

A Language-Based Approach to Security

Fred B. Schneider¹, Greg Morrisett¹, and Robert Harper²

¹ Cornell University, Ithaca, NY

² Carnegie Mellon University, Pittsburgh, PA

Abstract. Language-based security leverages program analysis and program rewriting to enforce security policies. The approach promises efficient enforcement of fine-grained access control policies and depends on a trusted computing base of only modest size. This paper surveys progress and prospects for the area, giving overviews of in-lined reference monitors, certifying compilers, and advances in type theory.

1 Introduction

The increasing dependence by industry, government, and society on networked information systems means that successful attacks could soon have widespread and devastating consequences. Integrity and availability now join secrecy as crucial security policies, not just for the military but also for the ever-growing numbers of businesses and individuals that use the Internet. But current systems lack the technology base needed to address these new computer-security needs [16].

For the past few years, we and others have been exploring the extent to which techniques from programming languages—compilers, automated program analysis, type checking, and program rewriting—can help enforce security policies in networked computing systems. This paper explains why this *language-based security* approach is considered so promising, some things that already have been accomplished, and what might be expected. But the paper also can be seen as a testament to research successes in programming languages, giving evidence that the area is poised to have an impact far beyond its traditional scope.

Section 2 discusses two computer security principles that suggest the focus on language-based security is sensible. Then, section 3 discusses the implementation of reference monitors for policy enforcement. Three language-based security paradigms—in-lined reference monitors, type systems, and certifying compilers—are the subject of section 4. Some concluding remarks appear in section 5.

2 Some Classic Principles

Work in language-based security is best understood in terms of two classic computer security principles [15]:

36 **Principle of Least Privilege.** Throughout execution, each principal should
37 be accorded the minimum access necessary to accomplish its task.

38 **Minimal Trusted Computing Base.** Assurance that an enforcement mech-
39 anism behaves as intended is greatest when the mechanism is small and
40 simple.

41 These principles were first articulated over twenty-five years ago, at a time
42 when economics dictated that computer hardware be shared and, therefore,
43 user computations had to be protected from each other. Since it was kernel-
44 implemented abstractions that were being shared, security policies for isolating
45 user computations were formulated in terms of operating system objects. More-
46 over, in those days, the operating system kernel itself was small and simple. The
47 kernel thus constituted a minimal trusted computing base that instantiated the
48 Principle of Least Privilege.

49 Computing systems have changed radically in twenty-five years. Operating
50 system kernels are no longer simple or small. The source code for Windows 2000,
51 for example, comprises millions of lines of code. One reason today's operating
52 systems are so large is to support basic services (*e.g.*, windowing, graphics,
53 distributed file systems) needed for the varied tasks they now perform. But an-
54 other reason is performance—subsystems are no longer isolated from each other
55 to avoid expensive context switches during execution. For example, the graphics
56 subsystem of Windows is largely contained within the kernel's address space(!)
57 to reduce the cost of invoking common drawing routines. So operating system
58 kernels today constitute an unmanageably large and complicated computing
59 base—a far cry from the minimal trusted computing base we seek.

60 Moreover, today's operating system kernels enforce only coarse-grained poli-
61 cies.

- 62 – Almost all code for a given machine is run on behalf of a single user, and
63 principals are equated with users. Consequently, virtually all code runs as a
64 single principal under a single policy (*i.e.*, a single set of access permissions).
- 65 – Many resources are not implemented by the operating systems kernel. Thus,
66 the kernel is unable to enforce the policies needed for protecting most of the
67 system's resources.

68 This *status quo* allows viruses, such as Melissa and the Love Bug, to propagate by
69 hiding within an email message a script that is transparently invoked by the mail-
70 viewer application (without an opportunity for the kernel to intercede) when
71 the message is opened.¹ In short, today's operating systems do not and cannot
72 enforce policies concerning application-implemented resources, and individual
73 subsystems lack the clear boundaries that would enable policies concerning the
74 resources they manage to be enforced.

¹ Because the script runs with all the privileges of the user that received the mes-
sage, the virus is able to read the user's address book and forward copies of itself,
masquerading as personal mail from a trusted friend.

75 Though ignored today, Principle of Least Privilege and Minimal Trusted
76 Computing Base, remain sound and sensible principles, as they are independent
77 of system architecture, computer speed, and the other dimensions of computer
78 systems that have undergone radical change. Traditional operating system in-
79 stantiations of these principles might no longer be feasible, but that does not
80 preclude using other approaches to policy enforcement. Language-based security
81 is one such approach.

82 **3 The Case for Language-Based Security**

83 *A reference monitor* observes execution of a target system and halts that sys-
84 tem whenever it is about to violate some security policy of concern. Security
85 mechanisms found in hardware and system software typically either implement
86 reference monitors directly or are intended to facilitate the implementation of
87 reference monitors. For example, an operating system might mediate access to
88 files and other abstractions it supports, thereby implementing a reference moni-
89 tor for policies concerning those objects. As another example, the context switch
90 (trap) caused whenever a system call instruction is executed forces a transfer of
91 control, thereby facilitating invocation of a reference monitor whenever a system
92 call is executed.

93 To do its job, a reference monitor must be protected from subversion by the
94 target systems it monitors. Memory protection hardware, which ensures that
95 execution by one program cannot corrupt the instructions or data of another, is
96 commonly used for this purpose. But placing the reference monitor and target
97 systems in separate address spaces has a performance cost and restricts what
98 policies can be enforced.

