

Attacking Malicious Code: A report to the Infosec Research Council*

Gary McGraw (Reliable Software Technologies)
and Greg Morrisett (Cornell University)

May 1, 2000

In October of 1999, the Infosec Research Council created a Science and Technology Study Group (ISTSG) focused on malicious code. The purpose of the Malicious Code ISTSG is to develop a national research agenda to address the accelerating threat from malicious code. The study is intended to identify promising new approaches to dealing with the problems posed by malicious code. This report is a byproduct of two meetings of Study Group members and their invited guests. Though this report was written by only two of study group members, we believe it represents an accurate distillation of the ideas and insights of all the participants. However, the authors are responsible for errors and oversights.

Study Group members include: Gary McGraw, Reliable Software Technologies, Chair; Avi Rubin, AT&T Research; Ed Felten, Princeton; Peter G. Neumann, SRI; Lee Badger, NAI Labs; Greg Morrisett, Cornell; Tim Teitelbaum, Grammatech; Virgil Gligor, University of Maryland; Tom Markham, Secure Computing; Jay Lepreau, University of Utah; Bob Balzer, ISI; Joshua Haines, Lincoln Labs; Roger Thompson, ICSA.net; Bob Clemons, NSA; Penny Chase, MITRE; Carl Landwehr, Mitretek; Brad Arkin, Reliable Software Technologies; Sami Saydjari, DARPA; Brian Witten, DARPA; and Dave Thompson, Mitretek. Guests who participated in the two day San Antonio workshop include: mudge, the l0pht; Crispin Cowen, Wirex; Fred Schneider, Cornell; Peter Lee, CMU; Richard Smith, pharlap; John Rushby, SRI; Dan Wallach, Rice University; Amy Felty, University of Ottawa; and David Evans, University of Virginia.

What is Malicious Code?

Malicious code is any code added, changed, or removed from a software system in order to intentionally cause harm or subvert the intended function of the system. Though the problem of malicious code has a long history, a number of recent, widely publicized attacks and certain economic trends suggest that malicious code is rapidly becoming a critical problem for industry, government, and individuals.

Traditional examples of malicious code include viruses, worms, Trojan Horses, and attack scripts, while more modern examples include Java attack applets and dangerous ActiveX controls.

- Viruses are pieces of malicious code that attach to host programs and propagate when an infected program is executed.
- Worms are particular to networked computers. Instead of attaching themselves to a host program, worms carry out programmed attacks to jump from machine to machine across the network.
- Trojan Horses, like viruses, hide malicious intent inside a host program that appears to do something useful (*e.g.*, a program that captures passwords by masquerading as the login daemon.)
- Attack scripts are programs written by experts that exploit security weaknesses, usually across the network, to carry out an attack. Attack scripts exploiting buffer overflows by “smashing the stack” are the most commonly encountered variety.

* The workshops on which this report is based were convened under the auspices of the Infosec Research Council (IRC), with members from U.S. Government organizations that sponsor and conduct information security research. Views expressed in the report are those of the authors and may not reflect those of the IRC, its members, or the organizations they represent.

- Java attack applets are programs embedded in Web pages that achieve foothold through a Web browser.
- Dangerous ActiveX controls are program components that allow a malicious code fragment to control applications.

Recently, the distinctions between malicious code categories have been bleeding together, and so classification has become difficult. Some concrete examples of malicious code are provided in Table 1. Note that recent versions of malicious code are really amalgamations of different categories.

