

DIMACS Center
Rutgers University

**Workshop on Software Security
January 6 - 7, 2003:
Report to the National Science Foundation
under Grant 0302708**

Report Date: March 27, 2003

DIMACS Center, CoRE Bldg. • Rutgers, The State University of New Jersey
96 Frelinghuysen Road • Piscataway, NJ 08854-8018 • USA
TEL: 732-445-5928 • FAX: 732-445-5932 • EMAIL: center@dimacs.rutgers.edu
Web: <http://dimacs.rutgers.edu/>

DIMACS is a partnership of Rutgers University, Princeton University, AT&T Labs - Research,
Bell Laboratories, the NEC Laboratories America and Telcordia Technologies.
Affiliate Members: Avaya Labs, IBM Research and Microsoft Research

The DIMACS Workshop on Software Security was held at the DIMACS Center on January 6 - 7, 2003. It was organized by Gary E. McGraw (Cigital, Chair), Edward W. Felten (Princeton University), Virgil D. Gligor (University of Maryland), and David Wagner (UC Berkeley).

The security of computer systems and networks has become increasingly limited by the quality and security of the software running on these machines. Researchers have estimated that more than half of all vulnerabilities are due to buffer overruns, an embarrassingly elementary class of bugs. All too often systems are hacked by exploiting software bugs. In short, a central and critical aspect of the security problem is a software problem. How can we deal with this?

The Software Security Workshop explored these issues. The scope of the workshop included security engineering, architecture and implementation risks, security analysis, mobile and malicious code, education and training, and open research issues. In recent years many promising techniques have arisen from connections between computer security, programming languages, and software engineering, and one goal was to bring these communities closer together and crystalize the subfield of software security.

A summary report, scheduled to appear in IEEE S&P issue, is Appendix 1. Appendix 2 provides a list of participants, Appendix 3 a program, Appendix 4 links to workshop presentations that are on the web, and Appendix 5 a booklet of talk abstracts. Further information can be found at the workshop website: <http://dimacs.rutgers.edu/Workshops/Software/>.

Appendix 1
Summary Workshop Report

From the Ground Up: The DIMACS Software Security Workshop

Computer system and network security is increasingly limited by the quality and security of the software running on constituent machines. Researchers estimate that more than half of all vulnerabilities are from buffer overruns, an embarrassingly elementary class of bugs.¹

life cycle, knowing and understanding common threats (including language-based flaws and pitfalls), designing for security, and subjecting all software artifacts to thorough objective risk analyses and testing.

By any measure, security holes in software are common, and the problem is growing. Figure 1 shows the number of security-related software vulnerabilities reported to CERT/CC over the past several years. The evidence points to a clear and pressing need for a disciplined approach to software security.

Before we continue, I need to establish some definitions. The “Software security definition” sidebar defines four of the most important terms we use when we discuss software security. I included them to reinforce the areas that each covers.

The Trinity of Trouble

Most modern computing systems are susceptible to software security problems. So, why is software security a bigger problem now than in the past? Three trends—a *trinity of trouble*—have heavily influenced the problem’s growth and evolution. (Interestingly, these three general trends also are responsible for the rise in malicious code.⁵)

Reliance on networked devices

First, growing Internet computer connectivity has increased attack-vector numbers and made it easier for hackers to launch attacks. This puts software at greater risk. We are connecting ever more computers, ranging from home PCs to systems that control critical infrastructures

GARY
MCGRAW
Cigital

Worse, more complex problems such as race conditions and subtle design errors wait in the wings for the buffer overflow’s demise. Software security problems will be with us for years, and hackers will continue to exploit systems via software defects. Clearly, a central and critical aspect of the computer security problem resides in software.

Software security—the idea of engineering software that continues to function correctly under malicious attack—is not new, but it is receiving renewed interest as reactive network-based security approaches such as firewalls prove to be ineffective. Unfortunately, today’s software is riddled with both design flaws and implementation bugs, which result in unacceptable security risks.

As security researcher Steve Bellovin puts it, “any program, no matter how innocuous it seems, can harbor security holes.” This notion is common knowledge, and yet developers, architects, and computer scientists only recently began systematically to study how to build secure software. In fact, the first books on software security and security engineering were published in 2001,²⁻⁴ and a body of literature on the subject has just begun to emerge. (See

Further reading sidebar.)

The DIMACS Software Security Workshop held in New Jersey this January (<http://dimacs.rutgers.edu/Workshops/Software>) explored issues such as security engineering, architecture and implementation risks, security analysis, mobile and malicious code, education and training, and open research issues. Many promising techniques have grown from connections between computer security, programming languages, and software engineering, and one workshop goal was to bring these communities closer together to crystallize the software security subfield. I report the workshop’s results in this article.

A new security paradigm

Internet-enabled software applications—especially custom applications—present the most common security risks we encounter today and are the targets of choice for malicious hackers. Software security is about understanding software-induced security risks and how to manage them. Good software security practice leverages good software engineering practices: thinking about security early in the software’s

(for example, the power grid), to enterprise networks and to the Internet. Furthermore, people, businesses, and governments increasingly depend on network-enabled communication such as email or Web pages provided by information systems. Things that used to happen offline now happen online.

Unfortunately, as these systems connect to the Internet, they become vulnerable to software-based attacks from distant sources. Attackers no longer need physical access to systems to exploit vulnerable software and shut down banking services and airlines (as shown by January's SQL Slammer worm).

