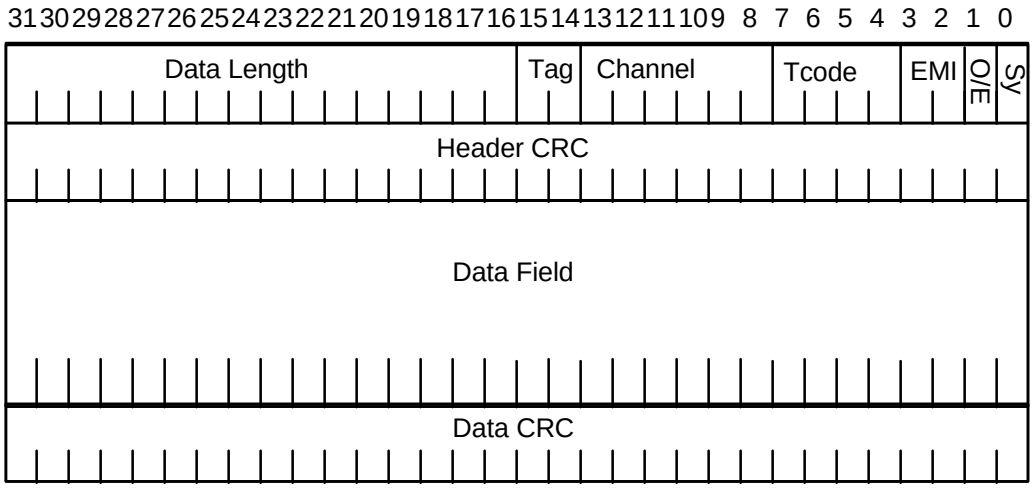

Table 2. IEEE 1394 Isochronous Packet Header

Table 3 contains the EMI value assignments:

value	EMI: Encryption Mode Indicator
00_2	Copy-freely
01_2	No-more-copies
10_2	Copy-one-generation
11_2	Copy-Never

Table 3. EMI Value Assignment

Table 4 contains the bit assignment of the Odd/Even bit:

bit	Odd/Even bit
0_2	Current content key is even value
1_2	Current content key is odd value

Table 4. Odd/Even Bit Assignment

Figure 5 shows an overview of content channel establishment and associated key management:

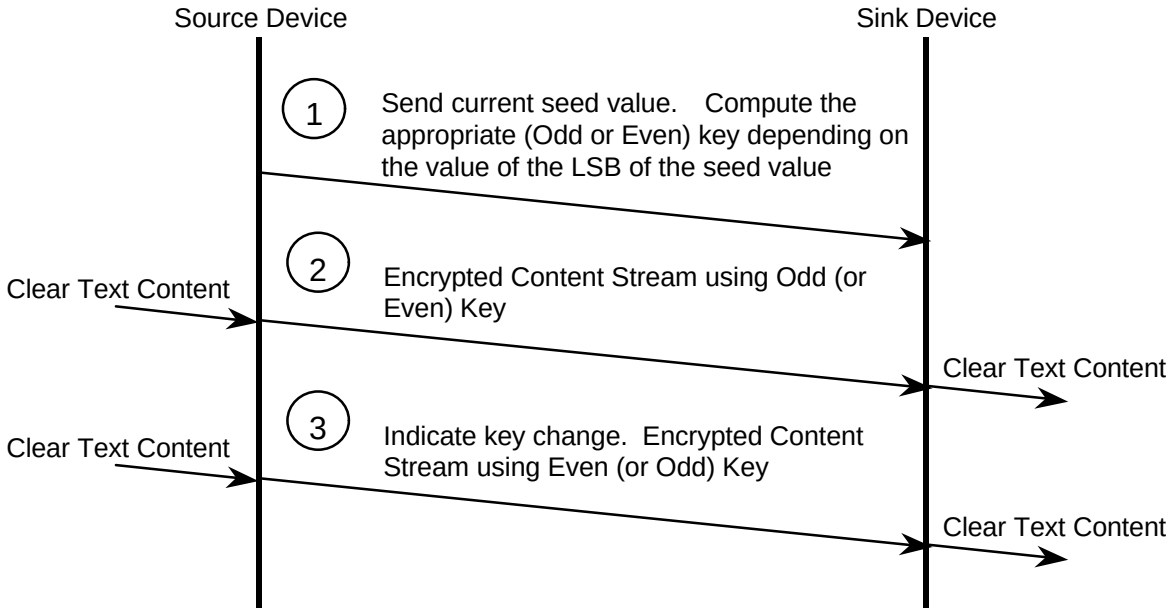

Figure 5. Content Channel Establishment and Management

Content keys are established between the source device and the sink device as follows:

1. When the source device starts sending the content, it generates a random number as an initial value of the seed of the content key. The initial seed is referred to as Odd or Even from its least significant bit.