- 99 – The performance cost results from the overhead due to context switches
100 associated with transferring control to the reference monitor from within
101 the target systems. The reference monitor must receive control whenever a
102 target system participates in an event relevant to the security policy being
103 enforced. In addition, data must be copied between address spaces.
- 104 – The restrictions on what policies can be enforced arise from the means by
105 which target system events cause the reference monitor to be invoked, since
106 this restricts the vocabulary of events that can be involved in security poli-
107 cies. Security policies that govern operating system calls, for example, are
108 feasible because traps accompany systems calls.

109 The power of the Principle of Least Privilege depends on having flexible and
110 general notions of principal and minimum access. Any interface—not just the
111 user/kernel interface—might define the objects governed by a security policy.
112 And an expressive notion of principal is needed if enforcement decisions might
113 depend on, among other things, the current state of the machine, past execution
114 history, who authored the code, on who’s behalf is the code executing, and so
115 on.

116 Language-based security, being based on program analysis and program rewrit-
117 ing, supports the flexible and general notions of principal and minimum access
118 needed in order to instantiate the Principle of Least Privilege. In particular,
119 software, being universal, can always provide the same functionality (if not per-
120 formance) as a reference monitor. An interpreter, for instance, could include not
121 only the same checks as found in hardware or kernel-based protection mech-
122 anisms but also could implement additional checks involving the application's
123 current and past states.

124 The only question, then, is one of performance. If the overhead of unadul-
125 terated interpretation is too great, then compilation technology, such as just-in-
126 time compilers, partial evaluation, run-time code generation, and profile-driven
127 feedback optimization, can be brought to bear. Moreover, program analysis,
128 including type-checking, dataflow analysis, abstract interpretation, and proof-
129 checking, can be used to reason statically about the run-time behavior of code
130 and eliminate unnecessary run-time policy enforcement checks.

131 Beyond supporting functionality equivalent to hardware and kernel-supported
132 reference monitoring, the language-based approach to security offers other ben-
133 efits. First, language-based security yields policy enforcement solutions that can
134 be easily extended or changed to meet new, application-specific demands. Sec-
135 ond, if a high-level language (such as Java or ML) is the starting point, then
136 linguistic structures, such as modules, abstract data types, and classes, allow
137 programmers to specify and encapsulate application-specific abstractions. These
138 same structures can then provide a vocabulary for formulating fine-grained secu-
139 rity policies. Language-based security is not, however, restricted to systems that
140 have been programmed in high-level languages. In fact, much work is directed
141 at enforcing policies on object code because (i) the trusted computing base is
142 smaller without a compiler and (ii) policies can then be enforced on programs
143 for which no source code is available.

144 **EM Security Policies**

145 A program analyzer operating on program text (source or object code) has
146 more information available about how that program could behave than does
147 a reference monitor observing a single execution.² This is because the program
148 text is a terse representation of all possible behaviors and, therefore, contains
149 information—about alternatives and the future—not available in any single ex-
150 ecution. It would thus seem that, ignoring questions of decidability, program
151 analysis can enforce policies that reference monitors cannot. To make the rela-
152 tionship precise, the class of security policies that reference monitors can enforce
153 was characterized in [17], as follows.

154 A *security policy* defines execution that, for one reason or another, has been
155 deemed unacceptable. Let EM (for Execution Monitoring) be the class of secu-
156 rity policies that can be enforced by monitoring execution of a target system and

² We are assuming that a reference monitor sees only security-relevant actions and values. Once the entire state of the system becomes available, then the reference monitor would have access to the program text.

157 terminating execution that is about to violate the security policy being enforced.
 158 Clearly, EM includes those policies that can be enforced by security kernels, refer-
 159 ence monitors, firewalls, and most other operating system and hardware-based
 160 enforcement mechanisms that have appeared in the literature. Target systems
 161 may be objects, modules, processes, subsystems, or entire systems; the execution
 162 steps monitored may range from fine-grained actions (such as memory accesses)
 163 to higher-level operations (such as method calls) to operations that change the
 164 security-configuration and thus restrict subsequent execution.

165 Mechanisms that use more information than would be available only from
 166 monitoring a target system’s execution are, by definition, excluded from EM. In-
 167 formation provided to an EM mechanism is thus insufficient for predicting future
 168 steps the target system might take, alternative possible executions, or all possi-
 169 ble target system executions. Therefore, compilers and theorem-provers, which
 170 analyze a static representation of a target system to deduce information about
 171 all of its possible executions, are not considered EM mechanisms. Also excluded
 172 from EM are mechanisms that modify a target system before executing it. The
 173 modified target system would have to be “equivalent” to the original (except
 174 for aborting executions that would violate the security policy of interest), so a
 175 definition for “equivalent” is thus required to analyze this class of mechanisms.

176 We represent target system executions by finite and infinite sequences, where
 177 Ψ denotes a universe of all possible finite and infinite sequences. The manner
 178 in which executions are represented is irrelevant here. Finite and infinite se-
 179 quences of atomic actions, of higher-level system steps, of program states, or
 180 of state/action pairs are all plausible alternatives. A target system S defines a
 181 subset Σ_S of Ψ corresponding to the executions of S .