Because access through a network does not require human intervention, launching automated attacks is relatively easy. The ubiquity of networking means more software systems to attack, more attacks, and greater risks from poor software security practice than in the past.

Easily extensible systems

A second trend negatively affecting software security is how extensible systems are. Extensible systems accept updates or extensions—sometimes referred to as mobile code—to evolve system functionality incrementally.⁶ The plug-in architecture of Web browsers, for example, 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, and applications such as word-processors, email 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, vendors must deploy software as rapidly as possible to gain market share. Yet the marketplace also demands that applications provide new features with each




Figure 1: Security-related software vulnerabilities reported to CERT/CC over several years. This figure shows that the problem is growing, however, we don't know the exact reasons for this growth (for example, more defective code, better hacking, more code, and so on).

release. An extensible architecture makes it easy to satisfy both demands by letting the base application code ship early and then shipping feature extensions later, as needed.

Unfortunately, the very nature of extensible systems makes it hard to prevent software vulnerabilities from slipping in as unwanted extensions. Advanced languages and platforms including Sun Microsystem's Java and Microsoft's .NET Framework are making extensibility commonplace.

Increasingly complex systems

A third trend is the unbridled growth of modern information systems' size and complexities, especially software systems. A desktop system running Windows/XP and associated applications depends on the proper functioning of the kernel as well as the applications to ensure that vulnerabilities won't compromise the system. However, XP consists of at least 40 million lines of code, and end-user applications are becoming at least as complex.

When systems become this large, bugs cannot be avoided. Figure 2 shows how Windows' complexity (measured in lines of code) has grown over the years. The point of

the graph is to emphasize the growth rate rather than the numbers. Note that the flaw rate tends to progress as the square of code size. Other factors that deeply affect complexity include whether the code is tightly integrated, the overlay of patches and other post-deployment fixes, and critical architectural issues.

The complexity problem is exacerbated by the use of unsafe programming languages such as C or C++, which do not protect against simple kinds of attacks, such as buffer overflows. In theory, we should be able to analyze and prove that a small program has no problems, but this task is impossible for even the simplest desktop systems, much less the enterprise-wide systems used by businesses or governments.

Proactive Security

Software security follows naturally from software engineering, programming languages, and security engineering. It fortifies computer security by identifying and expunging problems in the software itself by building software that proactively resists attack.

For decades, computer security literature has discussed some aspects

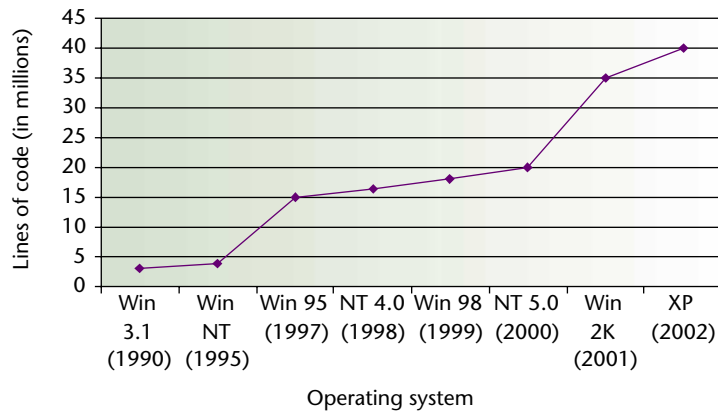


Figure 2. Growth of Microsoft's operating system code base from 1990 to 2002. These numbers reflect lines of code for all aspects of Windows, including device drivers. The striking growth rate is apparent.

of software security but only recently has software security been recognized as a clearly defined subdiscipline. It's still in its infancy.

Architecture versus implementation

Though Figure 1 shows that software vulnerability is growing, scientists have done little work to classify and categorize software security problems.

Security vulnerabilities in software systems range from local implementation errors (for example, use of the `gets()` function call in C and C++), through interprocedural interface errors (for example, a race condition between an access control check and a file operation), to much higher design-level mistakes (such as error-handling and recovery systems that fail insecurely or object-sharing systems that mistakenly include transitive trust issues).

These vulnerabilities define a large range based on how much program code we must consider to understand the vulnerability, how much detail regarding the execution environment we must know to understand the vulnerability, and whether a design-level description is best for

determining whether a given vulnerability is present. For example, we can determine that a call to `gets()` in a C and C++ program can be exploited in a buffer overflow attack without knowing anything about the rest of the code, its design, or the execution environment other than assuming that the user entering text on standard input might be malicious. Hence, we can detect a `gets()` vulnerability with good precision using a simple lexical analysis.

Mid-range vulnerabilities involve interactions among more than one code point. For example, precisely detecting race conditions depends on more than just analyzing an isolated line of code—it could depend on knowing about the behavior of several functions, understanding sharing among global variables, and knowledge of the operating system providing the execution environment.

Design-level vulnerabilities extend this trend further. Ascertaining whether a program has design-level vulnerabilities requires great expertise, which makes design-level flaw identification hard to do and particularly difficult to automate. Unfortunately, design-level problems ap-

pear to be prevalent and are at the very least a critical category of security risk in code. During the workshop, Microsoft reported that over 50 percent of the problems it encountered during its security push were architectural in nature.