2. The source device begins transmitting the content using the Odd or Even content key corresponding to the above reference of the initial seed to encrypt the content. The content key is computed by the source. A bit in the IEEE 1394 packet header is used to indicate which key (ODD or EVEN) is being used to encrypt a particular packet of content. If the initial seed is ODD, the Odd/Even bit in the 1394 packet header is set to Odd, otherwise it is set to Even.

   Upon receiving the seed, the sink device checks if the least significant bit of the seed matches the status of the Odd/Even bit. If both bits are identical, the sink computes the current content key. If those bits are different, it shows the key has been changed and the sink device computes the current content key. The source device prepares the next content key by computing the seed using the same process used for the initial calculation with exception that the seed is incremented.

3. Periodically, the source device shall change content keys to maintain robust content protection. To change keys, the source device starts encrypting with the new key computed above and indicates this change by switching the state of the Odd/Even bit in the IEEE 1394 packet header. The minimum period for change of the content key is defined as 30 seconds. The maximum period is defined as 120 seconds.

System Renewability

Renewability for Full Authentication

Devices that support full authentication can receive and process SRMs created by a licensing administrator and distributed via new content or new devices. System renewability messages can be updated from other compliant devices that have a newer list, from media with prerecorded content, or via compliant devices with external communication capability (i.e., over the Internet, phone lines or cable). This procedure should take place after completing the authentication and key exchange.

How System Renewability Messages Enter the Home

Figure 6 shows an example of how updated System Renewability Messages might be distributed into the home. In this example, the home device is shown as a PC. This device might receive updated renewal information over a network via a cable box or Digital Broadcast System set-top box. Another source of renewal information might be a new audio or video DVD disc purchased to run on a DVD player connected to the PC. Yet another source of renewal information might be a newly purchased DTV device that carries a newer SRM than the PC to which it is connected.

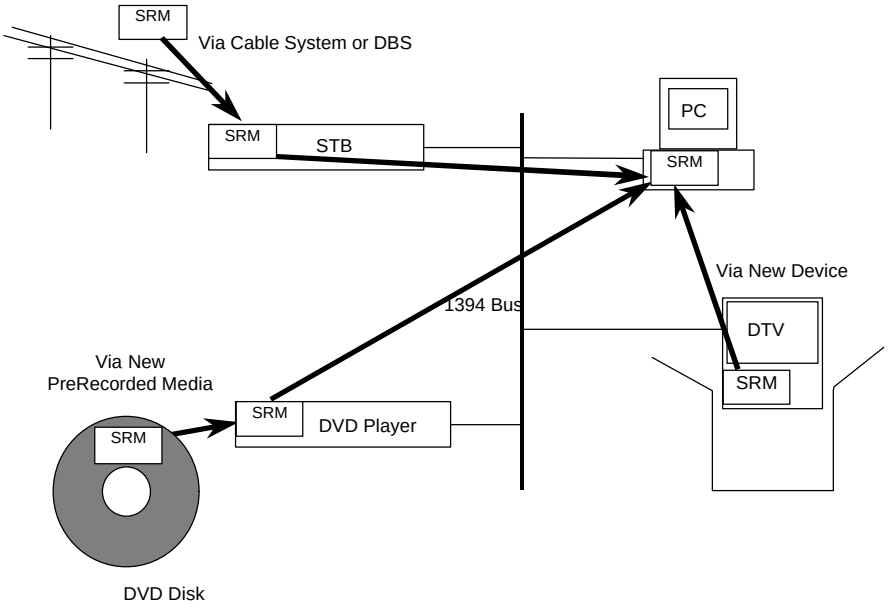

Figure 6. How System Renewability Messages Enter the Home

Revoking a Compromised Device

Figure 7 illustrates how a compromised device being used to circumvent digital protection might have its license revoked.