Malicious Code	Date	Category	Explanation
Love Bug	2000	Mobile code virus	The fastest spreading virus of all time used VB script and Microsoft Outlook mail to propagate. Caused an estimated \$10billion in damage.
Trinoo (and other dDoS scripts)	2000	Remote-control attack script	The highly-publicized denial of service attacks of February 2000 were carried out by remotely-planted agent programs.
Melissa	1999	Mobile code virus	The second fastest spreading virus of all time used e-mail to propagate. Infected over 1.2 million machines in a few hours.
Explore.Zip	1999	Mobile code worm	An e-mail borne worm that exploited problems in Microsoft Windows to propagate.
Happy99	1999	Virus	A widespread virus infecting Microsoft PCs.
CIH	1998	Virus	A particularly dangerous virus that attacks BIOS in PCs. Ran rampant in Asia before being contained.
Back Orifice	1998	Offensive code	Remote control program installed on Windows machines by crackers. Pervasive.
Attack scripts		Offensive code	Crackers called "script kiddies" download malicious code from the Internet and run it against any number of targets. Some expert must create and release the script to begin with. Widespread. Most common attack: buffer overflow.
ActiveX (scripting)	1997	Mobile code	Decried by security professionals, Microsoft's ActiveX system introduces grave security risks by relying on user's discretion and judgment.
Java Attack Applets	1996-1999	Mobile code	Attack applets placed on Web sites take advantage of flaws in the Java security model to carry out attacks. 17 known attacks.
Morris worm	1988	Worm	Released in 1988 by Robert Morris, Jr, this program crashed around 6000 computers (around 10% of the Internet at the time).
Thompson's compiler trick	1984	Trojan Horse	Ken Thompson introduced a Trojan Horse in a C compiler that compiled itself into future programs [Tho84].

A Growing Problem

Complex devices, by their very nature, introduce the risk that malicious functionality may be added (either during creation or afterwards) that extends the original device past its primary intended design. An unfortunate side effect of inherent complexity is that it allows malicious subsystems to remain invisible to unsuspecting users until it is too late. Some of the earliest malicious functionality, for example, was associated with complicated copy machines. Extensible systems, including computers, are particularly susceptible to the malicious functionality problem.

When extending a system is as easy as writing and installing a program, the risk of intentional introduction of malicious behavior increases drastically.

Any computing system is susceptible to malicious code. Rogue programmers may modify systems software that is initially installed on the machine. Users may unwittingly propagate a virus by installing new programs or software updates from a CDROM. In a multi-user system, a hostile user may install a Trojan Horse to collect other users' passwords. These attack vectors have been well known since the dawn of computing, so why is malicious code a bigger problem now than in the past? We argue that a small number of trends have a large influence on the recent wide spread propagation of malicious code.

Networks are Everywhere: The growing connectivity of computers through the Internet has increased both the number of attack vectors, and the ease with which an attack can be made. More and more computers, ranging from home PCs to systems that control critical infrastructures (e.g., the power grid), are being connected to the Internet. Furthermore, people, businesses, and governments are increasingly dependent upon network-enabled communication such as e-mail or Web pages provided by information systems. Unfortunately, as these systems are connected to the Internet, they become vulnerable to attacks from distant sources. Put simply, it is no longer the case that an attacker needs physical access to a system to install or propagate malicious code.

Because access through a network does not require human intervention, launching automated attacks from the comfort of your living room is relatively easy. Indeed, the recent denial-of-service attacks in February of 2000 took advantage of a number of (previously compromised) hosts to flood popular e-commerce Web sites with bogus requests automatically. The ubiquity of networking means that there are more systems to attack, more attacks, and greater risks from malicious code than in the past.

System Complexity is Rising: A second trend that has enabled widespread propagation of malicious code is the size and complexity of modern information systems. A desktop system running Windows/NT and associated applications depends upon the proper functioning of the kernel as well as the applications to ensure that malicious code cannot corrupt the system. However, NT itself consists of tens of millions of lines of code, and applications are becoming equally, if not more, complex. When systems become this large, bugs cannot be avoided. This problem is exacerbated by the use of unsafe programming languages (e.g., C or C++) that do not protect against simple kinds of attacks, such as buffer overflows. However, even if the systems and applications code were bug free, improper configuration by retailers, administrators, or users can open the door to malicious code. In addition to providing more avenues for attack, complex systems make it easier to hide or mask malicious code. In theory, we could analyze and prove that a small program was free of malicious code, but this task is impossible for even the simplest desktop systems today, much less the enterprisewide systems used by businesses or governments.

Systems are Easily Extensible: A third trend enabling malicious code is the degree to which systems have become *extensible*. An extensible host accepts updates or extensions, sometimes referred to as *mobile code*, so that the functionality of the system can be evolved in an incremental fashion. For example, the plug-in architecture of Web browsers makes it easy to install viewer extensions for new document types as needed. Today's operating systems support extensibility through dynamically-loadable device drivers and modules. Today's applications, such as word-processors, e-mail clients, spreadsheets, and Web-browsers support extensibility through scripting, controls, components, and applets. From an economic standpoint, extensible systems are attractive because they provide flexible interfaces that can be adapted through new components. In today's marketplace, it is crucial that software be deployed as rapidly as possible in order to gain market share. Yet the marketplace also demands that applications provide new features with each release. An extensible architecture makes it easy to satisfy both demands by

allowing the base application code to be shipped early, and by later shipping feature extensions as needed.