As an example, let's consider an error-handling and recovery system. Failure recovery is an essential aspect of security engineering, but it is complicated because it interacts with failure models, redundant design, and defense against denial of service. Understanding whether an error-handling and recovery system in an object-oriented program is secure involves ascertaining a global property spread throughout many classes in typical design. Error-detection code usually is included in each object and method, and error-handling code usually is separate and distinct from the detection code. Sometimes exceptions propagate up to the system level and the machine running the code handles them (for example, Java 2 VM exception handling), which makes determining whether a given error-handling and recovery design is secure quite difficult. This problem is exacerbated in transaction-based systems commonly used in e-commerce solutions where functionality is distributed among many different components running on several servers.

Other examples of design-level problems include object-sharing and trust issues, unprotected data channels (both internal and external), incorrect or missing access control mechanisms, lack of auditing/logging or incorrect logging; ordering and timing errors (especially in multithreaded systems), and many others. To make forward progress as a scientific discipline, software security must rigorously understand and categorize problems like these.

Building secure software is like building a house. We liken correct low-level coding (such as checking array references) to using refractory bricks instead of bricks made of saw-

Software security definitions

The reintroduction of basic terminology with a security emphasis can help clarify things when trying to categorize problems. I propose the following usage:

- **Defects.** Implementation and design vulnerabilities are defects. A defect is a problem that may lie dormant in software for years and then surface in a fielded system with major consequence.
- **Bug.** A bug is an implementation-level software problem. Bugs might exist in code but never execute. Though many software practitioners apply the term “bug” quite generally, I reserve use of the term to encompass fairly simple implementation flaws. Bugs easily can be discovered and remedied. Researchers have made significant progress in detecting security vulnerabilities stemming from low-level and mid-level implementation bugs. Tools include FIST,¹ ITS4,² Jscan,³ Splint,⁴ Metal,⁵ and PreFix/PreFast.⁶ These tools detect a wide range of low-level implementation bugs including buffer overflow vulnerabilities, format string bugs, resource leaks, and simple race conditions, all of which depend on only limited code analysis and knowledge of the external environment.
- **Flaw.** A flaw is a problem at a deeper level. Often flaws are much more subtle than simply an off-by-one error in an array reference or use of an incorrect system call. A flaw certainly is instantiated in software code, but is also present (or absent!) at the design level. For example, a number of classic flaws exist in error-handling and recovery systems that fail in an insecure or inefficient fashion. Automated

technologies to detect design-level flaws do not yet exist.

- **Risk.** Flaws and bugs lead to risk. Risks are not failures. Risks capture the probability that a flaw or a bug will impact the purpose of the software. Risk measures also take into account the potential damage that can occur. A very high risk is not only likely to happen, but also likely to cause great harm. Risks can be managed by technical and non-technical means.

References

1. A. Ghosh, T. O'Connor, and G. McGraw, “An Automated Approach for Identifying Potential Vulnerabilities in Software,” *Proc. IEEE Symp. on Security and Privacy*, IEEE CS Press 1998, pp. 104–114.
2. J. Viega et al., “ITS4: A Static Vulnerability Scanner for C and C++ Code” *Proc. Ann. Computer Security Apps. Conf.*, 2000.
3. J. Viega et al., “Statically Scanning Java Code: Finding Security Vulnerabilities,” *IEEE Software*, vol. 17, Sept./Oct. 2000, pp. 68–74.
4. D. Evans et al., “LCLint: A Tool for Using Specifications to Check Code,” *Proc. SIGSOFT Symp. Foundations of Software Eng*, ACM, 1994, pp. 87–96.
5. D. Engler et al., “Checking System Rules Using System-Specific, Programmer-Written Compiler Extensions,” *Proc. Symp. OS Design & Implementation (OSDI)*, Usenix, 2000, pp.
6. W. R. Bush, J.D. Pincus, and D.J. Sielaff, “A Static Analyzer for Finding Dynamic Programming Errors,” *Software Practice and Experience*, vol. 30, June 2000, pp.775–802.

dust. The kinds of bricks that we use are important to the house’s integrity, but even more important (if the goal is to keep bad things out) is having four walls and a roof in the design. The same thing goes for software: what system calls are used and how they are used is important, but overall design properties often count for more. In general, software security research has paid much more attention to bricks than to walls.

Software security areas of interest

Practitioners and scientists focus on the following areas of interest in software security:

- Reconciling security goals and software goals, software quality management in commercial practice
- Security requirements engineering
- Design for security, software archi-

itecture, and architectural analysis

- Security analysis, security testing, and the use of the Common Criteria
- Guiding principles for software security, case studies in design and analysis, and pedagogical approaches to teaching security architecture
- Software security education; educating students and commercial developers
- Auditing software; implementation risks, architectural risks, automated tools, and technology developments (such as code scanning, information flow, and so on)
- Common implementation risks: buffer overflows, race conditions, randomness, authentication systems, access control, applied cryptography, and trust management
- Application security; protecting code post production
- Survivability and penetration resis-

tance, type safety, and dynamic policy enforcement

- Denial-of-service protection for concurrent software
- Penetrate and patch as an approach to securing software
- Code obfuscation and digital content protection
- Malicious code detection and analysis

Much work remains to be done in each of these areas, but some basic practical solutions are becoming available in the market. Examples include using simple code-scanning tools during implementation review and reactive “application firewalls.”

