

# Chapter 13

## ‘Weird Machine’ Patterns

Sergey Bratus, Julian Bangert, Alexandar Gabrovsky, Anna Shubina,  
Michael Locasto and Daniel Bilar

1 **Abstract** You do not understand how your program *really* works until it has been  
2 exploited. We believe that computer scientists and software engineers should regard  
3 the activity of modern exploitation as an applied discipline that studies both the actual  
4 computational properties and the practical computational limits of a target platform  
5 or system. Exploit developers study the computational properties of software that are  
6 not studied elsewhere, and they apply unique engineering techniques to the challeng-  
7 ing engineering problem of dynamically patching and controlling a running system.  
8 These techniques leverage software and hardware composition mechanisms in unex-  
9 pected ways to achieve such control. Although unexpected, such composition is not  
10 arbitrary, and it forms the basis of a coherent engineering workflow. This chapter  
11 contains a top-level overview of these approaches and their historical development.

### 12 1 Introduction

13 When academic researchers study exploitation, they mostly concentrate on two  
14 classes of attack-related artifacts: “malicious code” (malware, worms, shellcode)  
15 and, lately, “malicious computation” (exploits via crafted data payloads containing  
16 no native code, a popular exploiter technique. These techniques have been discussed

---

S. Bratus (✉) · J. Bangert · A. Gabrovsky · A. Shubina  
Dartmouth College, Hanover, New Hampshire, USA  
e-mail: sergey@cs.dartmouth.edu  
<http://www.cs.dartmouth.edu/~sergey/>

M. E. Locasto  
University of Calgary, Alberta, Canada  
<http://pages.cpsc.ucalgary.ca/locasto/>

D. Bilar  
Siege Technologies, Manchester, New Hampshire, USA  
<http://www.siegetechnologies.com/>

17 in the hacker community since before early 2000s ([1] offers a brief history sketch).  
18 They were brought a decade later to the attention of academia by Shacham [2, 3] in  
19 2007. These artifacts are tackled with a variety of approaches, from machine learning  
20 on either captured payloads or execution traces to automatic construction of exploit  
21 payloads for specific targets. The goals of such studies also vary, from straightfor-  
22 ward detection of particular payloads in network traffic to finding and removing  
23 vulnerabilities or asserting that a vulnerability is not exploitable beyond a crash  
24 or denial-of-service. Considerable technical depth has been reached in all of these  
25 directions, yet we seem no closer to the ultimate goal of constructing trustworthy,  
26 non-exploitable software.

27 We argue that focusing on just these two classes of artifacts is a limiting factor  
28 in our understanding of exploitation (and therefore of how to prevent it). We believe  
29 that, as a means of making progress toward the goal of more fundamentally secure  
30 software, we must understand how exploitation relates to *composition*, which is  
31 fundamental to all modern software construction. Understanding the patterns of this  
32 relationship will expose new artifacts to study and indicate new technical directions.

33 We note that practical exploitation has long been about *composing* the attacker  
34 computation with the native computation of the target, allowing most of the target’s  
35 functions to proceed normally, without undue interference. We posit that such compo-  
36 sition is the source of the most powerful and productive concepts and methodologies  
37 that emerge from exploitation practice.

38 When researchers focus on attack artifacts alone, they frequently miss an impor-  
39 tant point of successful exploitation: the exploited system needs to remain available  
40 and reliably usable for the attacker.

41 In order to support this assertion and further discuss exploitation, we need to make  
42 an important terminological point. The word *hacking* is used to refer to all kinds of  
43 attacks on computer systems, including those that merely shut down systems or  
44 otherwise prevent access to them (essentially achieving nothing that could not be  
45 achieved by cutting a computer cable). Many activities labeled as “hacking” lack  
46 sophistication. In this chapter we focus on *exploitation* or, more precisely, *exploit*  
47 *programming*. We take *exploitation* and *exploit programming* to mean subverting the  
48 system to make it work for the attacker—that is, lend itself to being programmed by  
49 the attacker. Exploiters are less interested in causing BSODs, kernel panics, and plain  
50 network DOS attacks that merely result in a DoS on the target and cannot otherwise  
51 be leveraged and refined to take control over the system rather than disabling it.

52 Not surprisingly, preventing a disabling crash and subsequently “patching up”  
53 the target into a stable running state requires significantly more expertise and effort  
54 than, say, a memory-corrupting DoS. By achieving this state exploiters demonstrate  
55 a sophisticated understanding of the target platform,<sup>1</sup> often beyond the ken of its  
56 developers or even system programmers.

57 In this chapter we review a series of classic exploitation techniques from the per-  
spective of composition. Many of these techniques have been extensively described

---

<sup>1</sup> It also serves as an excellent teaching aid in advanced OS courses; see, e.g., [4].

58 and reviewed from other perspectives; however, their compositional aspect is still  
59 treated as ad hoc, and has not, as far as we know, been the subject of systematic analy-  
60 sis. Specifically, we regard composition as the basic unit of activity in an engineering  
61 workflow, whether that workflow is a traditional software engineering workflow or  
62 a workflow focused on engineering an exploit. We compare these workflows in the  
63 Sect. 2.