182 A characterization of EM-enforceable security policies is interesting only if
 183 the definition being used for “security policy” is broad enough so that it does
 184 not exclude things usually considered security policies.³ Also, the definition must
 185 be independent of how EM is defined, for otherwise the characterization of EM-
 186 enforceable security policies would be a tautology, hence uninteresting. We there-
 187 fore adopt the following.

188 **Definition of Security Policy:** A *security policy* is specified by giving a pred-
 189 icate on sets of executions. A target system S *satisfies* security policy \mathcal{P} if
 190 and only if $\mathcal{P}(\Sigma_S)$ equals *true*.

191 By definition, enforcement mechanisms in EM work by monitoring execution
 192 of the target. Thus, any security policy \mathcal{P} that can be enforced using a mechanism
 193 from EM must be specified by a predicate of the form

$$194 \mathcal{P}(\Pi) : (\forall \sigma \in \Pi : \widehat{\mathcal{P}}(\sigma)) \tag{1}$$

195 where $\widehat{\mathcal{P}}$ is a predicate on (individual) executions. $\widehat{\mathcal{P}}$ formalizes the criteria used
 196 by the enforcement mechanism for deciding to terminate an execution that would
 197 otherwise violate the policy being enforced. In [1] and the literature on linear-
 time concurrent program verification, a set of executions is called a *property* if

³ However, there is no harm in being liberal about what is considered a security policy.

198 set membership is determined by each element alone and not by other members
 199 of the set. Using that terminology, we conclude from (1) that a security policy
 200 must be a property in order for that policy to have an enforcement mechanism
 201 in EM.

202 Not every security policy is a property. Some security policies cannot be de-
 203 fined using criteria that individual executions must each satisfy in isolation. For
 204 example, information flow policies often characterize sets that are not proper-
 205 ties (as proved in [10]). Whether information flows from variable x to y in a
 206 given execution depends, in part, on what values y takes in other possible execu-
 207 tions (and whether those values are correlated with the value of x). A predicate
 208 to specify such sets of executions cannot be constructed only using predicates
 209 defined on single executions in isolation.

210 Not every property is EM-enforceable. Enforcement mechanisms in EM can-
 211 not base decisions on possible future execution, since that information is, by de-
 212 finition, not available to such a mechanism. Consider security policy \mathcal{P} of (1), and
 213 suppose σ' is the prefix of some finite or infinite execution σ where $\widehat{\mathcal{P}}(\sigma) = \text{true}$
 214 and $\widehat{\mathcal{P}}(\sigma') = \text{false}$ hold. Because execution of a target system might terminate
 215 before σ' is extended into σ , an enforcement mechanism for \mathcal{P} must prohibit σ'
 216 (even though supersequence σ satisfies $\widehat{\mathcal{P}}$).

217 We can formalize this requirement as follows. For σ a finite or infinite exe-
 218 cution having i or more steps, and τ' a finite execution, let

219 $\sigma[..i]$ denote the prefix of σ involving its first i steps
 220 $\tau' \sigma$ denote execution τ' followed by execution σ

221 and define Ψ^- to be the set of all finite prefixes of elements in set Ψ of finite
 222 and/or infinite sequences. Then, the above requirement for \mathcal{P} —that \mathcal{P} is *prefix*
 223 *closed*—is:

$$(\forall \tau' \in \Psi^- : \neg \widehat{\mathcal{P}}(\tau') \Rightarrow (\forall \sigma \in \Psi : \neg \widehat{\mathcal{P}}(\tau' \sigma))) \quad (2)$$

224 Finally, note that any execution rejected by an enforcement mechanism must
 225 be rejected after a finite period. This is formalized by:

$$(\forall \sigma \in \Psi : \neg \widehat{\mathcal{P}}(\sigma) \Rightarrow (\exists i : \neg \widehat{\mathcal{P}}(\sigma[..i]))) \quad (3)$$

226 Security policies satisfying (1), (2), and (3) are *safety properties* [8], proper-
 227 ties stipulating that no “bad thing” happens during any execution. Formally, a
 228 property Γ is defined in [9] to be a safety property if and only if, for any finite
 229 or infinite execution σ ,

$$\sigma \notin \Gamma \Rightarrow (\exists i : (\forall \tau \in \Psi : \sigma[..i] \tau \notin \Gamma)) \quad (4)$$

230 holds. This means that Γ is a safety property if and only if Γ can be characterized
 231 using a set of finite executions that are prefixes of all executions excluded from
 232 Γ . Clearly, a security policy \mathcal{P} satisfying (1), (2), and (3) has such a set of finite
 233 prefixes—the set of prefixes $\tau' \in \Psi^-$ such that $\neg \widehat{\mathcal{P}}(\tau')$ holds—so \mathcal{P} is satisfied
 234 by sets that are safety properties according to (4).

235 The above analysis of enforcement mechanisms in EM has established:

236 **Non EM-Enforceable Security Policies:** If the set of executions for a secu-
237 rity policy \mathcal{P} is not a safety property, then an enforcement mechanism from
238 EM does not exist for \mathcal{P} .

