Digital Rights Management

Hubris, history, hacks.

Yan Shoshitaishvili
UCSB Seclab

Overview

- Content Duplication
- Digital Rights Media - History
- Media-specific DRM
- MovieStealer
  - Design
  - Optimizations
  - Countermeasures
  - Results
  - Ethics and Legality

Content Duplication

The Situation

- Classical media model relied on difficulty of duplication
- In the modern age, copying is (nearly) effortless
  - VCRs
  - Tapes
  - Floppies
  - Internet
- Producers "lose" revenue for copied content

Filesharing

- VHS piracy was a minor annoyance
- Filesharing was BIG.
  - First generation: Napster
  - Second generation: Kazaa, Gnutella (Limewire), eDonkey
  - Third-generation: Bittorrent (Suprnova, isohunt, The Pirate Bay)
  - Also Usenet

Why?

- Motivations for piracy
  - finances
    - not paying for media
    - actually selling copied media
  - fun
    - challenging
    - social
  - archiving
  - interoperability

Mitigation

"Our media is all over the net. What can we do?"

- Several approaches, depending on type of media (video, software, etc) and desired control
- Cat and mouse game!

DRM Goals

- Many possible goals of DRM
  a. prevent copying (copy protection)
  b. prevent playback by unauthorized devices
  c. prevent playback by unauthorized users
  d. identify pirates

Digital Rights Management - History

VCR DRM

- Movie studios were concerned with easy movie copying
- Macromedia developed a method to scramble copied media for VCRs
- Takes advantage of differences between TVs and VCRs to scramble copy

- Reference: http://bit.ly/eWGUrF

VCR DRM - Bypass

- Hardware exists to strip out the scramble-causing data
- Bypass is rare due to specialized hardware

Software DRM

- Computer software was easily copied, leading to a perceived loss of profits by software makers
- Several approaches to copy protection

Software DRM - Education

- "Don't copy that floppy!"
- Bypass: no one cared

Software DRM - Possession-based

- Manual checks
  - "Type the third word in the second paragraph on page 4 of the manual."
  - Bypass: copy the manual
- Physical dongles
  - "Plug the dongle into the serial port to continue."
  - Bypass: serial port emulation
- CD/Floppy check
  - intentional bad sectors created by special process
- General bypass: software patching

Software DRM - Online

- Online activation
  - EA controversy
  - Software patching
- Require the user to be "always-on"
  - MMOs have this built-in
- Future: game streaming?

Media-specific DRM

Media DRM - Challenges

- Media is "dumb"
  - audio files don't execute code
  - attempts to change this end in tears
    - Sony DRM debacle http://en.wikipedia.org/wiki/Sony_BMG_copy_protection_rootkit_scandal
- In the old days: must be playable offline
- Solution: cryptography

Media DRM - CSS

- CSS - Content Scramble System
- Produced by the DVD Copy Control Association
- Encrypts DVD content
  - hides keys in a special area of the DVD to prevent copying

Media DRM - CSS Bypasses

- A group of people broke CSS in 1999
  - Most famous member: Jon "DVD Jon" Johansen
  - "DeCSS" used extracted key from software player
  - Legal insanity ensues
- CSS also found to be brute-forceable
  - 40-bit keys
  - with optimizations, several seconds on modern systems

Media DRM - HDCP

- "Trusted path" from media to TV
- The goal: never leave content unprotected
- The reality: not effective
  - re-encryption
  - master key leak (2010)

Media DRM - Streaming Services

- Rise of streaming services
  - Video (MS Playready, Adope RTMPE): Netflix, Hulu, Amazon
  - Audio: Spotify, Rhapsody
- Different requirements
  - Ok to require internet connection
- General approach: encrypt everything
  - encrypt media with "content key"
  - encrypt content key with "user key"

Digital Rights Management - Weaknesses

Cryptographic DRM schemes have three main weak points:

- Content keys
  - Too platform-specific.
- Analog hole
  - Suffers quality loss due to lossy encoding.
- Content sniffing
  - Our approach.