64 Since our focus is on composition, we do not distinguish between techniques used  
65 by *rootkits* vs. *exploits*. Rootkits are traditionally separated from other exploit-related  
66 artifacts such as exploits proper, “shellcode”, etc., since they are meant to be installed  
67 by the successful attacker who already attained the “root” level of privilege by other  
68 means. However, we note that such installation often involves means of composition  
69 that are only available to developers, not administrators however privileged; thus,  
70 composing parts of a rootkit with the system poses challenges due to lacking infor-  
71 mation and limited available context. The complexity of such challenges may vary,  
72 but they have the same nature as those faced by an exploit programmer, and indeed  
73 similar techniques are used to overcome them. In our discussion, we draw equally  
74 upon rootkit and exploit examples.

75 We posit that composition-centric analysis is required for designing defensible  
76 systems (see Sect. 4). The practical properties of composition in actual computer sys-  
77 tems uncovered and distilled by hacker research have often surprised both designers  
78 and defenders. We believe that the relevant methods here must be cataloged and  
79 generalized to help approach the goal of *secure composition* in future designs.

## 80 2 A Tale of Two Engineering Workflows

81 “Language design is library design.”

82 – B. Stroustrup

83 Hacking, vulnerability analysis, and exploit programming are generally perceived  
84 to be difficult and arcane activities. The development of exploits is still seen as  
85 something unrepeatable and enabled only by some unfortunate and unlikely com-  
86 bination of events or conditions. Almost by definition, something as imbued with  
87 arbitrary chance cannot or should not be an engineering discipline or workflow. Pop-  
88 ular perception casts these activities as requiring specialized cross-layer knowledge  
89 of systems and a talent for “crafting” input.

90 This chapter asserts that what seems arcane is really only unfamiliar. In fact,  
91 although it may be difficult to conceive of exploit development as anything other  
92 than fortunate mysticism, we argue that its structure is exactly that of a software  
93 engineering workflow. The difference emerges in the specific constructs at each  
94 stage, but the overall activities remain the same. A software developer engineers in  
95 terms of sequences of function calls operating on abstract data types, whereas an  
96 exploit developer engineers in terms of sequences of machine-level memory reads  
97 and writes. The first one programs the system in terms of what its compile-time API

98 promises; the other programs it in terms of what its runtime environment actually  
99 contains.

100 This section contains a brief comparison of these two engineering workflows. We  
101 do so to help give a conceptual frame of reference to the enumeration of exploit  
102 techniques and composition patterns detailed in Sect. 3.

103 The main difference between the two workflows is that the exploit engineer must  
104 first recover or understand the semantics of the runtime environment. In either case,  
105 programming is composition of functionality.

106 In the “normal” workflow of software engineering, the programmer composes  
107 familiar, widely-used libraries, primitive language statements (repetition and deci-  
108 sion control structures), and function calls to kick input data along a processing path  
109 and eventually produce the result dictated by a set of functional requirements.

110 In the exploit workflow, the reverser or exploit engineer attempts to build this  
111 programming toolkit from scratch: the languages and libraries that the software  
112 engineer takes for granted are not of *direct* use to the exploit developer. Instead,  
113 these elements define a landscape from which the exploit developer must compose  
114 and create his own toolkit, language primitives, and component groups. The first  
115 job of the vulnerability analyst or reverse engineer is therefore to understand the  
116 latent functionality existing in runtime environments that the software engineer either  
117 neglects or does not understand.

## 118 **2.1 The Software Engineer**

119 Based on functional requirements, a software engineer’s goal is to cause some  
120 expected functionality happen. In essence, this kind of programming is the task  
121 of choosing a sequence of library calls and composing them with language primi-  
122 tives like decision control structure and looping control structures. Data structures  
123 are created to capture the relevant properties of the system’s input; this structure  
124 usually dictates how processing (i.e., control flow) occurs.

125 A software engineer follows roughly this workflow path:

- 126 1. design and specify data types
- 127 2. design data flow relationships (i.e., an API)
- 128 3. write down source code implementing the data types and API
- 129 4. ask compiler and assembler to translate code
- 130 5. ask OS to load binary, invoke the dynamic linker, and create memory regions
- 131 6. run program according to the control flow as conceived in the source level.

132 In this workflow, we can see the software engineer engaged in: memory layout,  
133 specifying control flow, program construction, program delivery (loading) and trans-  
134 lation, and program execution. As we will see below, the exploit engineer engages  
135 in much the same set of tasks.

136 The software engineer’s goal is to bring order to a composition of procedures  
137 via compilation and assembly of machine code. One does this through tool chains,

138 design patterns, IDEs, and popular languages—the software engineer therefore does  
139 not need to relearn the (public) semantics of these operations every time he prepares  
140 to program.

141 These conventions are purely an effort–saving device aimed at increasing pro-  
142 ductivity by increasing the lines of code and features implemented in them. These  
143 patterns, tools, and aids reduce the level of thought required to emit a sequence of  
144 function calls that satisfy the functional requirements. They are an effort to deal with  
145 complexity. The goal of software engineers in dealing with complexity is to eliminate  
146 or hide it.

## 147 **2.2 The Exploit Engineer**

148 In contrast, exploit engineers also deal with complexity, but their goal is to manipulate  
149 it—expressiveness, side effects, and implicit functionality are a collective boon, not  
150 a bane. Any operations an exploit engineer can get “for free” increase his exploit  
151 toolkit, language, or architecture. A software engineer attempts to hide or ignore  
152 side effects and implicit state changes, but the very things encouraged by traditional  
153 engineering techniques like “information hiding” and encapsulation on the other side  
154 of an API become recoverable primitives for a reverser or exploit engineer.