239 One consequence is that ruling-out additional executions never causes an EM-
240 enforceable policy to be violated, since ruling-out executions never invalidates
241 a safety property. Thus, an EM enforcement mechanism for any security policy
242 \mathcal{P}' satisfying $\mathcal{P}' \Rightarrow \mathcal{P}$ also enforces security policy \mathcal{P} . However, a stronger policy
243 \mathcal{P}' might proscribe executions that do not violate \mathcal{P} , so using \mathcal{P}' is not without
244 potentially significant adverse consequences. The limit case, where \mathcal{P}' is satisfied
245 only by the empty set, illustrates this problem.

246 Second, Non EM-Enforceable Security Policies implies that EM mechanisms
247 compose in a natural way. When multiple EM mechanisms are used in tandem,
248 the policy enforced by the aggregate is the conjunction of the policies that are
249 enforced by each mechanism in isolation. This is attractive, because it enables
250 complex policies to be decomposed into conjuncts, with a separate mechanism
251 used to enforce each of the component policies.

252 We can use the Non EM-Enforceable Security Policies result to see whether
253 or not a given security policy might be enforced using a reference monitor (or
254 some other form of execution monitoring). For example, access control policies,
255 which restrict what operations principals can perform on objects, define safety
256 properties. (The set of proscribed partial executions contains those partial ex-
257 ecutions ending with an unacceptable operation being attempted.) Information
258 flow policies do not define sets that are properties (as discussed above). And,
259 availability policies, if taken to mean that no principal is forever denied use of
260 some given resource, is not a safety property—any partial execution can be ex-
261 tended in a way that allows a principal to access the resource, so the defining
262 set of proscribed partial executions that every safety property must have is ab-
263 sent. Thus we conclude that access control policies can be enforced by reference
264 monitors but neither information flow nor availability policies (as we formulated
265 them) can be.

266 4 Enforcing Security Policies

267 The building blocks of language-based security are program rewriting and pro-
268 gram analysis. By rewriting a program, we can ensure that the result is incapable
269 of exhibiting behavior disallowed by some security policy at hand. And by ana-
270 lyzing a program, we ensure only those programs that cannot violate the policy
271 are ever given an opportunity to be executed.

272 That is the theory. Actual embodiments of the language-based security vision
273 invariably combine program rewriting and program analysis. Today's research
274 efforts can be grouped into two schools. One—in-lined reference monitors—
275 takes program rewriting as a starting point; the other—type-safe programming
276 languages—takes program analysis, as a starting point. In what follows, we dis-
277 cuss the strengths and weaknesses of each of these schools. We then discuss
278 an emerging approach—certifying compilation—and how the combination of all

279 three techniques (rewriting, analysis, and certification) yield a comprehensive
280 security framework.

281 4.1 In-lined Reference Monitors

282 An alternative to placing the reference monitor and the target system in sepa-
283 rate address spaces is to modify the target system code, effectively merging
284 the reference monitor in-line. This is, in effect, what is done by software-fault
285 isolation (SFI), which enforces the security policy that prevents reads, writes,
286 or branches to memory locations outside of certain predefined memory regions
287 associated with a target system [20]. But a reference monitor for any EM secu-
288 rity policy could be merged into a target application, provided the target can be
289 prevented from circumventing the merged code.

290 Specifying such an *in-lined reference monitor* (IRM) involves defining [6]

- 291 – *security events*, the policy-relevant operations that must be mediated by the
292 reference monitor;
- 293 – *security state*, information stored about earlier security events that is used
294 to determine which security events can be allowed to proceed; and
- 295 – *security updates*, program fragments that are executed in response to security
296 events and that update the security state, signal security violations, and/or
297 take other remedial action (*e.g.*, block execution).

298 A load-time, trusted *IRM rewriter* merges checking code into the application
299 itself, using program analysis and program rewriting to protect the integrity
300 of those checks. The IRM rewriter thus produces a *secured application*, which is
301 guaranteed not to take steps violating the security policy being enforced. Notice,
302 with the IRM approach, the conjunction of two policies can be enforced by
303 passing the target application through the IRM rewriter twice in succession—
304 once for each policy. And also, by keeping policy separate from program, the
305 approach makes it easier to reason about and evolve the security of a system.

306 Experiments with two generations of IRM enforcement suggest that the ap-
307 proach is quite promising. SASI (Security Automata SFI Implementation), the
308 first generation, comprised two realizations [5]. One transformed Intel x86 assem-
309 bly language; the other transformed Java Virtual Machine Language (JVML).
310 Second generation IRM enforcement tools PoET/PSLang, (Policy Enforcement
311 Toolkit/Policy Specification Language) transformed JVML [6].

312 The x86 SASI prototype works with assembly language output of the GNU
313 gcc C compiler. Object code produced by gcc observes certain register-usage
314 conventions, is not self-modifying, and is guaranteed to satisfy two assumptions:

- 315 – Program behavior is insensitive to adding stutter-steps (*e.g.*, nop’s).
- 316 – Variables and branch-targets are restricted to the set of labels identified by
317 gcc during compilation.

318 These restrictions considerably simplify the task of preventing code for checking
319 and for security updates from being corrupted by the target system. In particular,
320 it suffices to apply x86 SASI with the simple memory-protection policy enforced
321 by SFI in order to obtain target-system object code that cannot subvert merged-
322 in security state or security updates.