Best practices and the state of the art

As practitioners become aware of software security’s importance, they adopt and evolve a set of best prac-

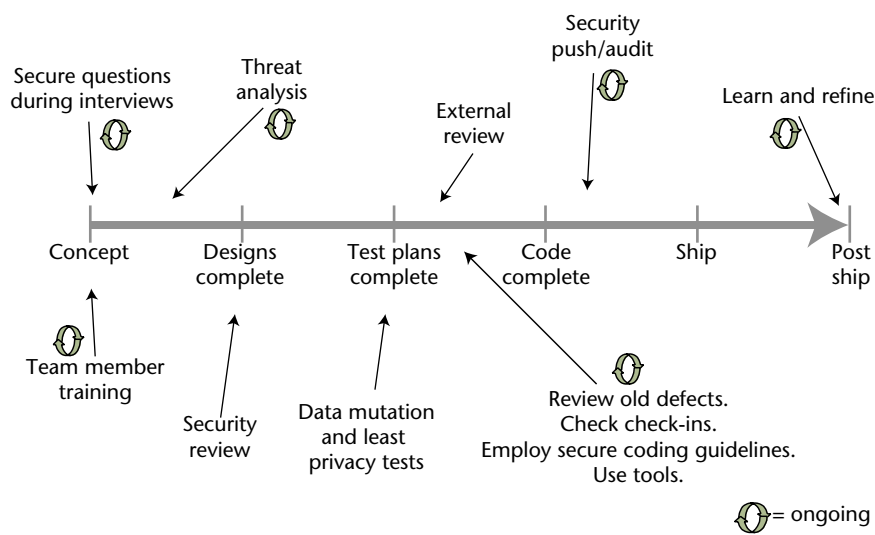


Figure 3. Microsoft has put the following software security process into place. Notice that security does not happen at one life cycle stage, nor are constituent activities “fire and forget.” The idea of constantly revisiting analyses at all levels is critical to software security.

tices to address the problem. Common approaches in practice today encompass training for developers, testers, and architects; analysis and auditing of software artifacts; and security engineering. Later we’ll see that Microsoft has carried out one noteworthy effort under the rubric of the Trustworthy Computing Initiative.

Training and education

Software security problem awareness is growing among researchers and security practitioners. However, the most important audience has in some sense experienced the least exposure. For the most part, software architects, developers, and testers remain blithely unaware of the problem. One form of best practice involves training software development staff on critical software security issues.

The most effective form of training begins with a description of the problem and then demonstrates its impact and importance. During the Windows security push in February and March 2002, Microsoft provided basic awareness training to all its developers. Beyond awareness,

more advanced software security training should cover security engineering, design principles and guidelines, implementation risks, design flaws, analysis techniques, and security testing. Quality-assurance personnel—especially those who perform testing—should have special tracks.

Today’s academic curricula do little to expose students in engineering and computer science to security. Introductory programming courses can and should cover basic security tenets, because software security education is critical to changing the way software is built. One excellent technique is to have groups of students analyze the work of their peers for security problems.

Analysis, auditing, and engineering

You cannot bolt security onto an existing program or even implement particular sets of “security functionality” as a complete solution. The most generic instantiation of this problem is the common, but flawed, over reliance on cryptography. A classic rejoinder heard from software

architects and developers when confronted with security is the claim, “It’s secure because we use SSL.” Though applied cryptography is an important weapon in the software security arsenal, it is insufficient for security. Put glibly, software security is not security software.

Because security is an emergent property of a complete system, we must take a holistic approach. We should weave in security throughout the complete software development life cycle:

- At the requirements level, we must consider security explicitly. Security requirements should not be an “add on,” and should take into account emergent characteristics of security, including explicit coverage of what should be protected, from whom, and for how long.
- At the design and architecture level, a system must be coherent and present a unified security architecture, document assumptions, and identify possible attacks. At this stage of development, risk analysis is a necessity—risks can be uncovered and ranked so mitigation can begin.
- At the code level, focus on implementation flaws, especially those that can be statically discovered in an automated fashion. Note that code review, although necessary (especially when problematic languages such as C are used), is not sufficient for software security. Finding and removing all implementation flaws in source code does nothing to address architectural problems.
- When it comes to testing, security is a special case. Security testing must encompass two strategies: testing security functionality with standard functional testing techniques and risk-based testing based on attack patterns and threat models. Note that security problems are not always apparent, even when directly probing a system.
- We must monitor fielded systems during use. Simply put, attacks will

happen, regardless of the strength of the design and implementation. Monitoring software behavior is an excellent defensive technique.

Successful security analysis involves two levels of system understanding: architectural risk analysis and implementation analysis. Doing only one provides a high likelihood of failure. Today, there is an over emphasis on code review, which alone cannot solve the software security problem.

Microsoft's Trustworthy Computing Initiative

Microsoft's chairman, Bill Gates released a memo in January 2002 that highlights the importance of building secure software to Microsoft's future. Its Trustworthy Computing Initiative has changed the way Microsoft builds software. To date, Microsoft has spent over \$200 million (2,000 person days) on its software security push.

Microsoft is focusing on people, process, and technology to tackle the software security problem. On the people front, Microsoft is training every developer, tester, and program manager in basic techniques of building secure products. Microsoft enhanced its development process to make security a critical factor in design, coding, and testing of every product, for example.

