WEIRD MACHINES AND REVISITING "TRUSTING TRUST" FOR BINARY TOOLCHAINS

JULIAN BANGERT, REBECCA SHAPIRO, SERGEY BRATUS

DARTMOUTH COLLEGE



OUTLINE

- Trust chains without bugs: what can go wrong?
- A chain of trust is a chain of parsers/loaders: how bug-less Babel breaks (badly!)
- Case studies:
 - Any input table is a program (recall 29c3)
 - ELF signing, Mach-O signing
 - ELF kernel loader vs. RTLD/ld.so

LANGSEC VS CHAINS OF TRUST

FURTHER REFLECTIONS ON TRUSTING TRUST

- Ken Thompson, "Reflections on Trusting Trust", 1984
 - (almost) 30 years ago
 - "You can't trust code that you did not **totally** create yourself"
 - invisible links in the chain-of-trust (..."well-installed microcode bugs"...)

BEYOND THE BUGS IN TRUSTING TRUST

- What if there were **no bugs** in any given piece of sw/hw link of the trust chain?
 - What if the code did exactly what the author **intended**, and
 - you can **trust** the author?
- Would we solve "trusting trust"?

HELL NO!



BECAUSE WE ARE IN BABEL



DIALECTS OF INPUT?



I CAN HAS UR TRUST CHAINZ?



CHAIN OF TRUST

- Chain: execution environments of increasing complexity and power (from boot to full OS ABI)
- Goal: no unexpected computation throughout
- Same code/data bytes interpreted (i.e., **executed**) by several consecutive environments
- Two kinds of trust in data or code:
 - input (code/data) can be checked for effects
 - input (code/data) was signed & been immutable since someone checked it for effects

TRUSTED BITS: HARD VS SOFT

- It's hard to statically find out what code does
 - So we "freeze" (sign, etc.) code



- But we can't freeze full binary images without impairing composition
 - Libraries, dynamic modules, ASLR,...
- So we add "tables" to drive composition/ mutability mechanisms



ANY TABLE IS A PROGRAM

- "Tables" drive computation that locates signed ABI sections & their signatures
- Tables are **bytecode** for automata in signature verifiers/loaders/parsers



"ANY INPUT IS A PROGRAM"

- (Meta)data is just a **program** for code that interprets it. [Hopefully, analyzable for effects]
- Any sufficiently complex input data is indistinguishable from **byte code** driving a VM
- Parser code for any sufficiently complex input format is indistinguishable from a VM for its inputs (= "byte code")
- Input validation is "runtime verification" of inputs as programs

WHAT CAN GO WRONG?

- Input not well-defined/recognized
 => code's assumptions about "checked"
 input will be violated (bug/vuln)
- Input well-formed but so complex there's no telling what it does
- Input is seen differently by different pieces of program/ toolchain



LIBERATING THE SOFT BITS



"LIBERATED SOFT BITS"



THE ELF/ABI CASE STUDY Ken Thompson's planted bug Compiler ld.so (RTLD) **DWARF** Relocator exceptions Linker ("dl-machine") Loader **#PF**

binfmt_elf

#DF

THE ELF/ABI CASE STUDY



THE ELF/ABI CASE STUDY



"WEIRD MACHINES"

- DWARF exception handling data + .eh_frame
 + Glibc = Turing machine (WOOT 2011)
- **Relocation** entries + dynamic **symbols** = Turing machine on process' address space
- GDT + IDT + TSS + page tables
 + # PF + # DF = Turing machine in ia32
 (WOOT 2013)
- More coming :)

VALIDATION IS VERIFICATION

- Tables are trusted when "valid" <=> drive computation as expected
- Validation of tables is static analysis of computations they induce on parsers & loaders
- Code that **interprets** ("executes") tables must be simple enough to allow trust via static analysis





THE MACHINES OF CODE SIGNING

CODE SIGNING

- Code signing -> primary trust evidence for binaries:
 - "trustworthiness from static measurements"
- Developer/distributer digitally signs bytes in binary
 - Integrity and attribution
- Easy to implement poorly
- It's not just an algorithm, it's a lifestyle
 - Key management
 - Program in memory =/= program on disk
 - It is merely *influenced* by what is on disk
 - Many "machines" involved in verification
 - Parsers, interpreters, validators

CODE SIGNING MACHINE COMPOSITION

Self-Operating Napkin



ON TRUSTING SIGNATURE VALIDATION

- Are our machines correctly implemented?
- Do we understand what our machines are capable of?
- Do different machines agree on how to parse / understand input?
- Do the tables carry correct and complete data?
- Can we trust **transformations** made after this static analysis?
- Enforcement?