323 The JVMML SASI prototype exploits the type safety of JVMML programs to
324 prevent merged-in variables and state from being corrupted by the target system
325 in which it resides. In particular, variables that JVMML SASI adds to a JVMML
326 object program are inaccessible to that program by virtue of their names and
327 types; and code that JVMML SASI adds cannot be circumvented because JVMML
328 type-safety prevents jumps to unlabeled instructions—these code fragments are
329 constructed so they do not contain labels.⁴

330 The type-safety of JVMML also empowers the JVMML SASI user who is formu-
331 lating a security policy that concerns application abstractions. JVMML instruc-
332 tions contain information about classes, objects, methods, threads, and types.
333 This information is made available (though platform-specific functions) to the
334 author of a security policy. Security policies for JVMML SASI thus can define per-
335 missible computations in terms of these application abstractions. In contrast,
336 x86 code will contain virtually no information about a C program it represents,
337 so the author of a security policy for x86 SASI may be forced to synthesize
338 application events from sequences of assembly language instructions.

339 Experience with the SASI prototypes has proved quite instructive. A refer-
340 ence monitor that checks every machine language instruction initially seemed
341 like a powerful basis for defining application-specific security policies. But we
342 learned from SASI that, in practice, this power is difficult to harness. Most x86
343 object code, for example, does not make explicit the application abstractions
344 that are being manipulated by that code. There is no explicit notion of a “func-
345 tion” in x86 assembly language, and “function calls” are found by searching for
346 code sequences resembling the target system’s calling convention. The author of
347 a security policy thus finds it necessary to embed a disassembler (or event syn-
348 thesizer) within a security policy description. This is awkward and error-prone.

349 One solution would be to obtain IRM enforcement by rewriting high-level
350 language programs rather than object code. Security updates could be merged
351 into the high-level language program (say) for the target system rather than
352 being merged into the object code produced by a compiler. But this is unattrac-
353 tive because an IRM rewriter that modifies high-level language programs adds
354 a compiler to the trusted computing base. The approach taken in JVMML SASI
355 seemed the more promising, and it (along with a desire for a friendlier language
356 for policy specification) was the motivation for PoET/PSLang. The lesson is to
357 rely on annotations of the object code that are easily checked and that expose
358 application abstractions. And that approach is not limited to JVMML code or

⁴ JVMML SASI security policies must also rule out indirect ways of compromising the variables or circumventing the code added for policy enforcement. For example, JVMML’s dynamic class loading and program reflection must be disallowed.

359 even to type-safe high-level languages. Object code for x86 could include the
360 necessary annotations by using TAL [11] (discussed below).

361 4.2 Type Systems

362 Type-safe programming languages, such as ML, Modula, Scheme, or Java, en-
363 sure that operations are only applied to appropriate values. They do so by guar-
364 anteeing a number of inter-related safety properties, including *memory safety*
365 (programs can only access appropriate memory locations) and *control safety*
366 (programs can only transfer control to appropriate program points).

367 Type systems that support type abstraction then allow programmers to spec-
368 ify new, abstract types along with signatures for operations that prevent unau-
369 thorized code from applying the wrong operations to the wrong values. For
370 example, even if we represent file descriptors as integers, we can use type ab-
371 straction to ensure that only integers created by our implementation are passed
372 to file-descriptor routines. In this respect, type systems, like IRMs, can be used
373 to enforce a wider class of fine-grained, application-specific access policies than
374 operating systems. In addition, abstract type signatures provide the means to
375 enrich the vocabulary of an enforcement mechanism in an application-specific
376 way.

377 The key idea underlying the use of type systems to enforce security policies
378 is to shift the burden of proving that a program complies with a policy from the
379 code recipient (the end user) to the code producer (the programmer). Not only
380 are familiar run-time mechanisms (*e.g.*, address space isolation) insufficiently
381 expressive for enforcing fine-grained security policies but, to the extent that
382 they work at all, these mechanisms impose the burden of enforcement on the end
383 user through the imposition of dynamic checks. In contrast, type-based methods
384 impose on the programmer the burden of demonstrating compliance with a given
385 security policy. The programmer must write the program in conformance with
386 the type system; the end user need only type check the code to ensure that it is
387 safe to execute.

388 The only run-time checks required in a type-based framework are those nec-
389 essary for ensuring soundness of the type system itself. For example, the type
390 systems of most commonly-used programming languages do not attempt to en-
391 force value-range restrictions, such as the requirement that the index into an
392 array is within bounds. Instead, any integer-valued index is deemed acceptable
393 but a run-time check is imposed to ensure that memory safety is preserved.

394 However, it is important to note that the need to dynamically check val-
395 ues, such as array indices, is not inherent to type systems. Rather, the logics
396 underlying today's type systems are too weak, so programmers are unable to
397 express the conditions necessary to ensure soundness statically. This is largely a
398 matter of convenience, though. It is possible to construct arbitrarily expressive
399 type systems with the power of any logic. Such type systems generally require
400 sophisticated theorem provers and programmer guidance in the construction of a
401 proof of type soundness. For example, recent work on dependent type systems [3,
402 21] extends type checking to include the expression of value-range restrictions

403 sufficient to ensure that array bounds checks may (in many cases) be eliminated,
404 but programmers must add additional typing annotations (*e.g.*, loop invariants)
405 to aid the type checker.