There are many possible ways to integrate software security practices into the development lifecycle. Figure 3 illustrates Microsoft's way by showing where "security cycles" can be applied during a standard software development life cycle.

Risk analysis, code review, security testing, and external review and testing have their places in the new process. Microsoft is building tools, including PRefix and PRefast for defect detection,⁷ to automate as many process steps as possible, and changing its Visual C++ compiler to detect certain kinds of buffer overruns at runtime. Microsoft also has begun thinking about measurement and metrics for security.

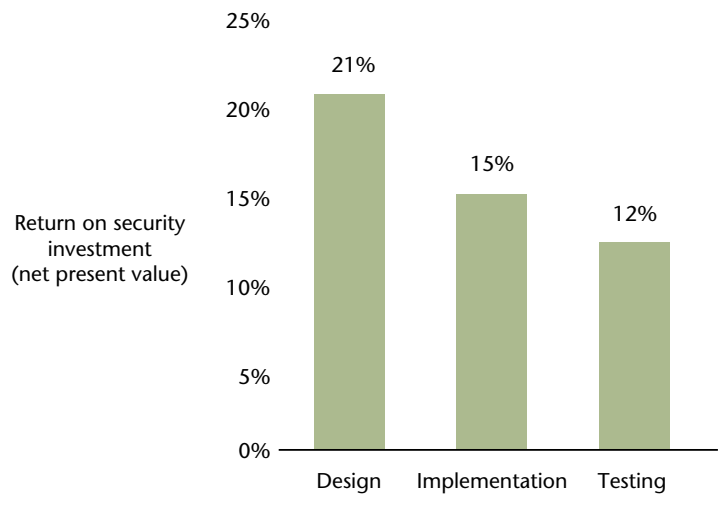


Figure 4. Phase return on investment (ROI) as measured by @stake over 23 security engagements. This study demonstrates the importance of design-level risk analysis as software security best practice.

The importance of measurement

Measurement is critical to the future of software security. Only by quantizing our approach and its impact can we answer such questions as: How secure is my software? Am I better off now than I was before? Am I making an impact on the problem? How can I estimate and transfer risk?

Measurement is one foundational approach critical to any science. As Lord Kelvin put it:

When you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meager and unsatisfactory kind: it may be the beginning of knowledge, but you have scarcely, in your thoughts, advanced to the stage of science.

The future belongs to the quants

We can begin to approach the measurement problem by recycling num-

bers from software literature. For example, we know that fixing software problems in the design stage is much cheaper than fixing them later in the life cycle. An often-quoted IBM study reports relative cost weightings as design = 1, implementation = 6.5, testing = 15, and maintenance = 100. We also know relative cost expenditures for life cycle stages: design = 15 percent, implementation = 60 percent, and testing = 25 percent. This range of numbers can provide a foundation for measuring software security impact.

Measuring ROI

A preliminary study (www.cio.com/archive/021502/security.html) reported by @stake (www.atstake.com) demonstrates the importance of concentrating security analysis effort at the design stage relative to the testing and implementation phases (see Figure 4). Separately, Microsoft reports that over 50 percent of the software security problems it finds are design flaws. We need more such analyses to advance the field.

Risk management calls out for quantitative decision support. We have work to do on measuring soft-

Further reading

In addition to the explicit references included in this article, I've collected a non-comprehensive bibliography of software security publications. This list is heavily biased towards recent publications. Use these references as a springboard for more reading.

- K. Ashcraft and D. Engler, "Using programmer-written compiler extensions to catch security holes," Proc. of the IEEE Security and Privacy Conference, IEEE CS Press, 2002.
- B. Arbaugh, B. Fithen, and J. McHugh, "Windows of Vulnerability: A Case Study Analysis," Computer, vol. 33, Dec., 2000, pp. 52–59.
- R. Bisbey II and D. Hollingsworth, "Protection Analysis Project Final Report," Technical Report ISI/RR-78-13, DTIC AD A056816, USC/Information Sciences Inst., May 1978.
- M. Bishop and M. Dilger, "Checking for Race Conditions in File Access," Computing Systems, vol. 9, Feb. 1996, pp. 131–152.
- C. Cowan et al., "Automatic Adaptive Detection and Prevention of Buffer-Overflow Attacks," Proc. 7th Usenix Security Symp., Usenix Assoc., 1998, pp. 63–77.
- M. Gasser, Building a Secure Computer System, Van Nostrand Reinhold, 1988.
- D. Larochelle and D. Evans, "Statically Detecting Likely Buffer Overflow Vulnerabilities," Proc. of Usenix Security Symp., Usenix Assoc., 2001.
- G. McGraw, "Software Assurance for Security," Computer, vol. 32, Apr. 1999, pp. 103–105.
- Miller, L. Fredricksen, and B. So, "An Empirical Study of the Reliability of Unix Utilities," Comm. ACM, vol. 33, Dec. 1990, pp. 32–44.
- B.P. Miller et al., Fuzz Revisited: A Re-Examination of the Reliability of Unix Utilities and Services," Technical Report CS-TR-95-1268, Univ. Wisconsin, Apr. 1995.
- F. Schneider, ed., Trust In Cyberspace, Nat'l. Academy Press. Washington, DC, 1998.

ware security and software security risk.