Unfortunately, the very nature of extensible systems makes it hard to prevent malicious code from slipping in as an unwanted extension. For example, the Melissa virus took advantage of the scripting extensions of Microsoft's Outlook e-mail client to propagate itself. The virus was coded as a script contained in what appeared to users as an innocuous mail message. When the message was opened, the script was executed, and proceeded to obtain email addresses from the user's contacts database, and then sent copies of itself to those addresses. The infamous Love Bug worked very similarly, also taking advantage of Outlook's scripting capabilities.

Defense against Malicious Code

Creating malicious code is not hard. In fact, it is as simple as writing a program or downloading and configuring a set of easily customized components. It is becoming increasingly easy to hide ill-intentioned code inside otherwise innocuous objects, including Web pages and email messages. This makes detecting and stopping malicious code before it can do any damage extremely hard.

To make matters worse, our traditional tools for ensuring the security and integrity of hosts have not kept pace with the ever-changing suite of applications. For example, traditional security mechanisms for access control reside within an operating system kernel and protect relatively primitive objects (e.g., files); but increasingly, attacks such as the Melissa virus happen at the application level where the kernel has no opportunity to intervene.

A useful analogy is to think of the computer and network security mechanisms of today like the walls, moats, and drawbridges of medieval times. At one point, these mechanisms were effective for defending our computing castles against isolated attacks, mounted on horseback. But the defenses have not kept pace with the attacks. Today, attackers have access to airplanes and laser-guided bombs that can easily bypass our antiquated defenses. In fact, attackers rarely need sophisticated equipment: because our kingdoms are really composed of hundreds of interconnected castles, attackers can easily move from site to site, finding places where we have left the drawbridge down. It is time to develop some new defenses.

In general, when a computational agent arrives at a host, there are four approaches that the host can take to protect itself:

1. **Analyze** the code and reject it if there is the potential that executing it will cause harm.
2. **Rewrite** the code before executing it so that it can do no harm.
3. **Monitor** the code while its executing and stop it before it does harm, or
4. **Audit** the code during executing and take policing action if it did some harm.

Today, software analysis is most likely to occur during the development cycle when source code is available. Even state-of-the-art software analysis techniques are nowhere near perfect, however. Software engineers working on their own systems often neglect to apply any analysis since the payoff appears to be limited, and development cycles are compressed by Internet time. Automated tools such as the open source security scanner ITS4 (see <http://www.rstcorp.com/its4>) may help. Primitive code analysis, including looking for particular patterns of system calls in an executable, has been incorporated into some commercially available security products. In the end, determining if an arbitrary piece of code will behave maliciously in advance of executing it is as difficult as the Halting problem. Thus, there will never be a complete solution.

Code rewriting is a less pervasive approach to the problem, but may become more important (see the next section). With this approach, a rewriting tool inserts extra code to perform dynamic checks that ensure bad things cannot happen. For example, a Java compiler inserts code to check that each array index is in bounds—if not, the code throws an exception, thereby avoiding

the common class of buffer overrun attacks. Rewriting can be carried out either at the application code level, or below that in subsystem functionality made available through APIs.

Monitoring programs, using a reference monitor, is the traditional mechanism used to ensure programs do not do anything bad. For instance, an operating system uses the page translation hardware to monitor the set of addresses that an application attempts to read, write, or execute. If the application attempts to access memory outside of its address space, then the kernel takes action (e.g., by signaling a segmentation fault.) A more recent example of an on-line reference monitor is the Java Virtual Machine interpreter. The interpreter monitors execution of applets and mediates access to system calls by examining the execution stack to determine who is issuing the system calls request. In this case, stack inspection is used as a policy enforcement mechanism.

If malicious code does damage, recovery is only possible if the damage can be properly assessed and addressed. Creating an audit trail that captures program behavior is an essential step. Several program auditing tools are commercially available.