406 Fundamentally, the only limitation on the expressiveness of a type system
407 is the effort one is willing to expend demonstrating type correctness. Keep in
408 mind that this is a matter of proof—the programmer must demonstrate to the
409 checker that the program complies with the safety requirements of the type
410 system. In practice, it is common to restrict attention to type systems for which
411 checking is computable with a reasonable complexity bound, but more advanced
412 programming systems such as NuPRL [3] impose no such restrictions and admit
413 arbitrary theorem proving for demonstrating type safety.

414 In summary, advances in the design of type systems now make it possible to
415 express useful security properties and to enforce them in a lightweight fashion,
416 all the while minimizing the burden on the end user to enforce memory and
417 control safety.

418 4.3 Certifying Compilers

419 Until recently, the primary weakness of type-based approaches to ensuring safety
420 has been that they relied on

421 ***High-Level Language Assumption.*** The program must be written in a pro-
422 gramming language having a well-defined type system and operational se-
423 mantics.

424 In particular, the programmer is obliged to write code in the high-level language,
425 and the end user is obliged to correctly implement both its type system (so
426 that programs can be type checked) and its operational semantics (so that it
427 can be executed). These consequences would have questionable utility if they
428 substantially increased the size of the trusted computing base or they reduced
429 the flexibility with which systems could be implemented. But they don't have
430 to. Recent developments in compiler technology are rapidly obviating the High-
431 Level Language Assumption without sacrificing the advantages of type-based
432 approaches. We now turn to that work.

433 A *certifying compiler* is a compiler that, when given source code satisfying
434 a particular security policy, not only produces object code but also produces a
435 *certificate*—machine-checkable evidence that the object code respects the policy.
436 For example, Sun's `javac` compiler takes Java source code that satisfies a type-
437 safety policy, and it produces JVMCL code that respects type-safety. In this case,
438 the “certificate” is the type information embedded within the JVMCL bytecodes.

439 Certifying compilers are an important tool for policy enforcement because
440 they do their job from outside the trusted computing base. To verify that the
441 output object code of a certifying compiler respects some policy, an automated
442 *certificate checker* (that is part of the trusted computing base) is employed. The
443 certificate checker analyzes the output of the certifying compiler and verifies that
444 this object code is consistent with the characterization given in a certificate.

445 For example, a JVMML bytecode verifier can ensure that bytecodes are type-safe
446 independent of the Java compiler that produced them.

447 Replacing a trusted compiler with an untrusted certifying compiler plus a
448 trusted certificate checker is advantageous because a certificate checker, including
449 type-checkers or proof-checkers, is typically much smaller and simpler than a
450 program that performs the analysis and transformations needed to generate
451 certified code. Thus, the resulting architecture has a smaller trusted computing
452 base than would an architecture that employed a trusted compiler or analysis.

453 Java is perhaps the most widely disseminated example of this certifying com-
454 piler architecture. But the policy supported by the Java architecture is restricted
455 to a relatively primitive form of type-safety, and the bytecode language is still
456 high-level, requiring either an interpreter or just-in-time compiler for execution.

457 The general approach of certifying compilation is really quite versatile. For
458 instance, building on the earlier work of the TIL compiler [19], Morrisett *et al.*
459 showed that it is possible to systematically build type-based, certifying compil-
460 ers for high-level languages that produce Typed Assembly Language (TAL) for
461 concrete machines (as opposed to virtual machines) [12]. Furthermore, the type
462 system of TAL supports some of the refinements, such as value ranges, needed
463 to avoid the overhead of dynamic checks. Nonetheless, as it stands today, the
464 set of security policies that TAL can enforce are essentially those that can be
465 realized through traditional notions of type-safety.

466 Perhaps the most aggressive instance of certifying compilers was developed
467 by Necula and Lee, who were the first to move beyond implicit typing annota-
468 tions and develop an architecture in which certificates were explicit. The result,
469 called Proof-Carrying Code (PCC) [13, 14], enjoys a number of advantages over
470 previous work. In particular, the axioms, inference rules, and proofs of PCC are
471 represented as terms in a meta-logical programming language called LF [7], and
472 certificate checking corresponds to LF type-checking. The advantages of using a
473 meta-logical language are twofold:

- 474 – It is relatively simple to customize the logic by adding new axioms or infer-
475 ence rules.
- 476 – Meta-logical type checkers can be quite small, so in principle a PCC-based
477 system can have an extremely small trusted computing base. For example,
478 Necula implemented an LF type checker that is about 6 pages of C code [13].

479 Finally, unlike the JVMML or TAL, PCC is not limited to enforcing traditional
480 notions of type safety. It is also not limited to EM policies. Rather, as long as
481 the logic is expressive enough to state the desired policy, and as long as the
482 certifying compiler can construct a proof in that logic that the code will respect
483 the policy, then a PCC-based system can check conformance.

484 4.4 Putting the Technologies Together

485 Combine in-lined reference monitors, type systems, and certifying compilers—
486 the key approaches to language-based security—and the sum will be greater

487 than the parts. In what follows, we discuss the remarkable synergies among
488 these approaches.