A Call to Arms

Much work remains to be done in software security. Some of it is basic and practical (for example, making software security part of the standard software development life cycle) and some of it is far beyond current capabilities (for example, automating software architecture analyses for security flaws). The US National Science Foundation suggests using the following 11 open questions as drivers for research. There is clearly overlap among these problems, but the list raises a large number of interesting subquestions. How do we

- Avoid building security flaws and security bugs into programs?
- Know when a system has been compromised?
- Design systems that can tolerate attack and carry out the intended mission?
- Design systems with security that can be reasonably managed?
- Provide reasonable protection of intellectual property?

- Support privacy enforcement technically?
- Get trustworthy computations from untrusted platforms?
- Prevent—withstand denial-of-service attacks?
- Quantify security trade-offs?
- Reveal and minimize assumptions in security system designs?
- Build programs and systems and know exactly what they will do and what they are doing?

We must give careful consideration to design for security. Given a set of principles and properties that we wish a system to have, we must identify design guidelines and enforcement rules. Open questions along this line of thinking include:

- Can principles be refined to guidelines?
- How can guidelines be reduced to rules that can be enforced statically?
- What technologies are suited for automated analysis?

Some concrete open research problems include explaining why

the software security problem is growing; quantifying, analyzing, and explaining bug/flaw categories; performing cost/benefit analyses to prove that early is good; untangling security software from software security at the requirements stage; exploring how to teach software security most effectively to students and professionals; and inventing and applying measures and metrics.

Software security is here to stay. In the near future, awareness, training, and process integration are likely to be a big focus. A market category will emerge with impact on traditional network security shops and software development organizations. As science begins to answer the questions posed here, the field will evolve in interesting ways. Because both theoretical and practical software security aspects remain, there's plenty of work ahead. Now is the time to get in on the ground floor. □

Acknowledgement

I acknowledge the thinking of all DIMACS Software Security Workshop participants, especially the program committee and the invited speakers: Ed Felten, Dan Geer, Virgil

Gligor, Michael Howard, Brian Kernighan, and David Wagner. Materials from the workshop are available at www.cigital.com/ssw.

References

1. D. Wagner et al., "A First Step Towards Automated Detection of Buffer Over-run Vulnerabilities," *Proc. Year of 2000 Network and Dist. Sys. Security Symp. (NDSS)*, 2000.
2. R. Anderson, *Security Engineering: A Guide to Building Dependable Distributed Systems*, John Wiley & Sons, 2001.
3. J. Viega and G. McGraw, *Building Secure Software*. Addison-Wesley, 2001; www.buildingsecuresoftware.com.
4. M. Howard and D. LeBlanc, *Writing Secure Code*, Microsoft Press, 2001.
5. G. McGraw and G. Morrisett, "Attacking Malicious Code: A Report to the Infosec Research Council," *IEEE Software*, vol. 17, Sept./Oct., 2000, pp. 33–41.
6. G. McGraw and E. Felten, *Securing Java: Getting Down to Business with Mobile Code*, John Wiley & Sons, 1999; www.securingsjava.com.
7. W. R. Bush, J. D. Pincus, and D. J. Sielaff, "A Static Analyzer for Finding Dynamic Programming Errors," *Software Practice and Experience*, vol. 30, June 2000.

Gary McGraw is Cigital's CTO, where he researches software security and sets technical vision for software quality management. He received a BA in philosophy from the University of Virginia and PhDs in cognitive science and computer science from Indiana University. He has coauthored four books: *Java Security* (Wiley, 1996), *Securing Java* (Wiley, 1999), *Software Fault Injection* (Wiley 1998), and *Building Secure Software* (Addison-Wesley, 2001) and more than 50 peer-reviewed papers. Contact him at gem@cigital.com; www.cigital.com.

Appendix 2
List of Participants

Workshop on Software Security List of Participants

Organizing Chair

Gary E. McGraw, Cigital

Organizers

Edward W. Felten, Princeton University
Virgil D. Gligor, University of Maryland
David Wagner, UC Berkeley

Participants:

Godmar Back, Computer Systems Laboratory
Lee Badger, DARPA
Lee Begeja, AT&T
Steven Michael Bellovin, AT&T
Ali Bicak, University of Maryland
Scott Robert Bourne, Rutgers University
Bill Cheswick, Lumeta Corporation
Crispin Cowan, WireX Communications, Inc.
Scott A. Crosby, Rice University
Drew Dean, SRI International
Jeffrey Dielle, Hewlett-Packard Managed Services
Dominic Duggan, Stevens Institute of Technology
David E. Evans, University of Virginia
Debin Gao, Carnegie Mellon University
Dan Geer, @stake
Santoshkumar Girish Nair, New Jersey Institute of Technology
Robert J. Hall, AT&T
Michael Howard, Microsoft Corp.
Trent Jaeger, IBM
Trevor Jim, AT&T
Danny Max Kaufman, Rutgers University
Lev Kaufman, Rutgers University
Gaurav S. Kc, Columbia University
Brian Kernighan, Princeton University
Larry Koved, IBM
Carl E. Landwehr, National Science Foundation
Ben Laurie, OpenSSL
Yow-Jian Lin, SUNY Stony Brook
Annie Liu, SUNY Stony Brook
Rebecca T. Mercuri, Bryn Mawr College
Lakshman Kumar Mikkavilli, Cisco Systems, Inc.
Afroze Naqvi, SIAST, Wascana Campus
Thomas Ostrand, AT&T
Jonathan D. Pincus, Microsoft
William W. Pugh, University of Maryland