Each of the basic approaches, analysis, rewriting, monitoring, and auditing, has its strengths and weaknesses, but fortunately, these approaches are not mutually exclusive and may be used in concert. Of course, to employ any of them, we must first identify what could be “harmful” to a host. Like any other computing task, we must turn the vague idea of “harm” into a concrete, detailed specification—a security policy—so that it can be enforced by some automated security architecture. Therein lies our greatest danger, for as we create the policy, we are likely to abstract or forget relevant details of the system. An attacker will turn to these details first, stepping outside our policy model to circumvent the safeguards.

Stick to Your Principles

To protect against this common failing, it is important to follow well-established security principles when designing security policies. One of the most important principles, first stated by Saltzer and Schroeder in 1975 [SS75], is the *Principle of Least Privilege*: a component should be given the minimum access necessary to accomplish its intended task. For example, we shouldn't give a program access to all files in a system but rather, only those files that the program needs to get its job done. This prevents the program from either accidentally or maliciously deleting or corrupting most files. Obviously, the fewer files that the program can access, the less the potential damage. Stated simply, tighter constraints on a program lead to better security.

Another important security principle is the *Principle of Minimum Trusted Computing Base*. The trusted computing base (TCB) is the set of hardware and software components that make up our security enforcement mechanisms. The Principle of Minimum TCB states that, in general, the best way to assure that your system is secure is to keep your TCB small and simple. Even in the mid 70's, operating system kernels were thought to be too large to be trusted. Those systems now seem small and tightly structured compared to today's widely used kernels composed of millions of lines of code.

In the next section, we give examples of currently deployed defenses for malicious code, focusing on their relative pros and cons. Unfortunately, the comparison shows that the pros are outweighed by the cons, largely because of a violation of the Least Privilege and Minimal TCB principles. We then discuss some promising technologies, identified by the research group, that are emerging from research labs.

Current Defenses

OS-Based Reference Monitors

Historically, mechanisms for security policy enforcement have been provided by the computer hardware and operating system. Address translation hardware, distinct supervisor and user-modes, timer interrupts, and system calls for invoking a trusted software base are used in combination to enforce limited forms of availability, fault containment, and authorization properties.

To a large degree, these mechanisms have proven effective for protecting operating system resources (e.g., files or devices) from unauthorized access by humans or malicious code. But the mechanisms work with a fixed system-call interface and a fixed vocabulary of principals, objects, and operations for policies. Only by incurring significant cost and usability penalties can that vocabulary be expanded. It rarely is. Currently, most desktop machines are configured as single-user so applications have complete access to the machine resources.

Scanning for Known Malicious Code

In the days before networking was rampant, malicious code mostly used the “sneaker net” as its vector. Viruses were spread from machine to machine by humans carrying floppy disks with infected programs on them. Perhaps the built-in limitations in the vector kept the number of viruses small. In any case, the limited number of viruses combined with the inefficiencies in the communication vector made possible the strategy of black listing.

Most commercial anti-virus products make use of a black listing strategy to this day. They rely on databases of virus “signatures” that are consulted when a new program arrives. Antivirus tools scan disks and sometimes e-mail looking for known viruses.

The limitations of this approach are obvious. Unknown malicious code will easily get by the simple defenses to carry out its dirty work. Until a new virus is contained by researchers and a signature entry is added to the database, it can run rampant. Recall both the Melissa virus and the Love Bug.

It should be clear that black listing by itself does not provide adequate security; in fact, all by itself, black listing is darn close to useless. It's too easy to make trivial changes to malicious code (a process that can be automated in the code itself) to thwart almost every black listing scheme.

Nevertheless, black listing is so cheap to implement that we believe it makes sense to do it anyway. Using a system that blacklists known malicious code may be worthwhile even if it only stops the occasional naïve attack.

Code Signing

Code signing as it is commonly used and (mis)understood today needs work. The idea itself is elegant and simple: a private key is used to sign code, both ensuring transmission integrity and enabling policy defined by trust in the signer. Unfortunately, a pervasive and common myth is that code signing signifies authorship or goodness. It does not. Encountering a piece of code that is cryptographically signed simply means that some private key was used to sign the code! From this we can reason about endorsement of the code by the person or organization who controls the private key in question.