MovieStealer - Design

MovieStealer - Intuitions

1. Decrypted content is accessible at some point in the program.
2. Media data is accessed in buffers.
3. Can differentiate between encrypted and encoded (compressed) buffers.
   - Specifically, encoded/compressed data has high entropy but low randomness, while encrypted data has high entropy and high randomness.
4. Can be used to locate the decryption point!

MovieStealer - Challenges

- Gigabytes of information
- Media players are *complex*
  - real applications
  - obfuscated
  - will not function with too much overhead
- Generality
  - We must choose the cases in which MovieStealer should be applicable

MovieStealer - Approach Overview

Goal: find the decrypted stream!

1. Loop detection.
2. Buffer detection.
3. Data paths.
4. Statistical analysis.
5. Content dumping.

Interlude - Basic Blocks

- Programs can be split into basic blocks
- BBs are a sequence of instructions that are always executed together

```
int x = getch();
int y = 2;
```

```
if (x == 2)
    printf("MATCH\n");
```

```
else
    printf("NO MATCH\n");
```

MovieStealer - Loop Detection

- Maintain call stack and basic block stack.
- Push block on entrance, pop on exit.
- If the same basic block is on the stack twice in a single function, we count it as a loop.

```
x = 10;

while (x > 0)
{
    printf("X is %d\n", x);
    x--;
}

printf("DONE\n");
```

MovieStealer - Loop Detection

Some crypto implementations might reuse the same loop for encryption and decryption.

```
void crypto_loop(void *key, void *in,
                 void *out, int len);

void encrypt() {
    crypto_loop("key", dec, enc, len);
}

void decrypt() {
    crypto_loop("key", enc, dec, len);
}
```

Solution:

Identify loops by the start address of their first basic block *and* the call stack.

MovieStealer - Buffer Detection

1. Instrument read and write operations.
2. Record target of each read and write. Each target is labeled as an *original buffer*.
3. These individual accesses are merged into *composite buffers*.
4. Composite buffers are merged.

Movie Stealer - Buffer Merging

The target of every read and write operation of a loop is labeled as an *original buffer*.

0x1000	Original buffer (size 4)
0x1004	Original buffer (size 4)
0x1008	Original buffer (size 4)
0x100c	Original buffer (size 4)
0x1010	Original buffer (size 4)
0x1014	Original buffer (size 8)
0x1018	

Movie Stealer - Buffer Merging

Two *original buffers* are merged into a *composite buffer* if they are adjacent and of the same size. Track element size.

0x1000	Composite buffer (element size 4)
0x1004	
0x1008	Composite buffer (element size 4)
0x100c	
0x1010	Original buffer (size 4)
0x1014	Original buffer (size 8)
0x1018	

Movie Stealer - Buffer Merging

An *original* and a *composite* buffer are merged if they are adjacent and the element sizes match.

<table>
<tr><td>0x1000</td><td rowspan="2">Composite buffer (element size 4)</td></tr>
<tr><td>0x1004</td></tr>
<tr><td>0x1008</td><td rowspan="3">Composite buffer (element size 4)</td></tr>
<tr><td>0x100c</td></tr>
<tr><td>0x1010</td></tr>
<tr><td>0x1014</td><td rowspan="2">Original buffer (size 8)</td></tr>
<tr><td>0x1018</td></tr>
</table>

MovieStealer - Buffer Merging

When no more original buffers can be merged, they are relabeled as composite buffers.