155 The main difference in the workflows is the preliminary step: you have to learn on  
156 a case by case or scenario by scenario basis what “language” or computational model  
157 you should be speaking in order to actually begin programming toward a specific  
158 functional end. Based on some initial access, the first goal is to understand the system  
159 enough to recover structure of “programming” primitives. The workflow is thus:

- 160 1. identify system input points
- 161 2. recapture or expose trust relationships between components (functions, control  
162 flow points, modules, subroutines, etc.)
- 163 3. recover the sequencing composition of data transformations (enumerate layer  
164 crossings)
- 165 4. enumerate instruction sequences / primitives / gadgets
- 166 5. program the process address space (prepare the memory image and structure)
- 167 6. deliver the exploit.

168 In this workflow, we can see the exploit engineer engaged in: recovering memory  
169 layout, specifying control flow, program construction, program delivery (loading)  
170 and translation, and program execution. We note that these steps may not (and need  
171 not be) sperabile: Unlike the software engineering workflow, the delivery of an  
172 exploit (i.e., loading a program) can be mixed up and interposed with translation of  
173 the program and preparation of the target memory space. Even though these activities  
174 might be more tightly coupled for an exploit developer, much of the same discipline  
175 remains.

176 Recent academic advances have the potential of automating (at least partially) the  
177 preparatory steps (1–4) the exploiter’s workflow. Holler’s *LangFuzz* tool automates



178 black-box fuzz testing of context-free grammar engines. It generates test cases from  
179 a given context-free grammar to exposes via fault generation inter-components' trust  
180 relations [5]. Caballero proposed and implemented the *Dispatcher* tool for automatic  
181 protocol reverse-engineering given an undocumented protocol or le format. Thus  
182 includes the structure of all messages that comprise the protocol in addition to the  
183 protocol state machine, which captures the sequences of messages that represent  
184 valid sessions of the protocol. As a proof of concept, his group managed to extract the  
185 grammar of Mega-D (a spam botnet), which sported an undocumented, encrypted  
186 Command & Control protocol [6]. The output of these tools can be repurposed  
187 for defenses. Samuels proposed a simple but clever approach against certain type  
188 confusion attacks through a generalizable annotated parse-tree-grammar scheme.  
189 Such annotated grammars can be converted to push-down automata from which  
190 input stress test can be derived [7]. The urgent need for such defenses is demonstrated  
191 by Shmatikov and Wang analysis of and attacks against AV parsers and undefined  
192 behavior in C language compilers, respectively [8, 9].

193 One major challenge exists for the exploit engineer: recovering the unknown  
194 unknowns. Although they can observe side effects of mainline execution or even  
195 slightly fuzzed execution, can they discover the side effects of “normally” dormant  
196 or latent “normal” functionality (e.g., an internationalization module that is never  
197 invoked during normal operation, or configuration code that has only been invoked  
198 in the “ancient past” of this running system)? This challenge is in some sense like  
199 the challenge a software engineer faces when exploring a very large language library  
200 (e.g., the Java class library API).

### 201 **2.3 A Simple Example of Exploit Programming**

202 Before we turn our attention to reviewing the composition patterns of hacking, we  
203 give a brief example of constructing an exploit as a programming task, to show the  
204 kind of workflow involved. The reader already familiar with such concept may skip  
205 directly to Sect. 3.

206 Assume a Program P that reads the contents of a file F into a buffer located on the  
207 stack and displays this content to standard output.

208 Note that our point here is not to say that stack-based buffer overflows are of inde-  
209 pendent or modern interest; rather, we use this scenario as the simplest illustration  
210 of exploitation as programming, where most readers likely already have a concep-  
211 tion of some of the issues in play, from both the software developer side (why this  
212 vulnerability exists: failure to check lengths) and the exploit engineer side (where to  
213 inject and how to structure shellcode)—it is popularly “understood” enough.

214 The simplest possible core task in the exploit engineer’s workflow in this scenario  
215 is to map input data to memory space. Concretely: where will the bytes that overwrite  
216 the return address and saved ebp land? And what should the content of my new return  
217 address be (i.e., the address of the file’s content on the stack)?

218 In this scenario, file F is simultaneously:

- 219 ● bytes in persistent storage
- 220 ● bytes traversing kernel space via the read (2) implementation (and callchain)
- 221 ● data for the program (i.e., contents of a buffer)
- 222 ● an overlay of some part of the process address space
- 223 ● bytecode for an automaton implicitly embedded in the target program (the so-
- 224 called “weird machine”, see Sect. 3) that will carry out malicious computation
- 225 ● shellcode for the CPU.

226 Understanding the layout of this memory is vital to actually constructing an input  
 227 file (i.e., bytecode) to program the exploit machine. In addition to understanding the  
 228 memory layout as a living artifact generated by both the compiler and the runtime  
 229 system (i.e., how the OS sets up the process address space and memory regions),  
 230 an exploit engineer must also understand other implicit operators, transformers, and  
 231 parsers co-existing in the *real* computational machine. He must understand where  
 232 they interpose on the processing path, how they are invoked, and what *actual* set of  
 233 transformations they have on the data (i.e., bytecode) as it journeys through the system  
 234 toward various resting places. He must ask: is my data filtered? Does it have a simple  
 235 transformation or encoding applied to it (e.g., Base64 encoding, `toupper()`)?  
 236 Do certain special bytes (i.e., NULL) truncate the data? Is the data copied to new  
 237 locations? Is the data reshuffled or reorganized internally?