In talking about code signing, many people make bad assumptions. These include: assuming that signed code is safe, treating signing as a binary indicator of goodness, assuming that goodness is compositional, and thinking that code that has been shown to be good in one environment will be good in all environments. Code signing is a useful technology, but these limitations are real.

The adoption of code signing has also been hampered by the lack of a Public Key Infrastructure (PKI). Very few PKI installations have been deployed, and those that have do not begin to approach Internet scale. Without a solid PKI, code signing will not become common.

Promising New Defenses

Software-Based Reference Monitors

Wahbe et al. suggested *software-based fault isolation* (SFI) as an alternative to the traditional hardware-based mechanisms used to ensure memory safety [WLAG93]. Their goal was to reduce the overhead of cross-domain procedure calls and providing a more-flexible memory-safety mode. Their basic idea is to rewrite binary code by inserting checks on each memory access and each control transfer to ensure that those accesses are valid. Schneider generalized the SFI idea to in-lined reference monitors (IRM) [Sch00]. With the IRM approach, a security policy is specified in a declarative language, and a general-purpose tool rewrites code, inserting extra checks and state that are used to enforce the policy. In principle, any security policy that is a safety property can be enforced, so the approach is quite powerful. For example, it can enforce any discretionary access control policy. The approach is also practical: Prototypes have been built at both Cornell and MIT [ES99,ET99,ES00]. One of the Cornell prototypes, PSLang/PoET, works for the Java Virtual Machine language and gives competitive performance for the implementation of Java's stack inspection security policy.

Type-Safe Languages

Type-safe programming languages, such as Java, Scheme, or ML, ensure that operations are only applied to values of the appropriate type. Type systems that support type abstraction allow programmers to specify new, abstract types and signatures for operations that prevent unauthorized code from applying the wrong operations to the wrong values. In this respect, type systems, like software-based reference monitors, go beyond operating systems in that they may be used to enforce a wider class of application-specific access policies. Static type systems have an additional attractive property: enforcement can be done offline through static type checking instead of each time a particular operation is performed. This allows the type checker to enforce certain policies that are difficult with on-line techniques. For example, Myers' JFlow [Mye99] extends the Java type system to enforce the policy that high-security data should never be leaked. Current research in type systems is aimed at eliminating more runtime checks (e.g., array bounds checks) or type-checking machine code (see for example [MWCG98]).

Proof-Carrying Code

Proof-carrying code (PCC), a concept introduced by Necula and Lee [NL96], is a promising approach for gaining high assurance of security in systems. The basic idea is to require any untrusted code to come equipped with an explicit, machine-checkable *proof* that the code respects a given security policy. Before executing the code, we simply verify that the proof is valid with respect to both the code and the policy. Because proof checkers can be quite simple (Necula's is about 6 pages of C code), it is easier to ensure that they are correct. And in principle, PCC can enforce any security policy—not just type safety—as long as the code producer can construct a proof. Necula and Lee have shown that such proofs can be constructed automatically for standard type-safety policies, if the code is generated by a compiler for a type-safe programming language. Unfortunately, going beyond standard notions of type safety cannot be performed automatically without either restricting the code or requiring human intervention. It is unlikely that programmers will construct explicit proofs. Thus an active area of research is how to integrate compilers and modern theorem provers to produce PCC.

Policy as Achilles' Heel

Thus far we have focused on technology solutions to the malicious code problem. To be sure, technology can be of service; but there is another critical aspect of the problem that remains to be addressed—the problem of policy.

In current forms, extensible systems do little to determine how a system will behave when extended in certain ways or, put another way, what a particular piece of code can and cannot do. In fact, today's computers are hyper-malleable and overly complicated. This greatly increases the malicious code risk. In the end, determining whether something malicious is happening requires first defining some policy to enforce.

When Policy Breaks Down

Clearly, the notion of policy is deeply intertwined with the concept of malicious code. Understood in terms of policy, the root causes of malicious code can be separated into two basic categories: 1) bad policy, and 2) incorrectly enforced policy.

Bad policy allows malicious code to do something malicious since policy does not forbid it. Even if policy is perfectly enforced by technology, the policy itself has to be well formed. Subcategories of bad policy include:

- misunderstandings of context, whereby policy makes no sense in the context where it was applied;
- inconsistency, whereby the policy is self-contradictory; or
- non-comprehensiveness, whereby policy fails to cover some situation or exists at the wrong level of abstraction.