489 ***Integrating IRM enforcement with Type Systems.*** Static type systems are
490 particularly well-suited for enforcing security policies that have been negotiated
491 in advance. Furthermore, enforcement through static checking usually involves
492 less overhead than a more dynamic approach. And finally, static type systems
493 hold the promise of enforcing liveness properties (*e.g.*, termination) and policies
494 that are not properties (*e.g.*, absence of information flow)—things that refer-
495 ence monitors cannot enforce. However, static type systems are ill-suited for the
496 enforcement of policies that depend upon things that can be detected at runtime
497 but cannot be ascertained during program development. Also, it may be simpler
498 to insert a dynamic check than to have a programmer develop an explicit proof
499 that the check is not needed. Consequently, by combining IRMs with advanced
500 type systems, we have both the opportunity to enforce a wider class of policies
501 and more flexibility in choosing an appropriate enforcement mechanism.

502 ***Extending IRM enforcement with Certifying Compilers.*** Program rewrit-
503 ing without subsequent optimization generally leads to systems exhibiting poor
504 performance. However, an IRM rewriter could reduce the performance impact
505 of added checking code by inserting checks only where there is some chance
506 that a security update actually needs to be performed. For example, in enforc-
507 ing a policy that stipulates messages are never sent after certain files have been
508 read, an IRM rewriter needn't insert code before and after every instruction. A
509 small amount of simple analysis would allow insertions to be limited to those
510 instructions involving file reads and message sends; and a global analysis might
511 allow more aggressive optimizations. Optimization technology, then, can recover
512 performance for the IRM approach.

513 But an IRM rewriter that contains a global optimizer is larger and more
514 complicated than one that does not. Any optimizations had better always be
515 done correctly, too, since bugs might make it possible for the security policy at
516 hand to be violated. So, optimization routines—just like the rest of the IRM
517 rewriter—are part of the trusted computing base. In the interest of keeping the
518 trusted computing base small, we should hesitate to employ a complicated IRM
519 rewriter.

520 Must an IRM architecture sacrifice performance on the altar of minimizing
521 the trusted computing base? Not if the analysis and optimization are done with
522 the lesson of certifying compilers in mind. An IRM rewriter can add checking
523 code and security updates and then do analysis and optimization to remove un-
524 necessary checking code, provided the IRM rewriter produces a certificate along
525 with the modified object code. That certificate should describe what code was
526 added everywhere and the analysis that allowed code to be deleted, thereby
527 enabling a certificate checker (in the trusted computing base) to establish in-
528 dependently that the output of the IRM rewriter will indeed never violate the
529 security policy of interest. Thus, the IRM rewriter is extended if ideas from
530 certifying compilers are adopted.

531 ***Extending Certifying Compilers with IRM enforcement.*** Certifying com-
532 pilers are limited to analysis that can be done automatically. And, unfortunately,
533 there are deep mathematical reasons why certain program analysis cannot be
534 automated—analysis that would be necessary for policies much simpler than
535 found in class EM. Must a certifying compiler architecture sacrifice expressive-
536 ness on the alter of automation?

537 In theory, it would seem so. But in practice, much analysis becomes possible
538 when program rewriting is first allowed. This is an instance of the familiar trade-
539 off between static and dynamic checks during type checking. For instance, rather
540 than verifying at compile time that a given array index never goes out of bounds,
541 it is a simple matter to have the compiler emit a run-time check. Static analysis
542 of the modified program is guaranteed to establish that the array access is never
543 out of bounds (because the added check prevents it from being so).

544 The power of a certifying compilers is thus amplified by the capacity to do
545 program rewriting. In the limit, what is needed is the means to modify a program
546 and obtain one in which a given security policy is not violated—exactly what
547 an IRM rewriter does. Thus, the power of certifying compilers is extended if
548 deployed in concert with an IRM rewriter.

549 **5 Concluding Remarks**

550 In-lined reference monitors, certifying compilers, and advanced type systems are
551 promising approaches to system security. Each allows rich instantiations of the
552 Principle of Least Privilege; each depends on only a minimal trusted computing
553 base, despite the ever-growing sizes for today’s operating systems, compilers,
554 and programming environments.

555 The idea of using languages and compilers to help enforce security policies is
556 not new. The Burroughs B-5000 system required applications to be written in a
557 high-level language (Algol), and the Berkeley SDS-940 system employed object-
558 code rewriting as part of its system profiler. More recently, the SPIN [2], Vino
559 [22, 18], and Exokernel [4] extensible operating systems have relied on language
560 technology to protect a base system from a limited set of attacks by extensions.

561 What is new in so-called language-based security enforcement is the degree
562 to which language semantics provides the leverage. The goal is to obtain inte-
563 grated mechanisms that work for both high-level and low-level languages; that
564 are applicable to an extremely broad class of fine-grained security policies; and
565 that allow flexible allocation of work and trust among the elements responsible
566 for enforcement.

567 **Acknowledgments**

568 The views and conclusions contained herein are those of the authors and should not
569 be interpreted as necessarily representing the official policies or endorsements, either
570 expressed or implied, of these organizations or the U.S. Government.

571 Schneider is supported in part by ARPA/RADC grant F30602-96-1-0317, AFOSR
572 grant F49620-00-1-0198, Defense Advanced Research Projects Agency (DARPA) and

573 Air Force Research Laboratory Air Force Material Command USAF under agreement
574 number F30602-99-1-0533, National Science Foundation Grant 9703470, and a grant
575 from Intel Corporation.

576 Morrisett is supported in part by AFOSR grant F49620-00-1-0198, and the National
577 Science Foundation under Grant No. EIA 97-03470.