ELF CASE STUDY

- Parsers
 - Signature and signature metadata
- Interpreters / translators
 - Binary -> hashes
- Validators
 - Validate certificates, signatures, hashes

ELF CODE SIGNING

- Executable signing implementations
 - bsign (Marc Singer)
 - elfgpg (Bart Trojanowski)
 - elfsign (skape)
 - SignELF (Joe Fox)
 - signelf (Vivek Goyal, proposed to kernel developers)
 - elfsign (Solaris)
 - ^ incompatible with each other
- Kernel module signing (3.7+, evolved over time)
 - DigSig (until 2009)

ÅRE OUR MACHINES CORRECTLY IMPLEMENTED?

- XML parsers (in the case of Mach-0)
- ASN.1 BER parsing (easy as pie, right?)
- Most written in C/C++ (...)

HOW POWERFUL ARE OUR MACHINES?

- 29c3 "The Care and Feeding of Weird Machines in ELF Metadata"
- Metadata-driven root shell backdoor in ELF and Mach-O
- LOCREATE (skape)
 - unpacker written in PE metadata

DROPPING A SHELL VIA AN EXECUTABLE'S METADATA

ping backdoor in ELF

Symbol table '.sym.p' contains 90 entries: Num: Value Size Type Bind Vis Ndx Name 0: 00000000060dff0 8 FUNC LOCAL DEFAULT UND Relocation section '.rela.p' at offset 0xf3a8 contains 14 entries: Type Sym. Value Sym. Name + Addend Offset Info 00000060dfe0 002d0000006 R_X86_64_GLOB_DAT 000000000000000 gmon_start_ + 0 00000060e9e0 004e00000005 R X86 64 COPY 00000000060e9e0 progname + 0 00000060e9f0 004b00000005 R_X86_64_COPY 00000000060e9f0 stdout + 0 00000060e9f8 005100000005 R X86 64 COPY 000000000060e9f8 __progname_full + 0 00000060ea00 005600000005 R X86 64 COPY 00000000060ea00 stderr + 0 00000060eb40 00000000005 R X86 64 COPY 00000060eb40 00000000001 R X86 64 64 0000000000000018 00000060eb40 000000000005 R_X86_64_COPY 00000060eb40 00000000001 R X86 64 64 000000000000018 00000060e218 0000000008 R X86_64 RELATIVE 0000000000401dc2

ping backdoor in Mach-O

>0016720: 7400 5f67 6574 7569 6400 5f73 6967 6e61 t._getuid._signa

>0016740: 005f 7373 6361 6e66 005f 6578 6563 6c70 ._sscanf._execlp >0016750: 0012 1212 125f 7374 7263 6872 005f 7374_strchr._st

CAN THESE BE TRUSTED?

```
bool is elf (char* pb, size t cb)
{
  if (cb < sizeof (HDR ELF32))
   return false;
  check byte sex (pb);
  HDR ELF32& header = *(HDR ELF32*) pb;
  if (memcmp (header.rgbID, "\177ELF", 4) != 0
        header.bitclass < 1
     || ( v (header.cbEntryProgram)
    && v (header.cbEntryProgram) != sizeof (PROGRAM ELF32))
     || v (header.cbEntrySection) != sizeof (SECTION_ELF32)
      v (header.iSectionNames) >= v (header.cEntrySection))
   return false;
 // *** FIXME: I don't recall why we need more than a header test.
(goes on to check section/program headers)
(from bsign)
```

PARSER DIFFERENTIALS

- PKCS 7 crytographic message
- ELF, and its multiple interpretations
 - Sections v. segments
- Multiple ways to locate a section
 - Is this the signature you are looking for?

(Kernel patch published by David Howells on 02 Dec, 2011)

DATA COMPLETENESS & CORRECTNESS

• How much of the file is signed?

```
/* only look at interesting sections */
if( !sname || s->shdr->sh_type == SHT_NULL ) {
    //|| s->shdr->sh_type == SHT_NOBITS ) {
    ES_PRINT("skipping null section\n");
    continue;
}
/* skip over the .pgptab and .pgpsig sections */
if( !strcmp( sname, ".pgptab" )
    || !strcmp( sname, ".pgpsig" ) ) {
    ES_PRINT(skipping internal section\n")
    continue;
}
(from elfgpg)
```

DATA COMPLETENESS & CORRECTNESS

We cannot sign the signatures, but they are loaded

```
// Include the ELF header, but with the number of sections set minus one,
// under the assumption that any binary having its checksum
// calculated will already have a signature header added to it.
// Yes, I can hear you screaming now. This makes my life easier. :P
//
// Note that elfsign, the tool, always creates the signature section before
// calculating the checksum.
elfHeader = melf_elfGetRaw(melf);
```

```
numSections = melf_elfGetSectionHeaderCount(melf);
sectionTableOffset = melf elfGetSectionHeaderOffset(melf);
```

```
melf_elfSetSectionHeaderCount(melf, numSections - 1);
melf elfSetSectionHeaderOffset(melf, 0);
```

```
(from elfsign)
```

WHAT ABOUT MACH-O CODE SIGNING?