238 When considering more complex examples, an exploit engineer must map and  
 239 understand other memory regions, including the heap and its management data struc-  
 240 tures, dynamic data embedded within various memory regions, and threading (con-  
 241 currency). Such things now present additional computational primitives to place in  
 242 context and understand. In some sense, the sum total composition of all these mech-  
 243 anisms is a compiler, translator, or interpreter for your bytecode—and you must first  
 244 understand how that compiler works. When you do, you can undertake the process of  
 245 writing and constructing bytecode appropriate to accomplishing your goal, whether  
 246 that is to drop shell, open a port, install a rootkit, exfiltrate a file, etc.

## 247 3 Patterns

### 248 3.1 Exploitation as Programming “Weird Machines”

249 Bratus et al. [1] summarized a long-standing hacker intuition of exploits as *programs*,  
 250 *expressed as crafted inputs, for execution environments implicitly present in the*  
 251 *target as a result of bugs or unforeseen combination of features* (“weird machines”),  
 252 which are reliably driven by the crafted inputs to perform unexpected computations.  
 253 More formally, the crafted inputs that constitute the exploit drive an input-accepting  
 254 automaton already implicitly present in the target’s input-handling implementation,  
 255 its sets of states and transitions owing to the target’s features, bugs or combinations  
 256 thereof.



257 The implicit automaton is immersed into or is part of the target’s execution envi-  
258 ronment; its processing of crafted input is part of the “malicious computation” —  
259 typically, the part that creates the initial compromise, after which the exploiter can  
260 program the target with more conventional means. The crafted input is both a pro-  
261 gram for that automaton and a constructive proof of its existence. Further discussion  
262 from the practical exploit programming standpoint can be found in Dullien [10],  
263 from a theory standpoint in Sassaman [11].

264 This perspective on exploit programming considers the exploit target as harboring  
265 a virtual computing architecture, to which the input data serve as *bytecode*, similar  
266 to, say, how compiled Java bytecode drives the Java virtual machine. In other words,  
267 the target’s input is viewed as an actual program, similar to how the contents of  
268 a Turing machine’s tape can be considered a program. Thus what is liable to be  
269 seen by developers as “inert data” such as inputs or metadata is in fact conceptually  
270 promoted to a vehicle of programming the target; in a sense, the exploiter treats the  
271 data as running and acting on the target program, not the other way around. Further  
272 discussion of this can be found in Shapiro [12].

273 In the following items, we focus on one critical aspect of the implicit exploit  
274 execution environments and the computations effected in them by exploit-programs:  
275 they must reliably co-exist with the native, intended computations both for their dura-  
276 tion and in their effects, while their composition is done in contexts more limited and  
277 lacking critical information as compared to the system’s intended scenarios. This is  
278 far from trivial on systems where state that is “borrowed” by the exploit computa-  
279 tion’s thread of control is simultaneously used by others. It involves dissecting and  
280 “slimming down” interfaces to their actual implementation primitives and finding  
281 out unintended yet stable properties of these primitives.

## 282 3.2 *Recovering Context, Symbols, and Structure*

283 To compose its computation with a target, an exploit must refer to the objects it  
284 requires in its virtual address space (or in other namespaces). In essence, except in  
285 the most trivial cases, a “name service” of a kind (ranging from ad-hoc to the system’s  
286 own) is involved to reconstruct the missing information.

287 Early exploits and rootkit install scripts relied on hard-coded fixed addresses of  
288 objects they targeted, since back then memory virtual space layouts were identical for  
289 large classes of targets.<sup>2</sup> As targets’ diversity increased, naturally or artificially (e.g.,  
290 OpenWall, PaX, other ASLR), exploits progressed to elaborate address space lay-  
291 out reconstruction schemes and co-opting the system’s own dynamic linking and/or  
292 trapping debugging.

---

<sup>2</sup> This fact was not well understood by most engineers or academics, who regarded below-compiler OS levels as unpredictable; Stephanie Forrest deserves credit for putting this and other misconceptions into broader scientific perspective.



293 Cesare [13] describes the basic mechanism behind ELF linking—based on little  
294 more that careful reading of the ELF standard. However, it broke the opacity and  
295 resulted in an effective exploit technique, developed by others, e.g., [14]. In [15]  
296 mayhem builds on the same idea by looking into the significance and priority of  
297 ELF’s `.dynamic` symbols. Nergal [16] co-opted Linux’s own dynamic linker into  
298 an ROP<sup>3</sup> crafted stack frame-chaining scheme, to have necessary symbols resolved  
299 and libraries loaded. Oakley [17] showed how to co-opt the DWARF-based exception  
300 handling mechanism.

301 Skape [18] takes the understanding of ELF in a different direction by showing how  
302 its relocation mechanism works and how that could be used for unpacking obfuscated  
303 Windows binaries. Recent work by Shapiro [12] demonstrated that the relocation  
304 metadata in ELF binaries is actually enough to drive Turing-complete computations  
305 on the Linux dynamic linker-loader. The ABI metadata in these examples serves as  
306 “weird machine” bytecode for the Turing machine implicitly embedded in the RTLD  
307 code.