Sandeep Radharkrishnan, Stevens Institute of Technology
Marcus J. Ranum, Ranum.com
Harshvardhan Ulhas Revankar, New Jersey Institute of Technology
Fred Roberts, DIMACS, Rutgers University
Algis Rudys, Rice University
John Charles Slimick, University of Pittsburgh at Bradford
Nicol So
John Steven, Cigital Inc.
Scott D. Stoller, Stony Brook University
John Viega, Secure Software
Dan S. Wallach, Rice University
Ronald Phillip Waterbury, AT&T
Elaine J. Weyuker, AT&T
Gregory Wright, AT&T
Rebecca Wright, Stevens Institute of Technology

Appendix 3

Program

DIMACS

*Center for Discrete Mathematics & Theoretical Computer Science
Founded as a National Science Foundation Science and
Technology Center*



DIMACS Workshop on Software Security

Monday, January 6th, 2003

Time	Agenda Item
8:00-9:00	Breakfast (DIMACS)
9:00-10:00	Gary McGraw, Cigital (Author of <i>Building Secure Software</i>) The Art and Science of Software Security
10:00-10:30	Group discussion
10:30-10:45	Morning break
10:45-12:00	Outrageous Opinions (to be submitted by attendees); <i>Examples to spark ideas:</i> <ul style="list-style-type: none">• The TCPA is good for everyone• Capabilities are the only way to go• Open source is a security panacea• Security does not matter to users at all• ... <i>Speakers and titles include:</i> <ul style="list-style-type: none">• Ed Felten: Nothing we do can improve security• Jon Pincus: Stop telling me I should be speaking Esperanto• Bill Pugh: It is time to abandon C and C++• Ben Laurie: TCPA and Palladium solve capabilities confinement• Dave Evans: Protecting Bits with Atoms (and Vices with Verses)• Steve Bellovin: We can't write secure programs• Crispin Cowan: Security and Open Source: the 2-Edged Sword
12:00-1:30	Lunch (DIMACS)
1:30-2:00	Breakout session Administration <ul style="list-style-type: none">• <i>Groups, group leaders, goals for sessions</i>
2:00-3:00	Invited talk: Michael Howard, Microsoft (Author of <i>Writing Secure Code</i>)

	The Microsoft Trustworthy Computing Initiative from the Inside	
3:00-3:30	Group discussion	
4:00-5:00	BREAKOUT: Security Engineering <ul style="list-style-type: none"> • Requirements • Architecture and design • Coding and Testing • Manageability Dave Wagner	BREAKOUT: On Architecture and Implementation <ul style="list-style-type: none"> • Design risks • Implementation risks • Technology Tradeoffs • Experience and expertise Gary McGraw
5:00-7:30	Dinner (on your own)	
7:30-10:00	Wine and cheese reception and poster session	

Tuesday, January 7th, 2003

Time	Agenda	
8:00-9:00	Breakfast (DIMACS)	
9:00-10:00	Invited talk: Brian Kernighan	
	Coding Excellence: Security as a Side Effect of Good Software	
10:00-10:30	Group discussion	
10:30-10:45	Morning break	
10:45-12:00	BREAKOUT: Security Analysis <ul style="list-style-type: none"> • Role of expertise • Auditing design • Auditing code • Security Testing Gary McGraw	BREAKOUT: Mobile code and Malicious Code <ul style="list-style-type: none"> • NET and Java • Web services • Modern malicious code Ed Felten
12:00-1:30	Lunch (DIMACS)	
1:30-2:30	Invited talk: Dan Geer, @stake	
	Software Security in the Big Picture: Repeating ourselves all over again	
2:30-3:30	BREAKOUT: Open Research Issues <ul style="list-style-type: none"> • Hard problems Virgil Gligor	BREAKOUT: Education and Training <ul style="list-style-type: none"> • Academia • Industry developers Dave Wagner
3:30-4:00	Break	
4:00-5:30	Workshop wrap-up	
	<ul style="list-style-type: none"> • Reports from breakout sessions • Program committee summary 	

Appendix 4

Links to Workshop Presentations

Presentation Links

Main Web page with Schedule:

<http://www.cigital.com/ssw/presentations.php>

Individual Talks:

Gary McGraw, Cigital
The Art and Science of Software Security
<http://www.cigital.com/ssw/presentations/gem/>

Ed Felten, Princeton University
Nothing we do can improve security
<http://www.cigital.com/ssw/presentations/felten/>

Dave Evans, University of Virginia
Protecting Bits with Atoms (and Vices with Verses)
<http://www.cigital.com/ssw/presentations/evans/>

Crispin Cowan, WireX Communications, Inc.
Security and Open Source: the 2-Edged Sword
http://www.cigital.com/ssw/presentations/dimacs_opensource/

Michael Howard, Microsoft
The Microsoft Trustworthy Computing Initiative from the Inside
<http://www.cigital.com/ssw/presentations/howard.ppt>

Brian Kernighan, Princeton University
Coding Excellence: Security as a Side Effect of Good Software
<http://www.cigital.com/ssw/presentations/kernighan/>

Dan Geer, @stake
Software Security in the Big Picture
<http://www.cigital.com/ssw/presentations/geer/>

BREAKOUT: Open Research Issues
<http://www.cigital.com/ssw/presentations/landwehr/>