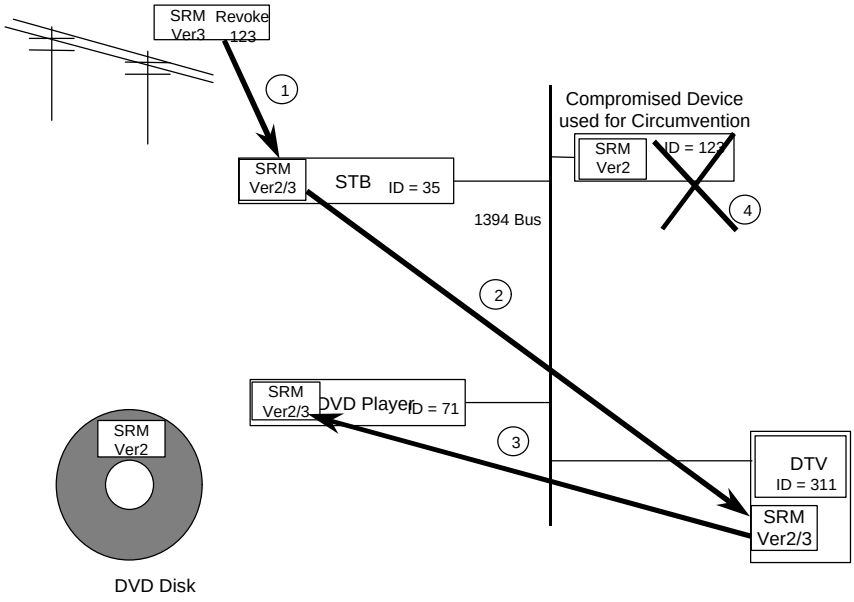

Figure 7. Revoking a Compromised Device

In this example, an illegal device 123 has been entered into the CRL in version 3 of the SRM. This new information enters the home over a network via a set-top box. The set-top box examines the version number of the new SRM. Since it is more recent than its own SRM version 2, it verifies the integrity of the new SRM using the licensing administrator's public key. If the SRM is valid and un-tampered with, the set-top box updates its SRM from version 2 to version 3.

The set-top box then passes the new SRM to a DTV device in the home when a cable movie is watched. The DTV device with SRM version 2 goes through the same verification procedure before updating its own SRM to version 3.

Next, a movie on a DVD disc with SRM version 2 plays on a DVD player. When the DVD movie is viewed on the DTV device, the DVD player notices that the DTV has a newer version of the SRM and requests that updated information. Again, the DVD player verifies the new SRM before accepting it as its own updated version. Device 123 has now been fully revoked in the home environment.

Similar procedures could be followed to rescind the revocation of device 123 in a future SRM, such as SRM version 4.

Renewability for Restricted Authentication

Source devices that support full authentication have the same renewability mechanism for sink devices that only support restricted authentication. Devices that only support restricted authentication do not have a renewability mechanism.

The Digital Transmission Licensing Administrator (DTLA)

The use of the DTCP specification and access to the intellectual property and cryptographic materials required to implement it will be the subject of a license. A license administrator, referred to as the **Digital Transmission Licensing Administrator (DTLA)**, is responsible for establishing and administering the content protection system based in part on the specification. Implementation of the DTCP specification requires a license from the DTLA.

While DTCP has been designed for use by devices attached to serial buses as defined by the IEEE 1394-1995 standard, the developers anticipate that it will be appropriate for use with future extensions to this standard, other transmission systems, and other types of content as authorized by the DTLA.

Figure 8 shows some potential DTLA operations that could result when the DTCP specification is implemented.

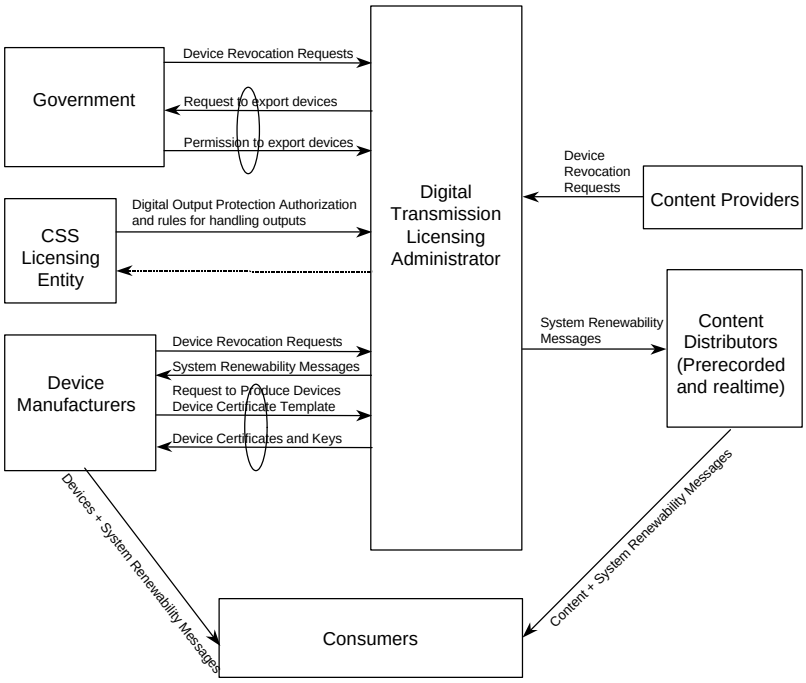

Figure 8. Potential Licensing Administrator Operations

For information regarding the DTCP specification and the DTLA, please send e-mail to:

dtla@dtcp.com