308 In all of the above detailed understanding of a mechanism comes before the insight  
309 of how an exploit could be built; in fact, once the mechanism is clear at the “weird  
310 machine” level, its exploitation use is almost an afterthought.

### 311 3.3 Preparing Vulnerable System State

312 Earlier classes of exploits leveraged conditions and configurations (such as memory  
313 allocation of relevant objects) present in the target’s state through all or most runs.  
314 Subsequent advancements such as Sotirov’s [19] demonstrated that *otherwise non-*  
315 *exploitable targets can have their state carefully prepared by way of a calculated*  
316 *sequence of requests and inputs for an exploitable configuration to be instantiated.*

317 This pattern of pre-compositional state-construction of targets is becoming essen-  
318 tial, as protective entropy-injecting techniques prevent setting up an effective “name  
319 service”. Recent examples [20, 21] show its applications to modern heaps (the former  
320 for the Windows low fragmentation heap), in presence of ASLR and DEP. Moreover,  
321 this method can target the injected entropy *directly*, by bleeding it from the target’s  
322 state (e.g., [22]).

### 323 3.4 Piercing Abstraction

324 Developers make use of abstractions to decrease implementation effort and increase  
325 code maintainability. However, abstractions hide the details of their implementation  
and as they become part of a programmers daily vocabulary, the implementation

---

<sup>3</sup> Which it pre-dates, together with other hacker descriptions of the technique, by five to seven years.

326 details are mostly forgotten. For example, few programmers worry about how a  
327 function call is implemented at the machine level or how the linking and loading  
328 mechanisms assign addresses to imported symbols.

329 Exploit engineers, however, distill abstractions into their implementation primi-  
330 tives and synthesize new composition patterns from them. Good examples of this are  
331 found in [16], who modifies the return addresses on the stack to compose existing  
332 code elements into an exploit, and the LOCREATE [18] packer which obfuscates  
333 binary code by using the primitives for dynamic linking.

### 334 **3.5 Balancing Context Constraints**

335 Wherever there is modularity there is the potential for misunderstanding: Hiding information  
336 implies a need to check communication.

337 A. Perlis

338 When a software architect considers how much context to pass through an inter-  
339 face, he has to balance competing constraints (see Sect. 4.2 in [23] for discussion  
340 of etiology and formalization sketch). Either a lot of context is passed, reducing  
341 the flexibility of the code, or too little context is preserved and the remaining data  
342 can no longer be efficiently validated by code operating on it, so more assumptions  
343 about the input have to be trusted. Exploiters explore this gap in assumptions, and  
344 distill the unintended side-effects to obtain *primitives*, from which weird machines  
345 are constructed [10, 24, 25]. We posit that understanding this gap is the way to more  
346 secure API design.

### 347 **3.6 Bit Path Tracing of Cross-Layer Flows**

348 When an exploiter studies a system, he starts with bit-level description of its contents  
349 and communications. Academic textbooks and user handbooks, however, typically  
350 do not descend to bit level and provide only a high-level description of how the  
351 system works. A crucial part of such bit-level description is the flow of bits between  
352 the conceptual design layers of the system: i.e. a binary representation of the data  
353 and control flow between layers.

354 Constructing these descriptions may be called the cornerstone of the hacker  
355 methodology. It precedes the search for actual vulnerabilities and may be thought  
356 of as the modeling step for constructing the exploit computation. The model may  
357 ignore large parts of the target platform but is likely to punctiliously describe the  
358 minutiae of composition mechanisms that actually tie the implementations of the  
359 layers together.

360 For example, the AlephOne Phrack article [26] famous for its description of stack  
361 buffer overflows also contained a bit-level description of UNIX system calls, which  
362 for many readers was in fact their first introduction to syscall mechanisms. Similarly,

363 other shellcode tutorials detailed the data flow mechanisms of the target’s ABIs (such  
364 as various calling conventions and the structure of libraries). In networking, particular  
365 attention was given to wrapping and unwrapping of packet payloads at each level  
366 of the OSI stack model, and libraries such as libnet and libdnet were provided for  
367 emulating the respective functionality throughout the stack layers.

368 What unites the above examples is that in all of them exploiters start analyzing the  
369 system by tracing the flow of bits within the target and enumerating the code units  
370 that implement or interact with that flow. The immediate benefits of this analysis  
371 are at least two-fold: locating of less known private or hidden APIs and collecting  
372 potential exploitation primitives or “cogs” of “weird machines”, i.e. code fragments  
373 on which crafted data bits act in predictable way.

374 Regardless of its immediate benefits, though, bit-level cross-layer flow descrip-  
375 tions also provide useful structural descriptions of the system’s architecture, or, more  
376 precisely, of the mechanisms that underly the structure, such as the library and load-  
377 able kernel functionality, DDKs, and network stack composition.