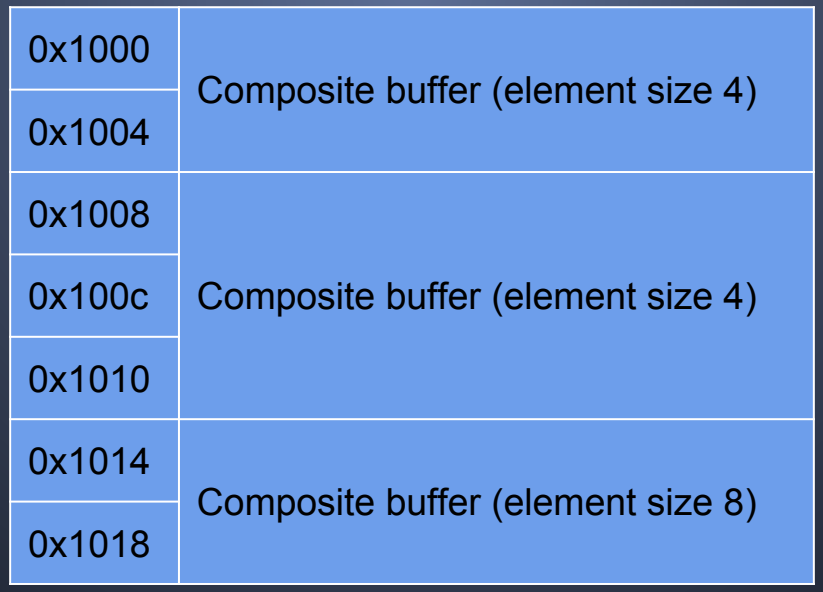

MovieStealer - Buffer Merging

Composite buffers are merged if:

- (distance) / (combined size) < 0.2
- Their element sizes are the same.

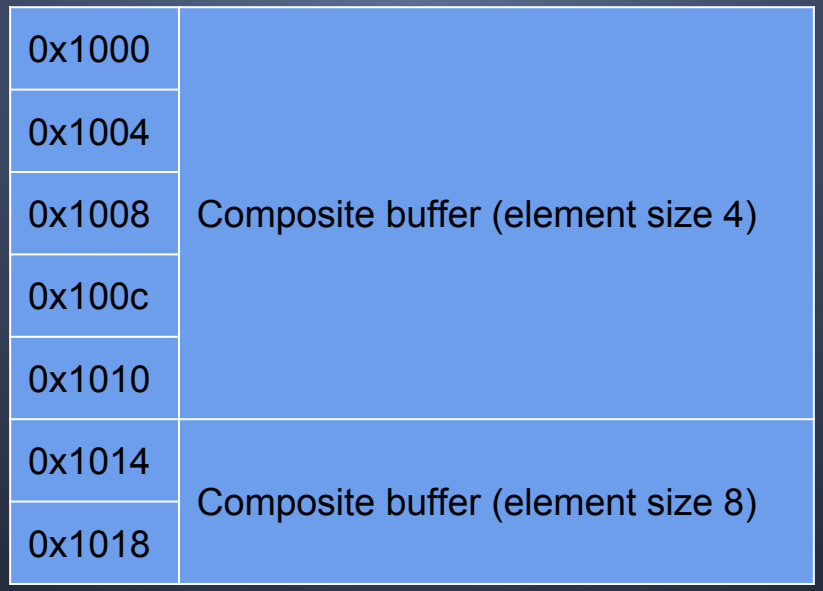

MovieStealer - Data Paths

- A data path consists of an input (a read buffer) and an output (a write buffer).
- Rather than track data flow, we create a data path for each combination of read/write buffers in a loop.
- Result: an over-approximation of the data flow in all loops of the application.

MovieStealer - Statistical Analysis

- The input and output of each data path is saved, and statistical analysis is performed on the aggregated data.
- We measure the difference in randomness and entropy across each data path.