Workshop wrap-up
http://www.cigital.com/ssw/presentations/gem_end/

Appendix 5
Booklet of Abstracts

ABSTRACTS

DIMACS Workshop on Software Security

January 6-7, 2003
DIMACS Center, CoRE Building, Rutgers University

Organizers:

Gary McGraw
Cigital
gem@cigital.com

Ed Felten
Princeton University
felten@cs.princeton.edu

Virgil Gligor
University of Maryland
gligor@umd.edu

Dave Wagner
University of California at Berkeley
daw@cs.berkeley.edu

Presented under the auspices of the Special Focus on Communication Security and Information Privacy.
See: http://dimacs.rutgers.edu/SpecialYears/2003_CSIP/

Supported by The National Science Foundation (NSF), the NJ Commission on Science and Technology, and also the DIMACS partners at Rutgers University, Princeton University, AT&T Labs - Research, Bell Labs, NEC Laboratories America, and Telcordia Technologies, and the DIMACS' Affiliate Members at Avaya Labs, IBM Research and Microsoft Research. In particular, DIMACS thanks Microsoft Research for a special contribution to support this workshop.

Steve Bellovin, AT&T

We can't write secure programs

Assertion: we cannot, in general, write secure programs. Security is a subset of correctness; correct programming is -- and will remain -- the oldest unsolved problem in computer science. **All** non-trivial programs, including firewalls, operating systems, and privileged or networked applications, are and will remain insecure. The challenge to security professionals is to design **systems** that will be "secure enough", despite the failure of many of the individual components.

Dan Geer, @stake

Software Security in the Big Picture

Security investment does not yet have the direct linkages to the creation of business value that other IT investments do, and executives cannot hope to show value via cost-benefit -- cost-effectiveness will be hard enough--but the future unequivocally belongs to the quants. Metrics that make security rational will be vulnerability-modeled and data-calibrated. They will rely on information sharing about incidents (frequency, dollars lost, whether the event was even caught...) and on log analysis that can separate the anomalous from the normal. The National Strategy already focuses on vulnerabilities rather than threats, and SDLCs should include security before software liability takes hold, as it soon will. None of this is surprising -- it is just what business maturity would predict.

Michael Howard, Microsoft

Trustworthy Computing - An Insider's View

In this presentation, Michael Howard will outline the overall Trustworthy Computing goals, as well as the short-term and long-term steps being taken to achieve the goal. He will also outline the tactical and strategic goals of the series of 'security pushes' at Microsoft, as well as the development life-cycle changes underway at the company.

Brian Kernighan, Princeton University

Coding Excellence: Security as a Side Effect of Good Software

Good programming languages are often thought a prerequisite for robust and secure software. Yet most of the time, language is secondary: sound design and good programming practices are much more important. Good programmers program well in any language, but no language can prevent a bad programmer from writing bad code. So while we wait for more perfect languages to be developed, and then be accepted by the majority of programmers, there is much that we can

do today to improve programming practice and thus improve the security properties of our programs.

Ben Laurie, OpenSSL

TCPA and Palladium solve capabilities confinement

A standard problem in distributed capabilities is the confinement problem. Once you have given someone a capability, you cannot prevent them from giving it to someone else. TCPA and Palladium provide a mechanism by which this can be solved. The protocol is left as an exercise for the reader, but it involves a private key closely held by the TCPA/Palladium hardware, a nonce, and, of course, the capability (or, more precisely, its Swiss number).

Gary McGraw, Cigital

The Art and Science of Software Security

Computer security researchers and practitioners have come to recognize the critical role that software plays in security. Software security is the art of proactively building software to be reliable and secure. By contrast, network security tends to emphasize a reactive law enforcement stance, and in many cases does not identify the root cause of security problems (bad software).

Making software behave is hard, and security subtleties only exacerbate the problem. Internet-enabled software applications, especially custom applications, present the most common security risk encountered today, and are the target of choice for real hackers.

This talk provides an introduction to the problem of software security. I discuss the magnitude of the problem and some of the root causes, which I call the trinity of trouble. I briefly discuss security engineering, security requirements, testing for security, and the idea of software risk management. I then propose some open questions meant to stimulate discussion and clarify some aspects of this exciting new field.

Jon Pincus, Microsoft

Stop telling me I should be speaking Esperanto!

Every time there's a buffer overrun, everybody shakes their head and sighs "if only people didn't program in C and C++ this wouldn't be an issue." It's obvious ... and, equally obviously, if only everybody spoke Esperanto, we'd all be able to understand each other. Can we just return to reality here? Today, virtually nobody writes system software (OSs, drivers, databases, web servers, browsers) in safe languages; perhaps, just like with Esperanto, there are other factors that mean the "obvious right answer" isn't actually right in practice. What will actually cause things to change (or, if it's not going to change, what should we do instead)? And since I'm

cynical enough to question the original Esperantists' assumption that a world language would bring us noticeably closer to world peace, why should I believe that changing languages will bring us noticeably closer to "software security"?

Bill Pugh, University of Maryland

It is time to abandon C and C++

It is difficult to write reliable and secure code in C and C++. Instead, any new projects should be performed in type safe programming languages such as Java, C#, Cyclone and CCured. While those languages do not eliminate security problems, they eliminate broad categories of them. Strong consideration should be given to trying to migrate security critical components of existing C and C++ applications to a type safe language.