378 For instance, the following sequence of Phrack articles on Linux rootkits is a great  
379 example of deep yet concise coverage of the layers in the Linux kernel architecture:  
380 *Sub proc\_root Quando Sumus (Advances in Kernel Hacking)* [27] (VFS structures  
381 and their linking and hijacking), *5 Short Stories about execve (Advances in Kernel*  
382 *Hacking II)* [28] (driver/DDK interfaces, different binary format support), and *Exe-*  
383 *cution path analysis: finding kernel based rootkits* [29] (instrumentation for path  
384 tracing). Notably, these articles at the cusp where three major UNIX innovations  
385 meet: VFS, kernel state reporting through pseudo-file systems (e.g., `/proc`), and  
386 support for different execution domains/ABI. These articles described the control  
387 and data flows through a UNIX kernel’s component layers and their interfaces in  
388 great detail well before tools like DTrace and KProbes/SystemTap brought tracing  
389 of such flows within common reach.

390 It is worth noting that the ELF structure of the kernel binary image, the corre-  
391 sponding structure of the kernel runtime, and their uses for reliably injecting code  
392 into a running kernel (via writing `/dev/kmem` or via some kernel memory corrup-  
393 tion primitive). In 1998, the influential *Runtime kernel kmem patching* [30] made the  
394 point that even though a kernel may be compiled without loadable kernel module  
395 support, it still is a structured runtime derived from an ELF image file, in which sym-  
396 bols can be easily recovered, and the linking functionality can be provided without  
397 difficulty by a minimal userland “linker” as long as it has access to kernel memory.  
398 Subsequently, mature kernel function hooking frameworks were developed (e.g.,  
399 *IA32 Advanced function hooking* [31]).

400 Dynamic linking and loading of libraries (shared binary objects) provide another  
401 example. This is a prime example of composition, implicitly relied upon by every  
402 modern OS programmer and user, with several supporting engineering mechanisms  
403 and abstractions (ABI, dynamic symbols, calling conventions). Yet, few resources  
404 exist that describe this key mechanism of interposing computation; in fact, for a  
405 long time hacker publications have been the best resource for understanding the  
406 underlying binary data structures (e.g., *Backdooring binary objects* [32]), the control  
407 flow of dynamic linking (e.g., *Cheating the ELF* [33] and *Understanding Linux ELF*

408 *RTLD internals* [34]), and the use of these structures for either binary infection (e.g.,  
409 the original *Unix ELF parasites and virus*) or protection (e.g., *Armouring the ELF:*  
410 *Binary encryption on the UNIX platform* [35]).

411 A similar corpus of articles describing the bit paths and layer interfaces exists for  
412 the network stacks. For the Linux kernel stack, the *Netfilter* architecture represents  
413 a culmination of this analysis. By exposing and focusing on specific hooks (tables,  
414 chains), Netfilter presents a clear and concise model of a packet's path through the  
415 kernel; due to this clarity it became both the basis of the Linux's firewall and a long  
416 series of security tools.

417 Not surprisingly, exploitative modifications of network stacks follow the same  
418 pattern as other systems rootkits. *Passive Covert Channels Implementation in Linux*  
419 *Kernel* [36] is a perfect example: it starts with describing the interfaces traversed  
420 on a packet's path through the kernel (following the Netfilter architecture), and then  
421 points out the places where a custom protocol handler can be inserted into that control  
422 flow, using the stack's native protocol handler interfaces.

### 423 3.7 Trap-Based Programming and Composition

424 In application programming, traps and exceptions are typically not treated as “first-  
425 class” programming primitives. Despite using powerful exception-handling subsystems  
426 (such as GCC's *DWARF*-based one, which employs Turing-complete bytecode),  
427 applications are not expected to perform much of their computation in traps or excep-  
428 tions and secondary to the main program flow. Although traps are obviously crucial  
429 to systems programming, even there the system is expected to exit their handlers  
430 quickly, performing as little and as simple computation as possible, for both perfor-  
431 mance and context management reasons.

432 In exploit programming and reverse engineering (RE), traps are the *first-class*  
433 *programming primitives*, and trap handler overloading is a frequently used techni-  
434 que. The target platform's trap interfaces, data structures, and contexts are carefully  
435 studied, described, and modeled, then used for reliably composing an exploit or a  
436 comprehension computation (i.e., a specialized tracer of debugger) with the target.

437 The tracing and debugging subsystems in OS kernels have long been the focus  
438 of hacker attention (e.g., *Runtime Process Infection* [37] for an in-depth intro to  
439 the `ptrace()` subsystem). Not surprisingly, hackers are the leading purveyors of  
440 specialized debuggers, such as *dumBug*, *Rasta Debugger*, and the *Immunity debugger*  
441 to name a few.

442 For Linux, a good example is *Handling Interrupt Descriptor Table for fun and*  
443 *profit* [38], which serves as both a concise introduction to the x86 interrupt system and  
444 its use on several composition-critical kernel paths, as well as its role in implementing  
445 various OS and debugging abstractions (including system calls and their place in  
446 the IDT). This approach was followed by a systematic study of particular interrupt  
447 handlers, such as the *Hijacking Linux Page Fault Handler* [39].