:06:58 PM kernel: CODE SIGNING: cs_invalid_page(0x1000): p=1994[GoogleSoftwareUp] clearing CS_VA :05:41 PM kernel: CODE SIGNING: cs_invalid_page(0x1000): p=2034[GoogleSoftwareUp] clearing CS_VA :04:24 PM kernel: CODE SIGNING: cs_invalid_page(0x1000): p=2088[GoogleSoftwareUp] clearing CS_VA





Saturday, December 28, 13

MACH-O CODE SIGNING DATA



signerInfos SignerInfos }

"You can't trust code that you did not totally create yourself" corollary: You can't trust code that you did NOT TOTALLY **LOAD** YOURSELF

PARSER DIFFERENTIALS



CHAIN OF TRUST = CHAIN OF PARSERS

- Parser differentials break chains of trust
- Two views of the same data => confusion
 - Android Master Key, http://saurik.com/id/{17,18,19}
 - Package structure seen differently by signature verifier & installer (Java vs C++)
 - Mach-O signed loading
 - Linux kernel module signing
 - Even Linux ld.so vs. kernel!

CASE STUDY: ELF ABI CHAIN

- How many **ELF parsers** in ABI toolchain?
 - Do they all see the same view of sections/ segments?
- Kernel's **binfmt_elf loader** vs userland **ld.so**

DOUBLE THE PARSING, DOUBLE THE FUN



PROGRAM HEADERS

- PT_LOAD maps bytes to address range
- Implementation: mmap()
- Evad3rs noted that (fixed) mmap replaces existing mappings
- **Order** of PT_LOAD is important, but that's not in the spec.
- Automaton, not just data

DOES THIS LOOK LIKE A PARSER TO YOU?

```
/*From binfmt elf.c, Linux 3.4, GPLv2*/
elf phdata = kmalloc(size, GFP KERNEL);
kernel read(bprm->file, loc->elf ex.e phoff, (char
*)elf phdata, size)
for (phdrs...) {
/*..., scans for PT INTERP and PT LOAD*/
if(phdr->p type == PT LOAD)
if (!load addr set) {
  load addr = (elf ppnt->p vaddr - elf ppnt-
>p offset);
  load addr set = 1;
NEW AUX ENT(AT PHDR, load addr + exec->e phoff);
                  /*^^^/ FAIL*/
```

LD.SO PHDR



THE BIRTH OF AN ELF

- 1. Kernel reads PHDR table into buffer
- 2. mmap PT_LOADS
- 3. Loads PT_INTERP (ld.so)
- 4. writes (addr of first PT_LOAD)+
 hdr.phoff to loader ("AUX vector")
- 5. ld.so looks at aux vector and processes PHDR table

PHDR BUG

- Works if program headers in **first segment**.
- Otherwise, points to some other memory
- We can "finger-paint" memory to our liking, and **ld.so** will use **different** PHDRS
- Hide your PHDRs, the reverse engineers are coming!

CRAFTED FILE



DEMO: .SO BACKDOOR

- Fun Fact: you can execute .so files (try / lib/libc.so.6)
- Kernel parser for exec(), different parser in ld.so for .so
- POC loads a different library (libevil.so) when loaded by kernel

REWRITING PROGRAMS IN MEMORY WITH LD.SO?

- We control **ld.so**'s idea of all relevant sections: **GOT**, dyn symbols, ...
- **Id.so** resolves (what it thinks are) symbols, writes (what it thinks is) **GOT**
- Now we can rewrite a loading program via only crafted .dynamic + library symbols

"TRUSTING TRUST" IN BABEL



- Trusting computers is not only about bugs! Bugs are part of a problem, but not by far all of it
- Complex data formats >> bugs
- There is no "chain of trust" in Babel!

SOLUTIONS FOR BABEL

- Squeeze complexity out of data until it stops being "code equivalent"
 - UEFI? Software package formats?
- Hobble unexpected computation by blocking implicit flows
 - (see our ELFbac TR http://elfbac.org/)
- Use new hardware security primitives to isolate parsers

THANK YOU

IEEE SPW 2014 LangSec workshop, May 18, 2014

Collocated with IEEE Secuity & Privacy Symposium 2014

http://spw14.langsec.org/