Incorrect policy enforcement allows code to do something malicious even if it is correctly forbidden by policy. In this case, correct technology-driven enforcement falls prey to poor policy creation and management. Subcategories of incorrect policy enforcement include:

- incorrect enforcement of safety policies;
- incorrect enforcement of liveness properties;
- incorrect enforcement of information flow.

Other subcategories may exist under incorrect enforcement as well.

Table 2 provides examples of malicious code understood in our policy-based framework.

BAD POLICY	Examples	INCORRECT POLICY ENFORCEMENT	Examples
Context misunderstood	<ul style="list-style-type: none"> • <i>Melissa (e-mail worms)</i> • <i>Morris worm (sendmail debug mode)</i> 	Safety properties	<ul style="list-style-type: none"> • <i>Thompson compiler trick</i> • <i>buffer overflows</i> • <i>guessable passwords</i>
Inconsistent	<ul style="list-style-type: none"> • <i>overly restrictive policy</i> 	Liveness properties	<ul style="list-style-type: none"> • <i>denial of service</i>
Non-comprehensive	<ul style="list-style-type: none"> • <i>Melissa (e-mail worms)</i> • <i>guessable passwords</i> 	Information flow	<ul style="list-style-type: none"> • <i>Javascript privacy hacks</i>

E-mail worms like Melissa fit into the *context misunderstood* box above because they are caused by the interaction between individual policy decisions made about separate parts of the system. Useful subsystems such as Javascript interpreters can cause problems if invoked in the incorrect context. For example, modern e-mail systems often include the ability to execute potentially-dangerous untrusted mobile code by default. This opens the door to malicious code.

E-mail worms like Melissa and the Love Bug also fit in the *non-comprehensive* box since policy often fails to cover systems like e-mail based Visual Basic code, which can be used maliciously.

The sendmail debug mode problem exploited in the Morris worm provides a good example of a *context misunderstanding* since something that made sense during development and installation (a debug capability) makes little sense in a fielded system. The added functionality came at the price of unnecessary risk.

Guessable passwords are a good example of *non-comprehensive* policy if they come about due to a lack of restrictions on user password choice. Without guidance, users tend not to behave in a secure fashion. Guessable passwords are a notorious security risk that is widely exploited by malicious code.

Thompson's compiler trick (a famous Trojan Horse to be built into the C compiler that made use of the login program [Tho84]) is an example of incorrect enforcement of *safety policy*. In this case, implicit policy assumes both that the compiler properly produces object code from source code and that the login program requires a correct password to be entered. In fact, the developer of the compiler can circumvent this implicit policy (which is not enforced technologically).

Distributed denial of service (dDoS) attacks are a clear example of *liveness policies* being incorrectly enforced. In this case, massive amounts of traffic are used to overwhelm the processing capabilities of a commerce server.

Addressing the malicious code problem requires the creation of sound policy and its enforcement through technology.

The Many Levels of Policy

System administrators and MIS security people think about policy in terms of user groups, firewall rules, and computer use. Security researchers steeped in programming languages think about policy in terms of memory safety and liveness properties. Government policy wonks think about policy in terms of rules and regulations imposed on users and systems. The problem is, all of these ways of thinking about policy are equally valid!

So how are we to set policy to combat malicious code? We believe the key is to focus on defining meta-level policies that system administrators work with naturally in terms of collections of lower-level enforcement mechanisms. This is no trivial undertaking.

Most of the technologies explored earlier in this article can be used to enforce particular aspects of software behavior. Many languages researchers, for example, consider the code safety problem "solved". Liveness and information flow properties are harder, but fairly clear research agendas exist to address the open issues. Of course, the terms *safety*, *liveness*, and *information flow* have technical meanings. Intuitively, a safety property states that a program will never perform a bad action, for some precisely defined notion of "bad". An example of a bad action is overflowing a buffer. A liveness property, on the other hand, states that a program will eventually perform some desired action or set of actions. For example, the property that a program will eventually release all of the memory that it allocates is a liveness property. Finally, information flow properties state that certain values or types of values will not be discernable to certain observers.

The problem is that low-level properties such as safety and liveness do not align nicely with what most security administrators think of as policy building blocks. Thus an open question is how to express reasonable security policy that can be directly transformed into technology enforcement solutions.