448 Overloading the page fault handler in particular has become a popular mechanism  
449 for enforcing policy in kernel hardening patches (e.g., *PaX*<sup>4</sup> and *OpenWall*<sup>5</sup>).  
450 However, other handlers have been overloaded as well, providing, e.g., support for  
451 enhanced debugging not relying on the kernel’s standard facilities—and thus not con-  
452 flicting with them and not registering with them, to counteract anti-debugging tricks.  
453 Since both rootkits (e.g., the proof-of-concept *DR Rootkit* that uses the x86 debug  
454 registers exclusively as its control flow mechanism) and anti-RE armored applica-  
455 tions (e.g., Skype, cf. *Vanilla Skype* [40]; also, some commercial DRM products).  
456 In particular, the *Rasta Debugger* demonstrates such “unorthodox debugging” trap  
457 overloading-based techniques.

458 Notably, similar trap overloading techniques are used to expand the semantics of  
459 classic debugger breakpoint-able events. For instance, *OllyBone*<sup>6</sup> manipulated page  
460 translation to catch an instruction fetch from a page just written to, a typical behavior  
461 of a malware unpacker handing execution to the unpacked code. Note the temporal  
462 semantics of this composed trap, which was at the time beyond the capabilities of any  
463 debugger. A similar use of the x86 facilities, and in particular the split instruction and  
464 data TLBs was used by the *Shadow Walker* [41] rootkit to cause code segments loaded  
465 by an antivirus analyzer to be fetched from a different physical page than the actual  
466 code, so that the analyzer could receive innocent data—a clever demonstration of the  
467 actual vs assumed nature of x86 memory translation mechanism. For an in-depth  
468 exploration of just how powerful that mechanism can be, as well as for background  
469 on previous work, see Bangert [42].

## 470 4 Conclusion

471 Exploit engineers will show you the unintended limits of your system’s functionality.  
472 If software engineers want to reduce this kind of latent functionality, they will have  
473 to begin understanding it as an artifact that supports the exploit engineer’s workflow.

474 Software engineers should view their input data as “acting on code”, not the other  
475 way around; indeed, in exploits inputs serves as a de-facto bytecode for execution  
476 environments that can be composed from the elements of their assumed runtime  
477 environment. Writing an exploit—creating such bytecode—is as structured a disci-  
478 pline as engineering “normal” software systems. As a process, it is no more arcane  
479 or unapproachable than the ways we currently use to write large software systems.

480 Yet, a significant challenge remains. If, as hinted above, we want to have a prac-  
481 tical impact on the challenge of secure composition, can we actually train software  
482 engineers to see their input parameters and data formats *as bytecode* even as they  
specify it? Even as they bring it into existence, where it is by definition partially

---

4 <http://pax.grsecurity.net/>

5 <http://www.openwall.com/Owl/>

6 <http://www.joestewart.org/ollybone/>

483 formulated, can they anticipate how it might be misused? We posit that this constant  
 484 and frequent self-check is worth the effort: Software engineers should familiarize  
 485 themselves with anti-security patterns lest preventable ‘weird machines’ arise in  
 486 critical applications.

## 487 References

- 488 1. Bratus S, Locasto ME, Patterson ML, Sassaman L, Shubina A. Exploit programming: from  
 489 buffer overflows to “weird machines” and theory of computation. *login*: Dec 2011.
- 490 2. Shacham H. The geometry of innocent flesh on the bone: return-into-libc without function calls  
 491 (on the  $\times 86$ ). In: Proceedings of the 14th ACM conference on computer and communications  
 492 security, CCS '07. New York: ACM; p. 552–561.
- 493 3. Roemer R, Buchanan E, Shacham H, Savage S. Return-oriented programming: systems, lan-  
 494 guages, and applications. *ACM Trans Inf Syst Secur.* 2012;15(1):2:1–2:34.
- 495 4. Dan R. Anatomy of a remote kernel exploit. [http://www.cs.dartmouth.edu/~sergey/cs108/2012/  
 496 Dan-Rosenberg-lecture.pdf](http://www.cs.dartmouth.edu/~sergey/cs108/2012/Dan-Rosenberg-lecture.pdf) (2011).
- 497 5. Holler C, Herzig K, Zeller A. Fuzzing with code fragments. In: Proceedings of the 21st USENIX  
 498 conference on security symposium, Security'12. Berkeley: USENIX Association; 2012. p. 38–  
 499 38.
- 500 6. Caballero Juan, Song Dawn. Automatic protocol reverse-engineering: message format extrac-  
 501 tion and field semantics inference. *Comput Netw.* 2013;57(2):451–74.
- 502 7. Samuel M, Erlingsson Ú. Let's parse to prevent pwnage invited position paper. In: Proceed-  
 503 ings of the 5th USENIX conference on Large-scale exploits and emergent threats, LEET'12,  
 504 Berkeley, USA: USENIX Association; 2012. p. 3–3.
- 505 8. Jana s, Shmatikov V. Abusing file processing in malware detectors for fun and profit. In: IEEE  
 506 symposium on security and privacy'12; 2012. p. 80–94.
- 507 9. Xi W, Haogang C, Alvin C, Zhihao J, Nickolai Z, Kaashoek MF. Undefined behavior: what  
 508 happened to my code? In: Proceedings of the Asia-Pacific workshop on systems, APSYS'12.  
 509 New York, USA: ACM; 2012. p. 9:1–9:7.
- 510 10. Dullien T. Exploitation and state machines: programming the “weird machine”, revisited. In:  
 511 Infiltrate conference, Apr 2011.
- 512 11. Sassaman L, Patterson ML, Bratus S, Locasto ME, Shubina A. Security applications of formal  
 513 language theory. Dartmouth College: Technical report; 2011.
- 514 12. Shapiro R, Bratus S, Smith SW. “Weird machines” in ELF: a Spotlight on the underappreci-  
 515 ated metadata. In: 7th USENIX workshop of offensive technologies. [https://www.usenix.org/  
 516 system/files/conference/woot13/woot13-shapiro.pdf](https://www.usenix.org/system/files/conference/woot13/woot13-shapiro.pdf). 2013
- 517 13. Cesare. S. Shared library call redirection via ELF PLT, Infection. Dec 2000.
- 518 14. Sd, Devik. Linux On-the-fly Kernel patching without LKM, Dec 2001.
- 519 15. Mayhem. Understanding Linux ELF RTLD internals. [http://s.eresi-project.org/inc/articles/elf-  
 520 rtd.txt](http://s.eresi-project.org/inc/articles/elf-<br/>
    520 rtd.txt) (2002).
- 521 16. Nergal. The advanced return-into-lib(c) Exploits: PaX Case Study. *Phrack Mag.* 2001;58(4).
- 522 17. Oakley J, Sergey B. Exploiting the hard-working dwarf: Trojan and exploit techniques with no  
 523 native executable code. In: WOOT. 2011. p. 91–102.
- 524 18. Skape. Lcreate: an anagram for relocate. *Uninformed.* 2007;6.
- 525 19. Sotirov A. Heap feng shui in javascript. In: *Blackhat*; 2007.
- 526 20. Redpantz. The art of exploitation: MS IIS 7.5 remote heap overflow. *Phrack Mag.* 68(12), Apr  
 527 2012.
- 528 21. Huku, Argp. The art of exploitation: exploiting VLC, a jemalloc case study. *Phrack Maga.*  
 529 2012;68(13).
- 530 22. Ferguson J. Advances in win32 aslr evasion, May 2011.