Stage	Input		Output	
	Entropy	Randomness	Entropy	Randomness
Download	*High*	*High*	*High*	*High*
Decrypt	*High*	*High*	*High*	*Low*
Decode	*High*	*Low*	*Low*	*Low*

MovieStealer - Statistical Analysis

- The Chi-Squared randomness test is used to measure randomness.
- Random data gives values ~1.0, while nonrandom data gives very high values.
- Care has to be taken to collect enough data to avoid false positives.
  - 800kb needed to avoid misclassifying non-random data as random.
  - 3.8kb needed to avoid misclassifying random data as non-random.

MovieStealer - Reconstruction

- Dumped data needs to be reconstructed.
- A reconstructor has to be implemented for each platform.
- Manual implementation process.

MovieStealer - Optimizations

MovieStealer - Optimizations

- The basic approach still has too much overhead for performance-demanding services to function.
- We developed several optimizations to improve speed.
- Two main categories:
  - Improved loop selection - optimally determine analyzation order.
  - Efficient loop analysis - quickly eliminate/confirm candidate loops.

MovieStealer - Order Optimizations

- On-Demand Instrumentation.
  - avoid analyzing startup code.
- Execution Frequency.
  - analyze most-frequently executed loops first
  - data streaming and decryption is the most common operation of a streaming media player.
- Instruction analysis.
  - Select loops likely to contain cryptographic code.

MovieStealer - Analysis Optimizations

- Bandwidth filtering
  - Eliminates loops that don't process enough data.
  - When streaming a media file of size S:

Stage	Input Bandwidth	Output Bandwidth
Download	*S*	*S*
Decrypt	*S*	*S*
Decode	*S*	*greater than S*

- Copying optimizations.
  - Avoid unnecessary data copying.
  - For writes, only copy data on loop exit.
  - For reads, copy immediately in case of overwriting.

MovieStealer - General Optimizations

- Callstack key.
  - Speeds up the callstack handling.
  - Keep a dword instead of a stack of function addresses.
  - On function entry, XOR function entry address onto callstack key.
  - On Function exit, XOR function entry address onto callstack key, cancelling it out.

MovieStealer - Experimental Results

Evaluation

- Three DRM platforms:
  - Microsoft PlayReady - used by Netflix for video streaming.
  - Adobe RTMPE - used by Amazon Instant Video and Hulu for video streaming.
  - Spotify's music protection.
- GPG for testing optimizations.

Results - GPG

GPG was used to quantify the effects of our performance optimizations, since the media players would fail to work without them.

Optimizations	Loops Instrumented	Seconds Elapsed
All	*7*	*31*
All but callstack key	*6*	*47*
Only instruction analysis	*10*	*49*
Only bandwidth filtering	*35*	*180*
Only execution frequency	*40*	*3480*

Results - DRM

All evaluated DRM platforms succumed to MovieStealer.

Optimizations	Loops Instrumented	Loops Traced	Buffers Identified	Seconds Elapsed
Netflix	*2274*	*58*	*80*	*110*
Hulu	*1529*	*46*	*14*	*281*
Amazon Video	*1258*	*35*	*6*	*146*
Spotify	*2305*	*224*	*60*	*536*

MovieStealer - Countermeasures

Countermeasures

- Several countermeasures are possible.
  a. Attacking the instrumentation.
     - intricate anti-debugging techniques
  b. Attacking the loop detection.
     - VM-ed loops to frustrate analysis
  c. Attacking the buffer detection.
     - non-consecutive buffer layouts
  d. Attacking the decryption detection.
     - pollute encrypted stream with nonrandom bytes
     - pollute decrypted data with random bytes
  e. Attacking the pirates.
     - watermarking

Ethics and Legality

Legality

- We believe this work to be legal under DMCA.
  - Consulted with UC counsel and the EFF

Ethics

- Responsible disclosure.
  - Contacted Microsoft, Spotify, Adobe, Amazon, and Hulu.
  - Microsoft, Spotify, and Adobe responded
    - Tested MovieStealer.
    - Confirmed DRM bypass.
    - Provided comments for the paper.
    - Encouraged publication.
- No tool release!

Question Time

Questions?