The answer is to understand policy as a layered set of abstractions. Some preliminary work exists (for example Netscape Navigator's approach to policy sets based on expected code behavior), but much work remains to be done.

Final Word

The malicious code problem will continue to grow as the Internet grows. The constantly accelerating trends of interconnectedness, complexity, and extensibility make addressing the problem more urgent than ever. As extensible information systems become more ubiquitous, moving into everyday devices and playing key roles in life-critical systems, the level of the threat moves out of the technical world and into the real world. We *must* work on this problem.

Our best hope in combating malicious code is creating sound policy about software behavior and enforcing that policy through the use of technology. An emphasis on one or the other alone will do little to help. Any answer will require a set of enforcement technologies that can be directly tied to policy set and understood by non-technical users.

References

[ET99] Evans D., and A. Twyman. Policy-directed code safety. In *Proceedings of the 1999 IEEE Symposium on Security and Privacy*, Oakland, CA, May, 1999. See also <http://www.cs.virginia.edu/~evans>.

[ES99] Erlingsson, U. and F.B. Schneider. SASI enforcement of security policies: a retrospective. In *Proceedings of the New Security Paradigms Workshop*, Ontario, Canada, September 1999.

[ES00] Erlingsson, U. and F.B. Schneider. IRM enforcement of Java stack inspection. In *IEEE Symposium on Security and Privacy*, Oakland, CA, May 2000, To appear.

[MF99] McGraw, G. and E. Felten. (1999) *Securing Java: Getting down to business with mobile code*. Wiley. Complete Web edition at <http://www.securingsjava.com>.

[Mye99] Myers, A.C. JFlow: Practical mostly-static information flow control. In *Proceedings of the 26th ACM Symposium on Principles of Programming Languages (POPL)*, San Antonio, TX, January 1999.

[MWCG98] Morrisett, G. D. Walker, K. Crary, and N. Glew. From System-F to typed assembly language. *ACM Transactions on Programming Languages and Systems*, 21(3):528-569, May 1999. See also <http://www.cs.cornell.edu/talc>.

[Nec97] Necula, G.C. Proof-carrying code. In *Proceedings of the 24th ACM Symposium on Principles of Programming Languages (POPL)*, pages 106-119, Paris, France, January 1997. See also http://www-nt.cs.berkeley.edu/home/necula/public_html/pcc.html

[Sch00] Schneider, Fred B. Enforceable Security Policies. *ACM Transactions on Information and System Security*, 2(4), March 2000.

[SESS96] Seltzer, M., Y. Endo, C. Small and K. Smith. Dealing with disaster: surviving misbehaved kernel extensions. In *Proceedings of the USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, pages 213-227, Seattle, WA, October 1996.

[SS75] Salzter, J.H., and M.D. Schroeder. The protection of information in computer systems. *Proceedings of the IEEE*, 9(63), September 1975.

[Tho84] Thompson, K. (1984) Reflections on Trusting Trust. *Communications of the ACM* 27(8). August 1984.

[WLAG93] Wahbe, R., S. Lucco, T. Anderson, and S. Graham. Efficient software-based fault isolation. In *Proc. 14th ACM Symposium on Operating System Principles (SOSP)*, pages 203-216, Asheville, NC, December, 1993.

Bios

Gary McGraw

Gary McGraw is the Vice President of Corporate Technology at Reliable Software Technologies where he pursues research in software security while leading the Software Security Group. He has served as principal investigator on grants from Air Force Research Labs, DARPA, National Science Foundation, and NIST's Advanced Technology Program. He chairs the National Infosec Research Council's Malicious Code Infosec Science and Technology Study Group. Dr. McGraw is a noted authority on mobile code security and co-authored both *Java Security* (Wiley, 1996) and *Securing Java* (Wiley, 1999). Dr. McGraw is currently writing a book entitled *Software Security for Developers* (2001).

Greg Morrisett

Greg Morrisett is an assistant professor in the Computer Science department at Cornell University. He is a Sloan Research Fellow, recipient of an NSF Career award, member of the IFIP Working Group 2.8 (Functional Programming), editor for the *Journal of Functional Programming*, and an associate editor for *ACM Transactions on Programming Languages and Systems*.