- 531 23. Bilar D. On callgraphs and generative mechanisms. *J Comput Virol*. 2007;3(4).  
532 24. Richarte D. About exploits writing. Core security technologies presentation 2002.  
533 25. Gera, Riq. Advances in format string exploitation. *Phrack Mag*. 2002;59(7).  
534 26. One A. Smashing the stack for fun and profit. *Phrack* 1996;49:14. <http://phrack.org/issues.html?issue=49&id=14>.  
535  
536 27. Palmers. Sub proc\_root auando sumus (Advances in Kernel hacking). *Phrack* 2001;58:6. <http://phrack.org/issues.html?issue=58&id=6>.  
537  
538 28. Palmers. 5 Short stories about execve (advances in Kernel hacking II). *Phrack* 2002;59:5. <http://phrack.org/issues.html?issue=59&id=5>.  
539  
540 29. Rutkowski JK. Execution path analysis: finding Kernel based rootkits. *Phrack* 2002;59:10. <http://phrack.org/issues.html?issue=59&id=10>.  
541  
542 30. Cesare S. Runtime Kernel kmem patching. 1998. <http://althing.cs.dartmouth.edu/local/vsc07.html>.  
543  
544 31. Mayhem. IA32 advanced function hooking. *Phrack* 2001;58:8. <http://phrack.org/issues.html?issue=58&id=8>.  
545  
546 32. Klog. Backdooring binary objects. *Phrack* 2000;56:9. <http://phrack.org/issues.html?issue=56&id=9>.  
547  
548 33. The Grugq. Cheating the ELF: subversive dynamic linking to libraries, 2000.  
549 34. Mayhem. Understanding Linux ELF RTLD Internals, 2002. <http://s.eresi-project.org/inc/articles/elf-rtld.txt>.  
550  
551 35. Grugq, Scut. Armouring the ELF: binary encryption on the UNIX platform. *Phrack* ; 2001;58:5. <http://phrack.org/issues.html?issue=58&id=5>.  
552  
553 36. Rutkowska J. Passive covert channels implementation in Linux Kernel. 21st chaos communications congress, 2004. <http://events.ccc.de/congress/2004/fahrplan/files/319-passive-covert-channels-slides.pdf>.  
554  
555  
556 37. (Anonymous author). Runtime process infection. *Phrack* 2002;59:8. <http://phrack.org/issues.html?issue=59&id=8>.  
557  
558 38. Kad. Handling interrupt descriptor table for fun and profit. *Phrack* 2002;59:4. <http://phrack.org/issues.html?issue=59&id=4>.  
559  
560 39. Buffer. Hijacking Linux page fault handler exception table. *Phrack* 2003;61:7. <http://phrack.org/issues.html?issue=61&id=7>.  
561  
562 40. Desclaux F, Kortchinsky K. Skype V. REcon. <http://www.recon.cx/en/f/vskype-part1.pdf> (2006).  
563  
564 41. Sparks S, Butler J. "Shadow Walker": Raising the bar for rootkit detection. *BlackHat*; 2005. <http://www.blackhat.com/presentations/bh-jp-05/bh-jp-05-sparks-butler.pdf>.  
565  
566 42. Bangert J, Bratus S, Rebecca S, Sean WS. The page-fault weird machine: lessons in instruction-less computation. In: 7th USENIX workshop of offensive technologies. Aug 2013. <https://www.usenix.org/system/files/conference/woot13/woot13-bangert.pdf>.  
567  
568