578 Harper is sponsored by the Advanced Research Projects Agency CSTO under the
579 title “The Fox Project: Advanced Languages for Systems Software”, ARPA Order No.
580 C533, issued by ESC/ENS under Contract No. F19628-95-C-0050.

581 References

- 582 1. B. Alpern and F.B. Schneider. Defining liveness. *Information Processing Letters*
583 21(4):181–185, Oct. 1985.
- 584 2. B. Bershad, S. Savage, P. Pardyak, E. Sirer, M. Fiuczynski, D. Becker, C. Cham-
585 bers, and S. Eggers. Extensibility, safety and performance in the SPIN operating
586 system. In *Proc. 15th ACM Symp. on Operating System Principles (SOSP)*, pages
587 267–284, Copper Mountain, Dec. 1995.
- 588 3. R. L. Constable *et al.* *Implementing Mathematics with the NuPRL Proof Develop-*
589 *ment System*. Prentice-Hall, 1986.
- 590 4. D. Engler, M. Kaashoek, and J. O’Toole. Exokernel: An operating system archi-
591 tecture for application-level resource management. In *Proc. 15th ACM Symp. on*
592 *Operating System Principles (SOSP)*, Copper Mountain, 1995.
- 593 5. U. Erlingsson and F. B. Schneider. SASI enforcement of security policies: A ret-
594 ropective. In *Proceedings of the New Security Paradigms Workshop*, Ontario,
595 Canada, Sept. 1999.
- 596 6. U. Erlingsson and F. B. Schneider. IRM enforcement of java stack inspection. In
597 *IEEE Symposium on Security and Privacy*, Oakland, California, May 2000.
- 598 7. R. Harper, F. Honsell, and G. Plotkin. A framework for defining logics. *Journal*
599 *of the ACM*, 40(1):143–184, Jan. 1993.
- 600 8. L. Lamport. Proving the correctness of multiprocess programs. *IEEE Transactions*
601 *on Software Engineering*, SE-3(2):125–143, March 1977.
- 602 9. L. Lamport. Logical Foundation. In *Distributed Systems-Methods and Tools for*
603 *Specification*, pages 119-130, Lecture Notes in Computer Science, Vol 190. M. Paul
604 and H.J. Siegart, editors. Springer-Verlag, 1985, New York.
- 605 10. J. McLean. A general theory of composition for trace sets closed under selective
606 interleaving functions. In *Proc. 1994 IEEE Computer Society Symposium on Re-*
607 *search in Security and Privacy*, pages 79–93, Oakland, Calif., May 1994.
- 608 11. G. Morrisett, D. Walker, K. Crary, and N. Glew. From System F to typed assembly
609 language. In *Proc. 25th ACM Symp. on Principles of Programming Languages*
610 *(POPL)*, pages 85–97, San Diego California, USA, January 1998.
- 611 12. G. Morrisett, D. Walker, K. Crary, and N. Glew. From System F to typed assembly
612 language. *ACM Transactions on Programming Languages and Systems*, 21(3):528–
613 569, May 1999.
- 614 13. G. C. Necula and P. Lee. Safe kernel extensions without run-time checking. In *Pro-*
615 *ceedings of Operating System Design and Implementation*, pages 229–243, Seattle,
616 Oct. 1996.
- 617 14. G. C. Necula. Proof-carrying code. In *Proc. 24th ACM Symp. on Principles of*
618 *Programming Languages (POPL)*, pages 106–119, Jan. 1997.

- 619 15. J. Saltzer and M. Schroeder. The protection of information in computer systems.
620 *Proceedings of the IEEE*, 9(63), Sept. 1975.
- 621 16. F. B. Schneider, editor. *Trust in Cyberspace*. National Academy Press, Washington,
622 D.C., 1999.
- 623 17. F. B. Schneider. Enforceable security policies. *ACM Transactions on Information*
624 *and System Security*, 2(4), Mar. 2000.
- 625 18. M. Seltzer, Y. Endo, C. Small, and K. Smith. Dealing with disaster: Surviving
626 misbehaved kernel extensions. In *Proc. USENIX Symp. on Operating Systems*
627 *Design and Implementation (OSDI)*, pages 213–227, Seattle, Washington, Oct.
628 1996.
- 629 19. D. Tarditi, G. Morrisett, P. Cheng, C. Stone, R. Harper, and P. Lee. TIL: A type-
630 directed optimizing compiler for ML. In *ACM Conf. on Programming Language*
631 *Design and Implementation*, pages 181–192, Philadelphia, May 1996.
- 632 20. R. Wahbe, S. Lucco, T. Anderson, and S. Graham. Efficient software-based fault
633 isolation. In *Proc. 14th ACM Symp. on Operating System Principles (SOSP)*, pages
634 203–216, Asheville, Dec. 1993.
- 635 21. H. Xi and F. Pfenning. Eliminating array bound checking through dependent
636 types. In *Proc. ACM SIGPLAN Conference on Programming Language Design*
637 *and Implementation (PLDI)*, pages 249–257, Montreal Canada, June 1998.
- 638 22. E. Yasuhiro, J. Gwertzman, M. Seltzer, C. Small, K. A. Smith, and D. Tang. VINO:
639 The 1994 fall harvest. Technical Report TR-34-94, Harvard Computer Center for
640 Research in Computing Technology, 1994.